Interruption of People in Human-Computer Interaction

by
Daniel Craig McFarlane

B.S. May 1987, Cornell University
M.S. May 1990, The University of Oregon

A Dissertation Submitted to the Faculty of
The School of Engineering and Applied Science of
The George Washington University
in partial satisfaction of the requirements for the degree of
Doctor of Science

August 31, 1998

Dissertation directed by

John Lee Sibert
Professor of Engineering and Applied Science

ABSTRACT

Just because something is new and clever does not necessarily mean that it is good. The utility of a technology must be evaluated strictly in terms of whether it actually helps people succeed. A new technology that is useful for isolated tasks may carry an invisible and costly side-effect like interrupting people. However, the side-effect of human interruption by machine only manifests itself in real world contexts where people normally perform several complex heterogeneous tasks in parallel. The telephone and email are examples. They are useful by themselves, but in a real work environment they also create annoying interruptions. A newer and more problematic example is the technology of semi-autonomous computer systems such as intelligent agents. These systems can be assigned to do useful things in the background while their human users work on other tasks. However, delegating a task requires supervising a task; and whenever an intelligent agent must initiate an interaction with its user it has to first interrupt them from whatever else they are doing. Interruption of people during human-computer interaction (HCI) is problematic because people have cognitive limitations that restrict their ability to work during interruptions. These human limitations for handling interruptions can cause people to make critical, even life-threatening mistakes. Unfortunately, no user interface design guidelines exist for directing solutions to this problem. In fact, no general theoretical tools exist for facilitating basic interdisciplinary investigations of this problem. This dissertation creates two fundamental interdisciplinary theoretical tools for addressing this problem: (1) a Definition of Human Interruption; and (2) a Taxonomy of Human Interruption. These tools are synthesized from the results of a comprehensive analysis of the existing literature of several relevant fields. The utility of these new tools was partially validated by demonstrations of their usefulness in facilitating two explorations of the problem of user-interruption in HCI. First, the Taxonomy of Human Interruption was used to structure an analysis and discussion of existing literature from several different disciplines relevant to the design of user interfaces for the human interruption problem. The taxonomy was shown to provide a valuable common ground for comparing works from different fields that are not obviously similar. Second, the Taxonomy of Human Interruption was used to guide an experiment with human subjects. It was used to formulate and operationalize a theory-based hypothesis about how to coordinate human interruption by machine during HCI. An experimental computer-based multitask was designed and built, and a human subjects experiment was conducted with 36 subjects (18 male and 18 female). The main hypothesis was conclusively supported. Subhypotheses received mixed support which indicates the existence of design trade-offs among the different methods for coordinating human interruption.

ACKNOWLEDGMENTS

Writing this dissertation has been a very strange process for me. I was officially constrained to work "alone" and not accomplish this project with a team. George Washington University (GWU) insisted that this project was an individual learning experience, and that the finished product, if successful, would serve as final proof of my suitability to be granted the degree of doctor of science. However, I did have an unofficial "team" of talented and persistent friends, colleagues, advisors, and family. These people went beyond all conventional limits in supporting me while I worked through this artificial process.

This dissertation project has made me painfully aware of how necessary and valuable other people are for doing research. Maybe that awareness is the main educational point of the dissertation.

My GWU advisor, John Sibert, and my first supervisor at The Naval Research Laboratory (NRL), Constance Heitmeyer, together with James Foley gave me the original opportunity to move to Washington, DC, and begin my dual role as doctoral student and employee. GWU gave me the opportunity for formal training, and NRL gave me the opportunity of funding and experience working with top people in my field.

Paying for a dissertation at GWU is a major accomplishment in and of itself. NRL has been both patient and committed to me and my research topic. They also gave me the necessary first-hand exposure to other people's research that I needed as background to identify the topic of this dissertation. James Ballas has made critical and large contributions to this work in funding and resources and in providing relevant experience. I hope the results of this dissertation can be transitioned into something useful for his work. Thanks to Alan Meyrowitz and Mike Shneier (Office of Naval Research) for direct funding of part of this project. Thanks also to Astrid Schmidt-Nielsen, Robert Jacob, Hellen Gigley, Constance Heitmeyer, Robert

Carter, and Randall Shumaker for extended opportunities of funding, resources, experience, and training. John Sibert and Walt Truskowski (NASA Goddard) provided funding to pay my experimental subjects $20 each.

A special thanks to all 43 subjects who volunteered for the pilot studies and for the experiment. They gave up about 2 hours and 30 minutes each to participate in the experiment (that makes a total of about 107.5 hours). This was a great sacrifice for many busy people, and made possible one of the core contributions of this work.

Thanks also to the members of my committee who have taken the time to thoroughly and professionally grill me. Their efforts in giving me advice and technical critique have made me rethink and rewrite portions of this dissertation for the better. The members of my examining committee were:

• Dr. John Lee Sibert (advocate)	George Washington University
• Dr. James A. Ballas	Naval Research Laboratory
• Dr. Rachelle Silverman Heller	George Washington University
• Dr. Robert J. K. Jacob	Tufts University
• Dr. Carol Dianne Martin	George Washington University

I have had the privilege of some patient and loyal friends who have helped me recognize my successes and recover a healthy perspective of my goals during setbacks. Manuel Pérez and Rudy Darken have been my backbone of stability during this project. Their persistent faith in me has helped me stay focused. I thank John Sibert for his friendly encouragement and for expressing a belief in my capability. Prof. Sibert, and the members of the GWU Human-Computer Interaction Research Group have given me consistent encouragement and helpful feedback in our regular meetings over the past several years. They have been my advisors, piers, and colleagues and have done me the favor of asking my advice. There has been nothing so healthy for me as trying to think about someone else’s problems.

The people who have gone the farthest beyond conventional limits of support are my family. My wife, Susan, has been unwaveringly supportive of my need to do this. I thank her for extensive work on editing, and for her patient tolerance of my need to spend so much time working. A special thanks to my children Katherine and Jared for telling me how great I am and for forgiving me when I work late. My father was my first example of how to get a doctorate. He and my mother and my brothers and sisters have constantly encouraged me to make my own decisions and follow through.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: INTRODUCTION

1.1 AIMS AND OBJECTIVES

This chapter introduces a recognized but currently neglected problem — interruption of human users in the context of human-computer interaction (HCI). Further, this chapter discusses why this problem is important and timely and describes the approach chosen for this dissertation research. After reading this chapter, the reader should be familiar with the topic of human interruption during HCI and understand the scope and important contributions of this dissertation.

1.2 OVERVIEW

Hansen (1971) proposes “Know thy user” as a useful maxim for HCI design. However, in the case of systems that support multiple concurrent user activities, user interface designers do not know their users. They instead must create ad hoc design solutions based only on their own intuition, which unfortunately is often inadequate. Many of these ad hoc efforts seem to have replaced Hansen’s design maxim of “Know thy user” with a new maxim of “Monopolize thy user.”

Computer system designers increasingly employ semiautonomous and multitasking technologies that provide systems with a degree of semiautonomy from direct user control for accomplishing complex tasks. These technologies are clearly useful in many domains. Semiautonomous and multitasking systems, however, have different user interface requirements than traditional, no-autonomy single task systems. The user must only intermittently interact with the semi-autonomous computer system, because some of the time, the system is working on its own. And multitasking systems require users to switch between concurrent or even simultaneous tasks.

This new style of intermittent interaction causes a dramatic increase in the side effect of user-interruption. When a semiautonomous system must communicate with its user, it must first interrupt the user from the other activity(ies) they are performing. The user's attention can become dominated by a single task; not because the task demands constant attention, but because the design of the system ignores peoples' cognitive limitations related to distraction and interruption.

Interrupting the user is not a "bad" thing. It is instead, just a necessary type of human-computer interaction. Interrupting people is, however, a complex topic and one for which users are especially sensitive to bad user interface design. The current literature recognizes the complexity and importance of this topic, however no one has yet published a comprehensive design solution for building user interfaces for systems that must interrupt their users.

This chapter makes a significant contribution by thoroughly introducing this important topic.

1.3 MOTIVATION

Current progress in computer technology increases the capability of building systems that allow people to perform multiple concurrent activities; however, people's cognitive capabilities are not progressing. Under certain conditions people are able to perform multiple tasks concurrently; however, they have cognitive limitations that make them vulnerable to mistakes and delays when they try to do more than one thing at a time. These cognitive limitations make people especially sensitive to bad user interface design.

If the user interface is designed in ways that accommodate people's cognitive limitations, then people will be able to successfully perform multiple concurrent computer-based activities. The question is, "what is good design for this problem?" There are three problems that hinder attempts to create successful user interface designs: (1) there is no general theory of human interruption that identifies the critical factors of the human interruption problem; (2) there are no general user interface design guidelines for this important problem of user-interruption by machine; and (3) people are exceedingly vulnerable to poor design solutions.

This HCI problem of computers interrupting their users is an unfortunate side effect of a powerful and increasingly popular kind of computer support — the ability of running useful computer activities in the background. For example, artificial intelligence technology can be applied to create “intelligent” semi-autonomous computer systems that provide valuable services for performing complex tasks. These intelligent systems can be useful things like intelligent decision aids, intelligent software agents, or autonomous robotic vehicles.

Semi-autonomous computer systems are particularly useful because they allow their human users to multitask, and multitasking is natural for people. People ordinarily perform two or more activities during the same time frame. They think in parallel and act serially (asynchronous parallelism (Edmondson 1989)). A computer user can delegate one or more tasks to intelligent software agents; start them running in the background and then go on to begin or resume some other activity, because background activities do not require constant user attention. The user must only intermittently interact with the computer system, because some of the time, the system is off working on its own. This new style of intermittent interaction causes a dramatic increase in the side effect of user-interruption. This side effect results from the problem that while the system is acting autonomously, the user goes on to start or resume something else (Maes and Wexelblat 1996). Whenever a semi-autonomous system must communicate with its user, it must first interrupt them from whatever other activity they have become involved with.

However, just because multitasking is a common and useful kind of human behavior that does not mean that people do it easily or reliably (Preece et al. 1994, p. 105). When people multitask they are susceptible to internal and external events that cause them to make mistakes. For example, a person can easily multitask while getting breakfast. They perform the following activities concurrently: cook a waffle; talk with another person; put away baking ingredients; watch TV news; heat syrup in a microwave; set the table; and load the dishwasher. Everything goes well until they hear on the TV that NASA's Mars Pathfinder has successfully landed on Mars and is sending back pictures. Five minutes later they realize they have made several mistakes: they burned the waffle, they put the milk jug in the dishwasher, they have completely forgotten the topic of their conversation, and they heated the syrup until it has swollen up like a balloon.

Computer systems support important multitasks where mistakes can be more expensive than funny, e.g., writing a report, collaborating with other people, projecting budgets, emergency 911 dispatching, flying an airplane, managing a nuclear power plant, or fighting a war. It is essential that user interfaces for systems that must interrupt their users be designed in ways to prevent expensive human errors and their costs. However, this is still an unsolved interface design question, and there are several current examples of computer systems with ineffective ad hoc solutions to this problem.

We know from previous investigations that people have cognitive limitations which restrict their ability to work during interruptions. These limitations can adversely affect people's performance on critical tasks. For example, an interruption of a commercial airline crew before takeoff contributed to their subsequent crashing of the plane. A Northwest Airline crew was preparing to fly out of Detroit Metropolitan Airport. The crew began the preflight checklist properly but were interrupted by Air Traffic Control before they verified the status of the airplane's flaps. The flaps were not down, as required. After the interruption by Air Traffic Control, the crew allowed other issues to distract them from resuming their checklist. Other distracters included confusion about which taxi-way to use because of a change in taxiing directions and delayed reports of weather and runway conditions. The crew took off without finishing their checklist. They never checked to see if the airplane's flaps were in the correct take-off position, and they were not. A flight emergency occurred shortly after takeoff. If the crew had understood their situation, they could have successfully become airborne without flaps. However, the crew had also received a windshear alert. When the emergency occurred, the crew mistakenly interpreted the problem as windshear instead of flap position and crashed the plane (National Transportation Safety Board 1988).

If the design of the user interface does not accommodate these limitations then the interface itself can affect people's performance on critical tasks. In cases where the cost of human error is high and the user interface for the computer system has not been designed to support graceful handling of interruptions and distractions then users must choose between risking costly mistakes or trying to force themselves to concentrate exclusively on a single task. Users may choose to limit their function to only one task at a time. The potential benefit of user multitasking with backgrounded autonomous aids is lost. The user's ability to multitask is suppressed not because one task inherently demands their constant attention, but because poor design of the user interface makes it difficult to switch between tasks without making errors -- the user interface design ignores peoples' cognitive limitations related to distraction and interruption.

> Quoted from the short story "Harrison Bergeron," by Kurt Vonnegut, from "Welcome to the Monkey House" (Vonnegut 1950).
>
> The year was 2081, and everybody was finally equal. They weren't only equal before God and the law. They were equal every which way. Nobody was smarter than anybody else. Nobody was better looking than anybody else. Nobody was stronger or quicker than anybody else. All this equality was due to the 211th, 212th, and 213th Amendments to the Constitution, and to the unceasing vigilance of agents of the United States Handicapper General.
>
> Some things about living still weren't quite right, though. April, for instance, still drove people crazy by not being springtime. And it was in that clammy month that the H-G men took George and Hazel Bergeron's fourteen-year-old son, Harrison away.
>
> It was tragic, all right, but George and Hazel couldn't think about it very hard. Hazel had a perfectly average intelligence, which meant she couldn't think about anything except in short bursts. And George, while his intelligence was way above normal, had a little mental handicap radio in his ear. He was required by law to wear it at all times. It was tuned to a government transmitter. Every twenty seconds or so, the transmitter would send out some sharp noise to keep people like George from taking unfair advantage of their brains.
>
> George and Hazel were watching television. There were tears on Hazel's cheeks, but she'd forgotten for the moment what they were about.
>
> On the television screen were ballerinas.
>
> A buzzer sounded in George's head. His thoughts fled in panic, like bandits from a burglar alarm.
>
> "That was a real pretty dance, that dance they just did," said Hazel.
>
> "Huh?" said George.
>
> "That dance — it was nice," said Hazel.
>
> "Yup," said George. He tried to think a little about the ballerinas. they weren't really very good — no better than anybody else would have been, anyway. They were burdened with sashweights and bags of birdshot, and their faces were masked, so that no one, seeing a free and graceful gesture or a pretty face, would feel like something the cat drug in. George was toying with the vague notion that maybe dancers shouldn't be handicapped. But he didn't get very far with it before another noise in his ear radio scattered his thoughts.
>
> George winced. So did two out of the eight ballerinas.

1.4 BACKGROUND

Researchers have observed that interrupting people causes side effects. Authors of human psychology have identified an effect they call the Zeigarnik Effect (Van Bergen 1968). The Zeigarnik Effect describes a finding that people have selective memory relative to interruption — people are able to recall the details of interrupted tasks better than the details of uninterrupted tasks. Results from many studies of the Zeigarnik Effect have produced somewhat inconsistent results. However, two findings seem conclusive: (1) interrupting people causes side effects, and (2) interruption of people is a complicated process.

These two results must be addressed in any research or development of user interfaces for computer systems that must interrupt their users. First, HCI design guidelines must be discovered for addressing the problems associated with the side effects caused by user-interruption by machine. Second, extreme care must be taken to control possible sources of confounding noise when designing investigations of human interruption.

Rubinstein and Hersh (1984) propose 93 guidelines for user interface design. Their guideline number 12, “interrupt with care” (p. 64), identifies user-interruption by computer as an important problem. To illustrate the usefulness of this guideline they say, “A system message announcing next month’s preventive maintenance schedule has no business appearing uninvited in the middle of a person’s edited text or command line. Computer interruptions of the user must be polite and occur only at places that don’t annoy or confuse” (p.64). Rubinstein and Hersh’s guideline, however, does not give any specific direction about how to make interruptions “polite” or how to schedule their occurrence for places that don’t “annoy or confuse.” Without these details, the only utility of guideline 12 is to identify the problem.

The Intelligent Control and Interface Design (ICID) research project at the Navy Center for Applied Research in Artificial Intelligence (NCARAI) is an example of how the capability of running intelligent decision aids in the background causes the unintentional side-effect of increasing user interruption (Ballas et al. 1996; Kushnier et al. 1996). ICID is an evolving research platform for investigating user interface design methods for building software tools to help Navy commanders perform better in command and control tasks (like managing the tactical air defense of an aircraft carrier battle group). Over the last 4 years, the ICID research team incrementally increased the capability of ICID by introducing new intelligent decision aids. First ICID included an intelligent decision aid that gave advice for the deployment and maintenance of a standard sector air defense. Next researchers added an aid that supports situational awareness by interactively deducing complex relationships between observed man-made objects and groups of objects in the environment. The ICID researchers are currently adding a new, intelligent decision aid that automatically deduces and alerts the user to occurrences of standard enemy attack patterns. The ICID team has observed that, while each of these additional decision aids provides a useful function, they each also place new interactional demands on the user. Each additional intelligent decision aid potentially interrupts and/or distracts the user in new and different ways making the design of the user interface more complicated.

Perse et al. identified user-interruption as a critical problem in the Navy's AEGIS combatant Integrated Survivability Management System (a combined combat and damage control system for the Navy's AEGIS cruisers) (Perse et al. 1991). They did an analysis of crew performance during extensive simulation tasks and found that interruptions significantly interfere with operator and mission performance. Perse et al. recommend that means be found to augment the AEGIS system to manage and reduce interruption of users.

The new internet "push" technology is a another example of a novel technology with the side effect of greatly increasing user-interruptions. Push is the name for a new internet technology that "pushes" information at the user. The user tells a push software agent what kind of information they want and the agent begins running in the background. The user then goes on to begin or resume another task. The backgrounded agent begins "pushing" information at the user. The result for the user is like having a personal 24 hour TV newsroom constantly sending stories of possible interest. Each new piece of information is a new source of interruption or distraction. A user's ability to perform their normal tasks could be drastically impaired because of all these "pushed" interruptions. HCI design methods must be discovered to construct user interfaces for push systems that allow users to manage this flood of interruption and distraction.

Technical improvements have historically resulted in side effects that must themselves be investigated. There are illustrative examples of this trend in noncomputer fields of technology. Transportation is a good example. Traffic at street intersections was not problematic until the invention and mass sale of automobiles. After cars became popular, people had to solve this now-serious side effect of greatly intensified traffic at street intersections. This problem was eventually solved with the invention and installation of traffic lights. The car itself is another good example of technical improvements causing important side effects. Original cars were slow and light. As car technology quickly improved, cars became both faster and heavier. However, an important side-effect emerged; as cars became faster and heavier, the severity of automobile accidents increased dramatically. This life threatening side-effect created a need for the invention of safety devices like steel unibody construction, seat belts, shoulder belts, crumple zones, air bags, high seat backs, etc.

Like street traffic and the automobile, improvements in computers sometimes cause previously inconsequential side effects to grow into important problems. Recent advances in semi-autonomy and multitasking have caused the HCI of user-interruption to become a critical problem.

1.4.1 Multitasking — People Performing Multiple Concurrent Activities

When people multitask with computers they do not do everything simultaneously. Miyata and Norman (1986) provide a useful theory-based classification of the different ways people manage the individual HCI activities of their multitasks (see also (Cypher 1986)). This classification can be used to describe the current state of action for each activity in people's multitasks. Miyata and Norman's categorization is especially useful because it can be used to describe individual HCI activities in a way that reveals the inherent timing dependencies of multitasks and people's strategies for concurrently accomplishing a the set of individual activities.

Miyata and Norman discuss activities in terms of people's cognition. A review of their ideas establishes a context for describing background literature relevant to the topic of this paper — HCI design methods to allow people to more successfully coordinate when they will perform externally initiated interruption tasks. Table 1 is a summary of Miyata and Norman's classification of multiple activities (p. 270-271).

Table 1 — State of Activities in HCI Multitask

Status of Activity with User	Meaning
(1) Current Activity	Actions for accomplishing this activity are being performed now.
(1.A) Foregrounded Activity	Current activity under the conscious control of the user.
(1.B) Backgrounded Activity	Current activity out side of the conscious control of the user.
(1.B.1) Internally Backgrounded Activity	Current activity under the subconscious control of the user.
(1.B.2) Externally Backgrounded Activity	Current activity under the control of some other agency.
(2) Suspended Activity	All activity suspended.

The important questions for describing the current state of action for an activity are: (1) is the activity currently being acted upon?; (2) if the activity is current, is it under the user's conscious control?; and (3) if the activity is current but not under the user's conscious control, is the user acting on it subconsciously or is some other entity acting on it?

The number of combinations of different activity states in a multitask is constrained because we can assume that there will always be one and only one activity as the foregrounded activity. When people are performing multitasks, they are always doing something; and people can only do one thing at a time in their conscious control (Davies et al. 1989). In the simple case of a dualtask (a multitask with two activities) there are only six possible states for the current status of the pair of activities. Table 2 identifies these six states.

Table 2 — Possible States of the Two Activities in a Dualtask

Dualtask State	State of Activity 1	State of Activity 2
1	Foregrounded Activity	Suspended Activity
2		Internally Backgrounded Activity
3		Externally Backgrounded Activity
4	Suspended Activity	Foregrounded Activity
5	Internally Backgrounded Activity	
6	Externally Backgrounded Activity	

The mirrored pairs can be combined into only three distinct dualtask activity pairs: (1) a foregrounded activity with a suspended activity; (2) a foregrounded activity with an internally backgrounded activity; and (3) a foregrounded activity with an externally backgrounded activity.

Interruption of the human by the machine is only a problem in two of the three dualtask activity pairs. In condition 1, a foregrounded activity with a suspended activity, there is no problem with external interruption of the user because suspended activities do not evoke interruption of current activities. The other two dualtask activity conditions each support different kinds of interruption, i.e., internal interruptions or external interruptions (Miyata and Norman 1986, p. 268-270). Condition 2, a foregrounded activity with an internally backgrounded activity, can be the context for internal interruptions. The person's subconscious cognitive processes that are acting on the internally backgrounded activity can initiate an internal interruption of the person's own current focus of attention. Condition 3, a foregrounded activity with an externally backgrounded activity, can be the context for external interruptions. The external entity acting on the backgrounded activity can initiate an interruption of the person during their foregrounded activity. (Note, external interruptions in HCI is the topic of this paper, so condition 3 is most relevant here.)

A dualtask example of a person driving a car (activity 1) while conversing with a passenger (activity 2) illustrates the three dualtask activity pairs. Suppose a person is driving a car as their foregrounded activity and they want to switch their conscious attention to the activity of conversing with a passenger. In condition 1, a foregrounded activity with a suspended activity, the driver would have to stop and park the car every time they wanted to begin or resume talking with the passenger. The activity of driving cannot be backgrounded, and so it cannot continue to be an current activity when the conversation activity is made the foregrounded activity. In condition 2, a foregrounded activity with an internally backgrounded activity, if the person is an experienced driver they can internally background the driving activity as a subconsciously controlled activity when they foreground the conversation activity. In condition 3, a foregrounded activity with an externally backgrounded activity, the person externally backgrounds the driving activity by starting the car's autonomous robotic driver. The person then lets go of the vehicle controls and foregrounds the conversation activity.

Computer multitasking is different than human multitasking. Mainstream personal computers (PC's) are built to multitask with preemptive multitasking schemes on a single CPU (Alford 1992). Grehan (1990) explains that preemptive multitasking is like Superboy playing a baseball game all by himself. He alternately zooms from position to position so quickly that he accomplishes the jobs of all the members of both teams.

PC's, of course, do not have "consciousness" so it's not useful to describe them as having conscious or subconscious activity like people. However, computer multitasking does have one thing in common with human multitasking — a PC with a single CPU can only do one task at a time like people's foregrounded conscious cognition. Unlike people, PC's cannot run simultaneously backgrounded current activities. The PC's CPU can only be used to do one thing at a time, but use a scheme of sharing the CPU's computing work by automatically switching the CPU's activity between all tasks — like the Superboy metaphor. Computers multitask maintenance is handled by a deterministic algorithm that ensures that each activity gets a "fair" share of CPU actions. Therefore computer multitasking is not vulnerable to the same kinds of errors as people's multitasking.

Although a single CPU cannot be used to simultaneously perform more than one thing at a time (like a person does when they internally background activities), a CPU can externally background activities. Peripherals devices can be designed to work as externally backgrounded activities (Minasi 1993). The keyboard is an example. This externally backgrounded function is: “accept keystroke events from the user.” PC’s CPUs are designed so that this activity is externally backgrounded to the keyboard device and its keyboard controller chip. The CPU does not poll the keyboard controller to see if a user has typed something. Instead, the keyboard and its controller are external entities. Whenever a user presses a key, the keyboard and its controller initiate a hardware interrupt request (IRQ) and send it to the CPU. These hardware interrupts can be considered external interruptions of the CPU received from external entities handling externally backgrounded activities.

1.4.1.1 Dualtask Activity Condition 1: A Foregrounded Activity with a Suspended Activity

User interface designers have created some useful ways of supporting users’ behavior of keeping only one current activity and switching focus between suspended activities. One useful solution is to provide users with reminders of suspended activities. Windowing systems have often been employed to support user multitasking because multiple windows graphically illustrate the different activities and present visible reminders of suspended activities. Several windowing systems have been used successfully, e.g., Xeorx Star, Macintosh, NeXTstep, Microsoft Windows, Openwindows, Motif. In a windowing system, there is usually one and only one active window, and the inactive windows act as reminders of other activities. Windowing systems should therefore help reduce human errors relative to people’s tendency to ignore activities that are out of sight, i.e., “out of sight out of mind.” (Note, even though it is more common for inactive windows to be used for suspended activities, inactive windows can also be used to support backgrounded activities.)

Preparing a financial report, for example, may require the user to perform two activities concurrently with two different computer applications, e.g., a word processor and a spreadsheet. While the user has the spreadsheet open as the active window (their foregrounded activity is generating a chart from data) the word processor sits idle in an inactive window (their suspended activity is writing a financial report). However, the inactive window is usually partially visible, and this visible presence serves as a constant reminder to the user of the existence of the suspended word-processor-supported activity.

Current windowing systems are not problem-free solutions as reminders for user-multitasking. There is no one-to-one correspondence between the number of windows and the number of human activities. One visible window can represent a tool that a person is using for more than one activity; and several visible windows can represent multiple tools that a person is using for a single activity (Cypher 1986). This mismatch means that the windowing system often does not directly support the user multitasking problem. Another issue is that windows are sometimes too effective as reminders. Inactive windows present a constant source of extraneous information that can distract users (Miyata and Norman 1986). Also, managing the layout of the active and inactive windows becomes a new activity itself that requires extra user time and effort (Holden and O'Neal 1992; Hsu and Shen 1992; Shneiderman 1992, p. 337).

1.4.1.2 Dualtask Activity Condition 2: A Foregrounded Activity with an Internally Backgrounded Activity

A person’s ability to internally background an activity depends on whether they have over-learned the activity through enough practice that it can be “automated.” With diligent repetition, some kinds of activities can become over-learned to the point where a person can perform the activity with their subconscious processes. Through hours of practice, most people become are able to automate activities like walking, driving automobiles, riding bicycles, and typing. Skilled typists are able to multitask the activities of touch-typing on a keyboard while composing ideas by internally backgrounding the typing activity while they simultaneously foreground the composition activity (Miyata and Norman 1986). Unskilled typists cannot internally background the typing activity, and therefore must suspend the activity of composing ideas whenever they need to type.

Interruption in HCI is not limited to computers interrupting people. Some kinds of intelligent computer systems present a continuous machine-initiated HCI dialogue or stream of information to the user. For example, Intelligent Computer-Assisted Instruction (ICAI) systems generally control the HCI in tutoring sessions with users. To be successfully “tutored” the user must maintain the foregrounded activity of attending to the ICAI system. However, people also multitask in ICAI tutorial sessions. They are able to internally background other supportive learning activities like verifying comprehension of presented information. If a person’s subconscious activity of verifying comprehension identifies a problem, the person’s subconscious can initiate an internal interruption. After the person consciously realizes the comprehension problem, they can command the machine to suspend the tutorial session and begin a new dialogue to help solve the learning problem. ICAI systems that allow people to interrupt them with questions are said to support “mixed-initiative dialogue” (Rickel 1989). Other types of intelligent systems that support this kind of mixed-initiative dialogue are: automated telephone voice systems (Potjer et al. 1996); surgery preoperative assessment expert systems (Hoogendoorn et al. 1991); and diagnostic expert systems (Anand and Lee 1989).

1.4.1.3 Dualtask Activity Condition 3: A Foregrounded Activity with an Externally Backgrounded Activity

Externally backgrounded activities (the topic of this paper) are the only source of external interruptions from computers in HCI. Whenever an external agent brings an externally backgrounded activity to a state that requires the user’s attention the agent must initiate an interruption of the user. Starting a print job is a common example of an activity that people externally background to an external entity. A user starts a print job for some electronic document and then goes on to begin another activity. The user does not have any conscious or subconscious involvement in the activity of printing the document. The computer operating system (OS) and the printer are the external entities working on the externally backgrounded print activity. If a problem occurs with the print job, like the printer is out of paper, the printer tells the computer OS and it initiates an external interruption of the user. Different operating systems interrupt users in different ways.

If the printer is out of paper, the Apple Macintosh OS 7.6.1 abruptly initiates an external interruption of the user and displays a modal dialogue box in the center of the screen obstructing the user’s view. All other user activities are automatically suspended by the computer, and the interruption activity of fixing the out-of-paper print problem is maintained in the foreground as the only current activity. The modal dialogue box with the interruption message has an "OK" button implying that the user should be able to acknowledge the print problem and then go back to finish whatever they were doing. However, that is not how it works. Clicking "OK" only causes the Macintosh OS to check if the printer is still out of paper; if so, it will repeat this cycle indefinitely. The computer forces the user to fix the print problem before allowing them to do anything else (therefore the Macintosh OS 7.6.1 should never be used to perform an activity that cannot be unexpectedly suspended).

People can employ intelligent software agents (another kind of external computer entity) to handle backgrounded activities. Current interactive agents, however, like current operating systems, are designed to interrupt their users under certain conditions. Rich (1996) proposes that users and their intelligent agents communicate using a metaphor of a shared working environment often used in computer-supported cooperative work (CSCW) systems. Rich employs the metaphor to create a user interface for a distributed multi-user system that manages the interaction and control of a shared application. People and their software agents are both represented with identity windows, and interaction with the shared application is represented in the same way for both people and graphical pointers (the system enforces a one actor at a time rule).

Human users background certain active tasks to agents. Agents in Rich's system are constructed to behave somewhat like other human collaborators. Therefore, there is some overlap between the users foregrounded activities and their externally backgrounded activities. Rich proposes two ways for the agent to interrupt. The first way is that it just grabs control of the interface away from the current actor (this is external interruption if the current actor is a human user). Then the agent communicates to the user with text messages and/or by driving the application interface with its pointer. Rich also proposes a second, "polite," method for an agent to interrupt a user. Instead of grabbing control, when the software agent is ready to interrupt it waves its hand-shaped pointer inside its identity window in an attempt to get the user's attention. It is then left to the user to decide whether to activate the agent and allow its interruption.

1.4.1.4 MIXED ACTIVITY CONDITIONS FOR MULTITASKS WITH MORE THAN TWO ACTIVITIES

There are obvious risks associated with people performing more than one activity at a time because of human vulnerability to error during multitasking. For important tasks, it would be much safer to limit all activity to a single foregrounded activity. People, however, need to and often enjoy doing things that are too complex to be done with serial activity. For example, flying an airplane is a complex multitask that requires several activities be accomplished concurrently. Airplane pilots have the difficult task of situational awareness (staying aware of the state of externally backgrounded activities) while performing focused work on a single foregrounded activity. Pilots must employ a mixed activity solution for accomplishing their several activities concurrently: they externally background several activities like measuring altitude and monitoring for fire; they suspend and later resume other activities like instructing the crew; and they internally background other activities like manually maintaining the airplane's attitude.

Situational awareness is an important activity that pilots have learned to internally background. While they perform a foregrounded activity they subconsciously maintain awareness of what they know about other backgrounded activities. When their subconscious decides that they no longer have adequate knowledge of the state of a backgrounded activity they initiate an internal interruption, check the status of that backgrounded activity, and then can resume the previously foregrounded activity.

People make mistakes while multitasking (Schneider and Detweiler 1988; Spelke et al. 1976). The Federal Aviation Administration (FAA) concluded that most errors made by air traffic controllers result from controllers' failure to maintain situational awareness (Redding 1992). Computer multitasking technology does not necessarily fix the problem. Computers can be built as useful tools for externally backgrounding activities; however, the user's success in a multitask depends on whether the design of the user interface augments people's cognitive weak spots. Adams and Pew (1990, p. 523) say, "Intentions notwithstanding, the inherent difficulties of the multitask situation are very often compounded by the introduction of automation. To maximize situational awareness, the dynamics and capabilities of such technologies must be designed with thorough respect for the dynamics and capabilities of human information-processing".

1.4.2 Apparent Trade-Off Between Speed on a Single Task and Coordinated Performance of Multitasks

System designers have traditionally chosen HCI methods that maximize the speed of getting information into and out of computers. They have employed the assumption that, if the design of a user interface increases the user's efficiency on a single task, then that user will be more productive overall. However, a user's speed on a single task is not necessarily a good predictor of their overall performance of a multitask.

It is possible that a user interface design that facilitates a person's fast performance of a single computer task may be the very same design that hinders their ability to perform multitasks. The design of a user interface for an intelligent software agent system is a good example. This user interface design must address the system's HCI requirement that it interrupt the user. From one perspective, the most efficient way to solve this problem is to use a method of immediately interrupting the user whenever needed; this design should facilitate the user's speedy input or output of information. This design may be superb if the user is performing only one task with one software agent. However, if the user is performing multiple tasks at the same time, this method of interruption may be counterproductive. For a human multitask environment, each computer system must be designed so that it does not monopolize its users and hinder them from performing other tasks. The user interface for an intelligent agent system must not be designed to seize the user's attention away from their other tasks (except in special critical situations, like warning a person of their immanent death, e.g., "There is a coolant leak in the nuclear reactor core!").

There is an old debate in the field of HCI about whether command-based interface designs or direct manipulation interface designs are better. This debate centers around a presumed trade-off between effort and speed. Authors have debated which end of the trade-off is more important — user effort and learning time or maximally efficient task performance. Card et al.

(1983) found support for a third possibility, i.e., that the presumed trade-off could be sidestepped altogether. They found that, for several kinds of computer systems, a direct manipulation design solution could produce a system that would be both easier to learn and faster for performing tasks than command-based design alternatives.

There is another apparent HCI design trade-off between speed and multitasking. It would seem that an HCI designer must choose between HCI designs which support users' efficient performance of single tasks and HCI designs which support users' performance of multitasks. Human cognitive limitations restrict peoples' ability to both perform focused work and maintain awareness of several tasks at the same time. Because of this human cognitive limitation, it would seem that computer system designers must decide to trade-off one kind of support for another. I assert that it is possible to discover a way to sidestep this apparent HCI design trade-off, in much the same way that Card et al. found a way to sidestep the "effort vs. speed" HCI design trade-off. It should be possible to design a computer system which will both support users' efficient performance of single task and, at the same time, manage user-interruption in ways necessary for multitasking.

To find such a "win-win" solution to the complex problem of human interruption during HCI researchers must have good theoretical tools. Unfortunately, there are none for this problem.

1.4.3 No Existing General Theoretical Tools

Practitioners need general design guidelines. However, authors in the field of HCI have not yet published generally useful user interface design principles and guidelines for managing the problems associated with user-interruption in intelligent computer systems.

In fact, currently, there is no common theoretical foundation to support research in human interruption. In the literature, authors from several different research domains each describe the interruption phenomenon only within the context of their particular field, without recognizing this phenomenon in other contexts. Some authors describe aspect(s) of interruption within their particular domain: how, when, why interruptions occur, or the observed effects and side effects of interruption. Other authors begin by implying vaguely that interruptions [whatever they are] are inherently "bad" and then propose specific methods for counteracting the implied problem within their particular domain.

Any general investigation of the problem of human interruption requires two theoretical tools — neither one of which exists yet. First, a general definition of human interruption is necessary for generalizing existing results in the literature from disparate fields. Second, a general taxonomy of human interruption is required for structuring analysis, literary survey, and empirical study.

User-interruption in HCI is not a solved problem; however, there are some partially useful sources in the current literature. Burton and Brown (1979) identify some design guidelines for building ICAI tutoring/coaching systems that must interrupt their users. Galdes and Smith (1990) improve on the ICAI design guidelines of Burton and Brown by analyzing the interruption behaviors of expert tutors. They postulate that expert human tutors should know best how to interrupt people. Cooper and Franks (1993) propose informal theoretical tools for investigating human interruption within the limited context of cognitive modeling.

Burton and Brown (1979) report on their effort to design a computer-based tutor for an ICAI system that teaches math skills. The computer-based tutor in their system is an intelligent aid that runs in the background and monitors peoples performance on math-learning games. The tutor is built to detect human learning errors and interrupt the user with attempts to help them overcome learning problems. Burton and Brown say that the design problem of when to interrupt is critical to the success of the ICAI system. They say that although interrupting students for coaching purposes is sometimes useful, "Every time the coach tells the student something, it is robbing him [or her] of the opportunity to discover it for himself. Many human tutors interrupt far too often" (Burton and Brown 1979, p. 15). Burton and Brown propose twelve design guidelines for determining when and how to interrupt the user. Their guidelines make user-interruption context sensitive. For example, "If a student is about to lose, interrupt and tutor him [or her] only with moves that will keep him from losing" (principle 4); and "Do not tutor on two consecutive moves, no matter what" (principle 6).

Galdes and Smith (1990) say that Burton and Brown's guidelines are useful but are not rigorous enough and need to be empirically validated. A more useful approach would be to observe how expert human tutors' interrupt their students and apply these interruption strategies to ICAI. Galdes and Smith analyze expert human tutors' teaching behaviors and identified these human tutors' successful interruption strategies. Galdes and Smith then present these identified strategies as design guidelines for building ICAI tutorial system that must interrupt people. These guidelines, like those of Burton and Brown's, say that timing of when to interrupt must be context sensitive.

Cooper and Franks (1993) propose an interesting definition and framework of human interruption. They say that creating general theoretical tools for researching human interruption is beyond the scope of their paper. However, they suggest an informal and non-general definition and framework of human interruption based on notions of people's cognitive limitations related to processing unexpected communication events. Cooper and Franks identify human interruption as a complex cognitive process that can be used as a formative example for designing cognitive models that combine both symbolic and connectionist concepts ("hybrid systems"). They suggest that human interruption can be defined as, "any disturbance to the normal functioning of a process in a system." Cooper and Franks identify the following useful dimensions of interruption in their framework: source, effects (degree and extent), content, applicability, duration, mechanism for recovery, and state space of the underlying system (Cooper and Franks 1993, p. 76-78). Their work is not general and its usefulness limited to informal research in the field of cognitive modeling; however, it is still interesting because no other such tools existed until this dissertation.

1.5 APPROACH

This document creates and partially validates the first generalizable theoretical tools for addressing the problem of human interruption. A broad and deep survey of current literature is conducted to analyze and identify a comprehensive collection of relevant theoretical constructs. This set of identified constructs is used to synthesize the first general definition of human interruption and an accompanying practical taxonomy of human interruption. The utility of these new theoretical tools is validated in part with a survey of HCI and a human subjects experiment.

The first validation effort illustrates the utility of the Taxonomy of Human Interruption as a tool that provides a literary framework. An extensive survey is reported of the published research about human interruption in HCI. This survey is structured with the Taxonomy of Human Interruption in a unique way that facilitates the generalization of previously disparate works. The second validation effort provides support for the claim that the Taxonomy of Human Interruption is useful for guiding general research. One factor of the taxonomy is used to guide the creation of a hypothesis and its operationalization into a detailed empirical study.

The creation and partial validation of these theoretical tools is a significant contribution because it provides the first general foundation for investigating the problems associated with human interruption. This first theoretical foundation makes it possible for future studies to discover general design guidelines for this user interface problem.

CHAPTER 2: SURVEY OF THEORETICAL CONSTRUCTS

2.1 AIMS AND OBJECTIVES

The aim of this chapter is to identify a comprehensive set of theoretical information about human interruption. A broad analysis of relevant existing theory must be accomplished to serve as a foundation for the objective of the next chapter, i.e., to synthesize some generally useful theoretical tools for the investigation of human interruption.

The current literature does not yet present a general and comprehensive theoretical model of human interruption; and building such a model is beyond the scope of this dissertation. It is postulated, however, that there do exist in the current literature sufficient theoretical constructs about human interruption to form a strong foundation from which to synthesize useful tools. The object of this chapter is to form such a foundation.

After reading this chapter, readers should understand the several individual theoretical constructs relevant to investigating the interruption of humans, and readers should understand how this set of available theory can serve as a foundation for building tools.

2.2 OVERVIEW

The literature contains many theoretical constructs relevant to human interruption. This chapter identifies and discusses them. This comprehensive set of theoretical constructs is composed as a theoretical foundation for researching questions about the who-what-where-when-why-and-how of human interruption. This theoretical foundation is also useful for the more narrow questions of HCI for user-interruption.

The next chapter (“Synthesis of the First Theoretical Tools,” pg. 117) uses the results of this analysis to synthesize a unifying definition of interruption, which establishes those theoretical constructs that are most significant and ubiquitous across different fields. The breadth and depth of this analysis and the resulting unified definition’s strict simplicity make the theoretical products of this dissertation powerful tools for guiding general research about human interruption.

In the following analysis, all theoretical constructs of interruption are categorized by the four things to which they must apply: (P) the people involved in the interruption; (T) the task(s) the person is attempting; (In) the interruption itself; and (C) the working context or environment. These four categories of theoretical constructs reveal the limited scope of the general definition created in the next chapter (pg. 121). This analysis of the interruption phenomenon is limited to the context of human interruption. It is postulated that these four categories are sufficient to address all relevant theoretical constructs of the interruption phenomenon within this context.

Authors of current literature have proposed useful theoretical constructs of interruption in each of the four categories of this analysis. In the first category, the people involved in the interruption, authors have discovered particular attributes of a person’s cognitive and physical structure that affect their behavior during and after interruption. These attributes represent important structural and behavioral characteristics of a person relevant to their interruption. In the second category, the task(s) the person is attempting, authors have identified aspects of tasks that are related to the user’s changes in performance during and after interruption. These task attributes represent important qualities of tasks that affect the outcome of interruption.

In the third category, the interruption itself, authors have discovered qualities of the interruption that affect how the people involved in the interruption behave. These qualities of the interruption represent significant theoretical constructs that are relevant to people’s performance during and after interruption. In the fourth category, the working context or environment, authors have discovered particular environmental characteristics that affect the outcome of an interruption. These characteristics represent important environmental influences on the interruption phenomenon.

This chapter examines several different domains of research from the current literature. Analysis proceeds one research domain at a time and systematically extracts the relevant theoretical constructs from each domain. This chapter is organized into subsections by the particular definitions of interruption employed by different fields of research. Within each subsection, this chapter discusses the relevant theories of a particular domain of research and identifies and explains the individual theoretical constructs promoted there.

The results of this chapter represent the theoretical information relevant to human interruption, and can be used as a foundation to synthesize powerful and generally useful theoretical tools for investigating the interruption of people.

This chapter makes a significant contribution by uncovering a large set of theoretical constructs relevant to the investigation of human interruption.

2.3 MOTIVATION

A comprehensive general model of human interruption is not published in the current literature. This chapter performs a comprehensive analysis to identify the existing relevant theoretical concepts, and this set of identified theory forms a strong foundation for the synthesis of useful theoretical tools (see the next chapter).

2.4 COLLOQUIAL MEANING

Domain Specific Definition of Interruption: a word of the English language.

It is useful to begin with the etymological perspective of the meaning and usage of the word, “interrupt,” in the English language. I quote an authoritative popular standard dictionary for a definition of interruption. *Webster's Encyclopedic Unabridged Dictionary of the English Language* is a widely respected standard source for definitions of English words.

> **interrupt,** *v.t.* 1. to cause or make a break in the continuity or uniformity of (a course, process, condition, etc.). 2. to break off or cause to cease, as in the middle of something: *He interrupted his work to answer the bell.* 3. to stop (a person) in the midst of doing or saying something, esp. by an interjected remark: *May I interrupt you to comment on that last remark?* —*v.i.* 4. to cause a break or discontinuance; interfere with action or speech, esp. by interjecting a remark: *Please don't interrupt.* [ME *interrupte*(*n*) < L *interrupt*(*us*) broken apart (ptp. of *interrumpere*), equiv. to *inter-* INTER- + *ruptus* broken; see RUPTURE] —interruptedly, *adv*. —interruptedness, *n*. —interruptible, *adj*. —interruptive, *adj*.
>
> —Syn. 1, 3. intermit. INTERRUPT, DISCONTINUE, SUSPEND imply breaking off something temporarily or permanently. Interrupt may have either meaning: *to interrupt a meeting.* To DISCONTINUE is to stop or leave off, often permanently: *to discontinue a building program.* To SUSPEND is to break off relations, operations, proceedings, privileges, etc., for a longer or shorter period, usually intending to resume at a stated time: *to suspend operations during a strike.* —Ant. 1, 2. continue (Random House 1989, p. 744).
>
> **interruption,** *n*. 1. the act or an instance of interrupting or the state of being interrupted. 2. something that interrupts. 3. cessation; intermission. [ME *interrupcio*(*u*)*n* < L *interruption-* (s. of *interruptio*)] (Random House 1989, p. 744).

These definitions and usage quotation examples help us understand what authors usually mean when they use the word interruption. I propose that this definition is useful because it describes what most English-speaking people believe to be common and obvious about the phenomenon of interruption. I distill from Webster's definition six purportedly common or obvious theoretical constructs about the interruption of people's activities (including interruption of their speech). I use a metaphor of water flowing through a ditch to illustrate the theoretical constructs of interruption, which I identify from the preceding definition. In this metaphor, interruption is what happens when the ditch is blocked and the water stops.

THEORETICAL CONSTRUCT(S) about people (Random House 1989, p. 744)**:**

P1. Human activities are continuous, fluid processes (like flowing water).

P2. Human activities have coherence over time (like the surface tension of water).

P3. People's actions are interruptible (in the same way that water can be divided).

THEORETICAL CONSTRUCT(S) about interruption (Random House 1989, p. 744)**:**

In1. An interruption is something that breaks the coherence of an activity and blocks its further flow (like dropping, or interjecting, a large rock into an irrigation ditch).

THEORETICAL CONSTRUCT(S) about people (Random House 1989, p. 744)**:**

P4. People can resume activities that have been interrupted once the interruption is removed (like removing the rock).

P5. People often use conventional protocols for interrupting each other's speech, for example, "May I interrupt you to comment on that last remark?" There are also protocols for interrupting all other kinds of human activities. (Also, I can use my metaphor of blocking an irrigation ditch with a rock to illustrate people's use of protocols. If I just drop the rock in the ditch, I will splash water and mud all over myself. Instead, I must use the protocol of slowly lowering the rock into place.)

2.5 MULTITASKING IN HCI

Domain Specific Definition of Interruption: an unanticipated request for task switching during multitasking.

Tsukada et al. (1994) present a practical discussion of how people get work done in computer-supported cooperative work (CSCW) office environments where a person is responsible for advancing several projects at the same time. It is unusual for a person to be engaged in only a single activity from start to finish to the exclusion of all other tasks. This behavior in which a person accomplishes two or more tasks within the same time period is called multitasking. Researchers in this field define interruptions as unanticipated requests for switching between different tasks during multitasking. (See also (Preece et al. 1994, p. 105).)

Tsukada et al. does not specifically address the interruption of a worker, however the authors make a useful distinction between a person's internal and external actions. Tsukada et al. says that people can multitask because they internally concern themselves with all their multiple tasks at once, in parallel, but externally act on only one task at a time. People multitask by frequently alternating their external efforts between each of their multiple tasks. The result is that a person can accomplish multiple tasks concurrently, as Figure 1 shows.

Tsukada et al. defines theoretical constructs about the cognitive and physical structure of the people involved in the interruption and about the requirements of tasks. These theoretical constructs are relevant to the interruption phenomenon.

THEORETICAL CONSTRUCT(S) about people (Tsukada et al. 1994):

P6. There is a useful distinction between people's internal efforts and their external efforts on tasks. People's internal and external efforts are related, but there is also some amount of independence between these two kinds of efforts. People's external efforts (observable behaviors) are dictated by their internal efforts (cognition), but not all of people's internal efforts become expressed as external efforts.

P7. People can exert external effort on only one task at a time.

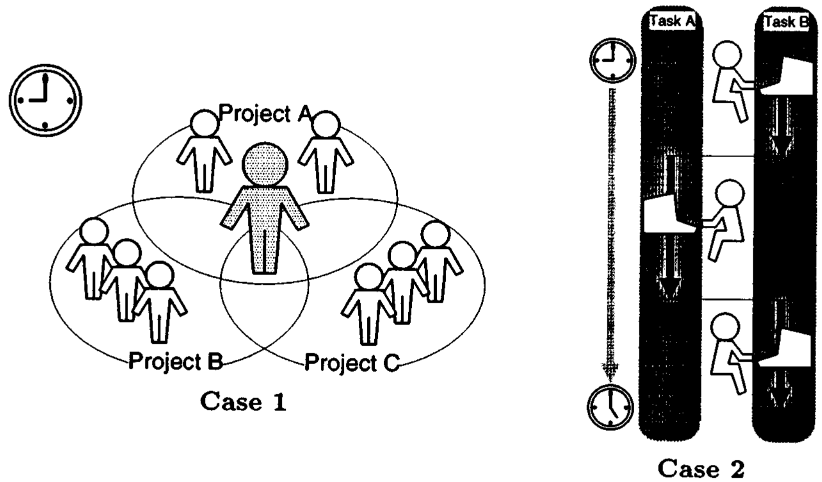

Figure 1.The left-hand diagram illustrates the parallel nature of a worker's internal effort in multitasking. The right-hand diagram illustrates the serial time-sharing nature of a worker's external efforts in multitasking. Reprinted from Tsukada et al., "The Multi-Project Support System Based on Multiplicity of Task," IEEE. © 1994 IEEE.

P8. People can switch their external effort from one task to another; i.e., a person can stop their external efforts on one task before it is completed and begin or resume their external efforts on another task.

P9. People can exert internal effort on multiple tasks at the same time — in parallel.

THEORETICAL CONSTRUCT(S) about tasks (Tsukada et al. 1994):

T1. It is not required that a task be accomplished all at once, but a task can be performed through the accumulation of many independent, noncontiguous efforts.

Card, Moran, and Newell (Card et al. 1983) present some of these same theoretical constructs (11 years before Tsukada et al.) in their book, *The Psychology of Human-Computer Interaction*. Card et al. only model a person performing one task at a time. However, I think their work is fundamental to this discussion of multitasking and user-interruption. Card et al. present a model of the structure and function of human cognition relevant to task execution. They call their model "The Model Human Processor" (Card et al. 1983, p. 24), as seen in Figure 2.

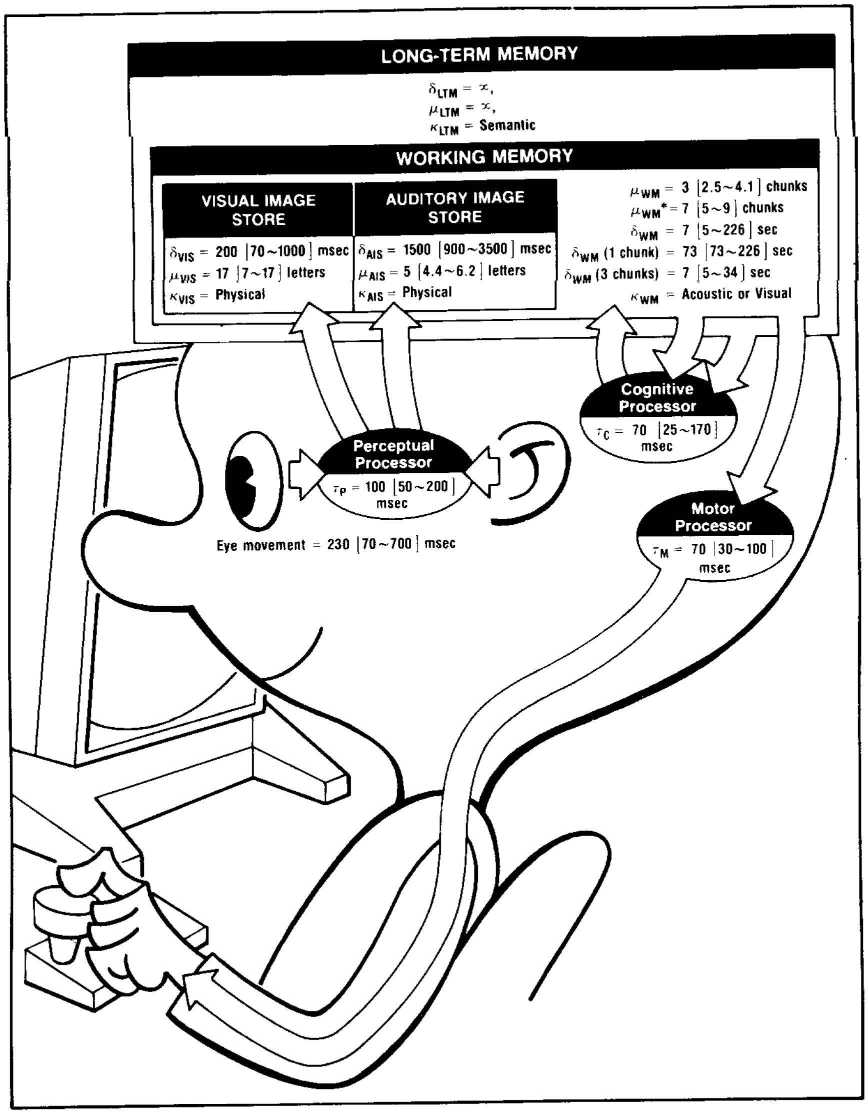

Figure 2.“The Model Human Processor — memories and processors.” From *The Psychology of Human-Computer Interaction* (p. 26) by Stuart K. Card, Thomas P. Moran, and Allen Newell, 1983, Hillsdale, NJ, Lawrence Erlbaum Associates, Inc., Copyright © 1983 by Lawrence Erlbaum Associates, Inc.

The Model Human Processor depicts human cognition with three parallel processors. These separate processes model human cognition in a way that allows a model of a human to perform three kinds of internal processing simultaneously. This model also limits a theoretical person to one external action at a time, because only one processor, the Motor Processor, controls external actions. Card et al. say, “the cognitive system is fundamentally parallel in its recognizing phase and fundamentally serial in its action phase. Thus the cognitive system can be aware of many things but cannot do more than one deliberate thing at a time" (Card et al. 1983, p. 42).[1]

Card et al. discretize actions at ~ 70 msec units. Actions are discretizable because the Motor Processor of the Model Human Processor is cyclical. This cyclic behavior of the Motor Processor divides its output into discrete units. Card et al. say that the cycle time of the Motor Processor is 70[30-100] msec (Card et al. 1983, p. 34). (This means the typical value is 70 msec and that the possible range is from 30 msec to 100 msec.)[2] The authors say that people perform all their motor behaviors merely with long chains of tiny 70 msec actions. The Model Human Processor allows researchers to quantify larger actions as sums of the different kinds of tiny 70 msec actions that comprise them, as seen in Figure 3.

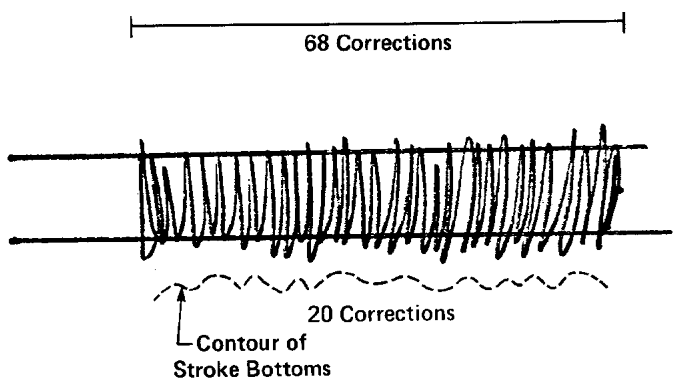

Figure 3. An illustrative example task from Card et al. People's observed behavior informally validates the Model Human Processor's cycle rate for the Motor Processor. From *The Psychology of Human-Computer Interaction* (p. 35) by Stuart K. Card, Thomas P. Moran, and Allen Newell, 1983, Hillsdale, NJ, Lawrence Erlbaum Associates, Inc., Copyright © 1983 by Lawrence Erlbaum Associates, Inc.

1. Compare this idea of parallel cognition and serial action to the theoretical constructs of interruption P7 (pg 28), and P9 (pg 29). This idea or theoretical construct is not unique to the Model Human Processor. In fact, several of the theoretical constructs that I discuss in Section 2 have sibling constructs in different domains of research. One contribution of this chapter is to find and gather these siblings together for examination so that the gereralizable part of the idea can be extracted.

2. Card et al. choose the specific numbers for their Model Human Processor capabilities from empirical studies in contemporary literature. For example, they set the cycle time of the Motor Processor at 70[30-100] msec. Card et al. describe the several other papers they used to find an average estimate for the Motor Processor speed at 70 msec, and the published extreme observations at 30 msec and 100 msec.

Figure 3 shows the results of one subject's moving a pen back and forth between two lines as fast as possible in 5 seconds. The subject made 68 lines in 5 seconds — thus 5000 msec divided by 68 actions equals about 74 msec per action. Card et al. say this observed behavior is illustrative, informal support for the speed of their Model Human Processor's Motor Processor cycle rate. The Model Human Processor allows us to say that this squiggly line drawn by a person in 5 seconds is actually the result of a chain of 68 tiny discrete actions (Card et al. 1983, p. 35).[3]

THEORETICAL CONSTRUCT(S) about people (Card et al. 1983):

P10. People can perform actions in parallel along three dimensions — perceptual, cognitive and motor, but within each of these three dimensions, people must perform actions sequentially.

P11. People discretize tasks cognitively. People hierarchically decompose large tasks into smaller tasks, and they continue this decomposition until subtasks are reduced into indivisible units of work. The size of these basic units of work is related to the cycle time of human's three cognitive processors.

P12. People perform large actions by executing chains of smaller discrete actions.

Card et al. also provide a modeling and analysis tool that supports the theoretical construct of interruption T1 (pg 29). They apply some of their ideas from the Model Human Processor to create a family of practical analysis tools called the GOMS (Goals, Operators, Methods and Selection rules) Models. The GOMS Models rely heavily on the idea that human actions are discretizable. GOMS can be employed to model how a person would perform a given task. The task is analyzed hierarchically into the subtasks that comprise it ("Goals"). These subtasks are modeled with chains ("Methods") of basic operations ("Operators") that must be

3. This observed 74 msec per action illustrates the cycle time of the Model Person's Motor Processor in isolation. The observed correction behaviors can be modeled with the Model Person with chains of behaviors in which the total time is the sum of the cycle rate of each processor employed. Each correction behavior represents a chain of one cycle each for the Model Person's Perceptual Processor, Cognitive Processor, and its Motor Processor. Twenty corrections in 5 seconds means 250 msec per corrective chain of behavior.

performed to accomplish them. The “Selection rules” are productions to simulate which chain a person would choose among alternatives to complete a subtask depending on the context. GOMS Models can be used to analyze tasks and make a priori, quantitative predictions of human performance.

Card et al. does not support our need to model multitasking or the interruption of the user by unanticipated requests for task switching. They have only intended the Model Human Processor and the GOMS Models to be used to model and analyze the event of a single user performing a single task, uninterrupted. The three processors of the Model Human Processor are intended only to model parallel cognition on a single task. The GOMS Models have no way of suspending the execution of one task and resuming another. I have included Card et al., however, because the Model Human Processor and GOMS Models represent early contributions of theoretical constructs about people and tasks that are relevant to the interruption phenomenon. They proffer the ideas that human actions can be discretized and that arbitrary chains of atomic actions can be composed to model complex actions. They also promote the idea that subtasks can be accomplished by dynamically selected chains of atomic operations.

More recently, the GOMS Models have been extended to model aspects of multitasking. John and Gray (1995) present a modified version of GOMS called CPM-GOMS. CPM-GOMS is Critical Path Method-GOMS (CPM also can stand for Cognitive, Perceptual, and Motor operations). CPM-GOMS further applies the structure of the Card et al. Model Human Processor by specifically employing the idea of three separate processors (Cognitive, Perceptual, and Motor). CPM-GOMS can be used to model human performance on tasks by first modeling subtasks with short chains of operations. These small chains must then be linked in sequence to model human performance on larger tasks.

These chains are executed by scheduling their respective operators on three separate time tracks — one per processor. Each operator needs to be scheduled on an appropriate processor (Cognitive, Perceptual, or Motor). Therefore, a small chain of operators that models a single subtask may have operators scheduled on each of the three processor schedules. The result of CPM-GOMS modeling is a three-part parallel schedule for the three processors of the model.

The schedule can be traced in parallel to estimate the time required for a person to perform the task.

However, the real improvement of the CPM-GOMS over the original GOMS comes from using the added flexibility of its improved control structure. After finishing the model of the entire task as a long succession of small chains, a researcher can begin using the flexibility of CPM-GOMS to improve the accuracy of the model. CPM-GOMS allows a researcher to collapse the final chain by interleaving the smaller chains that comprise it. As long as temporal dependencies are preserved, the operators from different subtasks can be interleaved within the processor schedule tracks. So, for example, if the researcher notices that at one point in its schedule, the Cognitive Processor is idle, waiting for the result of the Motor Processor, the researcher can collapse the Cognitive Processor's schedule so that it can work on an operator for a successive subtask while waiting. Later, when the Motor Processor has finished, the Cognitive Processor can resume executing the operators for the subtask it had begun.

John and Gray have implemented fundamental aspects of the structure of human cognition in CPM-GOMS. They apply the idea that people can do things while they wait for themselves to finish doing other things. This sounds strange, but since people are processing information in parallel on three processors (the Card et al. Model Human Processor — Cognitive, Perceptual, and Motor), they can do three things at a time. So if a person is only using one of their three processors to perform some task, they can use their two idle processors to do other tasks while they wait for themselves.

THEORETICAL CONSTRUCT(S) about people (John and Gray 1995):

P13. People can intermix their actions for different tasks because of their ability to act in parallel along three dimensions (cognitive, perceptual, and motor). While they wait for themselves to finish some basic processing along a single dimension, they can perform work on other, possibly unrelated, tasks within each of their other two processing dimensions.

P14. People can switch their actions between different tasks quickly and effortlessly.

Preece et al. (1994) say that although people are able to multitask, they have cognitive limitations that make them vulnerable to distraction.

> "While most people show great flexibility in coping with multitasking, they are also prone to distraction. On returning to a suspended activity, it is possible for them to have forgotten where they were in the activity. As a result they may not restart from where they left off but will recommence at a different point of entry. For example, pilots may think they have completed part of a procedure (such as a checklist) but in fact they have not done so. [See story on p. E-1 about the airplane pilots and an uncompleted checklist.] Alternatively, they may forget that they have already done something and repeat it. This most frequently occurs for routine procedures where knowledge for carrying out the various tasks has become largely automated. An everyday analogy is forgetting to salt the potatoes or adding the salt twice, if our routine procedures when cooking are interrupted by having to answer the phone" (Preece et al. 1994, p. 105).

Preece et al. (1994) say that distraction affects people's memories. Distraction diverts their attention and causes them to forget what they were doing. The combination of distraction and interruption can have especially bad consequences. The occurrence of an interruption is a circumstance when it is particularly important for a person to remember what they had been doing on the interrupted task. Since distraction affects memory, it can cause people to make serious memory errors when resuming the interrupted activity.

Webster's Encyclopedic Unabridged Dictionary of the English Language presents a useful colloquial definition of "*distraction*." I quote it here:

> **distract** *v.t.* 1. to draw away or divert, as the mind or attention: *The music distracted him from his work*. 2. to divide (the mind, attention, etc.) between objects. 3. to disturb or trouble greatly in mind: *Grief distracted him*. 4. to amuse; entertain; provide a pleasant diversion for: *I'm bored with bridge, but golf still distracts me*. 5. to separate or divide by dissension or strife. —*adj*. 6. Obs. distracted. [< L *distract*(*us*) (ptp. of *distrahere* to draw apart), equiv. to *dis*- DIS- + *tract*- (perf. s. of *trahere* to draw) + -*tus* ptp. suffix] —

distracter, *n*. —distractibility, *n*. —distractible, *adj*. —distractingly, *adv*. (Random House 1989, p. 417).

distraction *n*. 1. the act of distracting. 2. the state of being distracted. 3. mental distress or derangement: *That child will drive me to distraction*. 4. that which distracts, divides the attention, or prevents concentration: *The distractions of the city hinder my studies*. 5. that which amuses, entertains, or diverts; amusement; entertainment: *Fishing is his major distraction*. 6. division or disorder caused by dissension; tumult. [< L *distraction*- (s. of *distractio*) separation (Random House 1989, p. 417).

THEORETICAL CONSTRUCT(S) about people (Preece et al. 1994)**:**

P15. There exist certain stimuli, called distracters, that can affect people's attention and working memory outside of their conscious control and awareness.

THEORETICAL CONSTRUCT(S) about tasks (Preece et al. 1994)**:**

T2. There exist some nonwork activities that can distract people from work tasks.

THEORETICAL CONSTRUCT(S) about interruption (Preece et al. 1994)**:**

In2. There is an interaction effect between distraction and interruption. When a distraction is associated with an interruption, people become prone to make serious memory errors when attempting to resume interrupted tasks.

2.6 MULTITASKING IN LINGUISTICS

Domain Specific Definition of Interruption: an unanticipated request for topic switching during asynchronous parallelistic human-computer interaction.[4]

4. This definition is similar to the preceding definition regarding multitasking, however it comes from a field with different goals and theoretical concepts.

Edmondson (1989) in his paper titled "Asynchronous Parallelism in Human Behaviour: A Cognitive Science Perspective on Human-Computer Interaction," explains how a particular linguistic theory can be useful in HCI research of multitasking. He proposes that a prominent theory of linguistics about phonology, called autosegmental or nonlinear phonology, can provide useful concepts and formalism for HCI research about multitasking environments.

Edmondson says that people exhibit the readily identifiable and common behavior of "asynchronous parallelism." He says that this behavior can be observed in many different kinds of human activities, including human-human interaction (the domain of linguistics) and human-computer interaction (HCI). Edmondson says that asynchronous parallelism describes a human behavior in which a person does several things at once (parallelism), but they accomplish this by working on only one thing at a time while interleaving the execution of all the different activities (asynchronism). Other popular terms that refer to people's asynchronous parallelistic behavior are nonlinear, plurilinear, or interleaved behavior.

Edmondson does not directly address the interruption of the user. However, he does show how tools for addressing asynchronous parallelism can be generalized across domains. He postulates that human asynchronous parallelistic behavior will have similar cognitive requirements and limitations across different domains. This premise has two important implications: (1) the tool Edmondson proposes, nonlinear phonology, can be useful for addressing some issues of HCI in supporting people's multitasking, and (2) other tools created for other domains involving people's asynchronous parallelism (some of which directly address human interruption) can be generalized to the HCI domain.

THEORETICAL CONSTRUCT(S) about people (Edmondson 1989)**:**

P16. Asynchronous parallelism is a common and easily observed behavior exhibited by people in many and widely varied activities, including human-human interaction and human-computer interaction (see Theoretical Constructs P7 (pg 28), P8 (pg 29), and T1 (pg 29)).

Edmondson reports that linguist authors have researched asynchronous parallelism in people's language use and have published useful theoretical concepts and formalisms in the theory of nonlinear phonology. Edmondson says that HCI researchers can adopt and apply these products of linguistic theory to help them research human-computer interaction in domains where people exhibit asynchronous parallelism, e.g., the HCI of systems that support multitasking. The concepts proposed by these linguist authors are useful here. The formalism, however, is not useful in our attempt to define interruption.

People exhibit asynchronous parallelism in their human-human interaction. This behavior is possible because of the particular structure and function of human cognition. Edmondson's proposed linguistics implies several theoretical constructs of human cognition. I summarize the implied theoretical constructs here. People's linguistic abilities are provided by six general theoretical constructs of cognition:

1. People maintain and operate discrete units of linguistic expression at several levels of abstraction,
2. People cognitively prepare linguistic expressions of meaning (at each level of abstraction) as sequences of concatenated or interleaved units of linguistic expression,
3. People physically express meaning by acting out, one at a time in sequence, each discrete linguistic expression from their cognitive plan,
4. People self-monitor the success and appropriateness of their linguistic expressions while they are making them,
5. People use the information from their self-monitoring behavior to dynamically modify and change their cognitively prepared sequence of linguistic expressions as they continue to execute each discrete expression, and
6. People exhibit asynchronous parallelism in expressing meaning to other people — they cognitively prepare in parallel but express in sequence (asynchronously).

Card et al. propose the concept that people accomplish physical actions by sequentially performing long strings of discrete movements (see Figure 3 on pg 31, and Theoretical Constructs P10 on pg 32, P11 on pg 32, and P12 on pg 32) (Card et al. 1983). Edmondson is proposing this same idea for people's linguistic expressions. For example, people maintain and operate a set of discrete phonemes (a low level of verbal abstraction). When people want to convey meaning to others, they cognitively prepare (in parallel) a sequential list of phonemes and then sequentially speak each phoneme in the list (asynchronously). Edmondson says this same method of asynchronous parallelism is employed at many different levels of linguistic abstraction, e.g., phonemic, morphemic, syntactic, semantic, and pragmatic (rhetorical).

People are not computers doing batch processing on the expression of a list of linguistic units. Instead, people continually monitor the success and appropriateness of their own linguistic expressions. People self-monitor. People's interactivity is not suspended when they are in the process of expressing a sequential list of linguistic units. They can watch their own interactive progress and revise and rework their cognitive composition of planned linguistic actions when needed.

This ability to self-monitor and dynamically replan allows a person to accept interruptions while they are in the very act (in flagrante delicto) of expressing a sequence of linguistic units. For example, if they begin speaking a word and then receive a request for interruption, they can immediately stop speaking that word (without finishing it) and directly begin a totally new sequence of linguistic units. For example, I am at my home talking to Robert (my brother) about which kind of paper he should use to print his résumé. Suddenly, I notice that Kate (my 2-year-old daughter) is about to draw on herself with a marker. I can interrupt myself instantly — I do not even have to finish the word I am currently speaking to Robert. [Quote] "I think this other paper is mo/Kate, no! Markers are not for skin."

In summary, people's asynchronous parallelism is supported by a combination of the following cognitive theoretical constructs: discrete operators; parallel cognitive planning; sequential physical action; continuous monitoring of self and environment; and dynamic cognitive replanning.

Why is asynchronous parallelism so ubiquitous in human behavior? The theory of natural selection provides a useful answer. Our progenitors who could not behave with asynchronous parallelism did not tend to survive. For example, a group of people are standing in the open and talking to each other. Suddenly a pride of lions rushes out of the tall grass nearby. Those people who can behave with asynchronous parallelism IMMEDIATELY interrupt whatever they are saying and RUN. Those other people who cannot behave with asynchronous parallelism do not start running immediately, but instead must stand and finish what they were saying before being interrupted. These stalwart talkers ... tended not to reproduce.

THEORETICAL CONSTRUCT(S) about people (Edmondson 1989)**:**

P17. People maintain and operate discrete units of linguistic expression at several levels of abstraction, e.g., phonemic, morphemic, syntactic, semantic, and pragmatic.

P18. People cognitively prepare linguistic expressions (at each level of abstraction) as sequences of concatenated or interleaved units of linguistic expression.

P19. People physically express their cognitive plan for linguistic interaction by expressing linguistic units one at a time, in sequence.

P20. People monitor themselves and their environment while they are making linguistic expressions.

P21. People use the information from their self-monitoring and environmental-monitoring to dynamically modify and change their cognitively prepared sequence of linguistic expressions as they continue to execute sequentially ordered discrete expressions.

P22. People exhibit asynchronous parallelism — they cognitively prepare in parallel but physically express in sequence (asynchronously).

P23. People can successfully interact with each other (giving and receiving meaning) in an asynchronous parallelistic way. People can understand another person when that person physically expresses, one discrete unit at a time, a sequence of cognitively planned linguistic units.

Edmondson says that asynchronous parallelism in linguistics explains why the inclusion of one linguistic unit in a planned sequence can affect the expression of other linguistic units close by. For example, “Consider the word ‘construe’ as it is often pronounced by native speakers of English. The second syllable is frequently articulated with lip-rounding throughout, although the lip-rounding is only required as a feature of the vowel. What is happening is that the specification of lip-rounding is spreading back to influence the articulation of the syllable initial consonants (backward assimilation of rounding to the consonant cluster)” (Edmondson 1989, p. 6). Edmondson says that this type of behavior is evidence of people’s parallel cognitive planning before their sequential physical expression of linguistic units — asynchronous parallelism. If the linguistic units were cognitively planned sequentially instead of in parallel, then planning one linguistic unit would not influence the specification of preceding linguistic units.

2.7 MULTITASKING IN SITUATIONAL AWARENESS

Domain Specific Definition of Interruption: an event that threatens the delicate balance between situational awareness and focused activity, i.e., the reception of unpredictable new data.

Situational awareness is the product of deliberate divided attention. Some authors report research in domains where human situational awareness is critical to successful performance of person-machine systems. One good example is the person-machine system of a cockpit of a commercial aircraft. A pilot has several tasks to perform in concert while at the same time keeping an awareness of the current state of the plane and its outside environment. The pilot must infer the current situation from the information presented by over 400 separate gauges and instruments (Adams and Pew 1990). Situational awareness is essential because a pilot must make decisions that are context sensitive — the correct decision depends on the current state of the airplane.

The balance between situational awareness and focused activity is delicate because of peoples’ cognitive limitations. Adams and Pew emphasize the fragility of this balance for aircraft pilots, “an **interruption**, an oversight, a hasty inference, of a decision based on incomplete knowledge or information: under conditions of heavy workload or tight temporal pressure — any crew is vulnerable to each — could mean disaster" (Adams and Pew 1990, p. 519).

A person who has invested the effort to construct and maintain situational awareness has the advantage of already possessing the critical information in their heads when they are called upon to make important decisions. Situational awareness is indispensable in time-critical tasks. Flying an airplane is such a task. An airplane pilot sometimes must make emergency split-second decisions. In these emergency situations, a pilot does not have time to reconstruct knowledge of the current state of the airplane and its environment. This situational awareness cannot be reconstructed when needed in emergencies because it requires too much time— the pilot must read the 400+ gauges and make the necessary inferences to acquire awareness of the situation. It is dangerous for a pilot to allow themself to become interrupted or distracted from their responsibility to construct and maintain situational awareness. When an emergency occurs, the pilot either has the essential situational awareness or not. If yes, then they are ready to make good decisions. If no, then they are ready to make bad decisions.

There are several sobering examples of the possible costs of failure of aircraft pilots to maintain situational awareness. A commercial aircraft crashed in 1972 killing 99 passengers and crew members because none of the crew was aware of the airplane's altitude. In fact, none of the crew was even aware that no one was flying the plane. All of the crew members had become totally focused on solving the problem of a burned-out light bulb in the system that indicates the status of the landing gear. All three crew members were so focused on this minor problem that no one noticed that the autopilot had become disengaged. While they were working on the light bulb problem, the airplane gradually drifted down and crashed in the Florida Everglades. As the airplane gradually descended, an air traffic controller noticed on his radar screen that the aircraft was losing altitude and called the crew on the radio and asked, "how are things comin' along out there?" Because the crew was focused on the light bulb problem they probably assumed that the air traffic controller was inquiring about that, and they responded that everything was all right (Foushee and Helmreich 1988, pp. 194-195; National Transportation Safety Board 1972).[5]

The light bulb example is not unique. A commercial aircraft crashed in 1978 because the crew failed to maintain situational awareness, also because of a burned-out light bulb. The crew was not able to confirm that the landing gear was down and locked because of a burned-out light bulb, so they prolonged landing the plane while they tried to solve the problem. The captain would not attend to the fact that the plane's fuel was getting dangerously low, and the plane ran out of fuel and crashed several miles from the Portland, Oregon airport (Foushee and Helmreich 1988, pp. 194-195; National Transportation Safety Board 1979).

The stories of airplane accidents attest to the fact that maintenance of situational awareness is very difficult for people. This observed difficulty suggests something about human cognition. Shneiderman (1992, p. 84) comments on this human weakness in his summary of the relative capabilities of humans and current machines. He says that machines are better than people at monitoring prespecified, especially infrequent events, and that people are better than machines at sensing unusual and unexpected events. These theoretical constructs of human cognition often result in people being easily distracted or interrupted from monitoring tasks they are attempting, i.e., people are predisposed to fail at situation awareness.

THEORETICAL CONSTRUCT(S) about people (Shneiderman 1992)**:**

P24. People are not proficient at monitoring events; instead, they are very sensitive to detecting unusual and unexpected events.

Adams and Pew (1990) have written a paper in which they define and review the aspects of human cognition that are relevant to situational awareness in the person-machine system that is a commercial aircraft cockpit. In this paper titled, "Situational Awareness in the Commercial Aircraft Cockpit: A Cognitive Perspective," Adams and Pew say why interruptions are disruptive to the task of situational awareness. "To notice the occurrence of an event in any useful way, the pilot must immediately **interrupt** ongoing activities, at least to evaluate its significance, and establish the priority of its response implications. Resumption of the **inter-**

5. There is a good source of literature about peoples' cognitive vulnerability to becoming fixated on one topic to the exclusion of others. See references to the Einstellung phenomenon or psychic blindness, e.g., Lane and Jensen 1993.

rupted task must require thoughtful review of its status and may require repetition or reinitiation of one or more of its procedural components. Thus, the very reception of unanticipated data must always introduce an additional and disruptive element of workload. The design implications, especially for time-critical systems, should not be ignored" (Adams and Pew 1990, p. 523).

Adams and Pew detail the relevant task requirements of piloting an airplane. The combination of these requirements describes a task that is especially vulnerable to interruption. We can learn about the interruption phenomenon from a discussion of why this task is difficult. Adams and Pew (1990, p. 520) present four categories of task requirements: (1) there are several tasks that must be performed in concert, each demanding focused attention; (2) each of these several tasks can be both knowledge intensive and procedurally complex; (3) the demands of each of these several tasks are interleaved in time with the demands of other tasks in no predictable order (this leads to situations where "the urgency of executing one or more tasks is liable to peak at the very moment when information triggering, enabling, or urging completion of others is arriving"); and (4) the relevance of each piece of available information is not conveyed by its source or presentation context.

THEORETICAL CONSTRUCT(S) about tasks (Adams and Pew 1990)**:**

T3. Multitasks combining the following requirements are especially vulnerable to failure due to interruption: (1) several tasks simultaneously require the person's focused attention; (2) each of several tasks requires extensive cognitive memory and processing resources of the person; (3) the operations of the tasks must be performed in an unpredictably interleaved order; and (4) the relevance of each piece of available information is not apparent.

Adams and Pew present theoretical constructs of human cognition that are useful here. They say that the constructivist approach to perception is most useful in explaining why it is difficult for people to maintain situational awareness while performing other tasks. According to the constructivist approach, people make sense of new information by employing their own memories to tell themselves what they are seeing. In other words, people do not directly see the world as the images that fall on their eyes; instead they look out at the world through the amazing lens of their memories. The information people receive is usually incomplete and fraught with error and noise. People use their memories to fill in the gaps and allow themselves to make sense of their environments (Preece et al. 1994, section 4.1). Adams and Pew contend that when people maintain situation awareness, they compete with themselves for the cognitive resources they need to constructively perceive incoming information. People are at the same time trying to use these same memory resources to accomplish other tasks.

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990)**:**

P25. People must tap their cognitive memory resources when maintaining situational awareness because of their constructivist method of interpreting incoming information. This need for resources causes internal competition between tasks, because each task requires the same cognitive resources.

Adams and Pew say that the structure of human long-term memory has several useful implications for explaining the problems associated with situational awareness. They adopt a connectionist model of memory. This theory says human memories represent information as networks of basic theoretical constructs or concepts and links between them. In other words, each piece of information is represented with the set of its component parts, together with their relationships to each other. The structure of information is preserved by the interconnections in these networks, and the details are preserved by the primitive units of memory. Adams and Pew (1990, p. 521) say that the three most useful aspects of this theory of memory are: (1) the primitive units of memory are not duplicated and are relatively small in number (this means that all memories are just hierarchies of networks of links to the same basic set of primitive units); (2) memories function not only as records of information but also as the medium of perception and interpretation for new experiences (constructivist perception); and (3) the salience of memories increases with the frequency of their use, and since different memories share the same basic set of primitive units, the use of one memory affects the salience of other similar memories.

This connectionist theory of memory is useful in explaining why situational awareness allows pilots to make good decisions during emergencies. While the pilot is constructing and maintaining situational awareness, they are affecting the salience of other memories they have that are related to the particularities of the current situation. This explains why, when emergency strikes, the pilot is able to remember specific information very quickly that will help to solve a particular emergency. The memories of related background information, contingencies, exceptions, and conditional responses have been made more salient because of the frequency of activation of common memory primitives with the specific situation (Adams and Pew 1990, p. 521).

Connectionist theory can also be used to explain why experts are better than novices at both making good emergency decisions and maintaining situational awareness. An expert, by definition, has a much broader repertoire of relevant memories to help with constructivist perception than a novice. The expert uses this more capable and more efficient perception to maintain situational awareness better and more easily than a novice. A by-product of constructivist perception is that other memories that have common memory primitives become more salient. Therefore, memories the expert uses to construct perception become more salient because of this use. When an emergency occurs, the expert has many relevant memories easily available (salient); the novice has many fewer. This accessibility of relevant memories (the causes of similar emergencies and alternative viable solutions) allows an expert to make better emergency decisions than a novice.

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990)**:**

P26. People have connectionist long term memories — each piece of information is represented as a hierarchical collection of its component parts, together with a network that represents the interrelationships of the parts to each other. This memory structure has three characteristic properties: (1) the primitive units of memory are not duplicated and are relatively small in number; (2) memories are used to both store information and as a medium for perception (people have constructivist perception); and (3) the salience of memories increases with the frequency of activation of their constituent primitive memory units.

P27. While people use their memories to constructivistly create and maintain a situational awareness, they affect the salience of all their memories that are somehow related to the particularities of the current situation. This increased salience of related memories prepared during situational awareness efforts gives people fast and ready access to relevant information in unexpected emergencies.

Many classical theories of human memory propose two kinds of memory: short-term memory (working memory) and long-term memory. Classical theory of human cognition says that short-term memory is limited to seven plus or minus two items at a time (Miller 1956). This limitation seems to conflict with a connectionist theory that portrays human memory as being capable of sustaining many, possibly complex, memories active at the same time. Adams and Pew address this apparent weakness in connectionist theory. They adopt a theory of human memory that introduces structures which are useful for explaining human cognitive limitations within a connectionist framework.

Adams and Pew support a theory by Sanford and Garrod (1981) that says there are four different kinds of memory: two kinds of active memory (explicit focus and implicit focus), and two kinds of latent (currently inactive) memory (long-term episodic and long-term semantic). Sanford and Garrod say that each of these different kinds of memory has a different structure and function and that these differences in structure are useful in explaining human memory behavior and limitations. They have limited the scope of this theory to defining structures that influence the way human memories are retrieved. (Other aspects of memory are not addressed by this theory, e.g., formation of memories, use of memories in constructivist perception, and selection between competing memories.) Sanford and Garrod demonstrate the usefulness of this theory of memory structure in their domain of text comprehension. Therefore, if we can generalize from the text comprehension domain, we can use this theory to explain how memories are made more or less salient. This is a useful tool for talking about the interruption phenomenon, because we can use it to explain memory problems related to interruption events.

Explicit focus memory is active memory that has been the subject of classical memory studies to measure “short-term memory.” The explicit focus consists of a tightly limited number of tokens (or pointers), which refer to larger knowledge structures in long-term memory. So, explicit focus memory is like an index, and its tokens are the references listed there. The salience of tokens in this index are maintained dynamically. (The degree of salience determines how easily a memory can be recalled and used.) The salience of a token is determined by its recency of use and by its relevance to the current context (relevancy to the current context can be explained by a discussion of constructivist perception). Implicit focus memory is composed of the large, possibly complex, active memory structures referred to by the tokens in the explicit focus memory.

Classical studies of short-term memory suggest that there are only seven plus or minus two tokens in short-term memory (Miller 1956), however Sanford and Garrod's theory suggests a more useful model. Explicit focus can contain more than seven plus or minus two tokens, each with different salience. I suggest that classical studies have measured the number of tokens that a person's explicit focus memory can keep at maximum salience — seven plus or minus two. Maximum salience of tokens is required to allow a person rote recall of arbitrary information, as tested in classical studies. However, Sanford and Garrod's theory supports the idea that explicit focus memory also supports other tokens at partial salience. There are several studies that say that people are much better at deciding whether they know something when prompted with the information itself, than they are at recalling things by rote (Shneiderman 1992). This suggests that, although a person's explicit focus can only sustain seven plus or minus two tokens at maximum salience at any one time, they can keep many other tokens partially salient at the same time.

Sanford and Garrod's theory can be used to explain how an interruption can make it difficult to resume preinterruption tasks. If explicit focus memory can only support seven plus or minus two maximally salient tokens, then an interruption will displace some or all of the original seven plus or minus two tokens into partial salience. These original seven plus or minus two tokens may still be relatively easily available in explicit focus memory, but since they have been reduced in salience, it will take effort to reactivate them when the interruption has passed. Since the preinterruption task(s) tokens still have partial salience in explicit focus memory, some kind of external reminders could help a person reactivate these tokens to maximum salience when they resume their preinterruption task(s).[6]

Long-term episodic memory is the total collection of currently inactive memories that a person has built or accessed during their current working session. Long-term semantic memory is all the rest of a person's memories that have not been accessed or referenced in the current working session. Both kinds of these latent memories require considerable effort and/or strong-cueing to activate and use.

THEORETICAL CONSTRUCT(S) about people (Sanford and Garrod 1981):

P28. People have four different kinds of memory: two kinds of active memory — (1) explicit focus and (2) implicit focus; and two kinds of latent memory — (3) long-term episodic, and (4) long-term semantic. These differences in kinds of memory represent cognitive structural differences that address a trade-off between memory accessibility and memory extent. These four kinds of memory can be ordered by: accessibility 1, 2, 3, 4; and extent 4, 3, 2, 1. Explicit focus memory is readily accessible but can contain very little information, and long-term semantic memory is difficult to access but can contain huge amounts of complex information.[7]

P29. People's explicit focus memory can contain several tokens at once. People dynamically maintain a level of salience associated with each token, which determines the token's accessibility.

---

6. This idea of external reminders is supported by the research that says that people can recognize information more easily than they can recall it from rote. For example, people often find it useful to construct external physical reminders of things that they need to remember, e.g., tie a string around one's finger. Airline crews sometimes use a version of this idea to remind themselves to turn off the air conditioning units before lowering the flaps — they place an empty coffee cup upside down over the flap handle (Norman 1992, p. 167).

7. These four categories of human memory show a similar trade-off between accessibility and extent as four categories of computer memory, For illustration, we can pair human memories with computer memories: explicit focus as a bank of CPU registers; implicit focus as RAM; long-term episodic as hard disk; and long-term semantic as DAT (digital audio tape).

P30. People can perform unaided recall of information in their explicit focus memory only if the relevant token has maximal salience.

P31. People are only able to maintain maximal salience on seven plus or minus two tokens in their explicit focus memory at any one time.

P32. People can perform recognition, or aided recall, of information in their explicit focus memory if the associated token has less than maximal salience.

P33. People do not have direct conscious control over the salience they ascribe to tokens in their explicit focus memory. Instead, the level of salience of tokens is a side effect of cognitive processing.

Adams and Pew use Sanford and Garrod's theory to explain how the structure and function of human memory affects the way airplane pilots can direct attention during multitasking. Adams and Pew hypothesize that the salience of memory affects the constructivist perception process. They say that because memories are used as filters to perceive and interpret new information, the accessibility (salience) of those memories will affect the perception and interpretation process. We use the word accessibility to refer to the relative effort and reliability of activating a particular memory.

In a difficult task, like flying an airplane (which requires situational awareness and multitasking), the user must jostle their memory resources back and forth between many demands. The person must constantly change the salience of tokens in their explicit focus memory. Since the person cannot support maximal salience on all relevant tokens in their explicit focus memory, giving immediate attention to a task(s) will cause its tokens to dominate in salience over the tokens of other tasks. This variance in salience between the groups of tokens associated with different tasks affects the perception process related to those tasks.

People will perceive and interpret new information relevant to tasks that have high salience tokens with ease and accuracy. However, people will perceive and interpret new information relevant to tasks with lower salience tokens with difficulty and inaccuracy. New information relevant only to inactive memories can only be perceived and interpreted with great effort. In a situation where time is short, like landing an airplane, people will totally ignore or inaccurately perceive new information irrelevant to the immediate task. (See the story about the uncompleted checklist (p. E-1).

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990):

P34. Those memories people allow to be most salient in their explicit focus memory affect the ease and accuracy of their perception and interpretation of new information. A person will more easily and accurately perceive and interpret new information that is relevant to whatever they are currently acting on than new information that is relevant to other pending tasks.

Adams and Pew reaffirm theoretical constructs P7 (pg 28), P8 (pg 29) and T1 (pg 29). However, they color these theoretical constructs differently than Tsukada et al. They adopt the perspective of focused attention upon external tasks instead of the perspective of internal vs. external effort as Tsukada et al. P7 (pg 28) (people are limited to giving thoughtful, conscious attention to only one thing at a time), P8 (pg 29) (people accomplish complex multitasks by shifting attention from one task to another), and T1 (pg 29) (tasks can be accomplished by the accumulation of many independent noncontiguous efforts.

Adams and Pew elaborate on theoretical construct P9 (pg 29) (people internally attend to many things at the same time). They say what kinds of internal action a person can do in parallel. While a person performs one and only one external activity at a time, they simultaneously manage a queue that reflects the prioritization of other pending tasks. Adams and Pew use a metaphor of a queue to model this group of metainformation about pending tasks to imply that the task at the top of the queue will be executed next. A person orders their mental queue of pending tasks by the tasks' relative urgency in time and relevancy to the current context. A person must dynamically reorder this queue because the passage of time affects tasks' time requirements, and changes in situation affect tasks' relevancy and significance.

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990):

P35. People can simultaneously perform external actions on one (and only one) task at a time while constantly and simultaneously maintaining subtle dynamic metainformation about other pending tasks.

THEORETICAL CONSTRUCT(S) about tasks (Adams and Pew 1990)**:**

T4. Tasks have temporal requirements on their execution. This theoretical construct changes dynamically as time passes.

T5. The successful completion of a particular task has some level of importance or significance within a person's overall goals. This theoretical construct changes dynamically relative to changes in a person's environment and the execution of other tasks.

Maintenance of this internal metainformation is susceptible to human memory limitations and, therefore, requires effort and is prone to error. Adams and Pew say that a person's ability to successfully balance situational awareness and multitasking, e.g., flying an airplane, is dependent on their success at maintaining this metainformation. For people to coordinate situational awareness and multitasking, they must make good decisions about two things: (1) when to switch external effort between tasks and (2) what task to switch to next. These decisions rely completely upon a person's ability to maintain subtle metainformation about impending tasks.

A person's ability to maintain accurate metainformation about pending tasks depends upon their continual efforts to correctly perceive and interpret new incoming information. However, Adams and Pew remind us that processing new information takes much effort and can temporarily monopolize scarce memory resources. Although a person can maintain metainformation at the same time they exert focused attention upon some external task, they must stop their external activities to process new information. People cannot simultaneously perform external actions and process new incoming information. Adams and Pew (1990, p. 523) say that for people to process new incoming data, they must immediately interrupt whatever they are doing long enough to perceive the new information; then they may try to resume their interrupted activity.

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990)**:**

P36. The accuracy of metainformation that a person has maintained is positively related to their successful completion of multitasks. People use this metainformation when switching their external efforts between tasks. If a person has maintained accurate metainformation, then they will make good decisions about (1) when to switch focused attention to another task and (2) which pending task to act upon next.

P37. People must temporarily interrupt their focused attention on a task in order to switch their attention and refocus on the task of processing new incoming information (see P25 (pg 45)).

Adams and Pew report that mental shifts between topics or semantic domains have measurable costs to performance. Each time a person shifts their focus of attention from one task to another, they must expend time and effort and expose themselves to potential informational errors and biases (Anderson and Pitchert 1978; Bower 1982; Sanford and Garrod 1981; Schank et al. 1982).

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990)**:**

P38. Shifting focused attention from one thing to another has measurable costs in effort, time, and frequency of error.

Adams and Pew say that there is not yet a formal predictive theory that can accurately predict people's process of switching focused attention. However, even though there is not yet a refined theory, Adams and Pew say that there exists some useful information about how people switch focused attention. Further, they provide some terminology for describing the behavior of attention allocation. They say that it is useful to model people's process of focused attention switching with a probabilistic approach, instead of with a deterministic approach.

Adams and Pew (1990, p. 523) say that people have a variable degree of ease of switching their focused attention. Sometimes people will easily switch their attention between tasks, and at other times, they will have great difficulty switching between tasks. Adams and Pew propose that people are more likely to switch their attention when and to what is most easy.

THEORETICAL CONSTRUCT(S) about people (Adams and Pew 1990)**:**

P39. People have a variable degree of ease of switching their focused attention relative to time and relative to their multitask requirements (individual differences).

2.8 Management of Semiautonomous Agents

Domain Specific Definition of Interruption: a costly side effect of delegating tasks to intelligent agents.

Delegation is a method by which one individual commissions another individual to act on their behalf. This method has been a standard operating procedure for all hierarchical human organizations; and more recently, as a model for client-server computer systems. Delegation is popular because it has been shown to be useful for accomplishing certain kinds of complex tasks. Some authors of computer science employ the idea of delegation to address the problem of overloading a human user.

It is common for human users to become overloaded while trying to perform some kinds of multitasks on computer systems. Job requirements can exceed a person's cognitive resources. Authors have proposed constructing intelligent software agents that can accept requests for delegation. A user can commission these intelligent agents to perform tasks on their behalf. Authors hypothesize that this ability to delegate responsibilities to intelligent agents will allow human users to avoid cognitive overload and successfully perform their multitasks.

A discussion of delegation is relevant here because one of the costs of delegation is increased interruption. A person does not free themselves from responsibility when they delegate a task to an intelligent agent. Instead, they only trade one kind of responsibility for another. The person gives up the responsibility to do the task personally and accepts a new responsibility of supervising the performance of the task by an intelligent agent. These supervisory duties can be nontrivial.

Intelligent agents are usually constructed so that they are required to make reports and requests of their users. Since an agent is somewhat autonomous, its user is not required to focus attention on the agent while it is working. The human, instead, is allowed to concentrate on other tasks while the agent is working. This means that when the agent reports its progress or requests information from its user, it must first interrupt or distract its user from what he or she is currently doing. Thus, delegating a task does not totally free the user from cognitive demands related to that task (Kirlik 1993).

Kirlik (1993) authored a research paper in which he observed that the costs of delegating a task to a task-offload aid (an intelligent software agent) can sometimes outweigh the benefits. Kirlik reaffirms that people have a limited capacity to perform multitasks. And, because there are cognitive costs of delegation, it is possible for peoples' performance on a multitask to actually decrease if they begin delegating tasks to intelligent agents. It can sometimes take more effort to supervise an agent than to do the task without intelligent aid.

THEORETICAL CONSTRUCT(S) about tasks (Kirlik 1993)**:**

T6. Tasks can be delegated from one individual to another.

T7. The delegation of a task begets a new task of supervision. When a user delegates a task, they give up the responsibility to perform the task themselves, but they gain a new task of supervision.

THEORETICAL CONSTRUCT(S) about people (Kirlik 1993)**:**

P40. People have a limited capacity to perform multitasks. (This theoretical construct is somewhat superfluous, but I include it anyway, for completeness.)

P41. People have a limited capacity to perform supervisory tasks.

Kirlik says it is possible for people to reduce their workload and improve their performance on multitasks by delegating tasks to intelligent agents. However, the utility of delegation depends on the management strategy chosen by the person — people's managerial decisions affect the utility of their delegation decisions. "Of great importance is the strategy the operator develops for managing interaction with an aiding device. Human supervisory controllers have the capability and often the freedom to strategically manage their interaction with automation in an effort to keep both workload and system performance at acceptable levels" (Kirlik 1993, p. 222).

THEORETICAL Construct(S) about people (Kirlik 1993)**:**

P42. People can successfully use delegation to perform multitasks. People can divide their responsibilities for the tasks that comprise their multitask into two categories: (1) tasks they perform personally and (2) tasks they delegate and supervise the performance by others. This means that people can simultaneously coordinate performing tasks personally and supervising performance of tasks by others.

P43. People know several varied managerial techniques for supervising the performance of delegated tasks.

P44. People can make strategic managerial decisions dynamically when supervising the performance of delegated tasks.

The managerial strategy a person chooses is critical to the success of their delegation attempts. Sheridan (1988) wrote a useful paper titled, "Task Allocation and Supervisory Control," in which he presents a broad overview of factors that affect people's selection of appropriate managerial strategies for supervising intelligent computer aids performing delegated tasks. Supervisory control systems incorporate intelligent computer aids to enable their users to accomplish complicated physical control tasks. The intelligent computer aids in a supervisory control system are intelligent software agents. However, these two fields (supervisory control systems and intelligent software agents) have different terminology because they traditionally address different domains. Research reported about intelligent software agents usually addresses information processing tasks, whereas research reported about supervisory control systems usually addresses control of physical processes.

Sheridan (1988, p. 159) explains the function of supervisory control systems. "The human supervisor works through the computer to effect what needs to be done in the physical world. The computer is then seen as a mediator — communicating upward to the supervisor, communicating downward to the physical process, whatever it may be." Typical domains of application include: control of vehicles (aircraft, spacecraft, ships), control of chemical and electrical power generating plants, and control of industrial and other robotic devices.

THEORETICAL CONSTRUCT(S) about people (Sheridan 1988)**:**

P45. People can work through mediators. People can both act on tasks and perceive task performance at an abstract level through a mediator.

When people delegate tasks to intelligent aids, they also accept the costs of personally managing that delegation. One of these costs is potential interruption by the subordinate intelligent aid. Sheridan says there are three theoretical constructs that describe people's behavior while managing delegated tasks: (1) the kind of action the person is trying to accomplish with the intelligent aid; (2) the level of interaction abstraction provided by the intelligent aid; and (3) the degree of autonomy provided by the intelligent aid. A discussion of these theoretical constructs is important to our discussion of user-interruption, because it categorizes human behavior in supervisory control tasks (managing intelligent agents as they perform delegated tasks). This categorization of behavior gives us a theoretical tool for investigating the delegation process and the effects of interrupting human supervisors.

Sheridan says that there are twelve different categories of human supervisor functioning. These represent twelve different actions that people try to accomplish with intelligent computer aids. Sheridan further breaks down these twelve supervisory functions into five general classes. (Sheridan 1988, pp. 161-167) The following is partially quoted from Sheridan (1988, Table 1 p163).

(1) **Plan** [discover the function and effective use of an intelligent aid],

(1a) understand controlled process,

(1b) satisfice objectives,

(1c) set general strategy, and

(1d) decide and test control actions.

(2) **Teach** [provide an intelligent aid with the information it needs to perform a delegated task],

(2a) decide, test, and communicate commands.

(3) **Monitor Auto** [monitor automatic execution of the programmed actions],

(3a) acquire, calibrate, and combine measures of process state,

(3b) estimate process state from current measure and past control actions, and

(3c) evaluate process state; detect and diagnose failure or halt.

(4) **Intervene** [respond to a failure or halt condition],

(4a) if failure: execute planned abort, and

(4b) if normal end of task: complete.

(5) **Learn** [learn from current experience to use the intelligent aid better in the future],

(5a) record immediate events, and

(5b) analyze cumulative experience.

People interact with intelligent computer aids at different levels of abstraction. Sheridan says that there are different ways of controlling a process. The nonabstract way is to skip the intelligent aid and manually control the process oneself. The abstract ways are to delegate the control task to an intelligent computer aid and then interactively supervise that aid. People communicate with intelligent aids in one of three levels of abstraction: (1) knowledge-based (high-abstract); (2) rule-based (medium-abstract); or (3) skill-based (low-abstract). Sheridan describes these differences in interaction abstraction with the model shown in Figure 4.

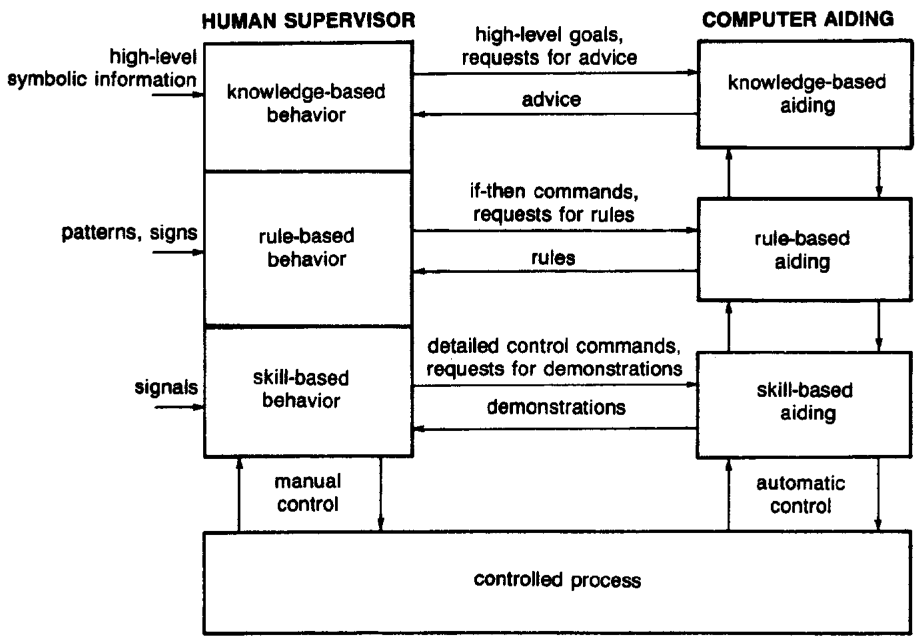

Figure 4.Sheridan's model of the different levels of interaction abstraction between a human supervisor and their intelligent computer aid. Reprinted from *Handbook of Human-Computer Interaction* by Thomas B. Sheridan, 1988, p. 168, from Elsevier Science — NL, Sara Burgerharstraat 25, 1055 KV Amsterdam, The Netherlands.

Sheridan's third theoretical construct of human supervisory behavior is the degree of autonomy provided by the intelligent aid. People choose different managerial strategies depending on the degree of autonomy exhibited by the intelligent computer aid. Sheridan classifies the autonomy of intelligent computer aids along two dimensions: capacity to act autonomously and accountability for actions. An intelligent computer aid has some degree of capability to automatically perform the different stages of a task (determine potential alternative approaches, select one approach to execute, implement the chosen approach, and inform the human of the results). The capacity of an aid determines: which parts of a task the human must do unaided; which parts of the task the aid does autonomously; and which parts the human and aid must perform cooperatively. An intelligent computer aid also has a degree of accountability toward its user. This defines who is given ultimate responsibility for accomplishing tasks — human or machine.

Sheridan (1988, p. 171) gives ten examples of supervisory interaction to illustrate his classification of automation. Each example shows an increasing degree of automation.

1. Human does the whole job up to the point of turning it over to the computer to implement.
2. Computer helps by determining the options.
3. Computer helps determine options and suggests one, which human need not follow.
4. Computer selects action and human may or may not do it.
5. Computer selects actions and implements it if human approves.
6. Computer selects action, informs human in plenty of time to stop it.
7. Computer does whole job and necessarily tells human what it did.
8. Computer does whole job and [if asked] tells human what it did.
9. Computer does whole job and [tells human only if it decides to].
10. Computer does whole job if it decides it should be done and, if so, tells human if it decides he [or she] should be told.

THEORETICAL CONSTRUCT(S) about people (Sheridan 1988):

P46. There are three factors that affect people's choice of managerial strategy when employing intelligent computer aids: (1) their purpose for employing the agent; (2) the level of abstraction with which they will interact with the agent; and (3) the degree of autonomy provided by the agent.

P47. People use intelligent computer aids for different purposes. These different purposes can be categorized into five general categories: (1) discover the function and effective use of an intelligent aid; (2) provide an intelligent aid with the information it needs to perform a delegated task; (3) monitor automatic execution of the programmed actions; (4) respond to a failure or halt condition; and (5) learn from current experience to use the intelligent aid better in the future.

P48. People interact at three different levels of abstraction with subordinate intelligent computer aids (knowledge-based level; rule-based level; or skill-based level).

THEORETICAL CONSTRUCT(S) about context (Sheridan 1988):

C1. Intelligent computer aids provide different degrees of autonomy along two dimensions: capacity to act autonomously and accountability for actions.

2.9 HUMAN-HUMAN DISCOURSE

Domain Specific Definition of Interruption: an example of human-human discourse that can be represented and analyzed with the theory of discourse analysis.

Authors have published useful theory of human-human interaction under the title of discourse analysis. This theory addresses the problem of modeling and analyzing occurrences of human-human interaction. The proposed theory of interaction can be extremely useful for addressing this difficult domain. Occurrences of human-human interaction are complex processes and, therefore, can be difficult to study. Interaction events are composed of all behaviors and their interrelationships, which happen across several different dimensions: abstraction, media, time, scale, and individual participants.

The theory of discourse does not address specific kinds of human-human interaction, such as events of people interrupting other people. Instead, this theory addresses the metalevel problem of how to represent and analyze the total complexity that is in every kind of human-human interaction. This theory is useful, because the patterns of human-human interaction are powerful models for studying human-computer interaction. Therefore, we can apply this theory of discourse to conduct detailed modeling and analysis of occurrences of people interacting with computers, and this includes the specific kind of human-computer interaction which is the topic of this dissertation — computers interrupting people.

Theory of discourse proposes to model human-human communication in context for the purpose of studying the communication of meaning between people. “Essentially, then, discourse analysis is an analysis of meaning but meaning seen not in the traditional philosophical or semantic sense of isolated concepts but rather the discourse analyst studies meaning as a construct of interaction, and he [or she] studies the various ways in which we create, relate, organize, and realize meaning in behavior” (Riley 1976, p. 2).

Riley (1976) in his paper titled “Discursive and Communicative Functions of Non-Verbal Communication,” presents a theory of human discourse that emphasizes the previously unappreciated importance of nonverbal communication acts. Riley says that spontaneous, authentic, face-to-face informal interaction between people can be modeled as a complex interdependent fabric of diverse communication acts. These communication acts are realized by a large variety of behaviors, e.g., language, tone of voice, gesture, posture, body movements, spatial orientation, physical proximity, eye contact, and facial expression (Riley 1976, p. 2). Riley says that it is important to model all these kinds of behaviors because people use them all together to express meaning, i.e., language by itself is not enough. This is the total context of communication.

Riley provides a way for all human dialogue behaviors to be represented within the same descriptive theory. This unifying theory embodies two useful concepts: (1) verbal behaviors do not deserve a special and separate status in discourse analysis and (2) that all discourse behaviors can and must be represented in a consistent way. This consistency and uniformity of representation allows us the freedom to express the complex interdependencies expressed in human interaction without being tied to the artificial and misleading "verbal or nonverbal" classification.

Riley uses a theoretical concept that can easily and uniformly model all kinds of dialogue behaviors — the communication act. A communication act is a basic unit of meaning, with which people attempt to express to other people. People **realize** their communication acts by a wide range of **behaviors,** which are **conveyed** along different **channels** of communication. Therefore, people can express a single meaning in any of several different ways — a particular communication act can be realized in different ways by different behaviors and conveyed along different channels. This distinction between meaning and its expression and its conveyance is useful for untangling the complexities of human interaction.

For example, speaking a word is a particular realization of a communicative act or a "speech act." Riley says, "a speech act is just one of the possible realizations of a communicative act: a shake of my head can communicate disagreement just as efficiently as the word 'No.' Indeed, so can the right intonation or key choice, so can facial expression and certain gestures. And of course this is an extremely crude example: the meaning of an act of communication is much more often the product or sum total of a head movement plus words plus intonation and key, plus facial expression, plus skeletal disposition, plus all the relevant situational features; meaning is the relationships, if you like, between all these features" (Riley 1976, p. 3).

THEORETICAL CONSTRUCT(S) about people (Riley 1976)**:**

P49. All aspects of people's communicative behavior are important and relevant to model and analyze the discourse and interaction between people. Speech alone is grossly insufficient.

P50. People interact with other people by the coordinated expression of interrelated sets of basic units of meaning — communication acts.

P51. People realize communication acts by various behaviors which are conveyed along different channels of communication. (Note as especially important this distinction between the meaning, behavior, and channels of conveyance.)

Riley expresses dissent from the popular assumption that verbal-vocal behaviors are relatively more important to communication modeling and analysis than other behaviors. Riley reviews the traditional categorization of human interactive behavior along the verbal and vocal dimensions. This categorization of human interactive behavior reveals three domains (or components) of linguistic research: verbal, paralinguistic, and kinesic (see Table 3). Riley says that the verbal component should not be esteemed more important than the other categories. He says that in a unified theory of discourse, these three categories must be of equal importance.

Table 3 (Riley 1976, p. 3) shows the traditional categorization of research into human interactive behaviors. People's behaviors are categorized along the verbal and vocal dimensions. I provide an example of each in parentheses.

Table 3 — Traditional Components of Discourse

	verbal	**nonverbal**
vocal	verbal component (speak the word “hello”)	paralinguistic component (speak the nonword “um”)
nonvocal	(“sign” a word of sign language)*	kinesic component (make the nonword gesture of pointing to an object)

* Riley leaves the nonvocal-verbal category vacant. It is irrelevant to this discussion because Riley states that this is not a useful categorization; however, for completeness there is at least one important example of verbal nonvocal behavior expressed by humans — sign language. Deaf people, fluent in sign language, are perfectly capable of communicating verbal information nonvocally.

Riley supports his assertion that kinesic behaviors are a critical part of discourse analysis by presenting a taxonomy of kinesic behaviors that classifies different expressions of communication acts. He shows that without including these previously devalued dialogue behaviors in our model, we cannot successfully model and analyze human interaction. Discourse analysis must include a representation of these communication acts. (I include a summary of Riley’s taxonomy of kinesic behaviors here because I believe that kinesic behaviors are an important class of behaviors for modeling and analyzing the human dialogue of interruption.

Riley’s taxonomy has three groups of kinesic behaviors: emblems, indices, and gestures. **Emblems** are behaviors that people consciously and intentionally perform to convey conventional, specific meanings that are easily expressed in words. For example, “thumbs up” or “the finger.” **Indices** are behaviors that people make to convey indexical information (information about the person making the behavior). There are three kinds of indices: psychological, social, and biological. Psychological indices are realizations of communicative acts about the state of the person’s psychological state (e.g., smiling, weeping, sweating, blushing). Social indices are realizations of communicative acts about a person’s social state (e.g., class, occupation). Biological indices are realizations of communicative acts about a person’s physical self (e.g., age, sex, health, fatigue).

Gestures includes all kinesic behaviors that are not easily categorized as emblems or indices. The following list contains different kinds of gestures.[8]

1. Kinematopoeia. These behaviors are realizations of communication acts to illustrate something. For example, holding one's hands far apart while speaking "I caught a fish this big."[9]

2. Deictics. These behaviors are realizations of communication acts to refer to something. For example, tipping one's head in the direction of an object while speaking "hey, look at that!"

3. Gestures having illocutionary force. These behaviors are physical realizations of communicative acts that convey content meaning in a conversation. These communication acts are used to convey many kinds of meaning:
   - agreement or disagreement, e.g., nodding or shaking one's head;
   - greeting, e.g., an eyebrow flash;
   - declining, e.g., placing one's hand over a cup when offered more coffee;
   - requesting, e.g., asking for the time by tapping one's wrist where a watch should be;
   - commanding, e.g., a policeman signaling traffic;
   - doubting, e.g., an appropriate facial expression; and
   - reporting ignorance, e.g., shrugging shoulders (Riley 1976, p. 9).

   [see footnote[10]]

4. Turn-taking signals. People use these gestures as realizations of communicative acts to regulate the process of turn giving and taking between the people in a discourse. For example, eye gaze is often used by people to indicate when they are prepared to relinquish a turn.

5. Attention signals. These behaviors are realizations of communication acts to regulate the conversants' attention. For example, eye gaze directed at someone outside of the current dialogue group.

6. Address signals. These behaviors are realizations of communication acts by a person to select and indicate their listeners. For example, alignment of head direction, eye gaze, and posture toward another person selects them as a listener.

Of all kinesic behaviors, there are three kinds of gestures that are particularly relevant to the phenomenon of interruption: turn-taking signals, attention signals, and address signals. Each of these physical behaviors is a realization of a person's communication act to regulate or manage the process of interaction at a meta level.

---

8. I tend to repeat myself about Riley's assertion that behaviors are realizations of communicative acts intended to convey some particular meaning. My repetition of this point is not good English, but I think it is necessary to remain clear about Riley's motivations.

9. My examples for kinesic behaviors are the common US English meanings of these gestures.

10. Gestures 4, 5, and 6 are kinesic behaviors that people make as realizations of communicative acts to regulate the interaction itself. These comprise metacommunication.

THEORETICAL CONSTRUCT(S) about people (Riley 1976):

P52. People have three different methods for expressing their communication acts: verbal, paralinguistic, and kinesic. (Note: I would add a fourth method here — a nonvocal verbal component. This component describes peoples' discourses by expressions of standard sign languages.)

P53. People have three different methods for kinesic expressions of communication acts: emblems, indices, and gestures.

P54. People make some communication acts at a metalevel of meaning in order to guide the process of interaction.

P55. People make gestures of turn-taking signals as realizations of communication acts to regulate turn-taking in their interaction with other people. Examples of these behaviors can be expressed with: eye gaze, change in speech timing, synchronized finishes of verbal and kinesic behaviors, creaky voice, low key voice, or cessation of kinesic behaviors.

P56. People make gestures of attention signals as realizations of communication acts to regulate the attention of all people involved in the interaction.

P57. People make gestures of address signals as realizations of communication acts to regulate who their listeners are and are not.

Riley (1976) asserts that kinesic dialogue behaviors comprise a critical part of the dialogue process and, therefore, must be included in dialogue analysis. The traditional distinction between verbal and nonverbal behavior (Table 3, p. 64), carries unnecessary confusion into discourse analysis. This distinction confounds the separation of meaning, expression, and conveyance. This confusion allows our irrational bias toward verbal behaviors to affect our ability to successfully model and analyze human interaction. It is much more useful to separate meaning from its expression and its conveyance. Therefore, the vocal/nonvocal, and verbal/nonverbal distinctions are inappropriate. Riley (1976, p. 4) says that the use of this inappropriate categorization has traditionally resulted in confusion between two different measures of people's discourse behaviors: (1) the degree of linguisticness and (2) the importance in communicative function.

Riley proposes a "unified or integrated model for the description of discourse." This unifying theory uses the concept of "communication act" as its basic unit of construction. This basic and uniform modeling tool affords representation and analysis of all discourse behaviors as potentially important to discourse at some level. Riley also proposes that all communication acts should be analyzed along three distinct levels: the realization level; the communicative level; and the discursive level. He says that these three levels capture the most useful categorization of discourse analysis. Riley (1976, p. 13) asserts these three levels of analysis are equally important for successfully analyzing dialogue.

The realization level of analysis addresses the mapping between all observable discourse behaviors (verbal, paralinguistic, kinesic) and the communication acts or meaning which those behaviors realize. This level of analysis is critical because it provides us with information about how people are realizing their communication acts. We discover how different behaviors are being used in concert to express meaning.

The communicative level of analysis addresses the illocutionary forces of communication acts (separate from their particular realization). This level of analysis provides us with information about the communicative intentions of the conversants, e.g., inviting, persuading, agreeing. Note that there is no one-to-one relationship between the illocutionary forces and the individual behaviors of realization.

The discursive level of analysis address people's attempts to regulate or manage the process of interaction at a metalevel, e.g., interactional tactics, turns, attention direction, address, relative distribution of utterances. Note also that there is no one-to-one relationship between discourse regulation and the individual behaviors of realization. We observe people's myriad discourse behaviors all woven together into a complex fabric. However, Riley's three levels of analysis allow us to extract specific information about two parts of people's communication acts: (1) their meaning — both illocution (communicative level) and dialogue regulation (discursive level) and (2) their chosen methods of expression (realization level). (Note: Riley relies on his "realization level of analysis" for extracting information about the "channels of conveyance" part of communication acts.)

THEORETICAL CONSTRUCT(S) about people (Riley 1976)**:**

P58. People intermix their expression of discourse behaviors so that their physical realization of meaning, their attempts at illocutionary force, and their attempts to regulate the process of interaction are all interwoven into a complex fabric of communication.

Riley's theory of discourse analysis gives us a powerful tool for analyzing human discourse relevant to this dissertation — the interruption of people. When we apply this theory, we see that for one communication act of interruption, it is possible to have more than one physical realization as discourse behavior. Indeed, it seems reasonable that there are a multitude of different behaviors and combinations of behaviors that can realize an interruption. For example, a person might interrupt another with the coordinated and synchronized expression of all the following behaviors together: speaking "excuse me please;" turning the head toward the other person; moving the eyes to make eye contact with the other person; reaching out an arm and hand to make a gesture similar to blocking the progress of something; moving closer to the other person; smiling; and then synchronized cessation of all movements and behaviors to indicate a change in turn to allow the other person to acknowledge this interruption request.

Riley's theory also allows us to differentiate between people's illocutionary forces and their meta dialogue discursive attempts. This is useful because people's behaviors intended to interrupt and those intended to convey meaning are intermixed. Riley's unified theory gives us a tool for separating those behaviors.

Riley's theory is useful in several ways, however it does not solve everything. There are some notable weaknesses or deficiencies to Riley's unifying theory. His theory does not provide structures for modeling interrelationships between individual communication acts or interdependencies between their behavioral realizations. Also Riley's theory does not provide temporal structures for modeling the coordination between communication acts or their behavioral realizations over time. This theory also does not provide useful methods to aid the difficult analysis task of discovering the abstract communication acts of human conversants' observable behaviors.

2.10 HUMAN-HUMAN DIALOGUE

Domain Specific Definition of Interruption: a common and normal part of human-human dialogue behavior.

Taylor and Hunt (1989) report the results of a workshop titled "Flexibility versus Formality." This workshop produced a discussion of requirements for a formalism for the design of multimedia human-computer dialogues. The workshop participants used human-human dialogue as a metaphor to identify requirements of human-computer dialogue. They began with the postulate that people's behavior in human-computer dialogue is similar to their behavior in human-human dialogue.

Taylor and Hunt say that some common human interaction behaviors are not well formalized as dialogue "turns." People frequently interrupt communication dialogue: (1) they interrupt themselves by breaking off their turn before completing a sentence, and (2) they interrupt other people by initiating a dialogue turn during another person's turn. Taylor and Hunt illustrate how frequently people interrupt themselves during human-human dialogue with the following real-life dialogue (Taylor 1989, p. 444) (the participants are discussing a forthcoming nasal operation):

> S: Do you know what they're doing?
>
> K: I think they take these poles and they just sort of (giggle) ...
>
> S: (giggle) violently knock your ...
>
> K: That's right. I think they're ... Basically it's like breaking inside, I think.

Taylor and Hunt say this example dialogue shows that interruption is a very common and normal part of dialogue. “In this little interchange, only three of the six potential sentences are completed, and to an uninvolved observer, the three broken ones do not appear to convey the necessary information. Obviously, however, from the viewpoint of the participants, the information is adequate. They seem to be quite happy with the interchange, which is experienced as well-formed. Interruption should be seen as an integral part of the dialogue process, not as some kind of irregularity that can be swept aside when analyzing ‘real’ conversation” (Taylor and Hunt 1989, p. 444).

Taylor and Hunt say that there are two other common dialogue behaviors that are difficult to model with the “turn” concept: sidechannel contributions and abort or emergency stops. Taylor and Hunt propose that a formalism for human-computer dialogue must have modeling structures to represent interruption, sidechannel contributions, and aborts as normal (first class) parts of dialogue. People express sidechannel contributions as feedback to the person they are attending. These expressions are provided to inform the communicator of the success and failure of their attempts to communicate. For example, while someone is speaking to me, I simultaneously communicate feedback to them — I say things like “uh-huh” or “yea” or make nonverbal gestures to mean that I hear and understand.

Sometimes a person aborts or abruptly quits a dialogue altogether. Taylor and Hunt say that this normal human behavior can cause serious problems in a human-computer dialogue if the computer has not been designed to support such behavior. A formalism for human-computer dialogue should include methods for explicitly dealing with the event of a person aborting the dialogue.

THEORETICAL CONSTRUCT(S) about people (Taylor and Hunt 1989)**:**

P59. People exhibit common and normal communication dialogue behaviors that are not well formalized as dialogue “turns.” These behaviors emphasize the dynamic nature of dialogue. Three important examples are: (1) interruption of self and others; (2) sidechannel contributions; and (3) abort.

2.11 Psychology of Human Attention

Domain Specific Definition of Interruption: the method by which a person shifts their focus of consciousness from one processing stream to another.

Davies, Findlay, and Lambert have written a paper titled, "The Perception And Tracking Of State Changes In Complex Systems" (Davies et al. 1989). They apply psychological theories of human attention in their research about the display design of interactive computer systems. Davies et al. address the task environment in which a user must maintain situational awareness of a complex, multiactivity process. Since they specifically address the user's need to switch attention between different monitoring tasks, some of the theoretical concepts they advocate are useful in our attempt to define the phenomenon of interruption.

People can execute several simultaneous cognitive processing streams. This allows people to perform cognitive processing on several topics at once. However, there is an important structural restriction. People's cognition supports only one principle processing stream. The remaining processing streams must be executed as subsidiary or peripheral streams. This restriction means that although people can execute several cognitive processes at once, they can perform only one activity (thought or action) at a time with conscious control and awareness. (Miyata and Norman (1986) present a similar excellent survey of the psychology of human attention and how it is useful for studying user multitasking in HCI.)

This theory has two interesting implications: (1) people are limited to one conscious activity at a time, and (2) people perform a large amount of cognitive processing outside of their conscious control and awareness. Davies et al. say that evidence of this theory is observable because information produced by people's subsidiary processing can dynamically influence their primary processing. "One example of this is evidenced by our ability to elicit changes in the orientation of focal attention in response to changes in the peripheral visual field. Such parallel processing is used dynamically. Studies of the reading process show that detailed textual information is received from quite a small region to where gaze is directed. However, less detailed information (word boundaries, initial letters or words, and so on) is also being simultaneously assimilated from more distant regions to facilitate eye guidance and to provide some preliminaries to more detailed analysis" (Davies et al. 1989, p. 511; Rayner 1983; Rayner 1992).

Davies et al. say that people have one focus of consciousness (a structure of their cognition). Whatever processing stream a person executes in their focus of consciousness becomes that person's principle processing stream. Therefore, a process is either principle or subsidiary depending on whether a person executes it inside or outside of their focus of consciousness. Human attention behavior can be usefully modeled as the result of the meta-activity of shifting processing streams into and out of a person's focus of consciousness. (Note that one important implication of this model is that processing streams continue to execute whether they are in the focus of consciousness or not. When a person switches their attention from one activity to another, the displaced processing stream continues to execute but out of consciousness.)

Davies et al. propose that interruption is the exclusive method by which a person switches processing streams into and out of their focus of consciousness. Further they say that "conscious human activity can be viewed as consisting of bouts of processing which are terminated at an 'interrupt'" (Davies et al. 1989, p. 512). This approach makes interruption a basic or pivotal concept for modeling the behavior of human attention.

Interruptions can be either internal or external. An internal interruption is a request by a subsidiary processing stream to be switched into the person's focus of consciousness. An external interruption is an event that triggers a subsidiary processing stream to request to be switched into the person's focus of consciousness. An external interruption might come from another person, ex., a telephone call or physical arrival; or an external interruption might come from the person themselves, ex., a physical reminder, like a sticky note, or an alarm clock buzzer.

THEORETICAL CONSTRUCT(S) about people (Davies et al. 1989)**:**

P60. People can execute several simultaneous cognitive processing streams.

P61. People have only one focus of consciousness. This is a unique cognitive structure that adds special support to one and only one processing stream at a time, i.e., conscious awareness and control.

P62. People have conscious awareness and control over only one activity (thought or action) at a time, i.e., whichever processing stream a person currently executes in their focus of consciousness. This is called the person's principle processing stream. The rest of a person's processing streams (those not executing in the focus of consciousness) are subsidiary and execute out of conscious control and awareness.

P63. People perform a large amount of cognitive processing outside of their conscious control and awareness.

P64. The information products of subconscious cognitive processing (subsidiary processing streams) can dynamically affect people's conscious cognitive processing.

P65. Human attention behavior is the result of the person's cognitive meta-activity of shifting processing streams into and out of their focus of consciousness.

P66. Subsidiary processing streams are not suspended but continue to actively process information out of conscious awareness and control.

P67. Interruption is the exclusive method by which a person switches processing streams into and out of their focus of consciousness.

The cocktail party phenomenon (Cherry 1953; Preece et al. 1994, p. 100) is a lucid example of people's cognitive attention behavior. The cocktail party is an environment that is over saturated with external events competing for attention. This is informational chaos. The senses of the people attending the party are overwhelmed with a tumult of incoming signals: many loud voices saying different things; a myriad of other sounds, noises and music; many attractive people wearing interesting clothing and jewelry; manifold physical gestures; multiple simultaneous eye contacts; people arriving, leaving and moving within the crowd; smells; tastes; and touches as people accidentally jostle each other (Preece et al. 1994, ch5.1).

People are fully capable of focusing on one stream of information amid such chaos. At first entering such an environment, people experience the chaos itself but can quickly become involved in one conversation with one group. They continue to experience the chaos, of course but they can extract one thread of human conversation from the chaos and pull it into their focus of conscious attention. They also simultaneously keep the rest of the chaos out of their focus of conscious attention.

However, it is clear from people's behavior that they process a tremendous amount of information subconsciously. For example, while a person consciously attends one conversation, they can notice the utterance of their own name spoken within some other distant conversation. They can then instantly switch their conscious attention from their current conversation to that other conversation where their name was spoken. Thus, while people consciously attend to one thing they are also simultaneously subconsciously attending to many many other things.

We can use the theoretical concepts of human attention to explain people's behavior at cocktail parties. This theory tells us that people can simultaneously process many streams of information but that only one stream can execute in a person's focus of consciousness. A person chooses one stream to execute in their focus of consciousness, and they begin consciously participating in one particular conversation of the cocktail party. They also continue to subconsciously work on several other subsidiary processing streams coming out of the chaos. At some point, one of their subsidiary processing streams recognizes something that it perceives as important to their conscious awareness, and a metacognitive activity occurs to switch that subsidiary processing stream into their focus of consciousness. The processing stream that is displaced from the focus of consciousness continues processing now as a subsidiary processing stream.

THEORETICAL CONSTRUCT(S) about people (Cherry 1953; Preece et al. 1994, p. 100)**:**

P68. People can select and focus on one stream of information amid dense informational chaos.

P69. People can extract several discrete streams of information simultaneously from dense informational chaos.

P70. People can subconsciously determine the relevancy and importance of the information they process in their simultaneous subsidiary processing streams. If needed they can perform the metacognitive activity of switching the important subsidiary processing stream into their focus of consciousness.

Preece et al. (1994, ch. 5.1) summarizes other useful concepts from attention theory — focused attention and divided attention. Focused attention describes human behavior in which a person tries to consciously attend to one information stream to the exclusion of all other competing stimuli. The cocktail phenomenon is an example of a context where people exhibit focused attention behavior. Divided attention describes human behavior in which a person is attempting to consciously attend to two or more things at the same time. Driving a car while participating in a conversation with a passenger is an example of divided attention behavior. The concept of divided attention does not imply that people have more than one focus of consciousness. Instead, it describes a kind of human cognition in which a person attempts to share their focus of consciousness between two or more processing streams by continuously alternating them into and out of their focus of consciousness.

Preece et al. (1994) say that people's metacognitive decision to switch from one processing stream to another can be either voluntary or involuntary. A person makes a voluntary attention switch when they make a conscious decision to switch from their current activity to something else. However, some kinds of events can cause a person to change their attention without conscious decision — involuntary attention switch. For example, the occurrence of a loud noise can cause an involuntary switch of attention.

THEORETICAL CONSTRUCT(S) about people (Preece et al. 1994, ch. 5.1)**:**

P71. People can maintain conscious awareness of two or more things at a time (divided attention) by continuously and alternately switching the relevant processing streams into and out of their focus of consciousness.

P72. A person's metacognitive decision to switch from one processing stream to another can be either voluntary or involuntary. This difference in behavior depends on people's environment.

P73. Information streams have characteristics each with an associated degree of salience.

2.12 A Metaphor of Cognitive Momentum

Domain Specific Definition of Interruption: something that extinguishes a person's cognitive momentum when they are performing concentrated work on a complex task.

Sullivan (1993) proposes an intuitive, nonscientific, definition of the interruption phenomenon. Sullivan used a metaphor of energy to express how interruptions are disruptive to performing work. He says that people build cognitive momentum in the performance of a task. An interruption, like a telephone call, can extinguish that momentum.

People have several different cognitive resources they use in concert to accomplish complex tasks. However, to accomplish demanding tasks, people must first exert concentrated effort to access and align these essential cognitive resources. There is an initial stage where people expend effort but do not accomplish any external portion of the task. People must first internally coordinate their cognitive resources before they can begin doing external work. This coordination of cognitive resources is a kind of cognitive momentum. People exert mental effort to get their cognitive resources into proper alignment once, and then that alignment propels itself forward as the person performs that task.

If person is interrupted after they have organized their cognitive resources, they lose their cognitive alignment. The start-up effort they invested is lost. Once the interruption finishes, they must regain their momentum before they will be able to resume doing work. They must recommence the process of gaining momentum.

THEORETICAL CONSTRUCT(S) about people (Sullivan 1993)**:**

P74. People have various cognitive resources that they must align in task-specific ways in order to use these resources in concert to accomplish complex tasks. In order for a person to accomplish a complex task, their various cognitive resources must be able to coordinate and cooperate with each other.

P75. People must exert concentrated effort over time to align their various cognitive resources. After a person has labored to align their cognitive resources, their alignment carries itself forward (like momentum) as they perform the task. Interruption, breaks a person's alignment; and they must exert much cognitive effort to re-align their resources before they can resume work on the task.

2.13 SOCIAL PSYCHOLOGY OF CONVERSATION

Domain Specific Definition of Interruption: a violation of people's conversational rights.

One way people interact with other people is through dialogue. People use dialogue to communicate their knowledge and wants to each other, however, they also use dialogue to affect their social relationships with each other. People assert their own worth or status when they interact with others. This potential for social influence is one reason why the United States of America has a "freedom of speech" amendment in its Constitution. Governments usually recognize the social power of communication and create, maintain, and execute laws to control or prevent control of the flow of that power.

People have many ways to assert social influence with dialogue — some of these ways address the "message," and some of these ways address the "medium." People can advance their social importance or status by making useful contributions to dialogue. These contributions provide useful knowledge or meaning to some topic of conversation. If a speaker makes a useful contribution, their audience may change its perception of the speaker's social value in a favorable way and begin to esteem the speaker more highly.

People are also able to wield social influence by directly affecting the dialogue process itself— without making any meaningful contribution to dialogue whatever. People will graciously allow themselves to be interrupted by someone whom they hold in high esteem, however this principle can also work in reverse. If a person with lower social status is allowed to interrupt the speaker, then the listeners, and/or the speaker themselves may positively change their belief in the interruptor's social worth. The listeners (and/or the interrupted speaker) may need to internally rationalize their own passivity at allowing the speaker to be interrupted. They may justify their passivity by saying to themselves, "I wouldn't have permitted the speaker to be interrupted unless the interrupting person were of high social status. Therefore I must actually have a higher opinion of the interrupting person than I had realized." The result is that the interruptor gains social prominence.

If we ascribe to the ideal that people's social status should depend directly upon their social usefulness, then we may conclude that it is "unfair" for a person to gain social standing by only affecting the dialogue process and not by making contributions of useful knowledge. Most people do implement this notion of "fairness" in conversation, and this application of "fairness" leads to the concept of conversational rights.

West (1982) says that people have conversational rights and that being interrupted while speaking is a violation of those rights. West addresses the topic of how women increase or decrease their social rank in the work environment by the way they react to violations of their conversational rights. West explains that there are individual differences between males and females in the way people perceive interruption. West makes a case that if a woman wants to gain social status in the workplace, she must address interruption events in ways that her male coworkers will appreciate.

THEORETICAL CONSTRUCT(S) about people (West 1982)**:**

P76. People combine dialogue with attempts to influence social relationships between each other. Dialogue between people serves as more than just a mode of communicating knowledge or needs. Dialogue also serves as a means for people to affect the social relationships between them.

P77. People can use dialogue to influence social relationships in two ways: (1) by making useful contributions of meaning and (2) by controlling the dialogue process itself through metadialogue regulation techniques.

P78. People recognize the concept of conversational rights. Interruption is a violation of conversational rights in which the interrupting person attempts to use a metadialogue regulation technique to affect the conversants' social standings.

P79. People sometimes need to internally justify their response to interruption. For example, suppose that one person interrupts another during a dialogue. If the interrupted speaker and/or the other listeners do not try to prevent this violation of the speakers conversational rights, they may need to later internally justify their passive behavior. This internal justification may have three possible outcomes: (1) they may convince themselves that the interrupting person has the high social status to merit such behavior; or (2) they may convince themselves that the interrupted person has the low social status to merit such behavior; or (3) both 1 and 2.

These theoretical constructs about people have significant social psychological implications for the design of human-computer interfaces. HCI system designers must be careful not to build computer systems that interrupt their users in ways that habitually violate their users' conversational rights. This kind of HCI design error would result in an ineffective and oppressive computer system.

2.14 Interactional Sociolinguistics of Politeness

Domain Specific Definition of Interruption: a face-threatening act.

People maintain social relationships with other people (see P76 on pg 79). An important way they do this is by deliberately controlling the way they express their communication acts. This deliberate control of expression is especially critical for interrupting people. Interruption is an inherently dangerous kind of communication act because of its potent social effect. Therefore, people very carefully package their communication acts of interruption so as not to damage their social relationships. This great care people take in constructing appropriate expressions of interruption is an important source of information for how we should create the HCI of user-interruption.

Brown and Levinson (1987) propose a theoretical tool for studying how people structure their communication acts for social effect. They say that people construct their communication acts in particular ways with which they intend to have specific social effects. Brown and Levinson propose a theory of interactional sociolinguistics that provides useful tools for studying how people construct their expression of communication acts for social effect.

> “We believe that patterns of message construction, or ‘ways of putting things,’ or simply language usage, are part of the very stuff that social relationships are made of (or, as some would prefer, crucial parts of the expressions of social relations). Discovering the principles of language usage may be largely coincident with discovering the principles out of which social relationships, in their interactional aspect, are constructed: dimensions by which individuals manage to relate to others in particular ways” (Brown and Levinson 1987, p. 55).[11]

---

11. Brown and Levinson’s theory explains that interruption is one kind of communication act (a kind of face-threatening act), and that the way the interruption is expressed determines its effect on social relationships.

Brown and Levinson propose an abstract model of a person they call the Model Person (MP). They employ the structure of their Model Person to explain systematic dialogue structures used by interacting people to influence social relationships.[12]

The Model Person is relatively simple — it consists of only a few structures and rules. "All our Model Person (MP) consists in is a willful fluent speaker of a natural language, further endowed with two special properties — rationality and face" (Brown and Levinson 1987, p. 58). Rationality gives the MP a predictable method of creating plans of communication actions to accomplish its social goals. Face endows the MP with two kinds of desires (or face-wants) ascribed by interactants to themselves and to one another: (1) the desire to be approved of (positive face) and (2) the desire to be unimpeded in one's actions (negative face).

The Model Person is designed to include more than just one person's own face-wants when they plan how to construct communication acts. The MP also incorporates the useful idea that a person's interlocutors (the other participants in a dialogue) also have the same kinds of face-wants themselves. Therefore the MP is useful for analyzing the construction of communication acts, because it supports careful examination of both the face-wants of the speaker and the face-wants of their interlocutors.

We can study the way people construct their dialogue for social effect by modeling human-human interaction as a dyad of MPs engaged in conversation. A dyad of MPs is a convenient model for exposing the conflicting social desires that people experience while interacting. The simple structure of an MP-MP dyad is sufficient to show that MPs' face-wants are interdependent — MPs depend on each other for the satisfaction of their face-wants. Brown and Levinson say that this simple interdependency is adequate to analyze all the subtlety of message construction for social effect observed in real human-human interaction.

---

12. Brown and Levinson's Model Person serves a similar function as the Card et al. Model Human Processor (Card et al. 1983). However each model addresses a different topic of investigation: Brown and Levinson's Model Person allows us to study how people construct and use patterns of dialogue structure to influence their social environment, and the Card et al. Model Human Processor allows us to study how people accomplish physical tasks through the processes of perception, cognition, and motor control.

People do not have direct control over the satisfaction of their face-wants. Face is also inconstant. Therefore, while people are interacting, they are also constantly attempting to maintain, enhance, and defend attacks against their stockpiles of positive and negative face. Face is a precious social commodity that is gained or lost during interactions but not manufactured in isolation.

THEORETICAL CONSTRUCT(S) about people (Brown and Levinson 1987)**:**

P80. People purposefully plan the construction of their communication acts in particular ways that they intend to have specific social effects (see Theoretical Construct P76 (pg 79)).

P81. People make rational decisions when they plan the construction of messages to accomplish their social goals.

P82. People have two kinds of "face" that represent their social desires (face-wants): (1) to be approved of (positive face) and (2) to be unimpeded in actions (negative face). People understand that other people have the same face-wants as themselves.

P83. People are interdependent for the satisfaction of face-wants. This interdependence originates from the fact that people are not able to independently satisfy their own face-wants but instead must depend on each other for face-rewards.

P84. While people interact, they constantly attempt to satisfy their social needs. They constantly apply their rational planning ability to enhance or defend their positive and negative face. There are two categories of people's continual social effort: (1) maintaining constant awareness of the influences of communication acts upon their face-wants and upon the face-wants of others and (2) dynamic adaptive rational planning of the structure of communication acts to exert appropriate social force (to satisfy both positive and negative face-wants). (Note: this new theoretical construct adds two more items to the growing list of things that people perform simultaneously while interacting.)

People sometimes design their communication acts in ways intended to attack each other's or their own face. Brown and Levinson call these attacks face-threatening acts (FTAs). People sometimes construct and execute FTAs as planned attacks upon face with which they intend to have some effect on social relationships. There are four categories of FTAs. These categories define to whom the threat is directed (the speaker or the listener) and the kind of face-threatened (positive or negative face) (Brown and Levinson 1987, p. 68). Table 4 illustrates these categories.

Table 4 — Face-threatening Actions by Who is Threatened

	speaker's face	**listener's face**
positive face	speaker degrades themself before listener	speaker degrades listener
negative face	speaker obligates themself to listener	speaker obligates listener to speaker

Categorization of face-threatening actions (FTAs) along two dimensions of threat: (1) who is being threatened and (2) what kind of face is being threatened.

Brown and Levinson say that interruption is a kind of FTA that intrinsically threatens both negative and positive face. Further, they define interruption as a blatant noncooperative action of discourse with which the speaker disruptively interrupts the listener's talk.

When the speaker interrupts the listener's talk, the speaker makes an FTA (face-threatening action) against the listener's negative face. By interrupting, the speaker implies that they do not intend to avoid impeding the listener's freedom of action. In fact, the interruption directly attacks the listener's freedom to continue talking. The speaker uses this same interruption FTA to also attack the listener's positive face. By interrupting, the speaker implies that they do not care about the listener's feelings, wants, needs, ... in effect, that the speaker does not want the listener's wants. Note, there are other kinds of dialogue behaviors (besides interruption) that intrinsically threaten both negative and positive face. These include: complaints, threats, strong expressions of emotion, and requests for personal information (Brown and Levinson 1987, pp. 65-67).

Interruption is an example of an FTA that threatens face by violating people's expected metadialogue regulation patterns for turn-taking. The theoretical constructs P53 (pg 66) to P57 (pg 66) of this chapter describe how people perform some communication acts solely to regulate the process of the interaction itself and do not convey meaning. Riley (1976) gave examples of three categories of behaviors that people use for metadialogue regulation — turn-taking signals, attention signals, and address signals. Brown and Levinson's (1987, p. 232) MP addresses the analysis of these behaviors as FTAs that act as violations of people's metadialogue regulation patterns of behavior. When people violate these regulation patterns, they impede their interlocutors' freedom to communicate (FTA of negative face), and they also imply that they do not care about their interlocutors' wants (FTA of negative face).

Face-threatening actions are not aberrations or infrequent outbursts of hostility — instead, they comprise a ubiquitous and essential part of normal human-human dialogue. People often need to disagree (threats to positive face) or arrange commitments (threats to negative face) between each other. The only way to accomplish these needs is to make FTAs. Most often, these disagreements or needed commitments are trivial, but people still must achieve them by making FTAs at each other. However, since the amount of "threat" in an FTA is relative to the significance of the disagreement (FTAs of positive face) or the needed commitment (FTAs of negative face), trivial FTAs carry trivial amounts of "threat." Brown and Levinson (1987, p. 13) claim that this theory of face and FTAs applies "universally" across all human languages and cultures.

Because FTAs are so common and because "threat" is a bad thing, regardless of the degree, people use many strategies to try to mitigate or hide the threats they make toward each other. These strategies are ways in which people plan the construction of their communication acts to provide redresses to those they threaten. Brown and Levinson (1987) call these strategies "politeness," and they present a useful taxonomy of politeness as a tool for investigating people's FTA behaviors.

It is useful to review their taxonomy here, with the exemplar of interruption as an FTA. I present examples of interruption for each category of politeness in Brown and Levinson's taxonomy (see Table 5 (pg 87)). This shows how people use different strategies to mitigate the negative effects of interruption.

Brown and Levinson's taxonomy of politeness has four basic categories of expression strategies: (1) bald-on-record, (2) positive politeness, (3) negative politeness, and (4) off record. The bald-on-record category describes the purposeful avoidance of any mitigating strategy (sometimes the most polite strategy is to be totally direct). The positive politeness category describes strategies people use to provide redress to threats to the listener's desire to be approved of (positive face). The speaker reduces the threat of their FTA by reducing the listener's cost of accepting the FTA and/or increasing the listener's ability to reject the FTA. For example, the speaker purposefully makes their own positive face obviously vulnerable. "Yes, I'm threatening your positive face, but as proof of my good intentions here, I make my own positive face vulnerable to counterthreats." Brown and Levinson say that positive politeness is the force behind familiar and joking behavior (see Figure 5).

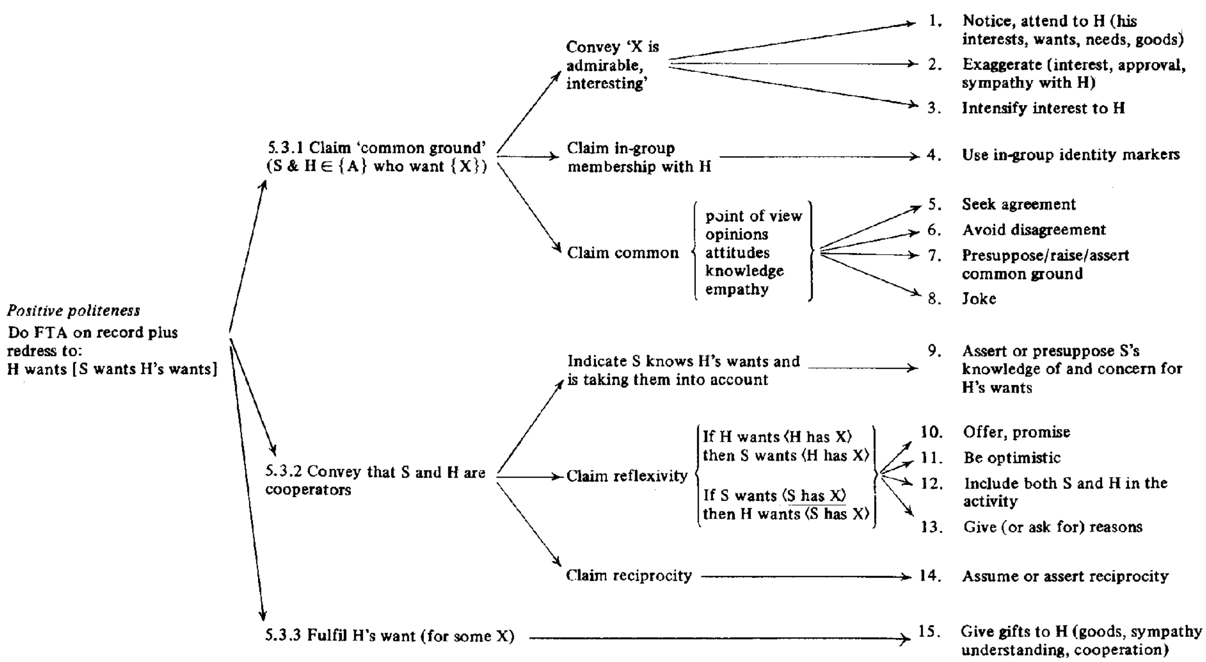

Figure 5.Positive politeness strategies. Reproduced from Brown and Levinson, *Politeness: Some Universals in Language Usage,* p. 102, © Cambridge University Press 1978, 1987.

The negative politeness category describes strategies people employ to provide redress to threats to the listener's desire to be unimpeded in their actions (negative face). The speaker reduces the threat of their FTA by reducing the listener's cost of accepting the FTA and/or increasing the listener's ability to reject the FTA. For example, the speaker proposes reciprocity (an FTA to the speaker's negative face). "Yes, I'm threatening your negative face, but as proof of my good intentions here, I make my own negative face vulnerable to counterthreats." Brown and Levinson say that negative politeness is the force behind respectful behavior (see Figure 6).

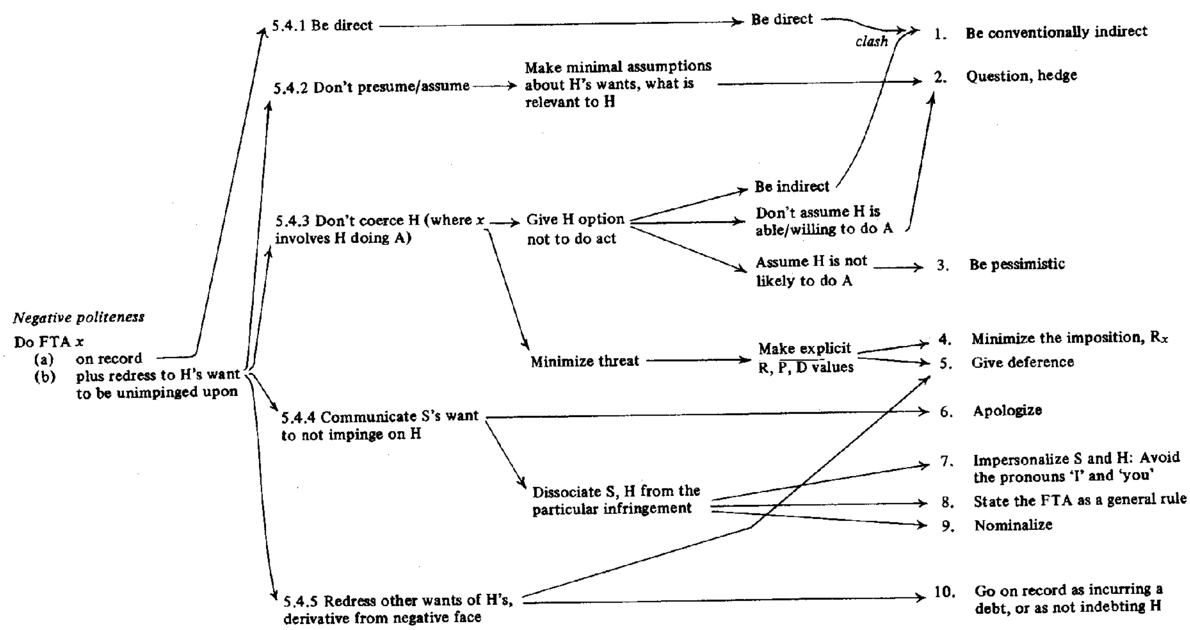

Figure 6.Negative politeness strategies. Reproduced from Brown and Levinson, *Politeness: Some Universals in Language Usage,* p. 131, © Cambridge University Press 1978, 1987.

The off-record politeness category describes strategies people use to create uncertainty about the existence of the FTA itself. This strategy embodies the idea that "people will feel less defensive if they are not sure they are being threatened." Off-record politeness allows the speaker to avoid counterattack by allowing them to claim that no FTA was made. Off-record politeness also allows the listener to ignore threatening requests by allowing them to claim that no FTA was made. When people use off-record politeness strategies, they plan the construction of communication acts in ways that do not have one clear interpretation of their communicative intention. This allows both speaker and listener the potential to hide behind the ambiguity of language if needed. "Maybe I'm making a face-threatening act to you, and maybe I'm not. You can respond to my implied FTA if you choose to — or not." Figure 7 shows the strategies for off-record politeness.

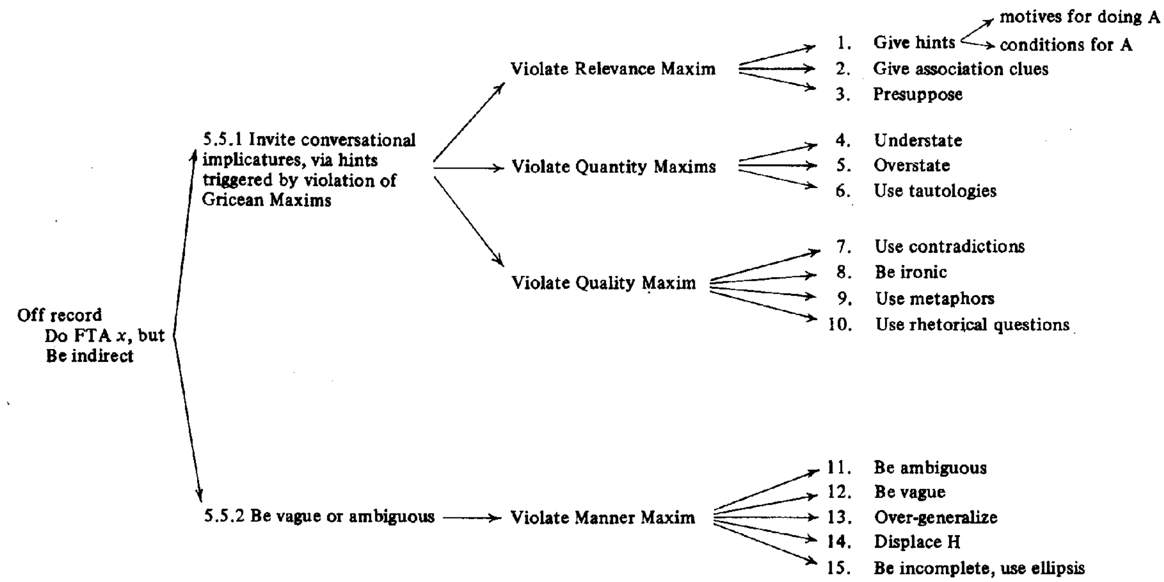

Figure 7.Off-record politeness strategies. Reproduced from Brown and Levinson, *Politeness: Some Universals in Language Usage,* p. 214, © Cambridge University Press 1978, 1987.

Table 5 is a taxonomy of polite interruption strategies. It employs Brown and Levinson's (1987, pp. 91-227) useful taxonomy of FTA (face-threatening action) threat-mitigation strategies (politeness) in the context of interruption. I have included examples of communication acts of interruption (interruption is an example of a kind of FTA that intrinsically threatens both positive and negative face) as examples for each strategy in the taxonomy.

Table 5 — Taxonomy of Polite Interruption

Category of Politeness	**Example**
I. Bald-on-record.	
[A] Be maximally efficient because other demands override any consideration of face-wants	"Hey, you are on fire!"
[B] Be maximally efficient, because the listener is anxious to communicate with the speaker	"OK, the doctor will see you now."

Table 5 — Taxonomy of Polite Interruption (Continued)

Category of Politeness	Example
II. Positive Politeness	
A. Claim common ground	
1. Notice, attend to the listener (their interests, wants, needs, or goods)	"Wow, your hair looks great that way! By the way, I want to talk to you about ...."
2. Exaggerate (interest, approval, sympathy with the listener)	"I had to interrupt you this moment; I thought I would die if I didn't get to talk to you."
3. Intensify interest to the listener	"Hey, your idea worked great! Let me tell you what happened ...."
4. Use in-group language or dialect	"Hey, pal, let's talk."
5. Seek agreement	"The Bulls did well last night. By the way, I want to talk to you about ...."
6. Avoid disagreement	"Excuse me. I agree with what you are saying, but ....."
7. Presuppose/raise/assert common ground	"I heard something that we need to talk about."
8. Joke	"Can I interrupt you from saving the world?"
B. Convey that speaker and listener are cooperators	
9. Assert or presuppose speaker's knowledge of and concern for the listener's wants	"I know you need to concentrate for several hours, so let me interrupt you now and leave you alone the rest of the day."
10. Offer, promise	"If you can give me a few minutes now, I'll try later to find a copy of that paper you want."
11. Be optimistic	[sit down opposite listener] "I've come to talk to you."
12. Include both the speaker and the listener in the activity	"Let's talk."
13. Give (or ask for) reasons	"Is there any reason I should not interrupt you now?"
14. Assume or assert reciprocity	"If you will let me interrupt you now, I will owe you one."
C. Fulfill the listener's want for something	
15. Give gifts to the listener (goods, sympathy, understanding, or cooperation)	"Here's that paper you wanted. By the way, may I interrupt you."

Table 5 — Taxonomy of Polite Interruption (Continued)

Category of Politeness	**Example**
III. Negative Politeness	
A. Be direct	
1. Be conventionally indirect	"Can you be interrupted?"
B. Don't presume/assume	
2. Question, hedge	"I need to interrupt you, if possible."
C. Don't coerce the listener	
3. Be pessimistic	"I don't suppose that I might interrupt you now."
4. Minimize the imposition	"I just need to interrupt you for a second."
5. Give deference	"Excuse me, Sir."
D. Communicate the speaker's want to not impinge on the listener	
6. Apologize	"Please forgive me for interrupting."
7. Impersonalize speaker and listener	"An interruption is necessary."
8. State the FTA as a general rule	[the speaker is wearing an usher's uniform] "Excuse me for interrupting, but it is this theater's policy that all patrons be asked to refrain from smoking."
9. Nominalize	"My interruption is regrettable."
E. Redress other wants of the listener	
10. Go on record as incurring a debt; or as not indebting the listener	"I would be very grateful if you would allow me to interrupt you now."
IV. Off record	
A. Invite conversational implicatures	
1. Give hints	[Clear throat loudly]
2. Give association clues	[Look at wrist watch in an exaggerated manner]
3. Presuppose	[Stare at the listener's face]
4. Understate	[Stand perfectly still and stare at nothing]
5. Overstate	[Jump up and down and wave arms]
6. Use tautologies	[Talking out loud to self] "Time is money."

Table 5 — Taxonomy of Polite Interruption (Continued)

Category of Politeness	Example
7. Use contradictions	[Speaker turns back to the listener exaggeratedly and fold arms]
8. Be ironic	[Talking out loud to self] "I know exactly what to do next."
9. Use metaphors	[Talking out loud to self] "I'm as confused as a hound dog at a tea party."
10. Use rhetorical questions	"I wonder if George [the listener] is busy right now."
B. Be vague or ambiguous: Violate Grice's Manner Maxim	
11. Be ambiguous	[Sigh loudly]
12. Be vague	"I'm here."
13. Overgeneralize	[Talking out loud to self] "Teams that don't coordinate often fail."
14. Displace the addressee	[Within hearing of the only expert on some topic who is busy working, interrupt someone else and ask them a question on the expert's topic.]
15. Be incomplete, use ellipsis	"I see you're ...."

THEORETICAL CONSTRUCT(S) about people (Brown and Levinson 1987)**:**

P85. People frequently design their communication acts in ways intended to threaten each other's or their own "face." People do this because they must make face-threatening actions (FTAs) to accomplish every kind of disagreement or commitment with others (regardless of how trivial).

P86. People's face-threatening actions (FTAs) can be usefully categorized along two dimensions of threat: (1) who is being threatened (speaker or listener), and (2) what kind of face is being threatened (positive or negative).

P87. People use many strategies to mitigate the face-threats they make to each other in their FTAs (face-threatening actions). These strategies are planned attempts to construct the expression of communication acts (i.e., FTAs) in ways that inherently provide redresses to those threatened (Brown and Levinson call these strategies "politeness").

P88. People use four basic kinds of politeness strategies: bald-on-record (when efficiency is paramount); positive politeness (give redress to threats of positive face); negative politeness (give redress to threats of negative face); and off record (the speaker gives themself an easy way of avoiding retaliation and speaker gives listener an easy way of rejecting the FTA).

THEORETICAL CONSTRUCT(S) about interruption (Brown and Levinson 1987)**:**

In3. Interruption is a kind of face-threatening action (FTA) that is intrinsically threatening to both negative and positive face. (Other intrinsic FTAs include: complaints, threats, strong expressions of emotion, and requests for personal information.)

Brown and Levinson's taxonomy of politeness is useful for researching human-human dialogue. It allows researchers to model and analyze the structure of people's communication acts and to infer people's intentions for social action. There are, however, interesting parts of human dialogue that are not included in this model and taxonomy. Two examples are: dialogue directed at oneself and other strategies for FTA construction besides mitigation of perceived threat.

People sometimes make communication acts that are clearly not intended to be perceived by anyone else. We can describe these communication acts as self-dialogue. Brown and Levinson say that people plan the structure of communication acts to accomplish their social intentions. What then, are people's social intentions when they communicate with themselves? Is it possible to analyze the structure of people's self-dialogue and find that they are trying to influence their own perceptions of themselves? If so, then the concept of "face" has another useful dimension — internal vs. external face. Brown and Levinson address external face, but people may also have internal face — self-esteem (positive internal face) and morals (negative internal face).

Brown and Levinson's interactional sociolinguistics theory is limited to addressing people's politeness strategies. Other authors address different human social motivations that influence people's construction of FTAs. West (1982) reports on research about how people construct their communication acts in order to gain dominance over coworkers in an office environment. It may be useful to follow Brown and Levinson's method and construct probable underlying social motivations that people employ to acquire dominance — a taxonomy of dominance strategies. Such a taxonomy would be another essential tool for performing general research into interactional sociolinguistics.

2.15 SIMULTANEOUS SPEECH IN LINGUISTICS

Domain Specific Definition of Interruption: a disruptive type of simultaneous speech.

When people interact, sometimes they talk at the same time (linguists describe this as "simultaneous speech"). Simultaneous speech can be problematic, because people are not good at comprehending more than one spoken message at a time (Pashler 1993; Schneider and Detweiler 1988). Most of the time people speak one at a time and avoid the problems of simultaneous speech. Sacks, Schegloff, and Jefferson (Sacks et al. 1978) propose a theoretical model for investigating how people coordinate turn-taking in conversations. Sacks et al. say that people engage in turn-taking behaviors as a necessary way of sharing the scarce verbal communication channel. "For socially organized activities, the presence of 'turns' suggests an economy, with turns for something being valued, and with means for allocating them affecting their relative distribution, as they do in economies" (Sacks et al. 1978, pp. 7-8). People usually act economically and coordinate turn-taking in order to most efficiently use the scarce verbal channel. Nevertheless, people also spice their dialogue with frequent small bouts of simultaneous speech — why?

People frequently do not engage in turn-taking; this results in more than one person talking at the same time (simultaneous speech). Interruption is a kind simultaneous speech, so it is useful for this discussion to examine the theory relevant to simultaneous speech. The current literature says it is useful to categorize simultaneous speech into three classes: (1) dialogue facilitation; (2) unintentional coordination errors in turn-taking; and (3) direct confrontation for control of the floor.

Back-channeling is a useful example of simultaneous speech for dialogue facilitation (Brennan 1990, p. 395; McCarthy and Monk 1994, p. 42; Pérez-Quiñones 1996, p. 114). Listeners usually provide feedback to their speaker of the success or failure of the communication processes. Listeners initiate brief simultaneous speech acts (or other kinds of brief communication acts) during the speaker's communication turn, e.g., acknowledging eye contacts, head nods, and "um humm's." These kinds of simultaneous speech tell the speaker whether the listeners are attending, hearing, and/or understanding the speaker's messages. (Note: these three states of listener understanding come from Clark's "States of Understanding" for grounding (Clark and Schaefer 1989; Pérez-Quiñones 1996, p. 89).)

Human dialogue is not error-free. Dialogue can breakdown in a myriad of ways. The fragility of the dialogue processes has created the need for some metacommunication acts, like back-channeling, for grounding. The regulation of turn-taking is yet another dialogue process that is itself susceptible to error. People sometimes begin talking simultaneously without intending to do so. When errors happen, people try to repair their dialogue. Sacks et al. (1978, pp. 39-40) report on such dialogue repair behaviors.

The third category of simultaneous speech describes interruption as a purposeful attempt to control the floor. There are different reasons that motivate people to interrupt for control — sometimes people intend to aid the speaker, and sometimes people intend to subvert the speaker. For example, Brown and Levinson (1987) (see theoretical construct P88 (pg 91)) acknowledge that sometimes people need to interrupt each other for cooperative reasons. If the speaker is about to tread on a snake, the listener can interrupt and yell, "Watch out!" (bald-on-record for maximal efficiency). However, people also sometimes use interruption to subvert the speaker. West (1982) describes the human behavior of using interruption with the intention of establishing or strengthening a dominance relationship with the speaker and/or other listeners. West calls occurrences of these subversive simultaneous communication acts “deep interruptions.”

West (1982) says that people react to being interrupted in strategic ways intended to minimize the subversive effects of the interruption. There are two basic categories of reactive strategies: (1) passively allow the interruption but draw attention to the turn-taking violation, and (2) actively fight off the interruption. If the speaker is interrupted, they can mark or draw attention to the interruption as an attack on their right to speak. They can do this by dropping out of the conversation suddenly, leaving a partially finished statement unfinished.

An interrupted speaker may instead try to directly fight off the subversive effects of an interruption by continuing to talk simultaneously throughout the interruption. West says that speakers sometimes stretch or repeat portions of their speech ad infinitum to continue talking and eclipse an interruption. Speakers may instead pretend to ignore the violation of their turn and continue without pause and finish their turn normally, i.e., just keep talking as if the interruption were not happening. Either of these fighting strategies, if executed successfully, seems to be effective in negating the subversive effects of interruption (West 1982).

THEORETICAL CONSTRUCT(S) about people (West 1982)**:**

P89. People sometimes speak simultaneously (“simultaneous speech”).

P90. Most speech communication is not simultaneous because of the problems humans experience in perceiving multiple verbal messages at the same time. People usually employ a metacommunication regulation process of turn-taking to avoid the problems of simultaneous speech.

P91. Although simultaneous speech does not constitute a majority of people’s speech communication, people do frequently spice their dialogue with simultaneous speech. There are three classes of simultaneous speech: (1) dialogue facilitation (for example, back-channeling); (2) unintentional coordination errors in turn-taking (for example, false starts); and (3) direct confrontation for control of the floor (for example, bald-on-record politeness or subversive attempts to establish dominance relationships).

P92. People possess defensive behavior strategies for negating the possible subversive effects of interruption. These strategies include: marking the interruption as a violation of one's right to speak; eclipsing the interruption; and ignoring the interruption.

2.16 Language Use in Linguistics

Domain Specific Definition of Interruption: a proposal for an entry into or exit out of a joint activity.

It is important to consider the interruption of people by the use of language. Clark (1996) in his book titled *Using Language*, presents an important and useful theory of human language. Clark proposes that for a theory of language to be useful, it must support the research of language within its complex context of use. Some other authors have tried to circumvent the complexities of language by asserting that language can be isolated from its context, but this is ineffectual. Language can not be separated from its context of use, i.e., who is using it (participants); what they are doing (joint activities); where they are using it (physical, psychological, and social contexts); why they are using it (joint purposes), and how they are using it (communication of meaning and coordination of the communication process itself).

People do things together. Clark says that the activities people perform together become more than the sum of individual behaviors. He calls these multiperson activities "joint activities," and he says that joint activities are composed of "joint actions."

> "A joint action is one that is carried out by an ensemble of people acting in coordination with each other. As simple examples, think of two people waltzing, paddling a canoe, playing a piano duet, or making love. When Fred Astaire and Ginger Rogers waltz, they each move around the ballroom in a special way. But waltzing is different from the sum

> of their individual actions — imagine Astaire and Rogers doing the same steps but in separate rooms or at separate times. Waltzing is the joint action that emerges as Astaire and Rogers do their individual steps in coordination, as a couple. Doing things with language is likewise different from the sum of a speaker speaking and a listener listening. It is the joint action that emerges when speakers and listeners (or writers and readers) perform their individual actions in coordination, as ensembles" (Clark 1996, p. 3).

Joint activities expose joint purposes. People use language to communicate and coordinate the accomplishment of their joint activity. However, the language use itself is only a means to an end. The relationship between language and joint activities can be illustrated with the metaphor of the relationship between a car and a person's need to get from point A to point B. In general, people drive cars to get somewhere — they do not just drive their cars around and around aimlessly. (Even "Sunday drivers" out for a rid, have the aim of seeing particular things, and they guide their cars purposefully toward some scenic road or other.) In a similar way, people do not just make communicative expressions without purpose. They talk, or gesture, or ..., as a means of accomplishing some joint purpose (Clark 1996).

Joint activities always include some amount of language use (language use is itself a kind of joint activity). This is why we must observe language within its context of use — people's language use is directly tied to the joint activities they are attempting.

> "Just as language use arises in joint activities, these [joint activities] are impossible without using language. Two or more people cannot carry out a joint activity without communicating, and that requires language use in its broadest sense. Yet whenever people use language, they are taking joint actions. Language use and joint activity are inseparable. The conclusion once again, is that we cannot understand one without the other. We must take what I call an action approach to language use, which has distinct advantages over the more traditional product approach" (Clark 1996, p. 29).

THEORETICAL CONSTRUCT(S) about people (Clark 1996)**:**

P93. People engage in "joint activities" with other people. The performance of joint activities is more than the sum of the individual behaviors of the participants. Individual behaviors become interwoven into a fabric of coordination created by the participants.

P94. People coordinate joint activities with language use. Since joint activities require coordination, participants must use language to perform every kind of joint activity.

P95. People perform joint activities to satisfy their joint projects. Joint projects are joint goals that people agree on. The language use that people exhibit while performing joint activities is only a means to their joint project ends — the language use itself is not meaningful in isolation.

THEORETICAL CONSTRUCT(S) about context (Clark 1996)**:**

C2. Language cannot be separated from its context of use (the participants; their joint activities; the physical, psychological, and social contexts; their joint purposes; and communication of meaning and coordination of the communication process itself).

This theory of language use and joint activities is useful for researching the topic of this dissertation — interruption of people by computers. Clark says that language use is itself a kind of joint action. He also says that joint activities have boundaries (entry, body, and exit) and that all entries and exits have to be engineered separately for each joint action (Clark 1996, p. 37). We can, therefore, define interruption as the occurrence of a person making communication acts with the intention of initiating the entrance into some joint activity.

To use Clark's theory, we must understand how it explains three aspects of joint activities relative to interruption: (1) the structure and dimensions of variation in joint activities; (2) the communicative mechanisms that people use to propose entrance into joint activities; and (3) coordination between participants of the entrance into joint activities.

2.16.1 The Structure and Dimension of Variation of Joint Activities

Clark (1996, p. 37) says that joint activities have flexible and dynamic structures: two or more joint activities can be performed simultaneously by the same participants; a single joint activity can be performed intermittently; joint activities may be divided, expanded, and/or contracted dynamically to accommodate changes to the group of participants (people can enter and exit from a joint activity already in progress). The support of simultaneity and intermittence of performance are particularly useful for researching interruption of a person interacting with a system of multiple intelligent software agents.

Another property of the structure of joint activities is that they are nested. “A joint activity ordinarily emerges as a hierarchy of joint actions or joint activities” (Clark 1996, p. 38). This idea is similar to the application of GOMS analysis by the hierarchical decomposition of goals into subgoals (Card et al. 1983). A high-level joint activity can be analyzed into its component low-level joint activities. For example, two people are in a store — a customer and a clerk. They attempt to accomplish the high-level joint activity of a business transaction. It is useful to decompose this joint activity into its component joint activities. These lower-level joint activities include: (1) both participants enter into a business transaction; (2) both participants settle the nature of the transaction (which products the customer will purchase); (3) both participants settle on the total cost of the transaction; (4) both participants make the exchange; and (5) both participants agree to end joint activity.

Clark enumerates dimensions of variability inherent in different types of joint activities. These dimensions are: scriptedness, formality, verbalness, cooperativeness, and governance. Scriptedness is the degree of prearrangement of specific behaviors (for example, participants in a marriage ceremony have preagreed on specific sequences of joint actions). Formality is the amount of agreed-upon behavioral protocols (for example, participants in a presidential debate are constrained by several preagreed upon protocols). Verbalness is the degree to which participants use speech in the joint activity (for example, participants in a fencing match use very little speech). Cooperativeness is the degree to which participants want the same outcome of the joint activity (for example, participants in a chess match want different outcomes). Governance is the degree to which control of the joint activity is shared among the participants (for example, the examiner has more control than the applicant in directing the joint activity of taking the driving test with the Department of Motor Vehicles) (Clark 1996, p. 30).

Clark presents a summary of the elements that affect people's performance of joint activities, Table 6.

Table 6 — The Structure of Joint Activities

element	**explanation**
Participants	A joint activity is carried out by two or more participants.
Activity roles	The participants in a joint activity assume public roles that help determine their division of labor.
Public goals	The participants in a joint activity try to establish and achieve joint public goals.
Private goals	The participants in a joint activity may try individually to achieve private goals.
Hierarchies	A joint activity ordinarily emerges as a hierarchy of joint actions or joint activities.
Procedures	The participants in a joint activity may exploit both conventional and nonconventional procedures.
Boundaries	A successful joint activity has an entry and exit jointly engineered by the participants.
Dynamics	Joint activities may be simultaneous or intermittent, and may expand, contract, or divide in their personnel.

Clark's general claims about the structure of joint activities. Reprinted from Clark, *Using Language*, pp. 38-39, © Cambridge University Press 1996.

THEORETICAL CONSTRUCT(S) about people (Clark 1996)**:**

P96. People's language use is itself a kind of joint activity.

P97. People propose entrance to and exit from joint activities to other people. (We call this interruption.)

P98. People’s joint activities have flexible and dynamic structures. Joint activities support the following kinds of flexibility: they can be performed simultaneously or intermittently; they can be divided, expanded, and/or contracted dynamically to accommodate changes to the group of participants; and they can be (and usually are) nested.

P99. There are dimensions of variability that affect people’s performance of joint activities: scriptedness, formality, verbalness, cooperativeness, and governance.

THEORETICAL CONSTRUCT(S) about context (Clark 1996)**:**

C3. The linguistic context affects joint activities: participants, activity roles, public goals, private goals, hierarchies, procedures, boundaries, dynamics.

One of Clark’s joint activity element structures is “boundaries.” He says that people’s joint activities have boundaries in time. Clark identifies stages of performance of joint activities: “entry,” “body” and “exit.” Entry is the stage where the participants of a joint activity transition from not being involved in the joint activity to being involved in the joint activity, i.e., the transition of individual people into joint activity participants. Body is the stage where participants perform the joint activity. Exit is the stage where participants of the joint activity transition from being involved in the joint activity to being not involved, i.e., the transition of participants of the joint activity back into individual people.

Clark says that transitions into and out of joint activities (entries and exits) are critical. “Entries and exits have to be engineered for each joint action separately. That makes entries and exits especially important features of joint activities” (Clark 1996, p. 37). One person (the speaker) proposes entrance into a joint project to another person (the addressee), who takes it up.

2.16.2 The Communicative Mechanisms that People Employ to Propose Entrance into Joint Activities

I use Clark's theory of language to define interruption as "the occurrence of a person proposing entrance to, or exit from, a joint activity to another person (or group of people)." This is directly useful to the subject of this dissertation — user-interruption in HCI. For HCI, this definition can be rewritten as "the occurrence of a computer proposing entrance to, or exit from, a joint activity to its human user(s)."

How do people make these interruptions? People propose entries and exits by using language to communicate their meanings to each other. Clark presents a theory of signals to support investigation of the basic process of communication meaning between people. What is most interesting is how people use the structure of joint activities and signaling to engineer joint activities. However, it is essential to discuss signaling theory before discussing how people actually coordinate entrance and exits of joint activities.

Clark (1996, p. 160) defines signals as "the presentation of a sign by one person to mean something for another." A signal is a process of conveying meaning between people. People deliberately choose a signaling class to form the expression of the communication acts they make. Each of the three signaling classes is composed of: (1) a kind of sign, (2) a method of signaling, and (3) a targeted cognitive process of sign interpretation within the listener.

This concept of signal is not limited to static representations of meaning but encompasses the total process of conveying meaning between participants. People signal other people — the speaker creates and expresses some composition of signs with the intention of conveying a particular meaning, and the audience receives those signs and interprets the meaning. Signs comprise signals and affect the audience by creating in their mind(s) particular ideas. Clark's signals are similar to the concepts of communication acts employed by Riley (1976) in Theoretical Constructs P50 (pg 63) and P51 (pg 63). Communication acts make clear distinctions between meaning, the expression of communicative behaviors, and the channels of conveyance. Clark's signals promote similar distinctions between meaning and its expression.

Signs evoke meaning in their human receivers. Clark categorizes signs into three classes: icons, indexes, and symbols. Clark (1996) says these categories reflect the three different ways that people are able to interpret meaning. Each class of sign expressed evokes a different kind of cognitive recognition process within its audience. Audiences interpret each class of sign by using different memory resources to construct meaning from the sign behaviors they perceive from speakers. Icon signs evoke perceptual memories. Index signs evoke memories of physical or temporal relationships. Symbol signs evoke memories of learned rules of association.

Clark's methods of signaling (Table 7) align the meaning a person intends to communicate with the particular cognitive process the listener will use to understand that meaning. Signaling processes and signs describe different ways people understand things, and the methods of signaling describe people's choice of communication expression to convey that meaning. The three classes of signaling methods are: demonstrate, indicate, and describe. These methods are each exclusively associated with a particular class of sign and sign process: people demonstrate with icon signs that evoke perceptual memories in their communication partners; people indicate with index signs that evoke memories of physical association in their communication partners; and people describe with symbol signs that evoke memories of learned rules of association.

Table 7 — Signaling Classes

	Description	**Indication**	**Demonstration**
Kinds of signs	symbol	index	icon
Cognitive processes of sign interpretation	learned rule association	physical or temporal association	perceptual similarity association
Methods of signaling	describing-as	indicating	demonstrating

Signal classes based on Clark's theory of signaling (Clark 1996, pp. 155-188). People create signals as trios of three interrelated structures: signs, cognitive processes within their audience, and methods of signaling.

Clark emphasizes the classes of signaling methods as easy means of discussing the different ways people signal each other. This approach is useful because each class of signaling method assumes a particular class of signal and a particular class of cognitive process for interpretation. For example, the describing-as method of signaling is always accompanied by the creation and expression of symbol signs which the audience interprets by using their memories of rules of association to constructivistly recognize the meaning conveyed. Conventional English language words like "whale," "opera" and "sing" are clear illustrative examples of symbols. Speakers use symbols when they describe things to their audience. Obviously, speakers often use symbols by uttering them as spoken words. However, there are other interesting ways of using symbols in description. Some physical gestures are expressions of symbols — there are a set of gestures that both speaker and audience have learned to associate meaning by rule. Clark calls these emblems. Examples are: the head nodding gesture means "yes" and the shoulder shrugging gesture means "I don't know." There are also auditory emblems that people have learned to associate nonword sounds with particular meanings, e.g., clapping means "I approve." Clark also describes a kind of gesturing he calls junctions, in which participants express the symbolic gesture simultaneously, e.g., shaking hands. (Clark 1996, pp. 161-164).

The indicating method of signaling is always accompanied by index signals that require the audience to apply their memories of physical or temporal associations. People signal for indication with their audience when they intend to create indices for the objects (including people), events, or states that they want to refer to. A person pointing their index finger at some object is a simple example of indication. Indexing has five requirements: (1) the index expressed so that all participants focus attention on it; (2) the index locates some object, event, or state in space and time; (3) the index presents a physical and/or temporal association; (4) the index is accompanied by a description that identifies a particular instance of the object; and (5) the speaker presumes that their audience can satisfy requirements 1-4 based on their current common ground (Clark 1996, pp. 164-172).

The demonstrating method of signaling is always accompanied by icon signals that require the audience to apply their memories of perceptual similarity. Clark (1996, p. 174) says, "the point of demonstrating is to enable addressees to experience selective parts of what it would be like to perceive the thing directly." Demonstration works because people are able to add an imaginary layer to their conversations. This imaginary layer affords people the opportunity to imagine that icons are the objects they represent. Clark says there are several ways of expressing demonstrations, including iconic gestures. Component gestures are selective depictions embedded within larger utterances. Concurrent gestures can be icons that a person makes while they utter something. Facial gestures often demonstrate what the speaker's face would look like if they were experiencing some meaningful reaction. Vocal gestures are icons people use to demonstrate meaning by making selectively depictive sounds (Clark 1996, p. 172-183).

Clark provides a classification table (Table 8) of examples of different, easily recognized communicative behaviors and where they fall within his signaling categorization theory (Clark 1996, p. 188).

Clark (1996) says that it is useful to describe people's signaling by clearly distinguishing the kind of signal, the instrument they use, and the depictive action they take. Table 9 shows how people can use all classes of signaling and all of the instruments of their bodies for interruption, i.e., to propose entrance into a joint project to another person(s). [Note: "<I>" stands for instrument and "<O>" for object.]

Table 8 — Signaling Methods and Instruments of Expression

	Method of Signaling		
Instrument	**Describing-as**	**Indication**	**Demonstrating**
voice	words, sentences, vocal emblems	vocal locating of "I" "here" "now"	intonation, tone of voice, onomatopoeia
hands, arms	emblems, junctions	pointing, beats	iconic hand gestures
face	facial emblems	directing face	facial gestures, smiles
eyes	winks, rolling eyes	eye contact, eye gaze	widened eyes
body	junctions	directing body	iconic body gestures

Examples of different kinds of communicative behaviors, and where they fit into Clark's signaling theory classification. Reprinted from Clark, *Using Language,* p. 188 © Cambridge University Press 1996.

Table 9 — Signaling for Interruption

Instrument	Depictive action	Example
Description		
voice	utter <O> with <I>	"ugh um" {clearing the throat loudly is a verbal emblem people have learned to mean "I need to interrupt you now"}
voice	utter <O> with <I>	"I need to interrupt you now" {words are symbols}
hands	gesture <O> with <I>'s	"[make the football game time-out signal by bringing the two hands together into a "T" shape]" {this is an emblem that people have learned to mean "I request an interruption"}
finger	gesture <O> with <I>	"[make the "shhhh" signal by placing an index finger vertically across closed mouth]" {this is an emblem that people have learned to mean "I request you to stop (or interrupt) your speaking"
Indication		
eyes	gazing at <O> with <I>	"[turn eyes to gaze at someone] Excuse me" {the eye gaze indicates who I want to interrupt}
finger	pointing at <O> with <I>	"[point at someone] Excuse me" {the finger pointing indicates who I want to interrupt}
arm	raising <I> in the visual field of <O>	"[a student raises their arm upwards toward the instructor during a lecture]" {the student indicates themself as the proposer of an interruption}
voice	identifying <O> with <I>	"I need to interrupt you now" {this utterance indicates "I," "you," and "now"; the rest of the utterance is describing-as}
Demonstration		
fingers	forming <I>'s into <O>'s shape	"[I put my fingers to my mouth as if I were whistling loudly]" {I interrupt you by demonstrating a loud whistle (whistle not actually made)}
voice	mimicking <O> with <I>	"Your mom said, 'get William off the telephone' [impersonation of mother's voice]" {I interrupt William with a demonstration of his mother interrupting him.}
arm	miming <O> with <I>	(you are across the room from me) "[I use one arm to demonstrate a pulling motion]" {I mean that I want to interrupt you from whatever you are doing and draw you to me}

Table 9 — Signaling for Interruption (Continued)

Instrument	**Depictive action**	**Example**
hand	forming <I> into <O>’s shape	(I want to interrupt you during your telephone conversation) “[I use my hand to make a cutting gesture between you and your telephone]” {I demonstrate cutting your phone cord so I can interrupt your conversation}

Clark says that people usually use composite signals in which they mix all three kinds of signals. The reader may ask, “If people actually signal with composite signals, why is it useful to employ Clark’s signaling theory which classifies signals into three separate categories?” My answer is that Clark’s signaling theory allows us to analyze people’s communicative behaviors in terms of their use of different kinds of cognitive resources. Even though a person signals with a mix of different signaling types, we can analyze how the speaker and their audience are using different cognitive resources to accomplish that signal — perceptual memories for icons, physical or temporal association memories for indexes, and rule association memories for symbols.

It is useful to show the utility of Clark’s signaling theory by employing it to analyze an example of interruption (a person proposing entrance into a joint project with another person). Thomas interrupts Miranda while she is talking on the telephone: [Thomas walks over near Miranda] [he turns his head and body toward her] [he fixes his eye gaze on her] [he reaches out his hand and makes a cutting gesture between her and the telephone] (while making his hand gesture, he utters) “Excuse me Miranda, I need to interrupt you now please, it’s urgent.”

Thomas’s behaviors of moving his body closer to Miranda, orienting his body, head, and eye gaze toward her are all indication signals. These behaviors draw Miranda’s attention to Thomas, they locate him in physical space. He specifies himself as a particular person making an interruption, instead of just an abstract person (description); and because of their common ground, he presumes that Miranda can understand his meaning. Thomas also uses utterance for indication signals. He indicates himself as the proposer of a new joint project with the word “I”; he indicates Miranda as the recipient with the words “Miranda” and “you”; and he indicates the time of his proposal with the word "now." We can analyze all these indication signaling in terms of how Thomas and Miranda use their cognitive memories of physical and temporal associations to accomplish the conveyance of index signs. Thomas's behavior of hand gesturing a cutting motion between Miranda and her phone is a demonstration signal. Thomas provides Miranda with an imaginary experience of severing her telephone conversation. We can analyze how Thomas and Miranda use their cognitive, perceptual, and similarity memories to accomplish the conveyance of this icon sign. Thomas's behavior of uttering words are description signals. His uttered sentence conveys symbol signs that Miranda must interpret. Thomas uses symbol signs to describe that he "needs" to "interrupt" her, "please." We can analyze how Thomas and Miranda use their cognitive rule association memories to accomplish the conveyance of these symbols signs.

THEORETICAL CONSTRUCT(S) about people (Clark 1996)**:**

P100. People's joint activities are bounded in time into entrance, body, and exit; and people must engineer each one of these stages of joint activities. (Note: theoretical construct P97 (pg 99) says that we can define interruption as the occurrence of one person proposing entrance to or exit from joint activities to other people.

P101. People convey meaning to each other by signaling. A signal is the means by which a speaker creates and expresses signs which cause their audience to interpret some intended meaning.

P102. People use three kinds of signals: description, indication, demonstration. These categories correspond to the three different ways that people are able to cognitively process meaning. Each kind of signal has (1) its own kind of signs of which it is composed, (2) its own cognitive process of interpretation in the audience, and (3) its own method of signaling.

P103. People convey meaning with description signals by referring to conventional systems of meaning — like words — that both speaker and audience have learned to associate by rule. Description (1) creates symbol signs, (2) requires the audience to use their memories of rules by association for interpretation, and (3) is expressed by the speaker with the describing-as method of signaling.

P104. People convey meaning with indication signals to identify or mark an object, event, or state for future reference. Indication (1) creates index signs, (2) requires the audience to use their memories of physical or temporal association for interpretation, and (3) is expressed by the speaker with the indicating method of signaling.

P105. People convey meaning with demonstration signals by enabling their audience to experience selective parts of what it would be like to perceive directly the original. Demonstration (1) creates icon signs, (2) requires the audience to use their memories of perceptual similarity for interpretation, and (3) is expressed by the speaker with the demonstrating method of signaling.

P106. People usually use composite signals in which they mix all three kinds of signals.

People create complex heterogeneous signals (theoretical construct P106 (pg 108)), but how do they decide which kinds of signals to combine? Clark says people choose the composition of their signals based on the dynamics of the pertinent joint activities. He proposes three dynamic factors of interaction that affect how people choose different kinds of signals to combine: purpose, availability, and effort (Clark 1996, pp. 186-187). First, joint projects have purposes that may suggest or even require specific types of signals. Clark says, "in so far as describing-as, indicating, and demonstrating serve different purposes, speakers' choices of composite must conform to their purposes" (Clark 1996, p. 187). The second factor — availability — affects peoples choices because some contexts do not allow some kinds of signals. For example, when people interact over the telephone, they do not have the option to choose many kinds of gestural indexing signals that depend on people's ability to see each other. The third factor — effort — reveals that different kinds of signaling afford the conveyance of different kinds of information. It may be possible to convey a single meaning in different ways, i.e., with alternate kinds of signaling. However, each alternate signaling type requires a different amount of effort to accomplish. For example, if I wanted to tell you how to tie a shoe, I might be able to do it with the describe-as method of signaling (I describe how to hold the two ends of the laces, and then how to wrap one around the other, and then ...), but I could accomplish it much easier by the demonstration method of signaling (I demonstrate by tying my shoe).

THEORETICAL CONSTRUCT(S) about people (Clark 1996)**:**

P107. People create complex heterogeneous signals (theoretical construct P106 (pg 108)) by selecting and intermixing different kinds of signals that best convey meaning in their dynamic joint activities. Three factors of joint activities affect people's selection of signal types: purpose, availability, and effort.

2.16.3 Coordination Between Participants of the Entrance into Joint Activities

People can construct appropriate signals to convey meaning within dynamically changing joint activities. The meaning people convey with signals allows them to accomplish their joint purposes. The most obvious kind of information people convey to each other is related to the subject of their joint activity, however people must also perform a significant amount of signaling just to coordinate the joint activity itself. Clark says that coordination is the key to joint activities. "What makes an action a joint one, ultimately, is the coordination of individual actions by two or more people. There is a coordination of both content, what the participants intend to do, and processes, the physical and mental systems they recruit in carrying out those intentions" (Clark 1996, p. 59).

The previous two subsections of this chapter presented the background information necessary for a discussion of the coordination of joint projects and its relevancy to interruption, i.e., (1) the structure of joint activities and (2) how people signal each other. Interruption is defined here as a proposal for entry into a joint activity; therefore, it is important to understand how Clark's theory addresses the coordination of such proposals and entries.

Clark (1996, p. 191) says, "a joint project is a joint action projected by one of its participants and taken up by the others." There are several aspects of a proposed joint project that must be coordinated between the proposer and their intended collaborators. Participants must coordinate on:

(1) agreement: whether or not there will be a joint project;
(2) who: the group of participants;
(3) what: the content of joint activity and the roles of participants;
(4) where: the location;
(5) when: times of state transitions (entry time, body time, and exit time);
(6) why: joint purpose; and
(7) how: state changes, and signaling channels (e.g. telephone, e-mail, or face-to-face).

People are not always successful at accomplishing the joint projects they desire. Clark (1996, p. 203) says there are four factors that specifically affect people's ability to coordinate entrance into joint projects. All of these factors must be successfully coordinated for people to commit themselves to enter into a joint activity. These factors are: identification, ability, willingness, and mutual belief. Since these four factors are requirements of entrance into joint projects, they are also the requirements for successful interruption. I use Clark's labeling convention of "*A*" for the person proposing a joint project, "*B*" for the person or people receiving the proposition, and "*r*" for the joint project itself.

1. Identification: both *A* and *B* must coordinate a joint understanding of the nature of *r*. This understanding comes from coordinating the seven aspects of a proposed joint project, as described in the preceding paragraph (agreement, who, what, where, when, why, and how).
2. Ability: both *A* and *B* must be able to fulfill their roles in *r*.
3. Willingness: both *A* and *B* must be willing to fulfill their roles in *r*.
4. Mutual belief: *A* and *B* must each believe that they both have successfully coordinated all three factors 1, 2, 3 and that they have now coordinated factor 4.

An illustrative example is useful here: Eric proposes entry into a joint activity to Barbara. Initially, Barbara is working alone on her computer; Eric walks into her office and interrupts.

> Eric:"Barbara, excuse me."
>
> Barbara:"Just a moment." [she keeps working for 10 seconds, then stops, and turns around], "Yes?".

Eric proposes a joint activity — that he and Barbara enter into a conversation (factor 1). Eric also shows that he is able and willing to do his part (factors 2 and 3). Barbara's response, "just a moment," conveys that she has identified Eric's purpose (factor 1) and that she will be willing and able to enter the proposed joint activity in a few seconds. Barbara continues to work for 10 more seconds, and then stops and turns around, and says, "Yes." At this point both Eric and Barbara accomplish factor 4. They enter the joint activity.

Barbara accepted Eric's proposal for entrance into a joint activity, but she altered one aspect of his original proposal. When Eric said, "Barbara, excuse me," Eric meant that he and she enter his joint activity "right now." Barbara altered the "when" aspect of Eric's original proposal.

People have alternative responses for answering propositions for entrance into a joint activities. Clark (1996, pp. 203-205) has identified four possible responses: take-up with full compliance; take-up with alteration; decline; and withdraw (see also Clark's discussion of the "Emergence of Conversations" and "Opening Sections" (Clark 1996, pp. 331-334)).

Table 10 discusses these alternative responses to entry propositions. To illustrate each possible response, I provide an example variant of Barbara's response to Eric's proposal (Clark 1996, pp. 203-205, pp. 331-334).

There is a subtle difference between Clark's "decline" and "withdraw." This reflects the idea that people in U.S. English culture maintain a subtle informal agreement not to ignore each other. This cultural agreement implies that once a person makes a proposition for a joint activity to other people, these recipients may already feel some commitment to at least decline the proposal. Clark has included the special case of “withdraw” to describe a kind of response in which people not only decline the proposed joint project but also decline their cultural contract to not ignore proposals in general.

Table 10 — Possible Responses to a Proposal for Entry into a Joint Activity

response type	definition	example
take-up with full compliance	recipient complies fully with entrance into joint project exactly as proposed	“[Barbara immediately stops whatever she is doing and turns around], yes?”
take-up with alteration	recipient agrees to enter into an altered version of the original proposal [recipient declines original proposal and counter-proposes with an altered version of the original]	Barbara says, “Just a moment. [she keeps working for 10 seconds, then stops, and turns around] Yes?”
decline	recipient declines to enter into the proposed joint project, because they are either unable or unwilling to comply	“[Barbara does not stop her work or turn around], sorry, I’m too busy right now”
withdraw	recipient withdraws entirely by responding with something completely unrelated	“[Barbara ignores the proposal: she does not stop her work, or turn around, or even respond]”

The four possible responses to a proposition for entry into a joint activity (Clark 1996, pp. 203-205, pp. 331-334).

The previous discussion about coordination (proposals for entry into joint actions and alternative responses by proposal recipients) is similar to Winograd and Flores' (1986) state transition diagram (Figure 8) for the possible states of people's conversation for action (Preece et al. 1994, p. 175).

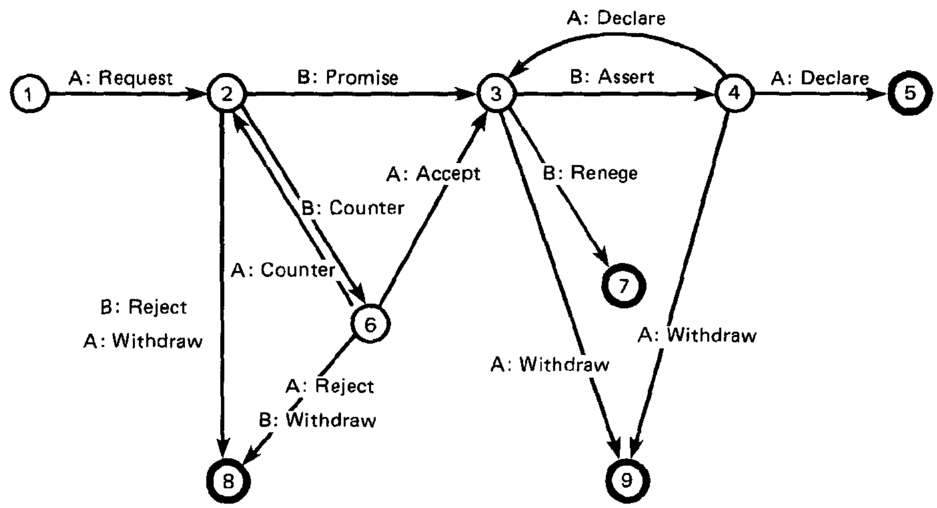

Figure 8.The basic conversation for action state transition diagram. Reprinted from Winograd and Flores (1986, p. 65). (Note that the states do not correspond to turn-taking, but only represent the different possible "states" for this type of conversation.)

If we can view a proposal for entrance into a joint activity as a kind of Winograd and Flores' "conversations for action," then we have a useful state transition diagram for analyzing people's joint project entrance coordination behaviors. Clark's alternative responses to entrance proposals would represent different paths through the Winograd and Flores' state transition diagram: "take-up with full compliance" could be a 1-2-3-4-5 path; "take-up with alteration" could be a 1-2-6-3-4-5 path; "decline" could be a 1-2-8 path; however "withdraw" has no comparative path. I propose that Winograd and Flores' state transition diagram for conversation for action has limited usefulness because of its failure to model the "withdraw" response option.

Coordination of joint activities takes time and effort to perform, so people use shortcuts whenever possible. Clark (1996, pp. 70-72) identifies two different kinds of coordination shortcuts. The first is explicit agreement. This is a planned future coordination that all participants commit to. For example, all participants can agree on the time, place, and format of their next meeting. This saves participants from having to dynamically coordinate part of the entrance into a future joint activity. The other kind of coordination shortcut is convention. Convention solves a recurring coordination problem by fixing a “standard” way to coordinate all future occurrences of a particular kind of coordination problem. Clark provides a useful example, “in America and Europe, placing knives, forks, and spoons on the table is a solution to the recurrent problem of what utensils to use in eating. In China and Japan, it is to place chopsticks” (Clark 1996, p. 70).

People use explicit agreement and convention to shortcut the coordination costs of interruption. For example, when people explicitly agree to a future meeting, they shortcut the need to interrupt each other to enter into that future joint activity. Also, if people can create a convention for meeting, they shortcut the coordination costs of interrupting each other for the entrance to all recurring meeting joint activities. For example, all participants agree to meet at a particular time and place on the second Tuesday of each month.

THEORETICAL CONSTRUCT(S) about people (Clark 1996)**:**

P108. People must coordinate several issues to enter a successful joint activity: *agreement* to enter; *who* will participate; *what* will be the content of joint activity and the roles of participants; *where* and *when* the joint activity will take place; the joint purpose (*why*); and *how* to signal and transition boundaries.

P109. There are four requirements for the successful coordination of entrance into joint projects (and therefore requirements for successful interruption): (1) identification — participants must understand the nature of the proposed joint project; (2) ability — participants must be able to fulfill their roles in the proposed joint project; (3) willingness — participants must be willing to fulfill their roles in the proposed joint project); and (4) mutual belief — participants must each believe that all participants have successfully coordinated requirements 1, 2, and 3 and then 4 itself.

P110. People can respond to propositions for entrance into joint activities in a number of different ways. Four possible alternative responses to the same proposition are: (1) take-up with full compliance (comply fully with proposition as-is); (2) take-up with alteration (comply with variant of original proposition, i.e., decline proposition and counter-propose an altered version of the original); (3) decline (decline by indicating inability or unwillingness to comply); and (4) withdraw (withdraw totally without responding to proposal).

P111. People in U.S. English culture maintain a subtle informal agreement not to ignore each other.

P112. Coordination of joint activities takes time and effort to perform.

P113. People can shortcut the costs of coordinating entrance into joint activities by employing explicit agreement (for single future joint activities) or convention (for recurring future joint activities).

CHAPTER 3: SYNTHESIS OF THE FIRST THEORETICAL TOOLS

3.1 AIMS AND OBJECTIVES

The aim of this chapter is to synthesize generally useful theoretical tools for the investigation of human interruption. No general tools exist yet for this research. This chapter uses the results of the preceding analysis chapter (“Survey of Theoretical Constructs,” pg. 23) to create a general Definition of Human Interruption and a practical Taxonomy of Human Interruption.

No general model of human interruption exists in the current literature, and building such a rigorous model is beyond the scope of this dissertation. However, the theoretical constructs identified in Chapter 2 (pg. 23) can serve as a strong theoretical foundation to synthesize theoretical tools that are generally useful for research about human interruption. Two tools are synthesized and presented here: (1) a general definition of human interruption with accompanying postulates and assertions and (2) a taxonomy of human interruption. This is a descriptive tool that exposes the important theoretical dimensions for analyzing and describing interruptions.

After reading this chapter, readers should be familiar with the theoretical tools presented here and should understand the claims of general utility of these tools and the tool’s potential for practical application into the investigation of human interruption.

3.2 OVERVIEW

This chapter proposes a useful and unique theoretical tool for the investigation of human interruption — a general unifying definition of human interruption. This chapter uses the term "unifying" for its general definition of human interruption because this unique theoretical tool is general enough to bridge the semantic boundaries between several different fields of research. The general definition presented here can unify HCI research about user-interruption by supporting the generalization of theoretical models and experimental results across domains.

The analysis of relevant theory from the current literature identified 126 theoretical constructs of interruption (113 about people, 7 about tasks, 3 about interruptions, and 3 about contexts). This set of identified theoretical constructs forms a useful theoretical foundation about human interruption. The component pieces of this foundation are derived from very broad overview of relevant research domains. The extensive breadth of the analysis makes the resulting set of theoretical constructs generally useful. This chapter uses this raw theoretical information to synthesize a general definition of human interruption and accompanying tools for practical investigation of human interruption.

This synthesis maximizes the general utility of the unifying definition of human interruption in four ways:

1. it uses the maximum breadth depth of analysis results from the preceding chapter;
2. it successfully identifies general common threads of theory across all the different theoretical perspectives from the various research fields analyzed;
3. it uses those general concepts to synthesizes a definition of interruption and accompanying theoretical tools; and
4. it excludes all nongeneral theoretical concepts from the synthesis of its general tools.

3.3 Motivation

Generally useful theoretical tools are necessary for the investigation of human interruption. No such tools exist. This chapter uses the results of the previous chapter (“Survey of Theoretical Constructs,” pg. 23) to synthesize a general, unifying definition of human interruption and an accompanying taxonomic tool for its practical application in research.

3.4 A General Unifying Definition of Human Interruption

Definition: *Human interruption is the process of coordinating abrupt change in people’s activities.*

Postulate 1: This abrupt change involves one or more of a person’s modes of activity: (1) cognition, (2) perception, or (3) physical action.

Assertion 1: This definition is most useful for investigating deliberate invocations of this process as attempts to cause meaningful effect(s).

Postulate 2: Interruption causes effects that are measurable with an acceptable level of measurement error.

Postulate 3: Interruption is accomplished with physical mechanisms in physical media. (Note, this supports postulate 2.) These physical mechanisms and media are identified in Chapter 2 (pg. 23) and are modeled with the theoretical constructs of interruption presented there.

Postulate 4: Causal relationships exist between the state of several important dimensions of the process of interruption and the effects of the interruption. These process dimensions are the significant veins of descriptive theory contained in the theoretical foundation created in Chapter 2 (pg. 23). These veins are the useful dimensions for investigating relationships between particular interruptions and their effects on people. (Note, see the Taxonomy of Human Interruption on pg. 127.)

Assertion 2: It is possible to intentionally affect the consequences of interruption. The values of the dimensions of the interruption process can be deliberately controlled in order to influence the outcome of interruption events, i.e., its costs and benefits.

Assertion 3: The effects of user-interruption in HCI are directly related to the particular design chosen for the user interface of the system. The design of the user interface directly affects the states of dimensions of the interruption process and, therefore, causally affects the results of interrupting the user.

Assertion 4: The theoretical taxonomic tool provided in this chapter is useful for applying the general unifying definition of interruption presented here. This taxonomy of human interruption is especially useful for investigating user-interruption in HCI.

This general unifying definition of human interruption is a tool. I assert that it is generally useful for the investigation of human interruption. Each part of this definition expresses some theoretical concept of human interruption in a way that is generally applicable.

3.4.1 Human Interruption: a Process of Coordination

Several theoretical constructs say it is useful to consider interruption as more than just a thing or a sign or token that people use. Interruption is an entire process. It involves the who, what, where, when, why, and how of each stage in the generation, transmission, reception, comprehension, and reaction to an interruption event. Other theoretical constructs say that interruption also includes how people coordinate transitions between the different stages of interruption event.

For example, the theoretical construct C2 (pg. 97, from Section 2.16, “Language Use in Linguistics”) says that “language cannot be separated from its context of use.” Therefore, any definition of interruption must include reference to the total process of interruption. Theoretical constructs P51 (pg. 63, from Section 2.9, “Human-Human Discourse”), and P80 (pg. 82, from Section 2.14, “Interactional Sociolinguistics of Politeness”), promote the usefulness of the idea of a “communication act” that embodies more than just the communication message.

A communication act is three things: a meaning, an expression, and a channel of conveyance. Another useful variant of the concept comes from theoretical construct P101 (pg. 107, from Section 2.16, “Language Use in Linguistics”) — signaling theory. It says that signals are more than just the signs they pass. Signals are signs, methods of signaling, and the cognitive processes of sign interpretation.

Interruption is also how people coordinate the interrelationships among the various parts of the interruption process. Theoretical construct P93 (pg. 97, from Section 2.16, “Language Use in Linguistics”) says that joint activities (like interruption) are accomplished with a “fabric of coordination created by the participants.” Theoretical construct P50 (pg. 63, from Section 2.9, “Human-Human Discourse”) says that communication acts are more than just their parts but must be coordinated by the people involved. Theoretical constructs P108 and P109 (pg. 114, from Section 2.16, “Language Use in Linguistics”) say that there are several elements and points of agreement that people must coordinate for their successful entry into a joint activity (i.e., interruption).

The definition of human interruption presented here has a good analogy in the event models of programming languages. Consider the Java programming language’s event model (AWT 1.0) (Chan and Lee 1997, pp. 458-492). This analogy is useful for illustrating the claim that human interruption is more than just the sign that carries that interruption.

Human interruption is somewhat like user interface events in Java. To investigate user interface events in Java (or to investigate human interruption), it is not useful to only look at Event objects. A productive investigation must examine all the parts of the event model and study how all those pieces coordinate in time. In Java, user interface events are part of a larger event model. Instances of Event objects are generated because of human users interacting with a Java program’s graphical component through an interaction device (e.g., a mouse or keyboard). After an Event object is generated, it is translated into the appropriate graphical context and delivered to the appropriate graphical component object. The event continues to be translated and delivered up the component hierarchy (visual hierarchy) until some component handles it or until it passes through the entire component hierarchy. The component objects in the component hierarchy receive Event objects and can be programmed with methods to handle those events. If programmed to do so, a component object's event-handler method(s) will be called and passed Event objects. A component object's event handler method identifies the Event objects it receives and uses their identity to choose and make appropriate reactions.

3.4.2 Human Interruption: Abrupt Change

The theoretical constructs identified in the analysis chapter do not provide a quantitative expression of how quickly interruptions happen (such a quantification would be very interesting). Instead, the word "abrupt" here distinguishes interruption from other kinds of gradual progressions of change that people experience in their activities. Theoretical construct P14 (pg. 34, from Section 2.5, "Multitasking in HCI") says that it is useful to model people's ability to quickly switch between different activities. Theoretical construct P21 (pg. 40, from Section 2.6, "Multitasking in Linguistics") says that it is useful to model people's cognitive behavior of dynamically modifying and changing their behavior while they are making it. The word abrupt in this chapter's definition of human interruption is meant to portray interruption as something that happens quickly and dynamically in real time.

3.4.3 Human Interruption: a Change in People's Activities

The simple intuitive explanation is that people have to be doing something before they can be interrupted to start doing something else. Theoretical constructs P65 and P67 (pg. 73, from Section 2.11, "Psychology of Human Attention") say that it is useful to model all of people's changes in focused attention as changes in the processing stream that currently is executing in the person's focus of consciousness. People perform activities that change. Theoretical construct P97 (pg. 99, from Section 2.16, "Language Use in Linguistics") says it is useful to model human-human behaviors, like interruption, as joint activities that have stages that change over time. People also use interruption to perform several simultaneous, intermittent, nested, and dynamically changing joint activities at the same time (theoretical construct P98 on pg 100, from Section 2.16, "Language Use in Linguistics").

3.4.4 Usefulness of the Definition for Practical Research

The general unifying definition of human interruption presented here affords great research questions. This power for motivating well-constructed research questions is one of the most useful contributions of this general definition. Here are some examples:

- What are the stages in the process of interruption?
- How do people coordinate the process of interruption?
- Of the people involved, who does what parts of the coordination?
- When do people accomplish the different stages and coordinations of interruption?
- Why are people changing theirs, and/or other people's activities, i.e., what is the meaning of an interruption?
- How and why do people carefully design their interruption expressions?
- What channels of conveyance do people use for interruption, and how does the channel used affect the interruption process?
- What kinds of human activities are changed through interruption?
- What are the effects or costs and benefits of interruption of human activities?

Investigation of these questions is directly supported by numerous theoretical constructs from the analysis chapter of this paper. Many of these theoretical constructs do not directly support the unifying definition itself, because they say how interruptions work but not what they are. However, the collection of theoretical constructs is broad enough to be a generally useful tool for both general definition and practical research.

The Taxonomy of Human Interruption (Table 11 on pg 129) is a practical synthesis of the theoretical constructs identified in the analysis chapter. This taxonomy describes and categorizes the useful dimensions of theoretical ideas for practical investigations of human interruption.

3.4.5 Generalizability of the Definition for Various Fields of Research

It is possible to create theoretical tools that allow researchers to share and generalize well-controlled research results across domain boundaries. I claim that the definition I present here is generally useful. I support this claim by showing how this general definition of human interruption can be applied to each one of the various research domains analyzed in the analysis chapter of this paper. I restate the general definition of human interruption ("human interruption is the process of coordinating abrupt change in people's activities") in terms of each field of research and give examples of how this general definition can help researchers form good questions.

Colloquial Meaning (see Section 2.4 on pg 25)
(The field of etymology views interruption as "a word of the English Language.")
The general definition presented here is a new, alternative definition for the common English meaning of interruption. It could help etymologists investigate potential improvements to existing dictionaries. It is possible that this general definition of interruption describes useful relationships of meaning that are not yet embodied in English dictionaries. For example, the general definition promotes the idea of interruption as a process of coordination.

Multitasking in HCI (see Section 2.5 on pg 28)
(This field views interruption as "an unanticipated request for task switching during multitasking.")
Interruption is the process of coordinating task switching in the human activity of multitasking in HCI. The general definition could help researchers form good questions about the process of coordinating task switching in multitasking and how people's multiple activities are affected by interruption.

Multitasking in Linguistics (see Section 2.6 on pg 36)
(This field views interruption as "an unanticipated request for topic switching during asynchronous parallelistic human-computer interaction.")
Interruption is the process of coordinating asynchronous parallelistic topic switching during the human activity of linguistic interaction in HCI. The general definition could help researchers form good questions about the process of coordinating topic switching in people's asynchronous parallelism and how people's parallel linguistic interactions are affected by interruption.

Multitasking in Situational Awareness (see Section 2.7 on pg 41)
(This field views interruption as "an event that threatens the delicate balance between situational awareness and focused activity, i.e., the reception of unpredictable new data.")
Interruption is the process of coordinating multitask switching and comprehension of new, incoming information during the human activity of situational awareness during multitasking.

The general definition could help researchers form good questions about the process of coordinating task switching and comprehension of new incoming information and how the multiple activities are affected by interruption.

Management of Semiautonomous Agents (see Section 2.8 on pg 54)
(This field views interruption as "a costly side effect of delegating tasks to intelligent agents.")
Interruption is the process of coordinating interaction with a semiautonomous agent during any other human activity. The general definition could help researchers form good questions about the process of coordinating intermittent interaction with a software agent and how people's agent management and their other activities are affected by interruption.

Human-Human Discourse (see Section 2.9 on pg 61)
(This field views interruption as "an example of human-human discourse that can be represented and analyzed with the theory of discourse analysis.")
Interruption is an example of a coordination process in the human activity of human-human discourse. The general definition could help researchers form good questions about the process of coordinating interruption in discourse and how people's discourse is affected by interruption.

Human-Human Dialogue (see Section 2.10 on pg 69)
(This field views interruption as "a very common and normal part of human-human dialogue behavior.")
Interruption is an example of a common coordination process in the human activity of human-human dialogue. The general definition could help researchers form good questions about the process of coordinating interruption in dialogue and how people's dialogue is affected by interruption.

Psychology of Human Attention (see Section 2.11 on pg 71)
(This field views interruption as "the method by which a person shifts their focus of consciousness from one processing stream to another.")
Interruption is the process of coordinating shifts of people's processing streams in their focus of consciousness during the human activity of attention. The general definition could help researchers form good questions about the process of coordinating shifts in people's focus of consciousness and how people's attentional activities are affected by interruption.

A Metaphor of Cognitive Momentum (see Section 2.12 on pg 76)
(This informal perspective views interruption as "something that extinguishes a person's cognitive momentum when they are performing concentrated work on a complex task.")
Interruption is the process of coordinating transitions between different human tasks. The general definition could help people form good informal questions about the process of transitioning between tasks and how shifting between activities affects people's cognitive momentum.

Social Psychology of Conversation (see Section 2.13 on pg 77)
(This field views interruption as "a violation of people's conversational rights.")
Interruption is the process of coordinating the control of the human dialogue process during the human activity of adjusting people's social status. The general definition could help researchers form good questions about the process of coordinating the control of people's social status and how people's activity of managing social status is affected by interruption.

Interactional Sociolinguistics of Politeness (see Section 2.14 on pg 80)
(This field views interruption as "a face-threatening act.")
Interruption is the process of coordinating FTAs (face-threatening acts) in ways to mitigate the severity of those threats during the human activity of linguistic interaction in social contexts. The general definition could help researchers form good questions about the process of coordinating FTAs and how people's activities of satisfying face-wants are affected by interruption.

Simultaneous Speech in Linguistics (see Section 2.15 on pg 92)
(This field views interruption as "a disruptive type of simultaneous speech.")
Interruption is the process of coordinating simultaneous speech during the human activity of human-human linguistic interaction. The general definition could help researchers form good questions about the process of coordinating simultaneous speech and how people's linguistic activities are affected by interruption.

Language Use in Linguistics (see Section 2.16 on pg 95)
(This field views interruption as "a proposal for an entry into or exit out of a joint activity.")
Interruption is the process of coordinating transitions between stages in people's human-human joint activities. The general definition could help researchers form good questions

about the process of coordinating transitions between stages in people's joint activities and how people's joint activities are affected by interruption.

3.5 Taxonomy of Human Interruption

The Definition of Human Interruption presented above is useful as a tool for forming meaningful questions about human interruption. This section present a practical new tool for answering those questions — The Taxonomy of Human Interruption.

Interruption of people is a complex process, and investigating human interruption can be difficult. The context of focus for this dissertation is the interruption of people in human-computer interaction. System designers can not do a good job of specifying system requirements with existing theoretical tools about human interruption. It is difficult to objectively describe all the significant parts of this process and their interrelationships. This difficulty makes it hard to create good user interface designs for systems that interrupt their users.

The taxonomy identifies the most useful dimensions of the problem. These dimensions or factors each describe a crucial aspect of the human interruption phenomenon that can stand alone and serve as a handle for gripping the problem from a useful perspective. They are each useful because they identify a uniquely different perspective for viewing human interruption. To be useful in this way, the description of a factor must not depend on other perspectives of other factors, but must be self-contained. These factors are orthogonal to each other.

Each factor of the taxonomy affords the application of some useful body of existing literature for addressing problems of human interruption. The factors were not chosen with some idealistic notion of how the problem should be broken into categories. Instead, the taxonomy was constructed from the theoretical foundation created in Chapter 2. The set of theoretical constructs identified and discussed in Chapter 2 expose the useful veins of existing theory that are relevant to this problem. The Taxonomy of Human Interruption is a highly concentrated summary of what is most generally useful from that theoretical foundation. Each factor of the taxonomy represents an independent perspective for looking at the problem from some foundation of existing work.

Table 11 — Taxonomy of Human Interruption

Descriptive Dimension of Interruption	**Example Values**
Source of Interruption	self [human]; another person; computer; other animate object; inanimate object. (See Table 12 on pg 132 for examples.)
Individual Characteristic of Person Receiving Interruption	state and limitations of personal resource (perceptual, cognitive, and motor processors; memories; focus of consciousness; and processing streams); sex; goals (personal, public, joint); state of satisfaction of face-wants; context relative to source of interruption (common ground, activity roles, willingness to be interrupted, and ability to be interrupted). (See Table 13 on pg 133 for examples.)
Method of Coordination	immediate interruption (no coordination); negotiated interruption; mediated interruption; scheduled interruption (by explicit agreement for a one-time interruption, or by convention for a recurring interruption event). (See Table 14 on pg 133 for examples.)
Meaning of Interruption	alert; stop; distribute attention; regulate dialogue (metadialogue); supervise agent; propose entry or exit of a joint activity; remind; communicate information (illocution); attack; no meaning (accident). (See Table 15 on pg 134 for examples.)
Method of Expression	physical expression (verbal, paralinguistic, kinesic); expression for effect on face-wants (politeness); signaling type (by purpose, availability, and effort); metalevel expressions to guide the process; adaptive expression of chains of basic operators; intermixed expression; expression to afford control. (See Table 16 on pg 136 for examples.)
Channel of Conveyance	face-to-face; other direct communication channel; mediated by a person; mediated by a machine; meditated by other animate object. (See Table 17 on pg 137 for examples.)
Human Activity Changed by Interruption	internal or external; conscious or subconscious; asynchronous parallelism; individual activities; joint activities (between various kinds of human and nonhuman participants); facilitation activities (language use, meta-activities, use of mediators). (See Table 18 on pg 138 for examples.)
Effect of Interruption	change in human activity (the worth of this change is relative to the person's goals); change in the salience of memories; change in awareness (metainformation) about activity; change in focus of attention; loss of willful control over activity; change in social relationships; transition between stages of a joint activity. (See Table 19 on pg 139 for examples.)

Note: see the following tables for examples of the categories of analysis afforded by this taxonomy.

The Taxonomy of Human Interruption has eight different factors, but a taxonomy can not contain everything. As a taxonomy, it identifies the important theory-based orthogonal perspectives for working on the problem, but there can be other important aspects of the problem. The topic of time is an example. Time can be an important in investigations of human interruption, but it is not included as a factor of the taxonomy. This because it is not orthogonal to the other factors. Time does not have meaning in and of itself, and is only meaningful within the context of other factors.

There are two ways in which time may be important for human interruption: (1) the timing of interruptions for causing transitions between different tasks in a multitask; and (2) the differential effects of varying duration of interruptions. These concerns are threads that run throughout all of the eight factors of the taxonomy and touch many aspects within the problem of human interruption. However important the thread of time is, it does not stand alone as a separate perspective.

The taxonomy is not just a list of important threads for the problem of human interruption. It is instead a compressed summary of theory-based perspectives. Time has no existing orthogonal theoretical foundation and so can not be included as a separate factor of the taxonomy. Each domain of relevant literature treats the concerns of time only in terms of its effects within a limited perspective. if there were a general theory of human time that explained the concepts of time for all human behavior then perhaps it would make sense to make time a separate factor of this taxonomy. This dissertation has not discovered such a theory in the current literature.

The Definition of Human Interruption and The Taxonomy of Human Interruption together are useful in three ways: analysis; prediction; and generalization. The factors of the taxonomy each provide a way to apply a particular perspective to problems of human interruption. If a problem can be grasped by the perspective of one factor of the taxonomy, then that factor's theoretical foundation can be brought to bear for analyzing the problem. For example, the taxonomy's factor "Human Activity Changed by Interruption" can be useful for analysis. If a problem of human interruption can be viewed from the perspective of human activity then the existing joint activity theory from linguistics (Clark 1996) can be invoked. Clark (1996) has worked out several useful mechanisms for investigating joint activities. These include signaling theory (see Table 7 on pg 102), and the structure of joint activities (see Table 6 on pg 99). The taxonomy provides a convenient way to use these mechanisms to analyze problems of human interruption.

The definition and taxonomy created in this chapter are also useful for prediction and generalization. The "Method of Coordination" factor from the taxonomy gives four examples. These are the four know solutions for coordinating human interruption. The Taxonomy of Human Interruption can be used to predict that there will be important differences in human behavior related to these different coordination methods. For example, the taxonomy affords a prediction that people's performance of computer-based multitasks will be differentially affected by each of the different interruption coordination methods. A user interface could be implemented with any one of the four methods as different solutions, and the taxonomy predicts that this design decision is critical. Chapter 5 is an empirical study to see whether this prediction is in fact useful in the real world.

Generalization is the task of combining interdisciplinary findings or mechanisms into something more useful than the pieces in isolation. Each factor of The Taxonomy of Human Interruption is an orthogonal perspective, however that does not mean that the factors do not interact. They share common threads. The thread of time is an example. A particular thread from one factor that is found to be important for a problem of human interruption can be used to bridge the gap to another factor of the taxonomy. Bridges between factors of the taxonomy allow investigators to combine the different theoretical foundations of the respective factors, and apply a new combined set of interdisciplinary tools that is more useful than any one could be in isolation.

The taxonomy can also facilitate generalization within a single factor. The theoretical foundations behind each factor of The Taxonomy of Human Interruption do not all represent single domains from the current literature. For each factor, the taxonomy brings together a set of relevant existing work under an umbrella of a single perspective. These umbrellas facilitate generalization of interdisciplinary findings that share common perspectives. For example, Chapter 4 is an in-depth discussion of the existing literature relevant to the "Method of Coordination" factor from the taxonomy. This factor serves as an umbrella to existing works from several different domains of research. Chapter 4 is facilitated by this umbrella structure of the taxonomy to discuss these interdisciplinary works in a combined way. The general discussion and comparison becomes something more useful to the problem of human interruption than any of the existing works in isolation.

3.5.1 Examples from the Taxonomy of Human Interruption

The following tables provide examples of interruptions for each of the dimensions in the Taxonomy of Human Interruption (Table 11):

Table 12 — Source of Interruption

Source of Interruption	**Example**
self [human]	While writing, I realize that I need to stop and rephrase my last sentence.
another person	Answer ringing telephone and hear "Hello Ms. Jones, you have been preapproved to receive our gold MasterCard!"
computer	While trying to save a document, the computer presents a modal dialogue box that says, "Not enough memory to complete last command. Please quit applications or close windows."
other animate object	Your dog walks over and presents its empty water bowl.
inanimate object	Avalanche.

Table 13 — Individual Characteristics of Person Receiving Interruption

Individual Characteristic	Example
personal resource	I am trying to remember a phone number. However, I am interrupted with something else that also requires my memory resources — the original phone number is lost.
sex	A female is interrupted by a male. She feels that her gender role does not give her the same right as a male to aggressively fight off the interruption.
goals	I am talking long distance to a hotel reservation clerk. We have the joint goal of getting me a room reserved and prepaid. My call-waiting clicks, telling me that someone is trying to interrupt me. I ignore the interruption — I decide that my joint goal with the hotel clerk is too important to risk failure because of an interruption.
face-wants	I perceive that I am very busy working. My interruptor understands that my negative face-wants (not to be impeded) are especially high now, so they design the expression of their interruption to provide me good redress for their threat to my negative face. “I don’t suppose that I might interrupt you now.”
context relative to source	I am much more willing to be interrupted by my supervisor at work than I am willing to be interrupted by a total stranger.

Table 14 — Method of Coordination

Method of Coordination	Example
immediate interruption	I am at the grocery store buying food at the check-out, and the all the store’s power goes out.
negotiated interruption	I am interrupted while writing. Stan walks into my office and says, “Excuse me, I need to talk to you.” I have four possible responses to Stan’s proposal for entry into a joint activity: accept, accept with alteration, reject, or withdraw. I respond, “Just give me a minute to finish my thought.... OK?”
mediated interruption	Sarah wants to interrupt the Chinese Commodities Office for information. She calls her secretary on the intercom, “Please call the Chinese Office and ask them for their current price on rice.”
scheduled interruption (explicit agreement)	“I’ll meet you for lunch tomorrow at 12 o’clock outside of Tony’s restaurant.”

Table 14 — Method of Coordination (Continued)

Method of Coordination	**Example**
scheduled interruption (convention)	"We'll meet in this conference room at 1:30 pm the first Monday of every month."

Table 15 — Meaning of Interruption

Meaning of an Interruption	**Example**
Alert	
divert attention, i.e., switch the processing stream in a person's focus of consciousness	A driver stopped at a traffic light and "beeps" their horn to divert the attention of the driver in front of them to the green traffic light.
warn	"Duck!"
announce the occurrence of an event	"Ladies and gentlemen, the show is about to begin."
Stop	
arrest perception	Turn off the lights
arrest cognition	Say to yourself, "It's only a movie."
arrest external behavior	In the SAT exam, "Time is up — everyone put down your pencil."
Distribute Attention	
multitasking	A person has adopted a policy of switching between tasks in order to perform more than one task at a time, "It's time to work on task 'x' for a while."
maintaining situational awareness	A person has adopted a policy of switching their attention between contexts in order to maintain awareness of several things at a time, "It's time to see how 'x' is doing."
Attack	
influence social relationships	Assert dominance over the current speaker in a social group by wresting the groups attention from the current speaker to oneself.
Metadialogue for Dialogue Regulation	
metacommunication for dynamically adjusting an activity to maintain appropriateness of efforts within a changing context	I'm telling a story to another person. I suddenly realize, based on my listener's facial expression, that I have told them this same story before. So I interrupt myself and apologize.

Table 15 — Meaning of Interruption (Continued)

Meaning of an Interruption	Example
facilitate speed of interaction	"The doctor will see you now."
request a turn in a conversation	I want to ask a question of an employee of a hardware store who is currently talking to another customer, "Excuse me for interrupting, may I speak with you next?"
dialogue regulation	Radio talkshow, "Excuse me, Mr. Jones, but we only have 30 more seconds of air time."
begin or end a conversation	"Excuse me, I need to talk to you."
Supervise Agent	
request for supervision or coordination by an agent to a supervisor	"Here is the report you requested."
request for delegation or coordination by a supervisor to an agent	"Calculate the Robertson's total federal taxes for 1996."
commanding	"Drop and give me 20 push-ups."
get progress report	"How is your dissertation coming?"
Manage a Joint Activity	
propose an entry into a joint activity and communicate the proposer's identity, ability, willingness, and need for mutual belief	Ted shows up in person at my office and asks, "Frank, I need to talk to you. Do you have a minute right now?"
propose an exit from a joint activity	"Excuse me, but I need to leave now to go to the dentist."
Reminder	
satisfy a prearranged interruption event of either explicit agreement, or convention	Jason and I prearranged a meeting in his office today at 2 p.m. I arrive in his office at 2 p.m., "Hi, shall we begin?"
Illocution	
education	Two students are talking during physiology class. The professor walks over and talks in their faces, "The human hand has three bones in each finger ...."
Accident (no meaning whatever)	

Table 15 — Meaning of Interruption (Continued)

Meaning of an Interruption	Example
passing in a hallway	Two pedestrians are approaching each other in a hallway. They accidentally both choose the same side of the hallway to pass each other. Their progress is interrupted until they can coordinate how to get past each other.

Table 16 — Method of Expression

Expression Of An Interruption	Example
Physical Expression	
verbal (vocal verbal)	Say, "Excuse me, please — I need to interrupt you now."
paralinguistic (vocal nonverbal)	Clear throat loudly, "ugh um."
kinesic (nonvocal nonverbal)	Make the time out "T" signal, like in a football game.
Expression for Effect on Face-Wants	
bald-on-record nonpoliteness	The airplane blasts a loud buzzer to alert the pilot of a fire (note: see also Table 5 on pg 87 for examples of politeness expressions).
positive politeness	"Hey, your idea worked great! Let me tell you what happened ...."
negative politeness	"I would be very grateful if you would allow me to interrupt you now."
off-record politeness	Stare at the person you want to interrupt.
Signaling Type (by purpose, availability, and effort)	
Description	Spoken words usually easy in face-to-face interaction — "I need to interrupt you now." (Note: see also Table 9 on pg 105 for examples of signaling expressions.)
Indication	In a noisy room and several yards away, indication can be easier than description — point at a person to get their attention and indicate that you want to interrupt them
Demonstration	In relating a message to someone, it can be much easier to demonstrate than to try to explain the speakers tone of voice — "Your mom said, 'get William off the telephone.'" [impersonation of mother's angry voice.]
Metalevel Expressions to Guide the Process	

Table 16 — Method of Expression (Continued)

Expression Of An Interruption	**Example**
back-channeling	While I talk to someone, they frequently say things like, “Yea,” “I see,” “Um hm.” These interruption tell me that my listener understands what I’m saying.
regulation of turn-taking	While I talk to someone, they say, “but....” This interruption tells me that the listener wants a turn to speak.
directing attention	I am asking Jane where she left my diskette. In the middle of my speech act, Jane points to a drawer.
select listener	While I am talking to Randy, Bill says, “Daniel” [my name]. Bill has interrupted me to select me as the listener for a joint project he wants to enter.
Adaptive Expression of Chains of Basic Operators	
express chains of basic operators — dynamically planned in parallel and expressed serially	I constantly monitor the dynamic effect caused by my communication act as I express it to my audience. I dynamically replan the chain of basic operators as I execute them to adapt my planed meaning and expression to conform to the changing context of interaction.
Intermixed Expression	
composite signals	Simultaneous: hand chopping gesture (as a demonstration signal to mean a request for interruption); “Please excuse me, will you hand me that?” (as a description signal); and eye gaze fixed to the object I want (as an indication signal).
Expression to Afford Control	
affordances for choices in responses to interruption	Express interruption to allow interruptee to choose among: accept, accept with alteration, reject, or withdraw.
affordances for defenses to interruption	Express interruption to allow interruptee to fight off interruption if they want to reject it. Allow them to: mark the interruption as a violation of their right to speak, eclipse the interruption, or ignore the interruption

Table 17 — Channel of Conveyance

Conveyance of an Interruption	**Example**
face-to-face	Walk up and present myself in person and say, “Excuse me.”
other direct communication channel	Call on an intercom system and say “Hi, this is Robert, I need to talk to you for a moment.”

Table 17 — Channel of Conveyance (Continued)

Conveyance of an Interruption	**Example**
mediated by a person	I'm giving a guest presentation at a new place, and I ask my host to quiet the audience because I'm ready to begin.
mediated by a machine	I leave a message for you on your answering machine
meditated by other animate object	A burglar has broken into my house — I send the dog to interrupt him.

Table 18 — Human Activity Changed by Interruption

Human Activity Changed by Interruption	**Example**
individual activities	I am drawing a picture of a landscape on a piece of paper.
internal	[drawing example, cont.] I do more than one internal activity simultaneously: I perceive my environment; I compare proportions of objects as I set them down on paper; I interpret the colors; I decide which aspects of the scene to emphasize; I day dream about a movie I saw.
external	[drawing example, cont.] I hold the paper in place with one hand and make a mark with the pencil in my other hand.
conscious	[drawing example, cont.] I consciously attend to the relative size of one object in the foreground.
subconscious	[drawing example, cont.] I subconsciously attend to many things: I listen for sounds of approaching animals; I monitor how much sun I'm getting; I evaluate the realism of the marks I make; I monitor how fatigued my back muscles are becoming; I shift my eye gaze to different areas of the scene.
asynchronous parallelism	[drawing example, cont.] I cycle the focus of my external actions to project my several internal activities one at a time. I accomplish this with a pattern of sharing my focus of consciousness among several processing streams.
Joint Activities	
human-human	I am at the check-out counter of the grocery store buying food. I create a joint activity with the check-out clerk.
human-computer (one of each)	I am writing an article on a microcomputer using word-processing software.
human-computer (more than one of each)	I work in a group of people on a network of computers using a digital video conferencing system and computer-supported cooperative work (CSCW) groupware.

Table 18 — Human Activity Changed by Interruption (Continued)

Human Activity Changed by Interruption	Example
Facilitation Activities	
language use	The act of using language is itself a kind of activity. The activity of language use allows us to perform other joint activities.
meta-activities	While I am interacting with another person, he or she and I are also performing the joint activity of maintaining common ground. Jim and I are performing the joint activity of social conversation during lunch. We make back-channeling communication acts during our conversation as a meta-activity to facilitate the success of our conversation.
use of mediators	The presidents of the U.S. and Russia are meeting. They employ translators to facilitate communication between English and Russian.

Interruptions can have good and/or bad effects. The worth of the effects of interruption are relative to the goals of the participants. For example, an airplane pilot has the goal of not crashing the airplane. If a pilot becomes completely focused on the task of fixing a broken light bulb, then an appropriate interruption would be good. An interruption could change the pilot's light-bulb-fixing activity and shift their attention to the airplane's altitude. This effect could save the lives of everyone on board. This interruption helps the pilot accomplish their goal of not crashing.

However, interruptions can also be bad. If a pilot is interrupted during the act of physically landing the plane, they could crash. An interruption of an accidental knock on the head could change the pilot's activity of landing the plane. This interruption prevents the pilot from accomplishing their goal of not crashing.

Table 19 — Effect of Interruption

Effect of Interruption	Example
change in human activity	An interruption substitutes one activity for another. The value of this substitution depends on whether it advances the participants' goals.

Table 19 — Effect of Interruption (Continued)

Effect of Interruption	Example
change in the salience of memories	An interruption reduces the salience of some memories and increases the salience of others. This change can help and/or hinder participants in accomplishing their goals.
change in awareness (metainformation) about activity	People maintain subtle metainformation about their activities. They use this metainformation to dynamically adjust their actions. However, if a person is interrupted and changes activities, they can become disconnected from metainformation about their previous task. This loss of awareness about the progress of the first task can result in a lag of dynamic behavior when resuming it.
change in focus of attention	An interruption switches a person's processing stream in their focus of consciousness. The worth of this change depends on the participants' goals.
loss of willful control over activity	People have "back doors" to their attention resources. An interruption can be expressed so as to immediately change a person's activity without their willful decision to allow it.
change in social relationships	Interruption can be a sign of social power. If a person allows themself to be interrupted, the interruptor can be perceived by all participants to have exerted social control.
transition between stages of a joint activity	Interruption is the mechanism for bridging the boundaries of joint activities. An interruption is a use of language to coordinate such a transition.

Two appendices are provided as aids for identifying domains of literature relevant to particular research contexts. Appendix A presents an index of domain perspectives of interruption and can be useful for identifying relevant fields of research. Appendix B presents an index of theoretical constructs of interruption and can be useful for identifying common concepts across domains.

3.6 SUMMARY

This chapter presents new general theoretical tools for researching human interruption -- a general unifying definition and taxonomy of human interruption. These generally useful and unique theoretical tools were synthesized from the significant and generally relevant theoretical constructs identified in an extensive analysis of current literature. The breadth and depth of this analysis, and the resulting unified definition's strict simplicity, make the theoretical products of this chapter powerful tools for guiding general research about human interruption. The general utility of these tools is maximized through the author's extensive analysis of each of the following domains of current literature: psychology, human cognition, linguistics, social psychology, socio-linguistics, cognitive modeling, human-computer interaction (HCI), artificial intelligence (AI), computer-supported cooperative work (CSCW), and computer science.

The general definition of human interruption is, "human interruption is the process of coordinating abrupt change in people's activities" (p. 67). Each part of this definition expresses an essential concept of human interruption in a way that is generally applicable. The four pieces of this definition say that human interruption is: (1) a process; (2) coordination; (3) abrupt; and (4) a change in people's activities. First, interruption is an entire process. It involves the who, what, where, when, why & how of each stage in the generation, transmission, reception, comprehension, and reaction of an interruption. Clark (1996) says that "language cannot be separated from its context of use." Therefore, any definition of interruption must include reference to the total process of interruption.

Second, interruption is coordination. Malone and Crowston (1994, p. 90) propose a general theory of coordination — "coordination is managing dependencies between activities." Interruption of people is a complex process composed of many subactivities with dependencies and interdependencies that must be managed. Third, interruption is abrupt. William Edmondson, in his paper about asynchronous parallelism in human behavior, says that people dynamically modify and change their behavior while they are making it (Edmondson 1989). Edmondson's article portrays interruption as something that happens quickly and dynamically in real time. Fourth, interruption is coordinating change. People have to be doing something before they can be interrupted to start doing something else. Herbert H. Clark, in his book "Language Use in Linguistics," says it is useful to model human-human behaviors [like interruption] as joint activities which have stages that change over time (Clark 1996).

The Taxonomy of Human Interruption identifies the most useful descriptive dimensions for investigating human interruption (p. 73): (1) source of interruption; (2) individual characteristic of person receiving interruption; (3) method of coordination; (4) meaning of interruption; (5) method of expression; (6) channel of conveyance; (7) human activity changed by interruption; and (8) effect of interruption.

CHAPTER 4:
A LITERARY FRAMEWORK

4.1 AIMS AND OBJECTIVES

The usefulness of the new Taxonomy of Human Interruption is validated in part by demonstrating its power in structuring a survey of current literature about human interruption during HCI. After reading this chapter, the reader should have experienced the utility of the Taxonomy of Human Interruption for structuring literary survey. The reader should also be familiar with the surveyed literature and understand how the different works interrelate and generalize.

4.2 OVERVIEW

It is useful to structure the survey of literature about user-interruption by computer in HCI according to the dimensions of the Taxonomy of Human Interruption. The Taxonomy of Human Interruption has a unique power for generalizing findings across different fields of research. One utility of a general theory is to help identify commonalities in research about similar topics. These identified dimensions of the human interruption problem can also aid in describing specific instances of user-interruption and in making hypotheses of user interface designs which may mitigate the costly effects of user-interruption. (The dimensions of the Taxonomy of Human Interruption are (p. 73): (1) source of interruption; (2) individual characteristic of person receiving interruption; (3) method of coordination; (4) meaning of interruption; (5) method of expression; (6) channel of conveyance; (7) human activity changed by interruption; and (8) effect of interruption.)

Background relevant to the taxonomy's dimension number 3, method of coordination, is the most relevant to the next chapter ("A Tool for Empirical Research," pg. 173) and will be discussed last.

A broad survey of literature about human interruption during HCI is brought together for the first time in a generalizable way. The Taxonomy of Human Interruption provides a common framework for discussion of diverse works. Without this taxonomy to structure this survey, the commonalities of many of these works would not be apparent.

4.3 MOTIVATION

The claimed utility of the new theoretical tools created in this dissertation must be validated.

4.4 SOURCE OF INTERRUPTION

The computer in HCI can be an *external source of user-interruption* (see section "Multitasking — People Performing Multiple Concurrent Activities" on pg 8, in Chapter 1 (pg. 1)). For example, a person can use the Maxims intelligent email agent (Lashkari et al. 1994; Maes 1994) to externally background the activity of repeating common email management patterns. Maxims learns its users email behavior patterns over time. After it gains experience, it calculates a confidence measure of the patterns it recognizes. If a confidence measure exceeds Maxims' "do-it" threshold then it automatically does it. If a confidence measure is below the "do-it" threshold but above the "tell-me" threshold, then Maxims initiates a message to the user with a suggestion for action. These actions performed by Maxims are potential user-interruptions.

Computers are only one example of an external source of interruption. Other sources can be other people, animals, or non-computer machines that a person uses to externally background activities. Also, some internal and external sources of interruption are not related to activities people have intentionally backgrounded, e.g., having a heart attack (internal interruption), and being struck by lightning (external interruption). People do not intentionally background their autonomic heart function or the weather.

Miyata and Norman's (1986) activity classes (described in Section 1.4.1 on pg 8) identify two different kinds of interruption: internal and external. These two kinds of interruptions come from two different sources of interruption: internal and external. Internal interruptions are side effects of internally backgrounded activities internal sources of interruption), and external interruptions are side effects of externally backgrounded activities (external sources of interruption).

4.5 Individual Characteristic of Person Receiving Interruption

People have individual differences in their ability to multitask while being interrupted. Some critical jobs, like public safety dispatch (911) and air-traffic control require the kind of people who can do these tasks reliably. Joslyn (1995) presents an empirically validated test called The Puzzle Game for predicting individuals' performance on the public safety dispatching task. The Puzzle Game is a computer game that has been designed to be as simple as possible but still require the user to do all the kinds of things that make public safety dispatching difficult. The game requires a subject to pack incoming simple geometric objects into bins by size, shape, and color. People must perform a few activities concurrently: make packing decisions; request specific information about non-packed shapes; and acknowledge receiving new shapes as they arrive. The arrival of new shapes causes user-interruptions, and information about non-packed shapes does not persist so subjects must remember the details of the packing job and the pending shapes.

Several factors of individuals have been found to affect their ability to multitask during interruptions. People's level of anxiety affects their ability to recall information about interrupted tasks (Husain 1987). People's ability to maintain a constant level of arousal affects their performance on an interrupted vigilance task (Cabon et al. 1990). Level of motivation affects: (1) people's ability to recall information about interrupted tasks (Atkinson 1953), and (2) people's tendency to resume interrupted tasks (Weiner 1965). Individuals have a degree of coordination ability which affects their ability to perform multitasks (Morrin et al. 1994). Children's individual differences in ability on conservation tasks (a child's ability to discern violations in conservation of amount) and reversal shift tasks (distinguish pattern transpositions) predict their multitask performance (Kermis 1977). People's level of apprehensiveness affects how often they initiate dialogue and how often they receive interruptions in human-human communication (Lustig 1980). People's sex affects their initiation and management of interruption in human-human communication (West 1982; Zimmerman and West 1975).

4.6 MEANING OF INTERRUPTION

Computer systems are built to interrupt their users for different reasons. Sometimes interruptions are supposed to act as reminders to help people not forget to resume activities they had suspended or backgrounded. The calendar application for the Macintosh named In Control (by Attain Corporation) initiates beeps that interrupt the user to **remind** them of scheduled meetings. Taylor and Hunt (1989) say interruption can mean **dialogue regulation**. They say that in human-human dialogue people interrupt each other as one means of regulating dialogue turns. Email applications initiate interruptions to **alert** the user of the existence of new messages.

Cars interrupt their users with beeps or even recorded voices to **warn** people when they accidentally leave the keys inside. **No meaning** is also a kind of meaning. Some interruption have no meaning other than as news that something has broken. Periodic failure of a communication channel (interruptions) have been observed to degrade Navy commanders' ability to make tactical decisions (Callan et al. 1990).

4.7 METHOD OF EXPRESSION

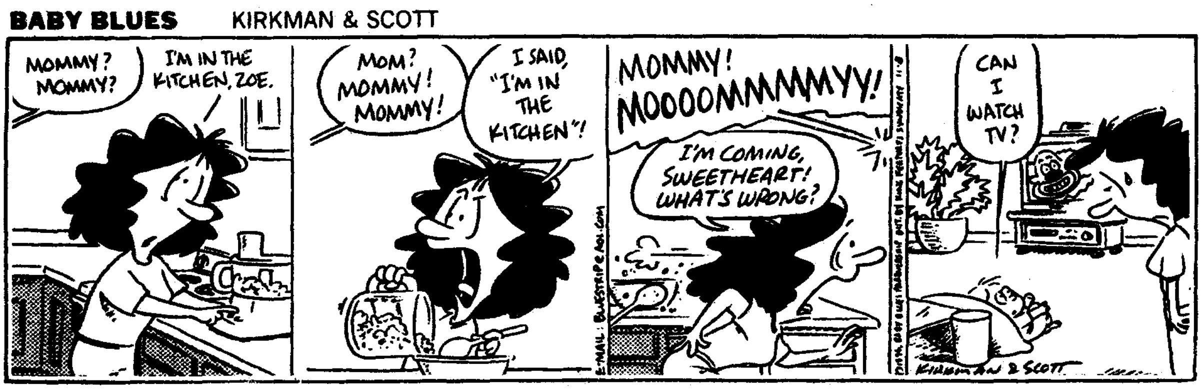

People carefully compose expressions for their communication acts because they know that other people interpret meaning from how the message is contrived. In fact a person's chosen expression for a communication message can be more important than the message itself. There are several significant aspects of expression. For example, gesticulation (people's body movements they act out as part of their communication acts) is an important source of information (Kendon 1972). People use expression of communication acts that direct the process of turn-taking in conversations (Duncan 1972; Duncan 1973). Expression also plays an important role in human-human interruption. Vocal amplitude, for example, directly affects people's ability to deny interruption in human-human communication (Morris 1971).

Investigations of tutoring have discovered the importance of expressing interruptions carefully. Galdes et al. (1991) say that deciding how to express interruptions is a critical activity for tutors, i.e., when to interrupt and what to say. They identify six factors that human tutors use to decide how to express interruptions of their human students. Galdes et al. discuss how to apply knowledge of the expressive behavior of expert human tutors to the design of computer-based tutoring systems. The six factors they identify are: (1) the tutor's goal for interrupting (to correct an error or to verify the existence of an error); (2) whether the student's error was caused by forgetting; (3) whether the student is likely to come across relevant information in the near future; (4) the tutor's expectation for normal knowledge level for this kind of student; (5) mutual understanding of appropriate contexts for making tutorial interruptions; and (6) whether the student seems likely to request help soon.

People expect that similarity in expression means similarity of meaning. This expectation is a strong influence on people's ability to multitask. The problem of psychic blindness or Einstellung (Lane and Jensen 1993) identifies a common error people make because of their expectations about similarity between expression and meaning for separate activities. Lane and Jensen say that using similar expressions for different problems can "blind" people to otherwise simple solutions. After people perform a few tasks that all require a common strategy for success, if they are presented with a different problem that looks like the previous ones they will often fail to discover an otherwise simple solution. McLeod and Mierop (1979) found that people have a difficult time switching between tasks in a multitask when the tasks require similar muscle movements for making responses. The design of the user interface (the expression of the computer system) should be contrived to provide people with changes in user interface techniques that correspond with changes in the foregrounded task of a multitask.

HCI researchers have investigated different ways of contriving the expression of system's user interfaces to support user multitasking and mitigate the costs of user-interruption by machine. Gillie and Broadbent (1989) found that the similarity between the interruption task and the current task, and the complexity of the interruption task directly affect the disruptiveness of interruptions. Storch (1992) found that interruptions expressed as on-screen messages were more disruptive to people performing a computer data-entry task than interruptions expressed as telephone calls or as human visitors. Harrison et al. (1995) demonstrated the utility of semi-transparency in windowing systems as useful for users for switching between concurrent activities. Spatial location can be an important expression choice for the user interfaces of interruption tasks. Osgood et al. (1988) compared interfaces which interrupted users with a set of numbers during a tracking task. They found people performed better when the interruption was expressed as a rapid display of numbers in the same location than when the interruption information was displayed at the same time but spatially distributed on the screen. Lee (1992) found that expressing the active window with an animated border instead of a static border reduced the number of times people became confused about which window was active when resuming a task after an interruption.

People's perception systems cannot resist certain expressions of visual interruptions (Müller and Rabbitt 1989). People have a reflexive mechanism and a voluntary mechanism in their perception system. If a visual signal is expressed as an abrupt change in lighting in a person's extrafoveal vision, then their reflexive mechanism automatically engages and their attention is automatically switched to focus on that spot.

Because of people's susceptibility to involuntary interruption, a "no expression" choice for expressing a user interface can sometimes be useful in supporting user control in multitasking. Marcus and Blau (1983) demonstrate the benefits of a user interface for an English composition task with an invisible display, i.e., subjects found it useful to not be able to see anything on the screen while they composed an essay on the computer.

4.8 CHANNEL OF CONVEYANCE

Not all kinds of communication channels support interruption. Chapanis and Overbey (Chapanis and Overbey 1974; Chapanis 1978) investigated the effects of interruption capability (whether a communication channels supports interruption) on human-human communication. Subjects were paired, and each pair was assigned to either the "free interchange condition" (subjects could interrupt each other), or the "restricted interchange condition" (subjects could not interrupt each other). Each subject in a pair was placed in a separate sound-proof room. Subjects were provided with various communication devices connected to their partner's room, and asked to cooperate together to solve a variety of joint projects. In the free interchange condition each subject had a button which allowed them to interrupt their partner's communication and seize control of the conversation. Subjects in this "free" condition could interrupt each other at any time. In the restricted interchange condition each subject had a button which allowed them, not to interrupt, but to release control of the conversation. Subjects in this "restricted" condition could never interrupt, but instead, had to wait until their partner manually released the communication channel.

In this study, Chapanis & Overbey recorded subjects' interactive communication for solving four different problems with interruption capability and without. They found that subjects under the two conditions solved problems in about the same amount of time, and used about the same number of words for task-related communication. However, subjects significantly differed in **how** they communicated. Subjects in the free condition (interruption capability) had more interchanges with fewer words per interchange; and subjects in the restricted condition had fewer interchanges with more words per interchange. It seems that in solving cooperative problems with another human, if people are able to interrupt and be interrupted (interruption capability) they will interrupt each other. Interruption capability causes people to solve problems with many short messages. However, if people cannot interrupt each other (no interruption capability), then they compensate by solving problems with fewer, but longer messages.

Karis (1991) also found that imperceptible inefficiencies in a communication channel can affect people's interruption behavior. He investigated the effect on human-human conversation of adding lag time (600 ms and 1200 ms delay between when a speaker said something and when their listener received it) in a wireless telecommunications channel. Karis found two results: (1) subjects did not notice the existence of the added lag and (2) the inclusion of delays increased the frequency with which people interrupted each other.

The channels of communication employed can affect the peoples interruption behavior. Ochsman and Chapanis (Ochsman and Chapanis 1974; Chapanis 1978) conducted an experiment in which they found that the existence of a voice channel in a communication system affects people's interruption behavior more than the existence of any other kind of communication channel. Their experimental design was very similar to the one described above from Chapanis & Overbey, i.e., paired subjects, in separate locations, are assigned to work cooperatively on some tasks by means of various communication systems. However in this experiment subjects are given two buttons instead of one for controlling the communication system. These two buttons allowed subjects to give and take control of the communication system from their partner: one button for taking control (interruption), and one button for giving control.

Ochsman & Chapanis then compared how subjects' behavior differed when provided with different combinations of communication channels. Subjects were treated with 10 different communication systems: voice + all others; voice + video; voice + handwriting; voice + typewriting; voice alone; handwriting + video; typewriting + video; handwriting + typewriting; handwriting alone; and typewriting alone. Ochsman & Chapanis found that when people have a voice channel in their human-human communication system, they take control (interrupt each other) much more frequently than they give control. However, when people do not have a voice channel they take and give control about equally often.

Taylor (1989) says that visual channels have a definite advantage over other channels for communicating spatial information. Taylor also says that pilots perform better using airplane cockpit user interfaces that include a speech/natural language capability (support for a voice channel) than with cockpit systems that do not. Computer initiated messages are better conveyed over the voice channel than the visual channel when the pilot is using their eyes for some other task. However, Taylor says that even though the voice channel is useful, if it is not implemented according to human requirements (its expression) it can cause the problematic side effect of user-interruption by machine. "Intelligent management of the priority, timing and repetition of speech is a distinctive characteristic of natural dialogues. The lack of these features in current speech technology leads to frequent ineffective machine-initiated transactions and undesirable interruptions which are difficult to ignore" (p. 265).

4.9 Human Activity Changed by Interruption

Some activities are more vulnerable to interruption than others. The frequency of interruption of group members in a computer-supported cooperative work (CSCW) activity affects group performance. Jessup and Connolly (1993) conducted an experiment with a group of subjects performing a brainstorming multitask on a CSCW system. Subjects had to alternate between the task of individually generating ideas and the task of sharing ideas with other group members, i.e., the individual task was intermittently interrupted with the group sharing task. Jessup and Connolly experimentally controlled and varied the frequency with which group members interacted. Three interruption frequencies were used: low, intermediate, and high. They found that groups that worked during intermediate and high frequency interruption performed better than groups that worked during low frequency interruptions; and that individuals in the high frequency interruption groups felt more stress (hurried, less able to concentrate, and more interrupted).

4.10 EFFECT OF INTERRUPTION

User-interruptions do not always cause critical effects. As described above, Chapanis and Overbey (1974) found that interruptions between team members did not affect their performance (measured as overall time required) on collaborative tasks. In this case, interruptions did not affect team performance time of collaborative tasks but only how subjects accomplished those tasks. Lee (1992) found that interruptions had no effect on people's probability of making unselected window errors while performing a multitask on in a windowing system.

Kreifeldt and McCarthy (1981) found (incidentally) that interruptions greatly increased the time subjects required to solve math problems with calculators. The main topic of their research was whether the "stress-tolerance" of user interface designs could be measured by using interruptions as the source of stress. They postulate that it is useful to not only evaluate the design of a user interface under normal task conditions but to also evaluate it under abnormal high-stress conditions. A user interface design's resistance to user performance degradation under stress is its degree of stress-tolerance. Kreifeldt and McCarthy chose user-interruption as a likely stressor because of its common familiarity as a source of stress in real world activities. They compared two different user interface designs that supported similar user performance levels under normal conditions: Reverse Polish Notation (RPN) style calculators and Algebraic Notation (AN) style calculators. They compared subjects performance solving math problems with each of these two different user interfaces under normal and under high-stress conditions (stress provided by injecting interruptions). They found that the RPN user interface design had much better stress-tolerance than the AN design.

Interruption also affects people's memory of their multiple activities in a multitask. Czerwinski et al. (1991) found that there is an inverse relationship between multitask similarity (similarity between foregrounded activity and interruption activity) and people's ability to remember information about the interrupted task after interruption. Subjects were given the primary task of monitoring a table of scrolling output from four space station expert systems (a simulation with mock data). The interruption task then replaced the primary task display with another (similar, or dissimilar) tabular monitoring task. All subjects performed the interruption task for five minutes, and then were asked to recall all possible information from the primary task. Czerwinski et. al. found significant support for their hypothesis that subjects who's interruption task was dissimilar to the primary task would remember more information about the primary task, than subjects who's interruption task was similar to the primary task.

Cellier & Eyrolle (1992) found empirical evidence about two kinds of effects of user interruption by machine. They observed that immediately after an interruption, peoples' task performance temporarily degrades in two ways: 1) they take longer to make decisions; and 2) they make more errors. Cellier & Eyrolle address the context of dualtasking where task-switching is externally controlled and unpredictable to the user. Subjects are given tasks of pattern matching on a one-row stream of mixed letters and numbers which scrolled from right to left. As the alphanumeric characters scrolled past, subjects use a rule to identify and mark particular patterns of two or three characters. Subjects are taught two pattern matching rules, and task switching occurs when subjects are interrupted and required to change from employing one rule and begin using the other.

Cellier & Eyrolle say that based on their observations of errors, there is a useful distinction related to user-interruption. They say there are two kinds of errors: "specific errors" which are the direct result of interruption; and "non-specific errors" which are unrelated to interruption. Cellier & Eyrolle observed three subgroups of specific errors: intrusions, confusions and omissions. Intrusions are errors where people incorrectly perform actions for the pre-interruption task after task-switching. Confusions are errors where people accidentally intermix actions from pre and post-interruption tasks. Omissions are errors where people fail to perform part of the post-interruption task.

4.11 METHOD OF COORDINATION

This final category of human interruption from the Taxonomy of Human Interruption (Table 11 on pg 129, see also (McFarlane 1997)) is most relevant to the topic of this dissertation. The method of coordination will be used in Chapter 5 (pg. 173) to form a theory-based hypothesis about user interface design for systems that must interrupt their users, i.e., that methods for coordinating human interruption can be applied to the design of the user interface to reduce the negative effects of external user-interruption on user performance of a multitask.

The Taxonomy of Human Interruption identifies four examples of methods that people use to coordinate the interruptions they receive: (1) immediate interruption; (2) negotiated interruption; (3) mediated interruption; and (4) scheduled interruption (or coordination by prearranged convention or explicit agreement). There are relevant papers in the current literature which address individual aspects of these four methods for coordinating user-interruption by machine. However, no paper has yet compared all four methods by applying them to the same context.

The earlier section (see Section 1.4.1 on pg 8) on Miyata and Norman's cognition-based activity theory (Multitasking — People Performing Multiple Concurrent Activities) describes the three dualtask activity conditions. Only dualtasks that fit into condition #3 (a foregrounded activity with an externally backgrounded activity) produce external interruptions. In the driving a car while talking example the person externally backgrounds the driving activity to the car's autonomous robotic driver when they begin or resume (foreground) their conversation. The robotic driver is an external entity that can initiate external interruptions of its user. The robotic driver will initiate these external user-interruptions whenever the driving activity (backgrounded) enters certain conditions, e.g., "nearly out of gas," or "flat tire," or "impact imminent!"

This example of a robotic driver system can illustrate the different possible user interface designs based on the Taxonomy of Human Interruption's four example methods of coordinating interruption. An "immediate interruption" solution would allow the robot to interrupt the person at any time in a way that insists that they immediately stop conversing and switch to the driving activity. A "negotiated interruption" solution would make the robot request the chance to interrupt, and then support a negotiation with the user. This would give the person control over when they dealt with the interruption — now or later or not. A "mediated interruption" solution would make the user interface so that the robotic driver could not directly interrupt the user. Instead, the robot would contact the person's mediator (e.g., intelligent interface agent) and request that the person be interrupted. The mediator would then determine when and how the user would be interrupted. A "scheduled interruption" solution would restrict the robot to interrupt on a prearranged conventional schedule, e.g., user-interruption must only take place on the hour, 15 minutes after the hour, the half hour, or 15 minutes before the hour. If the robot needs to interrupt the user at 10:07 P.M. then it would have to wait until 10:15 P.M. to initiate the user-interruption.

The "immediate interruption" solution would probably be best for this example robotic driver system. Driving activities are more important than conversation activities because driving errors have more serious consequences than conversation errors. The best user interface design would therefore attempt to minimize driving errors and ignore all other kinds of errors (including conversation errors). However, in general the different activities in other computer-supported multitasks are not so unequal and other coordination methods or combinations of them can create more successful solutions.

4.11.1 Immediate Interruption

Sometimes computer users can not coordinate interruptions they receive, but must deal with them immediately. Many of the detrimental effects of interrupting people are related to people's difficulty resuming the original task after handling the interruption. Authors of HCI research have investigated user interface design methods to support this error-prone activity of resuming previously interrupted tasks. Ballas et al. (Ballas et al. 1992; Ballas et al. 1992) discovered that user interface design significantly affected people's ability to recover interrupted tasks in the airplane cockpit. When automated activities unexpectedly fail and users must resume a previously automated activity (externally backgrounded) they experience a trouble-some initial decrease in performance called automation deficit. Ballas et al. found that employing the direct manipulation user interface design methods of low semantic distance and direct engagement allowed people to resume an externally backgrounded activity more successfully than text-based indirect design methods. Direct manipulation design methods put metainformation into user interfaces in ways that allow people to easily understand the structure and function of backgrounded activities (Shneiderman 1992).

The user interface can be designed to present information about the interrupted activity(ies) in ways that help people resume those activities more successfully than otherwise. The Notepad program (Cypher 1986) helps users not forget to resume interrupted activities by constantly displaying a list that reminds users of the existence of those interrupted activities. Other studies have investigated the utility of embedding information into the user interface to help people maintain awareness of the details of backgrounded tasks. Gaver (1989) proposes that people can gain important information from the sounds of backgrounded activities. For example, background sounds of the bottling factory floor were added to the CSCW team process control system for a remote and distributed team (Gaver et al. 1991). The previously unavailable factory sounds helped users maintain subconscious awareness of the various factory control activities they had externally backgrounded to floor workers. Robertson et al. (Robertson et al. 1993; Card and Robertson 1996; Rao et al. 1995) have successfully used peripheral information to help users maintain awareness of their location in information spaces, e.g., Cone Tree, Perspective Wall, Document Lens, Spiral Calendar, and the Hyperbolic Tree Browser. This awareness of location aids users when they must resume their backgrounded activity of navigation. Shneiderman (1992) promotes embedding location structure into menus of windowing systems for similar navigational reasons.

Smith and Hudson (1995) found that audio information can be added to CSCW systems to help people maintain awareness of the interruptibility of other team members. This is an “immediate interruption” design that helps people recover more easily from interruptions by allowing human interruptors make intelligent decisions about when to interrupt their coworkers. Smith and Hudson’s system allows people to eavesdrop on filtered versions of coworkers conversations to determine others’ interruptibility without invading their privacy. Coworkers speech is automatically reduced to non-speech signals which communicate only information about the speaker's tone of voice. This sound-based interface is less intrusive than similar video-based solutions for directly viewing coworkers to determine their interruptibility, e.g., (Li and Mantei 1992).

Gaver and Smith (1990) introduced action sounds (sonification of otherwise noiseless computer-based activities) into the CSCW system SharedARK for shared virtual environments. Users could hear not only sounds for their own actions but also the sounds for everyone else's actions too. Users found this useful for staying aware of each others activities and for locating people within the information space. Pedersen and Sokoler (1997) combine the CSCW group awareness ideas of video and audio access of team members activities with the idea of sonification for team activity awareness. Privacy is maintained by presenting only an abstraction of other team members physical and computer-based activities. Users see each other as abstract images doing abstract things. Pedersen and Sokoler have found that this is useful, but they say that building a natural and extensive abstract semantic language for activity is beyond the scope of their article.

Davies et al. (1989) discuss the merits of different user interface designs for interrupting people with reminders of backgrounded and suspended activities. Reminders help people recover from interruptions by reminding them of the existence, and sometimes the details of previously interrupted activities. Davies et al. apply theories of cognitive psychology and cognitive modeling to propose four categories of designs for reminders: normal switch, minimum switch, micro switch, and information at the fixation point. These four categories represent four different user interface designs for reminders that each require users to exert different cognitive effort to get state information about interrupted tasks. The four different designs differ in where the state information of the interrupted activity is available: (1) normal switch — off screen; (2) minimum switch — on screen but not in user's central viewing location; (3) micro switch — on screen and in the user's peripheral vision in such a way that does not require them to move their eyes to get the state information; and (4) information at the fixation point — on screen at the user's current eye fixation point. Davis et al. conclude that the inclusion of reminders is a useful design method for recovering from interruption. They also found support for their proposed categories by showing that people could more easily maintain awareness of the editing mode of a word processor when the mode information was conveyed by the cursor shape (information at the fixation point design) instead of in a separate window (minimum switch design).

From these above studies, it might seem that the best way help users recover from interruption would be to design the user interface to constantly present obvious reminders about the existence and state of interrupted activities; however, the problem is not that simple. Unfortunately, the constant portrayal of information about interrupted tasks can negatively affect people's performance on their foregrounded activity. Noy (1989) found that providing auxiliary displays for navigation-like secondary tasks in an automobile simulator caused a degradation in people's performance on the driving task. Nakagawa et al. (1993) found that monitoring the computer's handwriting recognition of live pen-based handwriting is a separate activity that distracts users and negatively affects their performance on pen-based interfaces.

One approach that does not depend on loading the display with information about backgrounded and suspended activities is to include tools that help users quickly review the state of an interrupted activity when attempting to resume it. Field (1987) compares two different user interface tools which allow people to review their interaction histories when resuming previously interrupted computer-based activities. Field presents some weak evidence that people can resume their primary task more easily after an interruption if they are provided with a selective retreat tool, and not a restrictive retreat tool. A selective retreat tool allows users to quickly see a complete history of their previous interaction with the information system. A person can use this tool when they try to resume a previously interrupted task by reviewing their interaction history, and "retreating," or jumping back to any of their previous contexts. The less powerful, restrictive retreat tool, does not show people their interaction context, and only allows them to "retreat," or jump back to the previous context, or to the main menu.

Malin et al. (1991) also say that the user interface should be designed to reorient users to previously interrupted activities when they try to resume them. If interruptions come from noncomputer sources, the machine is not necessarily able to detect when the interrupt happens. Malin et al. present a design that specifically allows users to suspend and resume activities. This way, users explicitly mark the occurrence of interruptions. The computer can then generate appropriate recovery support. Malin et al. also present a useful design to allow users to orient themselves to the current state of the system when they take over a task from a previous user. A simple log of relevant recent decisions is made easily available. This same design could be used to aid users in recovering from interruption.

4.11.2 Negotiated Interruption

Clark (1996) says that people normally negotiate human-human interruptions Unlike the “immediate interruption” method of coordinating interruption discussed above, people usually have choices of whether to allow the interruption and when to handle it. Clark says that in normal human-human language usage people have four possible responses to interruption: take-up with full compliance; take-up with alteration; decline; or withdraw (Clark 1996, p. 203-205, p. 331-334; see also this dissertation pg. 111). It is useful to design user interfaces in ways that take advantage of people's ability to negotiate interruptions. An external entity that initiates an external interruption may do so in a way that gives the user control. The interface could afford the user four options of when or whether to handle the interruption: (1) handle the interruption immediately (take-up with full compliance); (2) acknowledge the interruption and agree to handle it later (take-up with alteration); (3) explicitly refuse to handle the interruption (decline); or (4) implicitly refuse to handle the interruption by ignoring it (withdraw).

There are useful examples from commercial applications that support rudimentary negotiation of user-interruptions. Several email applications give users some level of control over when to read their incoming email messages. For example, when a new email message arrives, the program can get the user's attention by interrupting them with a signal notification like a beep and a modeless dialogue box. The user then can decide whether to allow the interruption immediately, or handle it later, or not.

One design approach is to present user-interruptions in ways that allow people to ignore them if they choose. Lieberman (1997) implements a version of this design in the Letizia autonomous interface agent. Letizia is an aid that runs in the background and makes recommendations of possibly related web pages to its user while they browse the world wide web. Letizia's interruptions do not directly interfere with users' web browsing activity. Instead, users are left to pursue their browsing activity with a normal browsing tool (i.e., Netscape), and the Letizia agent displays its suggestions in a separate but visible window. Letizia automatically loads web pages that it decides may be of interest to the user. Since these automatically loaded pages are displayed in a visible window, the user must see those changes in their peripheral vision. When Letizia initiates one of these interruptions, the user has their choice of four possible responses: (1) look at the Letizia window and decide to read that page immediately; (2) look at the Letizia window and decide to read that page later; (3) look at the Letizia window and decide not to read that page; or (4) ignore the Letizia window.

Oberg and Notkin (1992) investigated a similar design for interrupting users with error reports in a computer programming environment. Oberg and Notkin generated a Pascal editor with a dynamic code debugger that ran in the background. While people use the computer to edit their computer program the debugger continuously runs in the background. Whenever the debugger detects a programming error it interjects an error message within the code near the user's cursor position. Oberg and Notkin specifically chose a user interface design that would give users control over when or whether to address these interruptions. They created an interface that does not interfere with the user coding activity, but instead uses color to notify users of the locations of existing errors. The age of existing errors is represented by changing saturations of color and the rate of increasing saturation varies by categories of errors. The notification marker for "important" errors gets darker more quickly than those for "less important" errors. This error coding alerts users to the existence of errors, but does it in an unobtrusive way so they have control over when and whether to handle these interruptions. Oberg and Notkin do not formally compare their unobtrusive design with other more disruptive alternatives; however, they say that their anecdotal evidence endorses its usefulness.

Any design solution that implements the "negotiated interruption" method for coordinating user-interruptions must have a mechanism for getting users attention while they attend some other activity. Users must be notified of incoming interruptions because, if not, users cannot control when or whether to handle them. People's attention is vulnerable to certain kinds of stimuli (Müller and Rabbitt 1989). Shneiderman (1992, p. 80-81) says, "Since substantial information may be presented to users for the normal performance of their work, exceptional conditions or time-dependent information must be presented so as to attract attention." He presents the following techniques for getting users' attention: intensity; marking; size; choice of fonts; inverse video; blinking; color; color blinking; and audio. Preece et al. (1994, p. 100-108) also present techniques for guiding users' attention: structure information according to the perceptual laws of grouping; use spatial and temporal cues; color; and alert with flashing, reverse video, and/or auditory warnings." Visual movement within people's peripheral vision has also been found to be an effective attention-getting technique. Ware et al. (1992) found an inverse relationship between the velocity of moving iconic interruptions and people's response time in detecting and handling them.

Rich (1996) investigated the utility of using a moving hand-shaped icon as an attention getting technique for interaction with an intelligent agent. In one version of the agent interface the agent does not interfere with the user but must wave its "hand" to get the user's attention. Then the user has control over when or whether to allow the agent's interruption. People's attention is also susceptible to other people's eye gaze, i.e., people looking at each other. Kendon (1967) says that gaze-direction is one of the principle signals by which people manage interruption in human-human communication. For social reasons, people are predisposed to attend to any occurrence of another person looking at them.

Although it is useful in some ways to give users control over when and whether to handle interruptions this is not a total solution to the user-interruption problem. One effect of interruption is to disrupt peoples' memories of the details of pre-interruption tasks. It may seem reasonable to hypothesize that this negative effect is caused by people being caught off guard and being forced to begin the interruption task without first rehearsing that information of the pre-interruption activity critical to its successful resumption after interruption. If this hypothesis were justifiable, then a negotiation design solution would successfully avoid this negative effect of interruption. Unfortunately, this hypothesis has not yet been empirically supported. Gillie and Broadbent (1989) found that allowing users to review their foregrounded activity previous to handling interruption, did not necessarily help them recover that activity after interruption. They observed that the disruptive effects of interruption on peoples' memories are not caused by peoples' inability to rehearse their memories before handling interruption. They instead found that this negative effect of interruption on memory is caused by memory interference created by interruption tasks that are complex or similar to the pre-interruption task.

Katz (1995) found that negotiation design solutions have disadvantages and that users can sometimes prefer "immediate interruption" user interface designs. He compared two different interfaces for a kind of telephone Call Waiting called Caller ID on Call Waiting (CIDCW). CIDCW gives people information not only of the existence of incoming calls while they are talking with another person, but also of the new caller's name and phone number. Katz conducted an experiment that compared two different user interface designs for the telephone CIDCW system: (1) automatic interruption (an "immediate interruption" solution); and (2) user-controlled (a negotiated interruption solution). The automatic interruption interface caused immediate break of what the user could hear. A beep and then the information of the new caller (total 1.1 second) occluded what they could hear, then the audio connection with their original conversant was restored. The original conversant was unaware that a break had occurred. The user-controlled interruption interface announced the existence of a new call with a beep, and then the user had to press a button to hear the caller ID information. Katz found that subjects preferred the automatic interruption interface 3 : 1 over the user-controlled interface. They said that the user-controlled interface was much more disruptive of their telephone conversation than the automatic interface.

Katz says that the automatic interface and the user-controlled interface design solutions for CIDCW systems each have advantages and disadvantages. The advantages of the automatic interface are: (1) users do not need to take any action to receive caller ID data; (2) users do not have to learn anything new to use the interface; and (3) users do not have to formally break their conversations and excuse themselves to the other party to get the caller ID information. The automatic interface, however, has two disadvantages: (1) people's conversations can be unexpectedly suspended for a second; and (2) people know that they could be interrupted at any time regardless of what they are saying. The user-controlled interface has the advantage of not unexpectedly blanking out chunks of peoples' conversations or causing uncertainty in users' expectations. However, the user-controlled interface has the following disadvantages: (1) users might need to formally break their conversation to hear the caller ID information; (2) users have to learn a new interface; and (3) users have to take specific action, and therefore might postpone doing it so long that the new caller gets tired of waiting and hangs up.

4.11.3 Mediated Interruption

The White House Communications Agency (WHCA) provides the President of the United States (and his or her associates) the capability to make public speeches anywhere. There is a critical human interruption problem that can affect WHCA's ability to successfully announce the President and the other dignitaries at these public meetings. The WHCA employs a mediator to solve this problem. Whenever the President schedules a public speech, the WHCA sends a team in advance to prepare the site. They must either set up a public address system or contract one locally, arrange the President's special podium and Teleprompter, and prepare a ready communication link out. One of the WHCA team is designated to sit in a van out of sight and announce the President and the other VIPs. The introduction must be done right the first time because the professionalism of the introduction sets the stage for how the President will be received. The WHCA team plays Hail to the Chief (from a CD), and then the announcer says "Ladies and Gentlemen, the President of the United States, Queen Jane Janga Yyptemshephesmfttlaywa of the Kingdom of Flagmanistan, Senator Henry Joyce Jones from Virginia, ... If the announcer stammers, or mispronounces an important local dignitary's name, or fails to include someone, the crowd will notice and start laughing or worse.

The WHCA team prepares an introduction card for the announcer to read. However, when Air Force One actually lands mad chaos often begins. The introduction must be immediately changed or amended and lots of aids and dignitaries come swarming over the WHCA announcer trying to give important new instructions. WHCA has solved this problem by assigning another team member to mediate between the chaos and the announcer. The mediator allows the announcer to concentrate while still being accessible for last minute changes in a controlled way.

Adding a mediator to the user interface abstracts the HCI and is not always a good solution for an interruption problem. Delegating the interruption problem to a mediator begets a new task of supervising the mediator. Kirlik (1993) observed that the costs of delegating a task to a task-offload aid (like a mediator) can sometimes outweigh the benefits. It is possible for a poorly designed mediator to be more disruptive than the interruptions they broker.

Most research on computer-based mediators in the current literature tries to find ways to reduce the supervision costs by increasing the mediators ability to intelligently accommodate people's cognitive limitations. There are five main approaches: (1) predict people's interruptibility and use the results to intelligently time interruptions; (2) investigate new user interface methodologies for supervision; (3) automatically calculate users' cognitive workload and use the results for dynamic task allocation; (4) categorize different human and computer abilities and design supervisory control systems that exploit the different abilities of each; and (5) build and employ a cognitive model and use the results to guide user interface design process.

4.11.3.1 PREDICTING INTERRUPTIBILITY

People's degree of interruptibility (or their vulnerability to the effects of interruption) dynamically changes and is dependent on conditions of the person, their multitask, and the context. Miyata and Norman (1986) have identified several useful factors of human behavior that can be used to predict people's interruptibility: task dependency; relative priority; activity stages; user-specified interruptibility; and the difference between notification and description for reminding people of backgrounded activities. Related tasks in a multitask often have dependencies. If the computer can mirror the user's activities with a task model, then it can automatically determine when a backgrounded activity will be needed within the context of the foregrounded activity. Activities in a multitask may have different importance and the relative importance of the interruption task and the foregrounded task can be used to quantify users' interruptibility. People's activities can be decomposed into stages relative to human cognition (Norman 1986). People's interruptibility changes depending on the stage of their foregrounded activity. For example, people are more interruptible at the point where they transition between the last stage ("evaluation") and the formation of a new goal or intention (Miyata and Norman 1986, p. 278). People have a metacognitive awareness of their own interruptibility. This is why they sometimes turn off sources of interruption by shutting office doors, turning off telephones, or putting up "do not disturb" signs. There is a useful distinction between notification and description for reminding people of a backgrounded activity. People are more interruptible for a brief signal that announces the existence of an interruption than they are for the full interruption itself.

4.11.3.2 HCI FOR SUPERVISION

Intelligent interface agents are a kind of intelligent user interface (Chignell and Hancock 1988) that uses an anthropomorphic design to easily convey the idea of an intelligent assistant to the user (Lieberman 1997). One example is a telephone receptionist agent with an expert system to mediate all of a person's telephone calls (Gifford and Turock 1992). The agent makes it so that a user only has one telephone number and is accessible anytime anywhere on that one number. People sometimes use telephone answering machines or caller-id boxes as dumb versions of this kind of telephone mediation (Sullivan 1993). A straight forward use of this kind of mediation is for a user to allow the answering machine to record their messages when they are away from their telephone. However, people also sometimes use these mediators to screen their calls when they are present but unwilling to be interrupted except by specific people or topics. In computer-mediated communication (Bannon 1986) the "**talk**" facility in terminal-to-terminal communication is a source of interruption. While a user is typing a message, **"talk**" messages can intrude unexpectedly and interrupt. Bannon says that this is a poor user interface design solution. People know how to give subtle signals of their interruptibility (e.g., varying positions of a person's office door) and that this ability should be exploited for the design of systems that must interrupt people.

4.11.3.3 COGNITIVE WORKLOAD AND DYNAMIC TASK ALLOCATION

Automatic cognitive work load assessment (Gopher and Donchin 1986; O'Donnell and Eggemeier 1986) is another approach to reducing mediation costs. Authors who use the concept of work load ascribe to the idea that human brains are just another kind of machine, and that the load on these machine can be measured. In studies of work load, people are often viewed as a kind of component ("man in the loop") to be used in constructing important systems. Berger et al. (1988) propose a measure of work load to be used to dynamically change automated assistance on continuous control tasks. Bergeron (1968) investigated the measurement of work load on tasks similar to piloting a lunar lander. Kuperman and Perez (1988) analyzed a team system for Air Force bomber missions, and used workload measurements to identify crew task "chokepoints." The work load measure can be used to dynamically allocate decision tasks between a human decision maker and computer-based intelligent decision maker. When the user has a light work load then all decisions are allocated to them, but when they become overloaded then a computer-based decision maker is invoked and begins taking over some of the person's decision-making responsibilities. Authors base their dynamic allocation on different allocation theories: queuing theory (Chu and Rouse 1979; Rouse 1977; Walden and Rouse 1978); and optimal control theory (Millot and Kamoun 1988). Mouloua et al. (1993) found that adaptive function allocation improved people's ability to monitor for system failures in simulated airplane flights.

4.11.3.4 HUMAN FACTORS FOR SUPERVISORY CONTROL

Computers are sometimes built to control physical processes that people cannot or should not control directly. When such a system controls an important process it must be supervised by a person to ensure its performance. These systems support supervisory control (Moray 1986; Sheridan 1987) and embody a kind of mediation in which the computer serves as a mediator between a person and the physical world. Sheridan (1988) categorizes human functions (human supervisory activities) and proposes that these categories be used to discover the human attention requirements of the different supervisory activities.

4.11.3.5 COGNITIVE MODELING FOR MEDIATION

If a computer could magically know everything about what a user has done, is doing, and intends to do, then it could always interrupt the user when and how they would best want to be interrupted. If such a system could be built, the mediator would become invisible and require no user supervision like in ubiquitous computing (Preece et al. 1994, p. 149-151). This is an attractive and popular solution and there are several reports in the literature of attempts to build applied models of human cognition for use in dynamic management of user interfaces for systems that support user multitasking. The Pilot's Associate program is a good example (Hammer and Small 1995). Its designers tried to use applied user models and task models to automatically infer users intentions in a the multitask of a military single-seat aircraft tactical mission. Once the Pilot's Associate had predicted what the user would want next, it would interrupt them with "appropriate" information and activities. Attempts to build such a system have not been adequately successful because of the difficulty of accurately inferring users' intentions even within this limited task domain.

Authors have applied several theoretical domains to human cognitive modeling. Some approaches emphasize the idea that the human brain is an information processing machine: Schweickert & Boggs (1984) investigate the utility of modern variants of the single channel theory from computer science; Forester (1986) examines the usefulness of a multiple resource model of human information processing; and Soulsby (1989) evaluates the utility of control theory and estimation techniques. Some approaches postulate that human cognition employs rational mechanisms and therefore other rational models can be generalized to modeling people, e.g., Navon and Gopher (1979) investigate the utility of economic theories of resource allocation. The COGNET (COGnition as a NEtwork of Tasks) model is based on a network of local goals or tasks the person must pursue (Ryder and Zachary 1991; Zachary and Ross 1991). COGNET has been applied to military multitask user interface domains: anti-submarine warfare (Weiland et al. 1992; Zachary et al. 1988; Zubritzky et al. 1989); and anti-air warfare (Zachary et al. 1993). Other authors have created models of human attention to investigate user interface design for user multitasks: managing supervisory control multitasks

(Enstrom and Rouse 1977; Pattipati et al. 1983; Tulga and Sheridan 1980); and monitoring of graphically displayed information (Senders 1964).

With so many different modeling approaches to choose from, it would be very useful to have some guidelines on how to evaluate competing models. Wickens et al. (1989) evaluates the relative utility of five different cognitive models for a helicopter flight multitask. The five models are: the fourth generation of Human Operator Simulator (HOS) model; the PROCRU model; the WINDEX model; a task network model; and Wickens' multiple resource model. Wickens says that the coding of demand level (how task performance is affected by the performance of other active tasks) is the most important question for evaluating the utility of competing models.

4.11.3.6 INTERRUPTION BY PROXY

One interesting idea for mediation that has not been applied to user interface design is that of interruption by proxy. Salter (1988) describes a method for extracting information from human experts for building expert systems. A human expert's knowledge can be recorded covertly with a version of interruption analysis. An expert is observed doing what they do best. In normal interruption analysis, the investigator interrupts the expert whenever the expert is seen to make a significant decision, and the interviewer asks them about the details of their decision. However, interrupting experts has the detrimental side effect of stopping them from their normal operations. The researcher can avoid this problem by getting a second proxy expert. A second expert in the same field is brought in and observes the first expert with the investigator. Whenever the investigator needs to interrupt the first expert to get information, they instead interrupt the proxy expert, and the proxy explains the decision processes of the first expert.

The White House Communications Agency (WHCA) uses a form of interruption by proxy for controlling the Teleprompter while the President is speaking. When the President gives a speech, he [or someday she] performs at least two activities concurrently. He delivers a speech (the currently foregrounded activity) and at the same time reads the next part of the speech from the Teleprompter (an internally backgrounded activity). There is a third activity of manually scrolling the Teleprompter, however. The President's first two concurrent activities are so demanding that he has decided to not participate in the scrolling activity. The President has decided that the WHCA will totally automate the scrolling task with no possible interaction with him once he has begun speaking; they cannot interrupt him, and he cannot give them directions. The WHCA solution is for one of their team to pretend to be the President (a proxy President), and try to scroll the Teleprompter live as the real President gives his speech. Being the proxy President is a very difficult job for several reasons: (1) the WHCA does not get the speech from the President's staff until within 15 minutes of its delivery; (2) there are several technical problems involved in preparing the Teleprompter; (3) the President dynamically changes his rate of delivery and often makes unannounced deviations from the prepared text; (4) the Teleprompter control system allows the WHCA team to only see what the President is seeing; (5) and the WHCA Teleprompter controller is not even in the same room with the President. The WHCA solution saves the President from ever being interrupted with the scrolling activity; however, speech time is high stress time for the WHCA. One WHCA team member is the proxy President, and several other team members huddle around them trying to help with the task of anticipating what the President will want to see next.

4.11.4 Scheduled Interruption

If people had foreknowledge of the when-what-where-why-and-how of incoming interruptions they could plan their other activities to minimize the negative effects of interruptions. However, to be able to know about interruption before they happen, people would need some control over the initiation of those interruptions. In fact, this kind of foreknowledge would change the status of the activity from being an interruption into being something that is not an interruption — a planned event. The user interface design solution of "scheduled interruption" can provide users with the ability to transform some future interruptions into non-interruption planned activities by giving them a kind of control over when the interruption are initiated.

One form of this control over interruptions comes from studies of time management for organizational management of people's work time. Hall and Hursch (1982) found that time management training had a large and significant effect on subjects ability to spend more time each day performing high-priority tasks. Applying the time management techniques allowed people to avoid being constantly taken away from high-priority activities (and the negative effect of interruption). Before training, one subject (a university physicist) complained that he had no time for his high-priority activities because of constant interruptions by his students working in a nearby lab. Hall and Hursch observed that this subject's average time spent on high-priority activities increased from 28 minutes a day to 2 hours and 19 minutes a day following the time management training (the experiment was over an eight week period). The subject successfully applied the time management technique of creating a daily schedule which indicated his interruptibility during different time periods in the day. He posted this schedule on his door and "scheduled" these rules for conventional interruptions with his students (although these rules had to remain somewhat flexible because of his need to participate in students' ongoing research). For example, his schedule indicated that: 8-10 a.m. was his high-priority activity time and that only 5-second interruptions would be allowed; 10 a.m. - 12 p.m. was "Quick Problems" time and interruptions of 5-minutes or less would be allowed; 1-3 p.m. was open for meeting time on demand; and 3-5 p.m. was for completing tasks and no interruption would be allowed.

Other time management professionals also promote the usefulness of this technique of scheduling dedicated time each day for performing high-priority activities (Covey 1989; Des Jardins 1998). They have found it useful for people to plan and announce their precoordinated schedule for interruptibility. This technique can automatically change some kinds of would-be interruptions into ordinary planned activities.

Clark (1996) says that people are very familiar with two useful kinds of scheduling techniques for normal human-human activities — explicit agreement and convention (see this dissertation pg. 113). Explicit agreement is a technique that people use for prearranging the coordination of a one-time event like a meeting for lunch at a particular restaurant on a particular day and time. Convention is a technique that people use for prearranging the coordination of a recurring event like a group meeting that happens in the same place and time every week. These familiar and useful methods for coordinating interruptions should be useful for solving some HCI design problems for user-interruption.

4.12 CONCLUSION

The Taxonomy of Human Interruption is shown to be useful for structuring a survey of multi-disciplinary literature relevant to the problem of human interruption during HCI. The literature relevant to one factor of the taxonomy, “Method of Coordination,” is surveyed in exhaustive depth. The four main examples of this factor are discussed in great detail. These four examples of the “Method of Coordination” factor from the taxonomy represent the four general classes of design solution for an important aspect of user interface design.

CHAPTER 5: A TOOL FOR EMPIRICAL RESEARCH

5.1 AIMS AND OBJECTIVES

The usefulness of the new Taxonomy of Human Interruption is validated in part by demonstrating its power in guiding the formation and operationalization of an important hypothesis about the effects of different methods of coordination on users' performance in HCI. A detailed human subjects experiment is conducted. Note that this experiment is not a comprehensive validation or "proof" of The Taxonomy of Human Interruption. Such a "proof" is beyond the scope of this dissertation. Instead, this experiment provides some partial support for the claimed predictive power of the taxonomy. After reading this chapter, the reader should have experienced the utility of the Taxonomy of Human Interruption for guiding empirical research. The reader should also be familiar with the results of this investigation.

5.2 OVERVIEW

The Taxonomy of Human Interruption represents a practical repackaging of the results of most of the previous research work on the human interruption problem. It provides an outline and categorization of the important factors of human interruption reported in the literature. These factors are the significant and practically useful dimensions of the human interruption problem. Therefore, the taxonomy can guide useful investigations of how to design user interfaces for systems that must interrupt their users.

The “Method of Coordination” factor from the Taxonomy of Human Interruption is chosen as an appropriate factor for the purpose of demonstrating the utility of the taxonomy for guiding research. A new hypothesis is created about the effect of implementing different coordination methods in the user interface. It is hypothesized that a difference in implemented method of coordination will disparately affect users’ behavior on a computer-based dual task. The taxonomy is used as a guide for operationalizing a human-subjects experiment to test this hypothesis.

A human-subjects experiment with 36 volunteers is reported. Each subject performed a series of computer-based multitasks for about two hours as part of a repeated measures experiment. The results are found to support the hypothesis and therefore contribute to the validation of the Taxonomy of Human Interruption.

5.3 MOTIVATION

The claimed utility of the new theoretical tools created in this dissertation must be validated.

5.4 APPROACH

An experiment with human subjects establishes empirical support for the utility of the Taxonomy of Human Interruption. The previous chapters of this dissertation establish the scope and usefulness of this taxonomy in three ways: (1) “Introduction” (pg 1) justifies the need for this theoretical tool as an instrument for general investigations of human interruption in HCI; (2) “Survey of Theoretical Constructs” (pg 23) and “Synthesis of the First Theoretical Tools” (pg 117) (see (McFarlane 1997)) establish its interdisciplinary theoretical foundation; and (3) “A Literary Framework” (pg 143) illustrates its utility as a general framework for describing existing literature about human interruption from diverse fields. Beyond these contributions, the Taxonomy of Human Interruption is also useful for guiding empirical investigation of the design of user interfaces for systems that must interrupt their users. This additional use is demonstrated in part with an human subjects experiment.

Chapter 4, "A Literary Framework" (pg 143), reveals that the problem of user-interruption in HCI is a large and complex topic. The existing literature has uncovered many more research questions than it has answered. The approach of this chapter is obviously not to try to investigate all these unanswered questions. Instead, this chapter proposes and tests a single question or hypothesis. A hypothesis was carefully chosen with the aid of The Taxonomy of Human Interruption.

A major power of the Taxonomy of Human Interruption is that its factors can be used to frame hypotheses about user-interruption in HCI. Each factor of the taxonomy identifies a separate dimension of the process. These dimensions are useful for describing user-interruption in HCI because they each focus the analysis on an especially important issue. The taxonomy can therefore be used to focus the generation of hypotheses on the most important issues for empirical research. This paper supports the validity of the Taxonomy of Human Interruption by using it to frame a specific hypothesis based on one of the taxonomy's eight dimensions. One dimension has relatively higher potential for creating generalizable results and is a better fit to the limited scope of this research.

The taxonomy's dimension of "Method of Coordination" is the best choice among factors for guiding the formation of a hypothesis. Five of the dimensions of the taxonomy identify areas for investigation that lead to results that are less obviously generalizable across all user and task contexts. These five dimensions are: (1) Source of Interruption; (2) Individual Characteristic of Person Receiving Interruption; (3) Meaning of Interruption; (4) Human Activity Changed by Interruption; and (5) Effect of Interruption. These dimensions from the taxonomy are tied to the important issues of user-interruption related to specific user and task contexts, and are therefore less suitable as choices for maximizing the generalizability of the empirical results.

The remaining three dimensions of the taxonomy are: (1) Method of Coordination; (2) Method of Expression; and (3) Channel of Conveyance. While each of these dimensions addresses general topics, the issues they address have different complexity. The Method of Expression dimension represents a complex topic with a large number of critical variables.

The investigation of important hypotheses for this topic would be beyond the scope of this dissertation. The Channel of Conveyance dimension was rejected for the opposite reason, i.e., it is not complex enough. The remaining dimension of the taxonomy, Method of Coordination, was chosen because its complexity best fit the scope of this paper and because of its general applicability.

The Taxonomy of Human Interruption provides four examples for the method of coordination factor. These four examples are the four recognized basic categories of coordination methods for dealing with human interruption. These recognized methods are: (1) immediate interruption; (2) negotiated interruption; (3) mediated interruption; and (4) scheduled interruption.

A section of Chapter 4, “Method of Coordination" (pg. 154), provides an in-depth literature survey and discussion of these four methods of coordinating interruption relative to the problem of human interruption during HCI. These four recognized coordination methods represent the four basic solutions to the user interface design problem for dealing with human interruption. Different authors have identified different individual solutions. However, no author has addressed all four solutions in a single discussion, and no author has empirically compared the relative effectiveness of these four basic categories of solutions to the problem of human interruption in HCI.

This dissertation accomplishes both of these things: (1) the first general discussion of all four solutions (see “Method of Coordination" on pg 154, in Chapter 4), and (2) the first empirical investigation of the relative effectiveness of all four solutions in user interfaces for a computer-based interrupt laden multitask.

5.4.1 Main Hypothesis

Hypothesis (H_a): the particular method for coordinating user-interruption that is implemented in a user interface will affect users’ performance on interrupt laden computer-based multitasks.[13]

5.4.2 Subhypotheses

There are several sources in the literature that indicate the plausibility of the Main Hypothesis (see Section 4.11, "Method of Coordination," on pg 154). However, this body of previous work suggests different orderings in people's performance scores for the different methods of coordinating interruption. It is not the purpose of this dissertation to take sides on predicting the directions of these conflicting perspectives. However, it is useful to discuss these different views and determine which are supported by the results of the experiment reported here.

5.4.2.1 SUBHYPOTHESES 1 & 2

Interruption causes people to switch between tasks; however, there is an overhead cost in cognitive effort that must be met at each switch. The recognized problem of automation deficit (see Section 4.11.1, "Immediate Interruption," on pg 155; and Ballas et al. 1992; and Ballas et al. 1992) describes the phenomenon that people experience an initial decrease in performance each time they switch to a new task. Task switching has overhead costs and each user-interruption by machine causes potential task switching. This perspective, together with the Main Hypothesis, suggest the following subhypotheses.

Subhypothesis 1: the total number of task switches encountered by a user is affected by which method of coordinating interruptions is implemented by the user interface.

Subhypothesis 2: there is an inverse relationship between the total number of times people switch tasks while performing a multitask and people's performance on those tasks.

5.4.2.2 SUBHYPOTHESIS 3

The interactional sociolinguistic theory of politeness (see Section 2.14, "Interactional Sociolinguistics of Politeness," on pg 80; and Brown and Levinson 1987) says that the social commerce of face-wants are responsible for deterring people from interrupting each other on

---

13. **H_0:** the particular method for coordinating user-interruption that is implemented in a user interface will not affect users' performance on interrupt laden computer-based multitasks

impulse. Instead, people go to great lengths to carefully time and package how they interrupt each other so as to not infringe on each other's face-wants. Computers do not have face, and it may be argued that people do not have face-wants relative to their computers. Face is an inherently human-human social concept. From this perspective politeness is an irrelevant topic for the design of user interfaces, and computers should always implement the "Bald on Record" type of nonpoliteness. This perspective leads to the following subhypothesis.

Subhypothesis 3: the methods of coordinating interruptions that are most direct and immediate will create user interfaces that support higher user performance on multitasks than the methods of coordination that express interruptions with delayed timing or require interactive negotiation.

5.4.2.3 SUBHYPOTHESIS 4

People negotiate entrances into joint activities with other people (see Section 2.16, "Language Use in Linguistics," on pg 95; and Clark 1996). This process of negotiation becomes overlearned as people mature so that adults automatically expect it and can manage it without conscious effort. Interruption can be viewed as a request for entrance in to a joint activity. When people are interrupted they automatically expect to be able to use their highly developed skills of negotiation for arranging when, or if, they will handle the interruption. This perspective suggests the following subhypothesis.

Subhypothesis 4: the negotiation method for coordinating interruption will create user interfaces that support higher user performance on multitasks than the other methods of coordination.

5.4.2.4 SUBHYPOTHESIS 5

People are unreliable in their ability to uniformly perform all parts of multitasks. This is the recognized problem that drives the study of situational awareness (see Section 2.7, "Multitasking in Situational Awareness," on pg 41; and Adams and Pew 1990). People easily become immersed in performing single tasks in a multitask and tend to forget about their responsibility to other waiting tasks. Empirical evidence from the study of a proposed CIDCW (Call ID on Call Waiting) telephone interface (see Section 4.11.2, "Negotiated Interruption," on pg 159; and Katz 1995) found that when a user interface gives people the capacity to postpone handling interruptions, they sometimes ignore them too long. When people have direct responsibility for timing when to handle calls coming in on Call Waiting they can sometimes delay so long that the new caller hangs up. This perspective leads to the following subhypothesis.

Subhypothesis 5: the negotiation method for coordinating interruption will create user interfaces that result in users exhibiting more errors of omission on some parts of computer-based multitasks.

5.4.2.5 SUBHYPOTHESIS 6

People's interruptibility varies over time as they perform multitasks. Sometimes interruptions can have critically negative effects on people's performance, but other times those same interruptions cause little problem whatever. This disparity in interruptibility is related to the transition points people experience between stages of accomplishing activities (see Section 4.11.3, "Mediated Interruption," on pg 163; and Miyata and Norman 1986). If the user interface could automatically detect these points of increased interruptibility in users, then interruptions could be intelligently delayed until times when they would have least negative effect. This kind of intelligent mediation of when interruptions are presented to the user has the added benefit that users do not have the added burden of responsibility for negotiating their own interruption points. This perspective suggests the following subhypothesis.

Subhypothesis 6: the mediated method for coordinating interruption will create user interfaces that support higher user performance on multitasks than the other methods of coordination.

5.4.2.6 SUBHYPOTHESES 7 & 8

Uncertainty in when users receive feedback of their actions on computer-based tasks (response variability) has been found to be an important design variable in the design of user interfaces (Shneiderman 1992). People need to be able to predict when computers will accomplish the commands they are given. In human-human multitask environments, scheduling has been found to be the most useful time management strategy for increasing the predictability of interruptions (see Section 4.11.4, “Scheduled Interruption,” on pg 169; and Hall and Hursch 1982). This perspective leads to the following subhypothesis.

Subhypothesis 7: the scheduled method for coordinating interruption will create user interfaces that support higher user performance on multitasks than the other methods of coordination.

Subhypothesis 8: the total number of task switches encountered by a user will be less for user interfaces created with the scheduled method for coordinating interruptions than for user interfaces created with other interruption coordination methods.

5.5 METHOD

An experiment was designed to test the main hypothesis (and its subhypotheses). The Taxonomy of Human Interruption’s “Method of Coordination” factor identifies the four recognized categories of solutions for the addressing this problem.

This theoretical information was used to guide the operationalization of an experiment. The main hypothesis was tested by observing the relative effectiveness of each of the four interruption coordination methods identified by the Taxonomy of Human Interruption. These recognized methods are: (1) immediate interruption; (2) negotiated interruption; (3) mediated interruption; and (4) scheduled interruption. These specific methods from the taxonomy were used as the values for the independent variable. Each of these four basic solutions were implemented in different user interfaces to a interrupt laden multitask. The experiment had subjects perform the same multitask with different versions of the user interface. The dependent variable was subjects’ performance on the multitask while using the different user interfaces.

5.5.1 Subjects

36 volunteers were successfully run as subjects in this experiment (18 males and 18 females). Subjects had a median age of 21 (mean 24.7, min. 18, max. 47). All subjects were sampled from the general population of engineers and computer scientists living in the Washington, DC area. Most subjects (28 of 36) were students in the School of Engineering and Applied Science (SEAS) at the George Washington University. Subjects were recruited by an email broadcast to all SEAS students with email accounts. This recruitment message did not reveal the purpose of the experiment, but portrayed the experimental task as "fun" and "similar to a video game." The message also advertised that each volunteer would receive $20, and explained that volunteering would be a significant contribution to the investigation of an important problem.[14] Subjects were self-selecting from this broad population.

This method for population sampling is less than random and therefore not optimal. However, it was judged adequate because of the exceptional diversity of the GWU student population, and because of the universal motivation of the monetary reward.

5.5.2 Design

A single-factor within-subjects Latin square design was chosen as an appropriate design for this experiment. Six treatments were devised: four experimental treatments and two base-case control treatments. Each of the four experimental treatments represents one of the four methods for coordinating interruption identified in the Taxonomy of Human Interruption.

Each treatment manifests a different version of a user interface (the independent variable) for a computer-based interrupt laden multitask. The computer-based multitask was not varied between treatment conditions. Subjects' performance (the dependent variable) on the multitask is observed and recorded under the six treatment conditions.

---

14. Six subjects were employees of the Naval Research Laboratory (NRL) and volunteered to participate in this experiment during working hours. NRL's policies for human-subjects experimentation made it impossible to pay these subjects. Therefore, they were not paid and were not told of the $20 reward given to other subjects.

All subjects received all six treatments. However, each subject was assigned to one of six groups that define the Latin squares ordering of the presentation of the six treatments. The presentation of each treatment was divided into two contiguous trials to avoid the confounding influences of fatigue and boredom.

Each subject performed a total of 24 trials of the computer-based multitask. Each trial was 4 minute and 30 seconds long, and there was a brief rest period enforced to a minimum of 25 seconds between each session. Therefore the total time for a subject to complete the experimental task was about 2 hours. For all subjects, the first 12 trials were practice (~1 hour) and the second 12 trials were experiment (~1 hour).

Subjects received the same Latin squares ordering of trials on the practice trials as they did on the experimental trials. For example, subjects assigned to the Latin squares order Group-2 received their 24, 4 minute 30 second trials of treatments in the following order (from left to right):

Table 20 — Trial Orderings of Treatment Pairs for Subjects in Group-2

practice trials	experimental trials
2 2 4 4 1 1 6 6 3 3 5 5	2 2 4 4 1 1 6 6 3 3 5 5

5.5.2.1 INTERNAL VALIDITY

Internal validity is the appropriateness of using the results to support the conclusions. The internal validity is important because it determines to what degree the experiment measures what it is supposed to measure. As with any human subjects experiment, there are many possible sources of confounding influence that can destroy internal validity. The survey of relevant literature presented in Chapter 4 (pg. 143) indicates that research about human interruption is especially sensitive to problems of confounding influence. This experiment implemented the following measures as an attempt to control or nullify all major confounds to internal validity.

5.5.2.1.1 Control of General Practice Effects

The following controls were implemented to address possible sources of confounding influences due to general practice effects (change in subjects performance caused by increasing exposure to the experimental context).

• Subjects performed one hour of practice trials before they began their experimental trials. One hour of practice was chosen as appropriate based on observations of pilot subjects. During pilot subjects testing it was observed that different subjects learn the experimental multitask at different rates. Some subjects could therefore have gained sufficient practice to overcome the confounding effects of learning in less than one hour. However, the practice time had to be chosen to accommodate the lowest common denominator of subject learning rate to insure that every subject had received sufficient practice to overcome the confounding effects of learning.

• Subjects received practice trials very similar to the experimental trials. This design avoids possible confounding effects caused by differences in subjects learning behaviors for different contexts.

• A counterbalanced grouping scheme was chosen to negate possible learning effects. In a repeated measured study there is no way to avoid the possible effects of order of presentation of treatments on subjects' performance. However, implementing a counterbalanced grouping scheme provides a way of separating main effects from any order effects. A digram-balanced Latin squares ordering was chosen as the best ordering design for counterbalancing this repeated measures study with six treatment conditions. A digram-balanced counterbalance ordering is a version of a Latin squares design in which each condition precedes and follows all other conditions exactly once (Keppel 1991, p. 339; Wagenaar 1969). Subjects were randomly assigned to one of six groups.

Table 21 — Counterbalanced Treatment Order — Digram-Balanced Latin Squares

Group	Treatment Condition Orders					
1	1	2	3	4	5	6
2	2	4	1	6	3	5
3	3	1	5	2	6	4
4	4	6	2	5	1	3
5	5	3	6	1	4	2
6	6	5	4	3	2	1

• Treatment sessions, or trials, were limited to 4 minute 30 second. Observations made during pilot subjects testing found that trials must be kept under about 5 minutes in length to avoid subject fatigue and boredom as a source of confounding influence. 25 second (minimum) rest periods were imposed between trials.

• The experimental task was designed as an engaging video game to control for subject boredom.

• The multitask for every trial for every subject was randomly scheduled. This randomization was implemented to control for the possible confound of subjects learning to predict and therefore anticipate multitask events.

5.5.2.1.2 Control of Differential Carryover Effects

The following controls were implemented to address possible sources of confounding influences due to differential carryover effects (change in subjects performance caused by interference between treatments).

• Rest periods were imposed between treatments to allow subjects time to recover from the effects of the previous treatment.

• A mask was used to block the display of the dualtask during the rest period to allow subjects to more easily forget the previous treatment.

• Consistency of task instructions was maintained by having consistent on-screen instructions presented before each trial. Also, subjects received detailed written instructions for the multitask that they were able to keep during the experiment for reference.

• Consistency of protocol was maintained by having the experimenter follow a written script for administering the experiment to each subject.

5.5.2.1.3 Control of Other Possible Confounds

The following controls were implemented to address other possible sources of confounding influence.

• A within-subjects design was used to control for individual differences in subjects.

• The multitasks presented during the different trials were created so as to be as uniform as possible but still have a randomized, unpredictable schedule of events. Each trial had the same number of events for each multitask, and the random scheduling was constrained to a fixed frequency distribution. This ensured an unpredictable multitask with a uniform level of difficulty between trials.

• Subjects were asked if their vision was corrected to 20/20, and given a screening test for color perception (Ishihara 1996). This controlled for the possible confound of subjects not being physically capable of performing the visual-based multitask activities.

• Subjects rights were explained and a mandatory consent form was introduced to control for the possible confound of abnormal anxiety or stress in subjects. Subjects were told that they were free to withdraw at any time for any reason without penalty.

• The experimental environment was constrained as much as possible to provide subjects with a non-distracting setting. The environment was also kept consistent in terms of comfort, lighting, and experimental apparatus.

• Instructions to the subjects did not mention performance levels or encourage extraordinary effort. This instruction was a control for abnormal anxiety in subjects. Instructions to the subjects explained that the purpose of this experiment was to compare different user interface designs and not the subject's personal abilities. Subjects were instructed that there is no such thing as "good" performance or "bad" performance, but that they should try to maintain a consistent level of effort throughout the experiment.

• Subjects were encouraged to ask questions during the practice period. This technique was a control for confusion in subjects about multitask requirements.

• Subjects were told that the experimental dualtask has been determined to be below normal risk and judged totally safe by the GWU Human Subjects Protection Committee. This explanation was a control for the confound of abnormal apprehensiveness in subjects.

• Multitask interaction has been restricted to the barest minimum to reduce possible sources of confounding influences. The interaction was limited to visual displays (no sound) and all subjects' multitask behaviors were limited to keyboard key presses. No supplementary computer feedback was provided to inform subjects of their performance.

5.5.2.2 External Validity

External validity is the appropriateness of generalizing conclusions to real-world domains. The external validity is important because it determines to what degree the experimental context is similar enough to a real-world context to allow generalization of the results.

This experiment is a pioneering investigation of basic science, and so it was decided that external validity would be sacrificed to some degree in favor of increased internal validity. However, some measures were taken to increase the external validity of this experiment where possible.

- Subjects were sampled from a large, and culturally and racially diverse population.
- Self-selection of subjects was constrained to equal numbers of male and female subjects.
- Incentives ($20) and the guise of entertainment (the game-like task) were employed as attempts to increase subjects' motivation and attentiveness on the experimental multitask.
- The experimental settings were chosen to be typical of normal computer-based office work. This "normal" context, however, was constrained to "not-so-normal" closed-door, uninterrupted work.

5.5.3 Multitask

An interruption laden multitask was created as an appropriate testbed for this experiment. All subjects performed this same multitask on all trials. This multitask provided a common benchmark for comparing subjects' relative performance on the six different treatment conditions.

The multitask is a dualtask (a two-task multitask) composed of a continuous game task and an intermittent matching task. The game task is loosely based on a video game by Nintendo Corporation called "Fire" that was originally released in 1980 & 1981 as a version of the Nintendo Game & Watch product series[15] (Nintendo 1980 & 1981).[16] The matching task is loosely based on the matching tasks used in experiments of the Stroup Effect (Stroup 1935).

The dualtask is conceptually simple and yet can be very difficult for people to perform. The results of pilot studies determined that this dualtask has some useful characteristics as a testbed for investigations of human interruption during HCI.

15. The original Nintendo Game & Watch named "Fire" was re-released in 1997 as a part of the Nintendo "Game & Watch Gallery" game cartridge for Nintendo Game Boy and Nintendo Pocket Game Boy (Nintendo 1997).
16. All software used in this experiment is original -- no software was copied from Nintendo.

5.5.3.1 GAME TASK

The game task required subjects to control the movement of cartoon style stretcher bearers that must catch other game characters as they fall from a building and bounce them three times into a waiting truck. Each falling character must be successfully caught and bounced three separate times at three different locations. If a character is missed at any of the three bounce points then they are lost. The original Nintendo game scenario included medical stretcher bearers catching babies that jumped from a burning building and bounce them into an ambulance. This game scenario has been altered from the original for two reasons: (1) to be less emotionally charged; and (2) to be more obviously relevant to this author's Navy sponsor. The new scenario includes Marine stretcher bearers catching diplomats jumping from an overrun U.S. embassy and bouncing them into a military truck.

The game task is trivial when game characters jump one at a time. However, when more than one game characters jump in quick succession it becomes a difficult game of juggling, and subjects had to move the stretcher bearers back and forth between the three jump points in quick irregular sequences to keep all the game characters in the air at the same time. See Figure 9 below.

Figure 9.Game Task: help the diplomats escape the overrun embassy by moving the Marine stretcher bearers to catch and bounce them three times into the truck.

This game task is an example of a apparently simple task that can be very difficult to perform. The task has the following important characteristics: appealing as a game for humans; discrete subtasks; simple yet powerful; and constrained scheduling of subtasks.

5.5.3.1.1 Appealing as a Game for Humans

It was important to measure how subjects perform when they are highly motivated and actively concerned for the outcome of the multitask. However, for the purpose of this experiment, subjects could not be asked to perform a critical real-world task in a critical real-world context. A multitask had to be found that could motivate subjects to perform as if the outcome mattered, but that would not affect subjects negatively when they failed. Beyond the prohibitive costs of failure in many real-world settings, it was important that this experiment not pose any psychological, emotional, or physical risk to subjects.

The Nintendo game "Fire" was chosen as a model for the game task because it is both compelling and utterly simple. "Fire" was and is a popular game, and despite its obvious simplicity it has been proven to engage people in motivated play for hours at a time.

A pilot study was conducted with three volunteers to set the complexity of the game to an appropriate level. The game had to be contrived so that it was complex enough to expose subjects' vulnerability to interruption, but simple enough not to cause subjects to despair of performing well. Through testing with pilot subjects it was discovered that 59 characters jumping over a 4 minute and 30 second trial was appropriate. This complexity ensured that no subject could ever save 100% of the jumping characters. This complexity also ensured that although subjects could not save 100% of the jumpers, they did not despair of the possibility of saving 100% of the jumpers.

The results of the pilot studies also revealed the need to have two different levels of complexity for the practice trials. As with Nintendo's "Fire," while users are learning the game and the controls, they need to have a simplified introduction version of the task. An introductory period of easy play gives subjects time to learn everything they need to know in a low stress context. It was concluded from pilot studies that if subjects were not given an easy introductory period their ability to learn the multitask would be negatively affected. The number of jumpers for the 24 trials for each subject was as follows (except for the "Basecase - matching task only" treatment which has no game task). Note, the cells of the table contain the pairs of trials for each of the six treatment conditions.

Table 22 — Number of Jumpers for Each Trial

practice trials					
first	second	third	fourth	fifth	sixth
38 59	38 59	38 59	38 59	38 59	38 59
experimental trials					
first	second	third	fourth	fifth	sixth
59 59	59 59	59 59	59 59	59 59	59 59

5.5.3.1.2 Discrete Subtasks

This game task was a good choice for experimentation because it is both continuous and discrete. It runs continuously, so the user does not get a break in their responsibility for performing the task. However, the task is composed of several individual subtasks. Saving each jumping character is a separate discrete subtask. This convenient subtask composition allows observations of peoples' behaviors to be easily broken down into discrete units.

The fact that this game task is composed of individual jumping characters provides the following useful task characteristics.

- Subject's performance on completed subtasks can be easily classified as success or failure.
- Subtasks require subjects to make more than one time-sensitive decision over time.
- Subtasks do not require constant attention from subjects, only a few well-timed actions.
- Subtasks continue long enough so that interruption in the middle is not impractical.
- Subtasks require a small but significant amount of situational awareness for successful completion. This insures that there is an overhead cost for resuming the game task after performing the matching task.
- Subtasks can be individually scheduled in a randomized way to prevent predictability.

5.5.3.1.3 Simple Yet Powerful

The game task is extremely simple. One subtask by itself can be easily accomplished with predictable actions. It is only the interactions of more than one randomly intermixed subtasks that requires dynamic unpredictable problem solving. This arrangement allows the overall complexity of the game task to be conveniently manipulated. Complexity here is defined in terms of how demanding the task is on the person performing it. The overall complexity of the game task is determined by two aspects of how subtasks are scheduled: (1) the total number of jumpers scheduled to jump in a given interval of time; and (2) the distribution of when subtasks begin within that time interval.

The game task provides a convenient experimental platform without introducing excessive noise. The following simplifications are implemented.

- All subtasks (individual jumping characters) require exactly the same amount of time to handle and require exactly the same decisions. For each subtask, the time it takes from its start jump until after its third (and last) bounce is about 13.7 seconds. After a character has been successfully bounced its third time it is on the screen about another 3.2 seconds until it falls safely into the military truck (total time on the screen for a saved jumper is 13.7 + 3.2 = 16.9 seconds).

- Each subtask is completely independent. Errors made while performing one subtask do not automatically cause errors on other subtasks.

- The game task is easily learned both in concept and in action. Only two keyboard keys control all game task interaction, and these two keys are spatially mapped to the left and right movements of the stretcher bearers.

- Keyboard keys were not mapped to more than one meaning.

- Feedback was kept to the minimum possible. There are many kinds of possible feedback that could have been implemented to give users additional information about the multitask or about their performance. These additional feedback stimuli could have been powerful sources of confounding influence on subjects' performance. Therefore the following kinds of feedback were not implemented: sound; performance scores; animation of secondary events; alerts of impending events; and information of the state of the hidden task.

5.5.3.1.4 Constrained Scheduling of Subtasks

Subjects each performed 24 trials (36 subjects * 24 = 864 total trials). Each of those 864 trials had a unique and randomly determined scheduling of when subtasks began. However, unconstrained randomization would have introduced unnecessary variance in the complexity of the game task across trials. A constrained randomization scheme was implemented to avoid this source of noise. Scheduling was constrained to a constant frequency distribution.

Each subtask was scheduled to begin (jump off the building) by a specified time interval since the last subtask began. A standard array of time intervals was used to constrain the randomization of the scheduling. Every 4 minute 30 second experimental trial had 59 subtasks (except the “Basecase - matching task only” condition). An array of standard time intervals was calculated such that each interval had a unique time and all 59 were in a linear progression starting at 0 msec. A standard increment was calculated so that the intervals increase linearly, and also so that the sum total of all the intervals ensured that the game will last 4 minutes and 30 seconds. The total of all the intervals equals 270,000 msec (4 minutes and 30 seconds) minus ~16,900 msec (time for the last jumper to reach the truck if saved). The following figure shows the msec values for the 59 standard intervals.

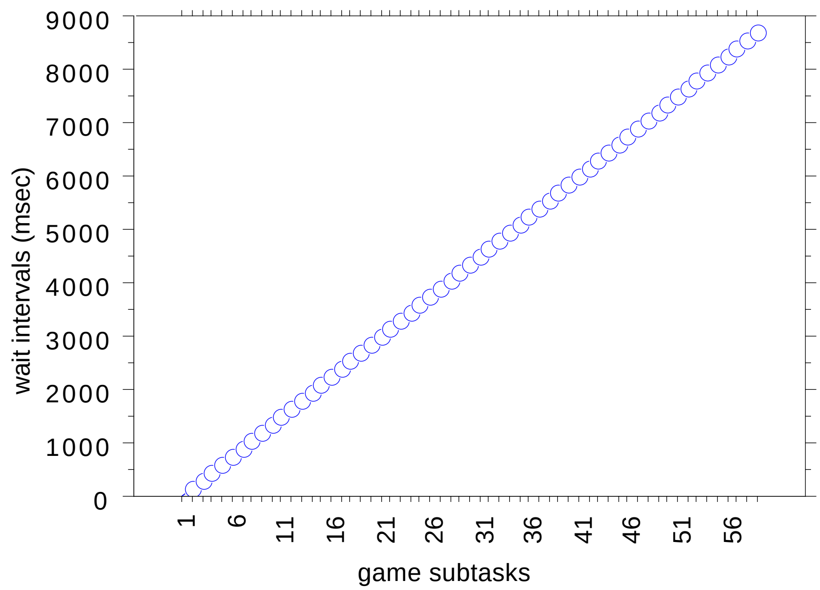

Figure 10.The 59 standard wait intervals (in msec) used to schedule every trial of the game task. Each subtask (jumping diplomat) is scheduled to begin a certain number of msec after the beginning of the previous subtask.

A histogram of the standard intervals shows a uniform distribution of interval times in the frequency domain (see figure below).

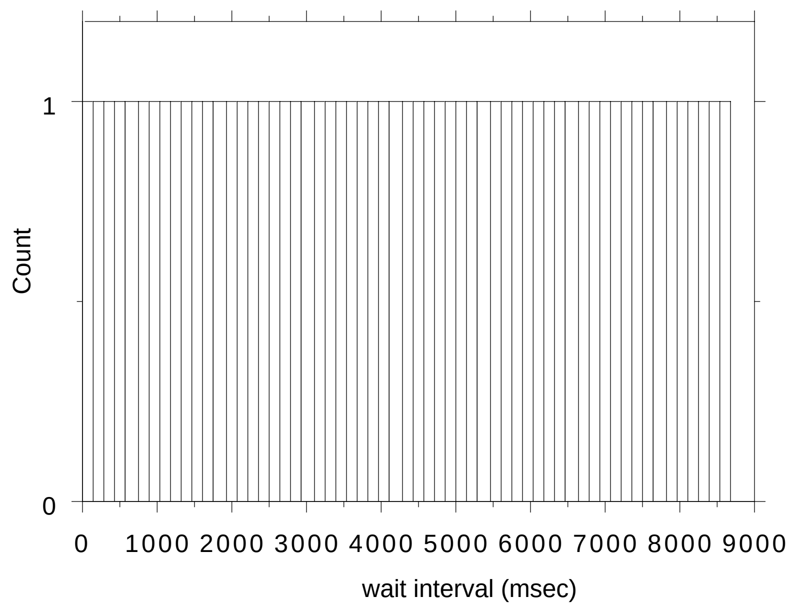

Figure 11.The standard intervals from Figure 10 shows a uniform distribution in the frequency domain.

Each of the 864 trials in this experiment achieved an unpredictable game by randomly re-sorting this standard array of time intervals. Resorting did not affect the frequency domain of the scheduling intervals. Therefore, although each of the 864 trials had a randomly different game, each also has exactly the same frequency distribution of scheduling intervals. For example, Figure 12 below shows the randomly re-sorted schedule of the 59 subtask time intervals of one trial for one subject.

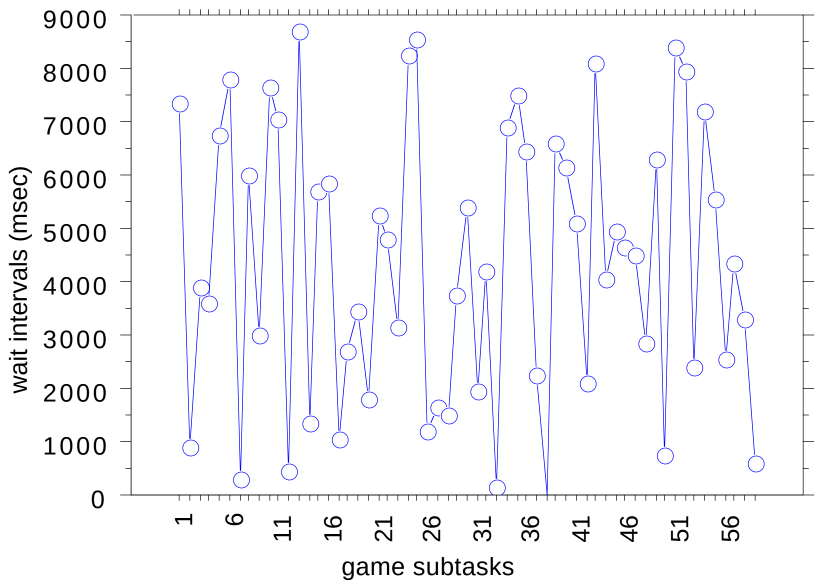

Figure 12.An example of randomly re-sorted standard intervals (from Figure 10) to create the unpredictable scheduling of the 59 subtasks (jumping diplomats) of a particular trial for a particular subject.

5.5.3.2 MATCHING TASK

The second task of this dualtask is an intermittent matching task loosely based on the matching tasks reported in investigations of the Stroup effect (Stroup 1935). The interruption task required subjects to make a matching decision either based on color or shape. When matching task events occurred they totally obscured subjects' view of the game task. The presentation of the matching task preserved the game-like context of the game. It was presented as a pager-like device.

Subjects were presented with a colored shape at the top of the pager window, and instructed to choose one of the bottom two colored shapes according to the matching rule displayed in the center. The matching rule instructed subjects to either "Match by shape" or "Match by color." See Figure 13.

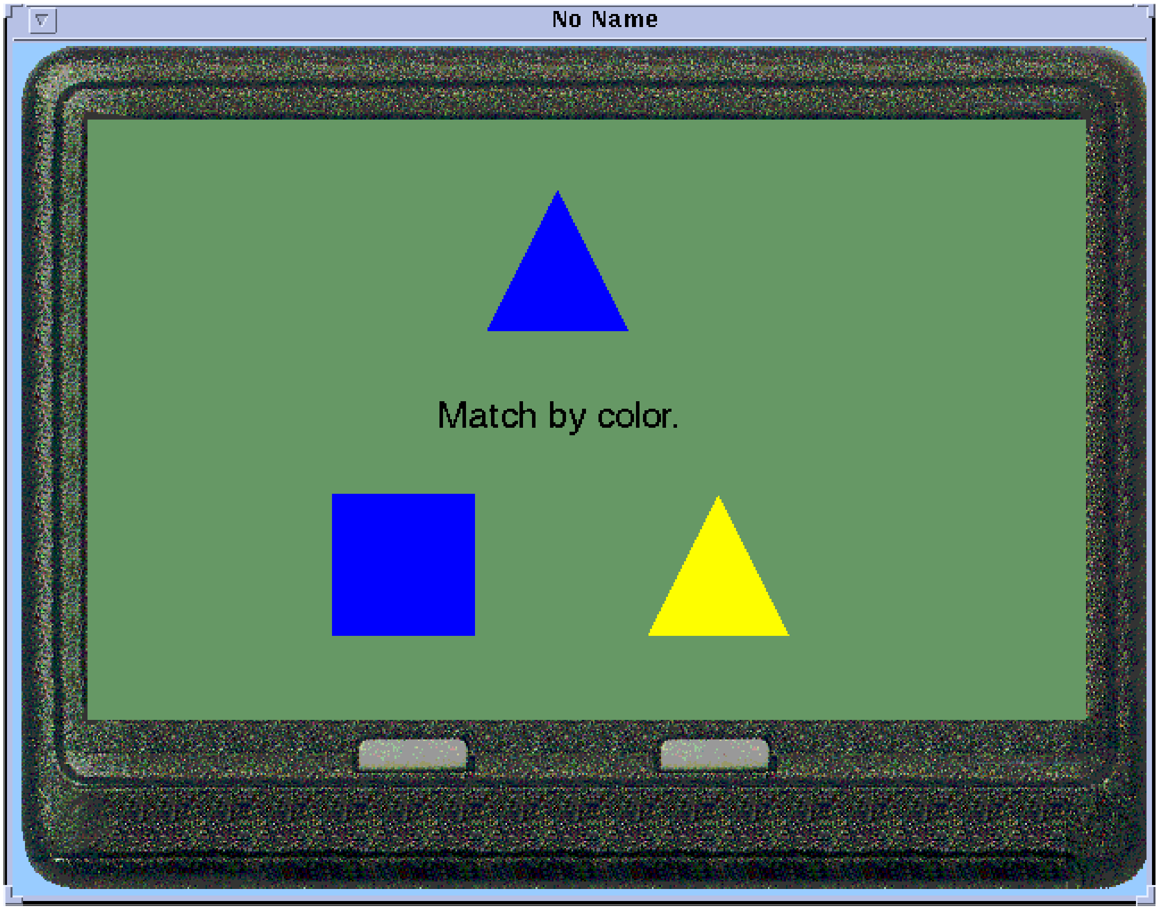

Figure 13.Matching Task: choose which of the bottom two colored shapes matches the top colored shape according to the displayed matching rule.

This matching task is conceptually simple, but deceptively difficult to perform. Pilot studies found that people were not able to learn to automate this task to their subconscious processing. Each individual matching task required subjects to focus their attention long enough to make a conscious decision. This required conscious decision, however, is minimally brief.

This task was useful for this experiment because of its following characteristics.

- The individual matching tasks were completely independent and subjects' choices were easily judged right or wrong.

- The graphic nature of the task corresponded with the graphic nature of the game task.
- It was easily learned.
- The left/right choice was conveniently mapped to a left/right keyboard selection.
- The matching choices required a relatively consistent amount of subjects' time. This consistency allowed subjects to be able to predict with some accuracy how long it would take them to accomplish a single matching task.
- It implemented no priority scheme.
- Matching tasks had to be done one at a time from a FIFO (first-in-first-out) queue.
- It did not allow interruptions of interruptions.

The same constrained randomization scheme that was used to schedule the subtasks of the game task was used to schedule the individual matching tasks. Through testing with pilot subjects it was discovered that 80 matching tasks within a 4 minute and 30 second trial were appropriate. Pilot studies also revealed the need to have a simplified introduction version of the matching task for the first trial of each pair of practice trials. The number of matching tasks for the 24 trials for each subject was as follows (except for the "Basecase - game task only" treatment which had no matching task). [Note, the cells of the table contain the pairs of trials for each of the six treatment conditions.]

Table 23 — Number of matching tasks per trial

practice trials					
first	second	third	fourth	fifth	sixth
40 80	40 80	40 80	40 80	40 80	40 80
experimental trials					
first	second	third	fourth	fifth	sixth
80 80	80 80	80 80	80 80	80 80	80 80

5.5.3.3 PERFORMING THE DUALTASK

All six versions of the user interfaces for the dualtask have some commands in common. All dualtask interaction is performed by single-handed and same-handed keyboard key presses. Subjects performed the game task by pressing the "Delete" and "Page Down" keys with one hand to control the back and forth movement of the stretcher bearers. Subjects performed the matching task by using the same hand to press the "Insert" and "Page Up" keys to choose either the left or right shapes.

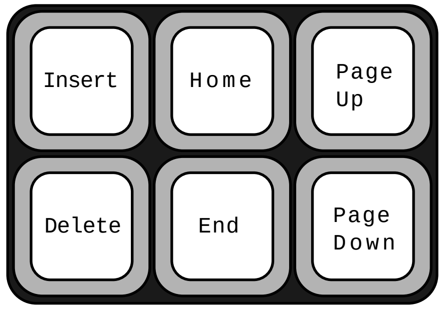

Figure 14. The keyboard keys used for performing the experimental multitask. The "Delete" and "Page Down" keys can be used to move the stretcher bearers left and right in the game task. The "Insert" and "Page Up" keys can be used to make left and right matching choices in the matching task. The "Home" and "End" keys are only relevant to the "negotiated interruption" condition.

5.5.4 Treatments

Each of the six treatment conditions employed a different user interface for performing the common dualtask described above. The game task continued to run without possibility for pause regardless of whether subjects could see it or whether it was occluded by the matching task.

5.5.4.1 TREATMENT 1

A "basecase — game task only" control condition implemented the game task in isolation. Subjects received no interruptions of matching tasks.

5.5.4.2 TREATMENT 2

A "basecase — matching task only" control condition implemented the matching task in isolation. The matching task's pager-like background was displayed the entire trial and matching tasks appeared within it.

5.5.4.3 TREATMENT 3

The “immediate interruption” treatment condition implemented the “interrupt immediately regardless of situation” strategy for coordinating user-interruption. When a matching task was scheduled to occur it was immediately presented. This interrupted the user and the matching task’s pager display was imposed in the same window and totally obscured the game task display. Users had to then perform the matching task before they can go back and resume the game task. If other matching tasks arrived while a subject was performing a matching task, they were queued in a FIFO manner. The user had to perform all queued matching tasks before they could resume the game task.

5.5.4.4 TREATMENT 4

The “negotiated interruption” treatment condition gave users control over when they would handle interruptions. When a matching task occurred, an announcement was immediately flashed of an empty pager and then the game task display resumed. This announcing flash lasted 150 msec, and was verified as an appropriate length through pilot testing. This announcement notified users of the existence of a waiting interruption task. The users then had to decide when to begin the queued matching task. Subjects could use the “Home” and “End” keys at any time to bring the matching task to the foreground or push it to the background. If more than one matching task was queued, users did not have to perform all of the queued tasks together, but instead could switch back and forth between the game task and the queued matching tasks at will. When a user completed the last queued matching task the game display was automatically resumed.

5.5.4.5 TREATMENT 5

The “mediated interruption” treatment condition automatically calculated a function of users’ workload. This user interface was an implementation of the idea of trying to time user-interruptions when subjects were not busy doing the game task. When the metric of workload was high the mediator queued up the arriving matching tasks, and then when the metric of workload changed to low the mediator interrupted users with the stack the queued matching tasks. User had to then perform all the queued matching tasks before they could return to the game task. When the user was busy (high metric for workload) no notification of the arrival of interruptions was presented.

5.5.4.6 TREATMENT 6

The "scheduled interruption" treatment condition saved up matching tasks without notifying users of their arrival like in the mediated condition. However, unlike the mediated condition users were interrupted with the queued matching tasks on a prearranged schedule. This schedule was a statically fixed cyclical time interval and users were informed of this interval before they began the treatment. This scheduled interruption interval was set at 25 seconds between the beginnings of the onset of interrupting subjects with queued matching tasks.

5.5.5 Apparatus

All subjects performed the computer-based dualtask on a Hewlett Packard OmniBook 5700CTX laptop computer running the Windows95 operating system. This computer had a 166 Pentium CPU with MMX, 32Mb of RAM, and a 2.9Gb hard disk. The built-in monitor was a 12.1 inch backlit liquid-crystal XGA display with 1024 X 768 pixel resolution and 16 bit color. Subjects used an externally attached, extended keyboard (Dell brand). The computer-based dualtask was displayed in a single 640X480 pixels window on the laptop's built-in color monitor. This window was located in the top left corner of the screen.

The laptop sat on a box four and a three quarter inch high on a table top in front of subjects. The box was added to bring the screen up to a comfortable viewing height, and the box was painted off-white so as not to be a source of distraction. The external keyboard sat on the table top directly in front of subjects. Subjects were seated on a padded chair typical of the kind used by office workers.

The experimental software was written in Java 1.1 on a Sun Sparcstation 20 using Sun's JDK 1.1.3 and 1.1.5 development tools and the XEmacs 19.14 editor. The software was then moved to the HP laptop and ran under Sun's JDK 1.1.5 for Windows95. Not porting changes were necessary. The game task is implemented with sprite-based double-buffered frame animation running at 20 frames a second. All animation and subtask scheduling was implemented with a multi-threaded approach.

5.5.6 Procedure

Subjects were run one at a time through the experiment. The experimenter followed a written script to ensure that the treatments were administered to each subject in a consistent way. This script dictated the following order for the different parts of the experiment.

1. Greeting and introduction.
2. Verify that subjects meet the minimal requirements to participate in this experiment. Subjects were asked: (1) "do you have normal color vision corrected to 20/20?" (2) "can you read English?" (3) "can you press keyboard buttons with one hand?" and (4) "are you 18 years old or older?"
3. Acquire subjects' signature on a consent form that explains their rights. See Appendix: Consent Form, pg. 271.
4. Administer a standard color test (Ishihara 1996).
5. Administer the Entrance Questionnaire. See Appendix: Entrance Questionnaire, pg. 273.
6. Give subjects the written "Instructions for Subjects" and ask them to read it. See Appendix: Instructions for Subjects, pg. 277.
7. Administer the 24 trials of the computer-based dualtask. Dualtask events and user input events were unobtrusively recorded by computer throughout the experiment. All trials were also video taped.
8. Administer the Exit Questionnaire. See Appendix: Exit Questionnaire, pg. 283.
9. Give subjects the "Debriefing Statement" (See Appendix: Debriefing, pg. 289) and the $20.

Each subject spent about 2 hours and 30 minutes participating in this experiment.

5.6 RESULTS

Observations recorded during this experiment are used to empirically compare the four known solutions to the problem of human interruption. This experiment has recorded observations of how people behave during interruptions. Internal validity and reliability were emphasized in the design and execution of this experiment. These data can now be analyzed with confidence to determine whether reality supports theory.

The main hypothesis for this experiment predicts that "the particular method for coordinating user-interruption that is implemented in a user interface will affect users' performance on interrupt laden computer-based multitasks" (see Section 5.4.1 on pg 176). Five different measures of users' performance have been chosen as appropriate for testing this hypothesis: (1) number of jumpers saved on the game task; (2) number of switches between game task and matching task; (3) number of matches done wrong; (4) number of matches not done; and (5) the average age of matching tasks before they are completed, i.e., the average time from the scheduled onset of each matching task until its is actually completed.

There are many other interesting performance measures that could have been analyzed. However, these other categories of data were judged to be secondary to the main purpose of this experiment — partial validation of the Taxonomy of Human Interruption as a useful tool for empirical research. The five performance measures chosen represent those that are most obviously appropriate for testing the main hypothesis and its subhypotheses. Analysis of other performance measures, although interesting, is left for future work.

Only data collected for the 12 experimental trials (not the 12 practice trials) are included in these analyses. Also, the data from each pair of treatment trials are combined for this analysis. It is postulated that any effects related to differences between trial 1 and trial 2 for each treatment will not affect the results of a combined analysis. Observations for each pair of treatment trials are summed. The following table contains an example.

Table 24 — Example of combined scores for number of jumpers saved for all 12 experimental trials of a single subject

Raw scores					
base-game	base-match	immediate	negotiated	mediated	scheduled
41 44	0 0	26 19	33 35	23 24	19 20
Combined scores					
base-game	base-match	immediate	negotiated	mediated	scheduled
85	0	45	68	47	39

5.6.1 Support for the Main Hypothesis

The following results support the main hypothesis (see Section 5.4.1 on pg 176) and permit H_o to be rejected. Significant evidence is found for each of the five performance measures to support the main hypothesis. Some significant secondary effects were found for two of the five performance measures. These secondary effects are examined and found not to confound support for the main hypothesis.

Four kinds of analysis are reported for each of the five measures of performance (1) a table of descriptive statistics; (2) a box plot; (3) an ANOVA; and (4) if warranted, a post-hoc analysis. An alpha level of .05 is used to make decisions of significance.

The descriptive statistics report the mean, std. dev., std. error, count, min., max. and median. The box plots display marks at the 10th, 25th, 50th, 75th, and 90th percentiles of the variable. The boxes, therefore, contain the center 50% of values with the center line at the median. The outer brackets enclose 80% of values, with 10% outliers graphed to the left and 10% outliers graphed to the right (Abacus Concepts 1996, pg. 185). The ANOVA uses a simple Latin squares model as defined in Bruning and Kintz (1987, pg. 85-93).

Post-hoc pair-wise comparisons are done for factors determined significant by ANOVA. A least squares means table analysis is the most appropriate post-hoc analysis technique because of the Latin squares repeated measures design of this experiment. The SuperANOVA Manual says that the correlation among observations from a repeated measures design renders post-hoc means separation test invalid for within-subjects tests (Abacus Concepts 1989, pg. 198). Therefore post-hoc tests like the Games-Howell and the Fisher PLSD are inappropriate. The least squares means table analysis is a “graphs of means” post-hoc test recommended by the SuperANOVA Manual as appropriate for the repeated measures model used in this experiment (Abacus Concepts 1989, pg. 204-205). A corrected alpha level is used to reduce the likelihood of finding significant pair-wise differences by chance because of the large number of comparisons. The corrected alpha level is calculated by dividing the .05 alpha level by the number of pairs being compared. For example, a corrected alpha level of .005 (.05 / 10) would be used to determine significance of pair-wise comparisons on a post-hoc analysis involving 10 pairs.

5.6.1.1 NUMBER OF JUMPERS SAVED

Table 25 — Descriptive Statistics of Number of Jumpers Saved (118 possible)

	Mean	Std. Dev.	Std. Error	Count	Min	Max	Median
Total	46.0	25.858	1.759	216	0	95	49
Baseline - Game Only	79.6	8.264	1.377	36	62	95	80.5
Baseline - Match Only	0	0	0	36	0	0	0
Immediate	47.5	9.918	1.653	36	30	70	48
Negotiated	60.2	13.474	2.246	36	21	83	59.5
Mediated	49.4	9.031	1.505	36	27	65	49.5
Scheduled	39.0	8.175	1.362	36	22	53	38.5

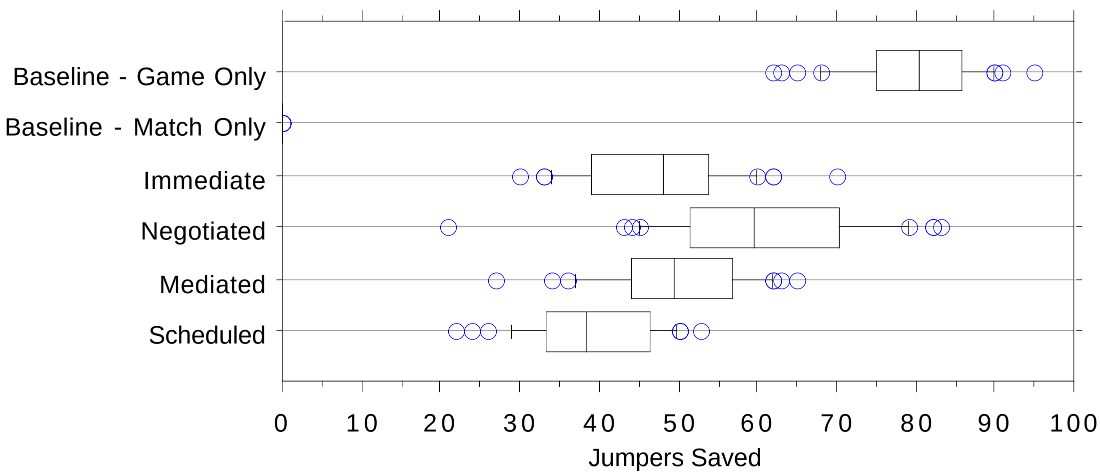

Figure 15. Box plot for number of jumpers saved, by treatment condition.

Table 26 — Simple Latin Squares ANOVA for Number of Jumpers Saved

Source	SS	df	ms	F	p
Total	143759.704	215			
Between Subjects	10120.370	35			
Groups(IO_b)	4311.981	5	862.396	4.454	.004*
$Error_b$	5808.389	30	193.613		
Within Subjects	133639.333	180			
Interruption (I)	126378.204	5	25275.641	742.730	<.0001*
Order (O)	597.593	5	119.519	3.512	.005*
Interrupt. X $Order_w$ (IO_w)	1558.926	20	77.946	2.290	.003*
$Error_w$	5104.611	150	34.031		

Note: a “*” indicates significance based on an alpha of .05.

Table 27 — Least Squares Means Table of Main Effect for Number of Jumpers Saved

	Vs.	Diff.	Std. Error	t-Test	P-Value
Baseline - Game Only	Immediate	32.111	1.351	23.767	.0001*
	Negotiated	19.389	1.351	14.351	.0001*
	Mediated	30.194	1.351	22.348	.0001*
	Scheduled	40.583	1.351	30.038	.0001*
Immediate	Negotiated	-12.722	1.351	-9.416	.0001*
	Mediated	-1.917	1.351	-1.419	.1586
	Scheduled	8.472	1.351	6.271	.0001*
Negotiated	Mediated	10.806	1.351	7.998	.0001*
	Scheduled	21.194	1.351	15.687	.0001*
Mediated	Scheduled	10.389	1.351	7.689	.0001*

Note: a "*" indicates significance based on a corrected alpha of .005 (.05 / 10).

The ANOVA found a significant main effect. The ANOVA also found other significant effects for this measure of performance that are unrelated to the main hypothesis. These other significant effects are examined and it is determined that none of these secondary effects presents a confound for the main effect. Therefore these data can be said to support the main hypothesis.

5.6.1.1.1 Order Effect

A significant order effect was discovered ("Between Subjects - Groups (IO_b)"). It is unlikely that these between subjects differences by subject group could have happened by chance. This means that the order in which subjects were presented the treatments affected their overall performance on saving jumpers. This order effect was unexpected. Three major efforts were taken to control for the effects of treatment presentation order: (1) the Latin squares ordering scheme; (2) the lengthy practice on all six treatments; and (3) the written instructions given to all subjects before they began the dualtask that describes all treatment conditions. The fact that a significant order effect is exhibited despite these controls is an indication that something important is affecting subjects.

This order effect may be interesting; however, it is not part of the main hypothesis of this experiment. This analysis only needs to show that the order effect does not confound the main effect. The following line chart shows the average performance for each of the six between-subjects groups for the different treatment conditions. There is a clear order effect; however, it does not affect the overall pattern of subjects' performance across treatment conditions.

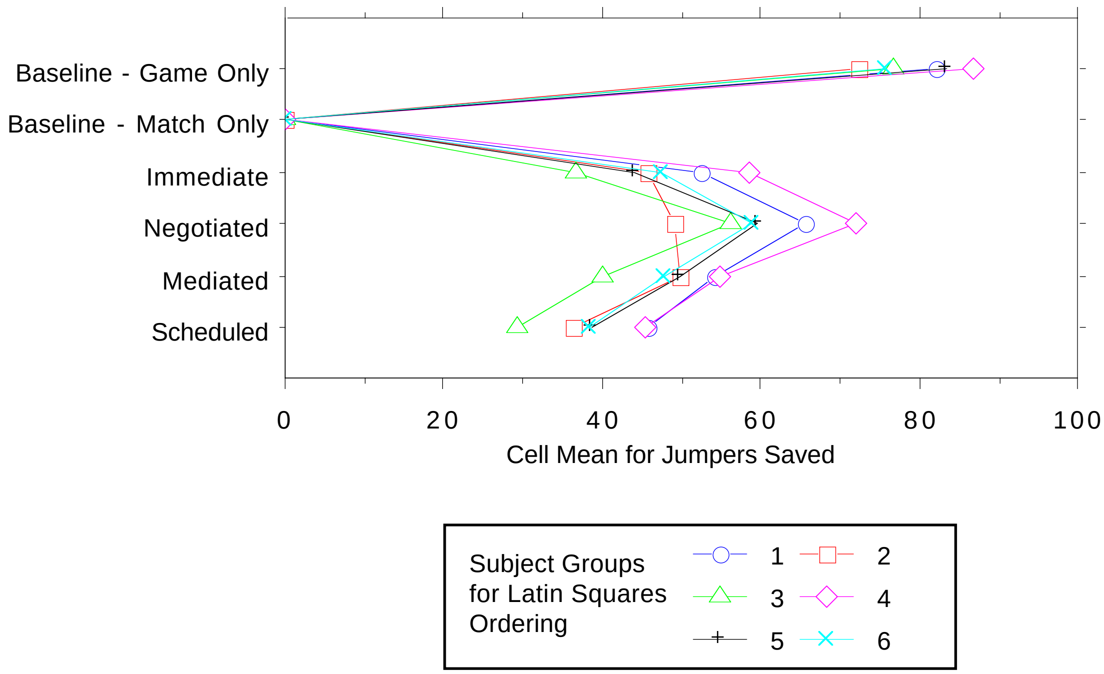

Figure 16.Line graph of jumpers saved for the different treatment conditions, split by between-subjects groups 1-6 (the Latin squares ordering groups).

A post-hoc pair-wise analysis is useful in explaining this order effect.

Table 28 — Least Squares Means Table of Between Subjects Order Effect for Number of Jumpers Saved

	Vs.	Diff.	Std. Error	t-Test	P-Value
1	2	7.778	3.280	2.372	.0243
	3	10.167	3.280	3.100	.0042
	4	-3.056	3.280	-.932	.3589
	5	4.028	3.280	1.228	.2290
	6	5.306	3.280	1.618	.1162
2	3	2.389	3.280	.728	.4720
	4	-10.833	3.280	-3.303	.0025*
	5	-3.750	3.280	-1.143	.2619
	6	-2.472	3.280	-.754	.4568
3	4	-13.222	3.280	-4.032	.0004*
	5	-6.139	3.280	-1.872	.0710
	6	-4.861	3.280	-1.482	.1487
4	5	7.083	3.280	2.160	.0389
	6	8.361	3.280	2.549	.0161
5	6	1.278	3.280	.390	.6996

Note: a “*” indicates significance based on a corrected alpha of .0033 (.05 / 15).

This post-hoc analysis of the order effect can be used to rank the six orders. This ranking is (by descending number of jumpers saved): 4 1 5 6 2 3. The only significant pair-wise differences between groups are between group 4 and group 2, and between group 4 and group 3. The “Counterbalanced Treatment Order” Table (see Table 21 on pg 183) shows that the treatment conditions were order balanced for the different groups. Therefore, the most obvious explanation for a significant order effect is that the particular treatment condition that subjects saw first differentially affected their process of constructing strategies for performing the dual-task. The powerful effect of first treatments is common and recognized in psychology as the “primacy effect” (Aronson 1995).

Subjects between groups saw different treatment conditions first. Using the “4 1 5 6 2 3” rank ordering, subjects from the six groups saw the following treatment conditions first: (group 4) negotiated interruption; (group 1) basecase — game only; (group 5) mediated interruption; (group 6) scheduled interruption; (group 2) basecase — matching only; and (group 3) immediate interruption.

This rank ordering can be explained in terms of decreasing amount of “intelligence” implemented in the method for coordinating interruptions. Subjects in group 4 saw the negotiated interruption treatment condition first. The negotiated solution makes the user directly responsible for deciding when to handle interruptions, and people are the most intelligent method for coordinating interruptions. Subjects in group 3 saw the immediate interruption treatment condition first. The immediate solution implements no intelligence in deciding when to interrupt. The other groups between 4 and 3 in the ranking lie in a decreasing order of intelligence implemented in their respective solutions.

It seems that subjects formed rigid task strategies based on whatever treatment they saw first. Another reasonable explanation for this order effect is degree of perceived control. If subjects felt they were in control of when to handle interruptions then they formed more successful strategies than if they felt that they had no control. Subjects who saw the negotiated solution first performed best, and subjects who saw the immediate solution first performed worst.

The one hour of practice given to all subjects was intended to negate any primacy effect from order of treatment conditions. However, it appears that subjects form rigid task strategies during the first treatment condition. The primacy effect on these task strategies is not negated by one hour of practice.

It is asserted that this order effect is the result of a stubborn primacy effect regarding subjects’ perception of degree of intelligence implemented in the first treatment condition they encountered, and that this does not pose a confound to the main effect.

5.6.1.1.2 General Practice Effect

A significant effect was found on within subjects order (“Within Subjects - Order (O)”). This is likely due to a general practice effect. Individual subjects are being affected by one or more effects like learning, boredom, or fatigue. The following line chart shows the average performance of all subjects over their 6 experimental treatments in the order they did them.

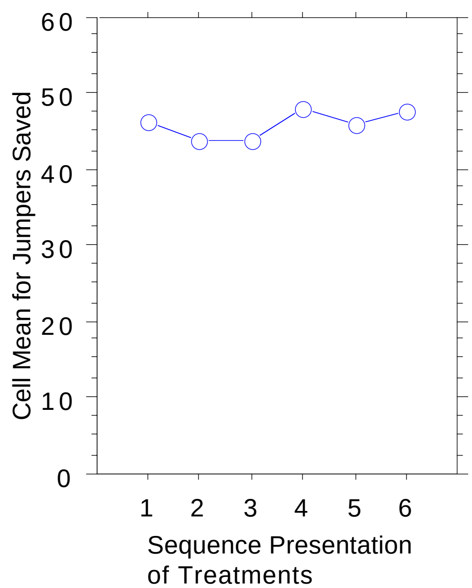

Figure 17.Line graph of the average number of jumpers saved, by the sequential order in which subjects performed the dualtask trial pairs.

The within-subjects order effect results from a slight increase in subjects performance. It is postulated that this increase is the result of some continued learning on the game task, and that this does not pose a confound to the main effect.

5.6.1.1.3 Interaction Effect

A significant interaction effect was found between interruption (main effect) and the within subjects order ("Within Subjects - Interrupt. X $Order_w$ (IO_w)"). This means that there is an interaction between the main effect (I_w) and the general practice effect (O_w). Treatment conditions affected subjects' performance differently depending on whether they received them first, second, third, fourth, fifth, or sixth. The following line chart shows the average number of jumpers saved by the sequential order subjects did the trial pairs, split by the treatment conditions.

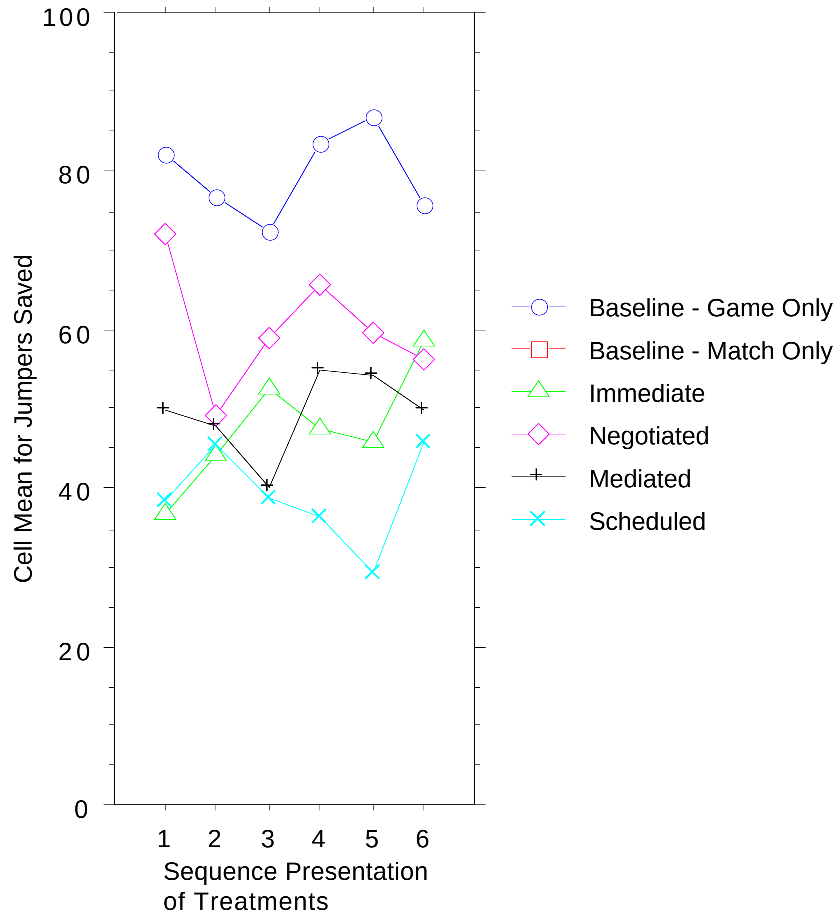

Figure 18.Line graph of the average number of jumpers saved, by the sequential order in which subjects performed the dualtask trial pairs, split by treatment condition.

There are a few crossovers, however, the differences do not dramatically affect the overall pattern. The within subject order (sequence) differentially affected subjects performance on different treatment conditions. This is not surprising given the observed between subjects order effect and the fact that different task strategies are more appropriate for different treatment conditions. Ideally, people would learn a different set of task strategies for accomplishing each of the different treatment conditions. However, there seems to be a large primacy effect that affects subjects ability to learn appropriate sets of task strategies for each different treatment condition.

It is postulated that this interaction effect is the result of interference from a primacy effect, but that this does not pose a confound to the main effect.

5.6.1.2 Number of Switches Between Game Task and Matching Task

Table 29 — Descriptive Statistics of Number of Task Switches

	Mean	Std. Dev.	Std. Error	Count	Min	Max	Median
Total	81.2	89.786	6.109	216	0	298	40
Baseline - Game Only	0	0	0	36	0	0	0
Baseline - Match Only	0	0	0	36	0	0	0
Immediate	249.6	17.524	2.921	36	194	274	253.5
Negotiated	110.1	60.548	10.091	36	22	298	109.5
Mediated	88.0	11.216	1.869	36	67	117	87.5
Scheduled	39.6	1.022	.170	36	36	40	40

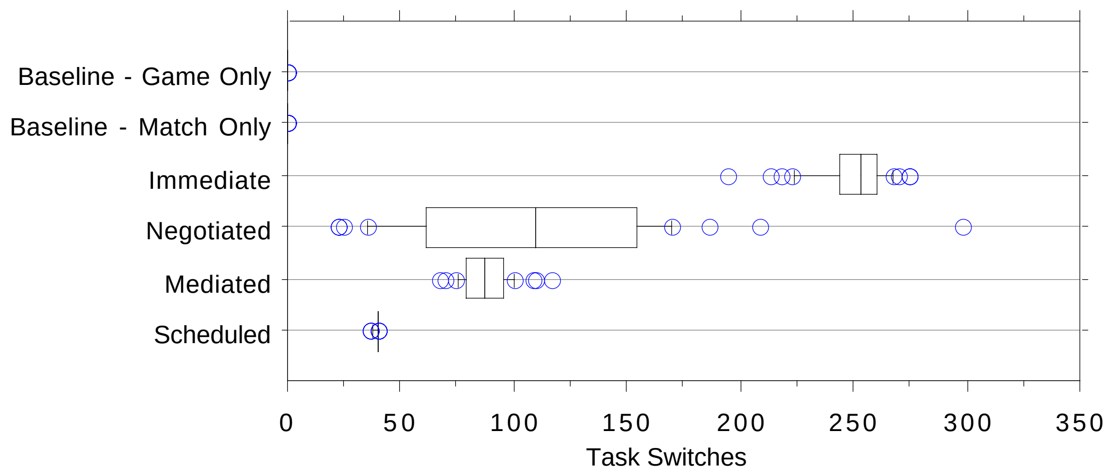

Figure 19. Box plot for number of task switches, by treatment condition.

Table 30 — Simple Latin Squares ANOVA for Number of Task Switches

Source	SS	df	ms	F	p
Total	1733247.333	215			
Between Subjects	27404.000	35			
Groups(IO_b)	10371.722	5	2074.344	3.654	.011*
$Error_b$	17032.278	30	567.743		
Within Subjects	1705843.333	180			
Interruption (I)	1589748.500	5	317949.700	577.490	<.0001*
Order (O)	5103.167	5	1020.633	1.854	.106
Interrupt. X $Order_w$ (IO_w)	28405.944	20	1420.297	2.580	.001*
$Error_w$	82585.722	150	550.571		

Note: a “*” indicates significance based on an alpha of .05.

Table 31 — Least Squares Means Table of Main Effect for Number of Task Switches

	Vs.	Diff.	Std. Error	t-Test	P-Value
Immediate	Negotiated	139.528	6.762	20.635	.0001*
	Mediated	161.583	6.762	23.896	.0001*
	Scheduled	210.000	6.762	31.057	.0001*
Negotiated	Mediated	22.056	6.762	3.262	.0016*
	Scheduled	70.472	6.762	10.422	.0001*
Mediated	Scheduled	48.417	6.762	7.160	.0001*

Note: a "*" indicates significance based on a corrected alpha of .0083 (.05 / 6).

Note the existence of variability in the number of task switches in the "Immediate" and "mediated" conditions. Although there are always the same number of matching tasks to perform in every trial, the number of task switches is not controlled. If a matching task occurs while a subject is already performing another matching task, then the new task is queued and presented immediately after they make a choice. No task switch occurs. There are two sources of variability for these matching task overlaps: (1) random scheduling of when the matching tasks occur; and (2) differences between subjects in how quickly they perform single matching tasks.

The ANOVA found a significant main effect. However, like the ANOVA for "jumpers saved" performance measure, this ANOVA also found other significant effects that are unrelated to the main hypothesis. These other significant effects are examined and it is determined that none of these secondary effects presents a confound for the main effect. Therefore these data can be said to support the main hypothesis.

5.6.1.2.1 Order Effect

A significant order effect was discovered ("Between Subjects - Groups (IOb)"). This means that the order in which subjects were presented the treatments affected their overall task switching behavior. This order effect was as unexpected as that found for "jumpers saved" performance measure. The following line chart shows the average performance for each of the six order groups of subjects for the different treatments. There is a clear order effect that affects the pair-wise relationship between the "negotiated" treatment condition and the "mediated” treatment condition. However, the order effect does not confound the overall significance of the main effect.

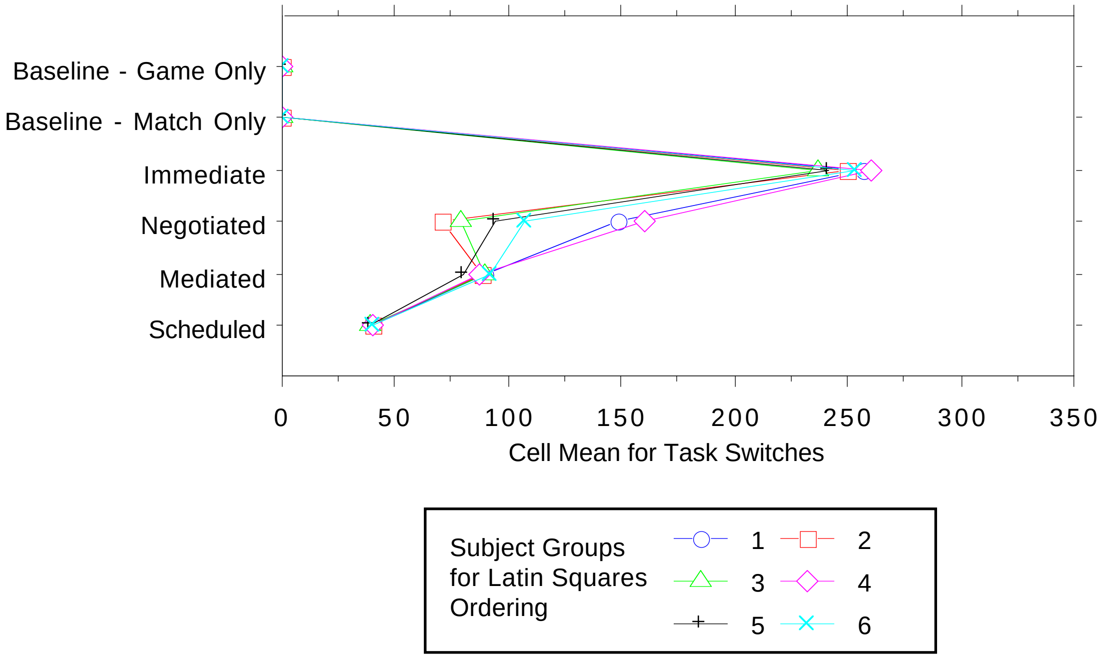

Figure 20.Line graph of task switches for the different treatment conditions, split by between-subjects groups 1-6 (the Latin squares ordering groups).

A post-hoc pair-wise analysis may be useful in explaining this order effect.

Table 32 — Least Squares Means Table of Between Subjects Order Effect for Number of Task Switches

	Vs.	Diff.	Std. Error	t-Test	P-Value
1	2	17.861	6.311	2.830	.0082
	3	19.861	6.311	3.147	.0037
	4	1.056	6.311	.167	.8683
	5	15.306	6.311	2.425	.0215
	6	8.889	6.311	1.409	.1693
2	3	2.000	6.311	.317	.7535
	4	-16.806	6.311	-2.663	.0123
	5	-2.556	6.311	-.405	.6884
	6	-8.972	6.311	-1.422	.1654
3	4	-18.806	6.311	-2.980	.0057
	5	-4.556	6.311	-.722	.4760
	6	-10.972	6.311	-1.739	.0923
4	5	14.250	6.311	2.258	.0314
	6	7.833	6.311	1.241	.2241
5	6	-6.417	6.311	-1.017	.3174

Note: a “*” indicates significance based on a corrected alpha of .0033 (.05 / 15).

This post-hoc analysis of the order effect can be used to rank the six orders. This ranking is (by descending number of task switches): 1 4 6 5 2 3. The are no significant pair-wise differences between groups. This rank order is almost identical to the rank order for the “jumpers saved” performance measure (see pg. 206): 4 1 5 6 2 3. From this similarity of ranks it is reasonable to conclude that the significant between subjects order effects found for both “jumpers saved” and “task switches” result from the same primacy effect. Subjects’ perceived degree of intelligence implemented in the first treatment condition they encounter affects their task strategies related to saving jumpers and switching between tasks of the multitask. The one hour of practice was insufficient for negating this primacy effect.

It is postulated that this order effect is the result of a stubborn primacy effect regarding subjects’ perception of degree of intelligence implemented in the first treatment condition they encountered. This primacy effect clearly effects the pair-wise difference between the “negotiated" treatment condition and the "mediated" treatment condition. However, the order effect does not confound the overall significance of the main effect

5.6.1.2.2 Interaction Effect

A significant interaction effect was found between within subjects treatment and order ("Within Subjects - Interrupt. X Order$_w$ (IO$_w$)"). Treatment conditions affected subjects' performance differently depending on whether they received them first, second, third, fourth, fifth, or sixth. The following line chart shows the average performance of all subjects over their 6 experimental treatments in the order they did them split by the treatment conditions.

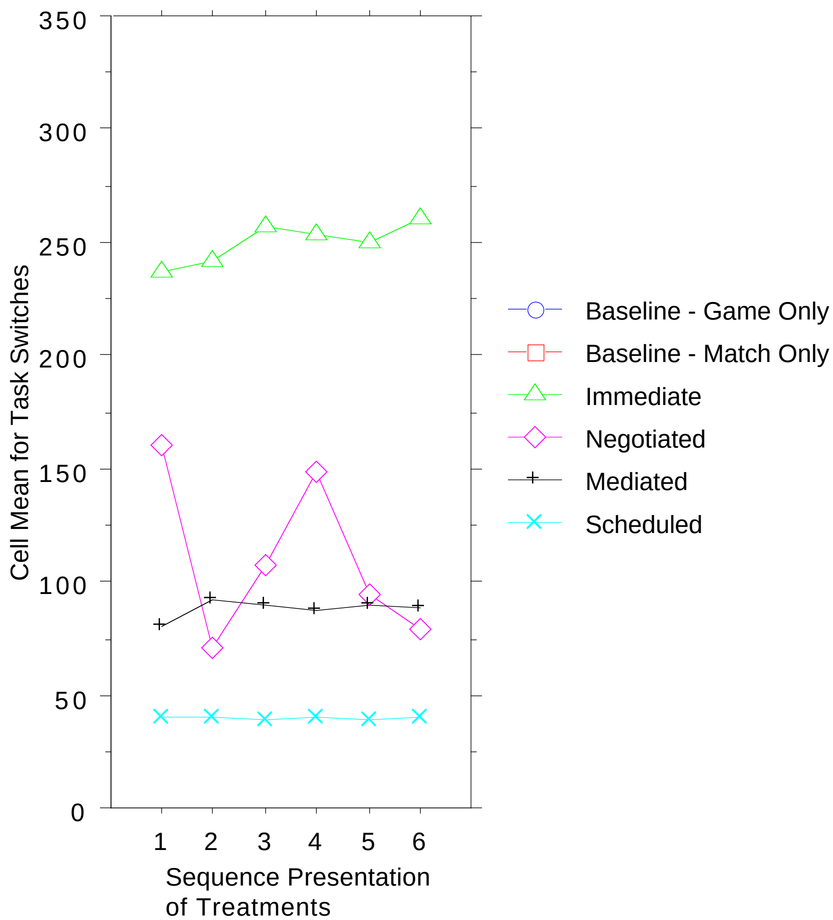

Figure 21.Line graph of the average number of task switches, by the sequential order in which subjects performed the dualtask trial pairs, split by treatment condition.

There is a strange crossover between the “negotiated” solution and the “mediated” solution. However, the overall pattern is largely unaffected by this interaction effect. This observed interaction effect is similar to that seen for the “jumpers saved” performance measure. There seems to be a large primacy effect that influences subjects ability to learn appropriate sets of task strategies for each different treatment condition.

It is postulated that this interaction effect is the result of interference from a primacy effect, but that this does not pose a confound to the main effect. Pair-wise comparison between the “negotiated” condition and the “mediated” condition, however, are confounded and are therefore not meaningful.

5.6.1.3 NUMBER MATCHED WRONG

Table 33 — Descriptive Statistics for Number Matched Wrong (160 possible)

	Mean	Std. Dev.	Std. Error	Count	Min	Max	Median
Total	8.0	9.606	.654	216	0	63	5
Baseline - Game Only	0	0	0	36	0	0	0
Baseline - Match Only	6.1	5.841	.974	36	1	29	4.5
Immediate	12.7	10.996	1.833	36	1	53	9
Negotiated	9.8	9.400	1.567	36	1	52	7.5
Mediated	10.2	10.868	1.811	36	0	63	7.5
Scheduled	9.4	10.165	1.694	36	0	50	7

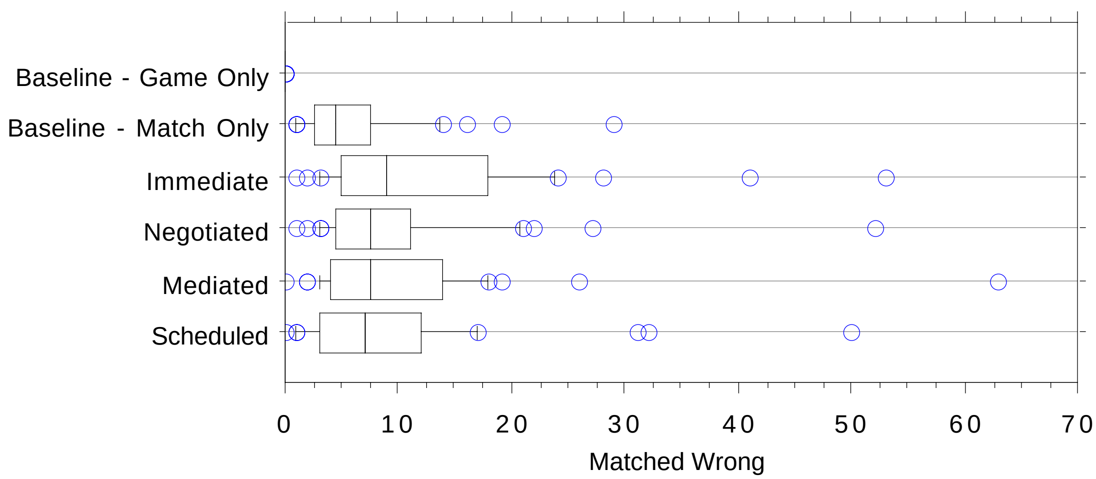

Figure 22. Box plot for number matched wrong, by treatment condition.

Table 34 — Simple Latin Squares ANOVA for Number Matched Wrong

Source	SS	df	ms	F	p
Total	19840.884	215			
Between Subjects	11686.384	35			
Groups(IO_b)	2375.412	5	475.082	1.531	.210
$Error_b$	9310.972	30	310.366		
Within Subjects	8154.500	180			
Interruption (I)	3572.245	5	714.449	28.375	<.0001*
Order (O)	113.301	5	22.660	0.900	.483
Interrupt. X $Order_w$ (IO_w)	692.093	20	34.605	1.374	.143
$Error_w$	3776.861	150	25.179		

Note: a "*" indicates significance based on an alpha of .05.

Table 35 — Least Squares Means Table of Main Effect for Number Matched Wrong

	Vs.	Diff.	Std. Error	t-Test	P-Value
Baseline - Match Only	Immediate	-6.556	.941	-6.963	.0001*
	Negotiated	-3.611	.941	-3.835	.0002*
	Mediated	-4.056	.941	-4.308	.0001*
	Scheduled	-3.222	.941	-3.422	.0008*
Immediate	Negotiated	2.944	.941	3.127	.0022*
	Mediated	2.500	.941	2.655	.0090
	Scheduled	3.333	.941	3.540	.0006*
Negotiated	Mediated	-.444	.941	-.472	.6377
	Scheduled	.389	.941	.413	.6803
Mediated	Scheduled	.833	.941	.885	.3779

Note: a “*” indicates significance based on a corrected alpha of .005 (.05 / 10).

The ANOVA found a significant main effect, therefore these data can be said to support the main hypothesis.

5.6.1.4 Number of Matches Not Done

Table 36 — Descriptive Statistics for Number of Matches Not Done (160 possible)

	Mean	Std. Dev.	Std. Error	Count	Min	Max	Median
Total	4.3	7.088	.482	216	0	52	2
Baseline - Game Only	0	0	0	36	0	0	0
Baseline - Match Only	3.0	1.183	.197	36	2	7	3
Immediate	1.2	.941	.157	36	0	3	1
Negotiated	11.3	14.395	2.399	36	0	52	3.5
Mediated	2.6	1.626	.271	36	0	6	3
Scheduled	7.6	1.360	.227	36	4	10	8

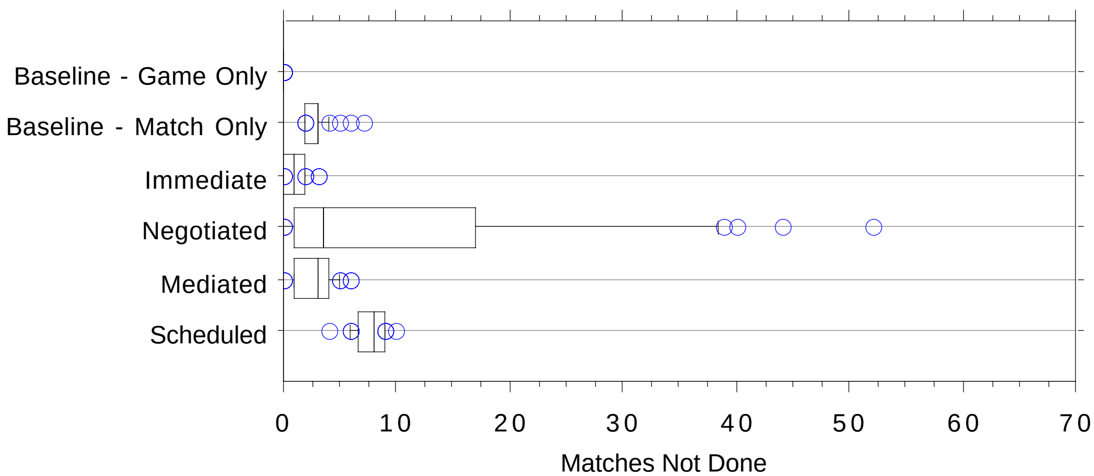

Figure 23. Box plot for number of matches not done, by treatment condition.

Table 37 — Simple Latin Squares ANOVA for Number of Matches Not Done

Source	SS	df	ms	F	p
Total	10801.958	215			
Between Subjects	1324.458	35			
Groups(IO_b)	205.486	5	41.097	1.102	.380
$Error_b$	1118.972	30	37.299		
Within Subjects	9477.500	180			
Interruption (I)	3311.930	5	662.386	18.735	<.0001*
Order (O)	185.931	5	37.186	1.052	.390
Interrupt. X $Order_w$ (IO_w)	676.444	20	33.822	0.957	.518
$Error_w$	5303.194	150	35.355		

Note: a "*" indicates significance based on an alpha of .05.

Table 38 — Least Squares Means Table of Main Effect for Number of Matches Not Done

	Vs.	Diff.	Std. Error	t-Test	P-Value
Baseline - Match Only	Immediate	1.806	1.533	1.177	.2414
	Negotiated	-8.278	1.533	-5.398	.0001*
	Mediated	.361	1.533	.235	.8142
	Scheduled	-4.611	1.533	-3.007	.0032*
Immediate	Negotiated	-10.083	1.533	-6.575	.0001*
	Mediated	-1.444	1.533	-.942	.3481
	Scheduled	-6.417	1.533	-4.184	.0001*
Negotiated	Mediated	8.639	1.533	5.634	.0001*
	Scheduled	3.667	1.533	2.391	.0184
Mediated	Scheduled	-4.972	1.533	-3.242	.0015*

Note: a “*” indicates significance based on a corrected alpha of .005 (.05 / 10).

The ANOVA found a significant main effect, therefore these data can be said to support the main hypothesis.

5.6.1.5 Average Age of Matching Tasks When Completed

Subjects did not always complete all the matching tasks before the end of trials (see the "number of matches not done" performance measure). It was decided that those unfinished matching tasks would be included in the calculation of the average time, and that their assigned age would be the time from their scheduled onset until the end of the trial.

Table 39 — Descriptive Statistics for Average Match Age When Completed (in msec)

	Mean	Std. Dev.	Std. Error	Count	Min	Max	Median
Total	11854.369	15793.234	1074.593	216	0	109476.375	5067.312
Baseline - Game Only	0	0	0	36	0	0	0
Baseline - Match Only	3121.764	837.445	139.574	36	2076.000	5083.375	2893.125
Immediate	2871.451	745.829	124.305	36	1931.375	4924.875	2604.750
Negotiated	31074.729	27655.589	4609.265	36	4834.000	109476.375	23032.500
Mediated	13730.181	1287.709	214.618	36	9417.125	15647.750	13808.812
Scheduled	20328.090	1006.949	167.825	36	17732.500	22158.625	20243.625

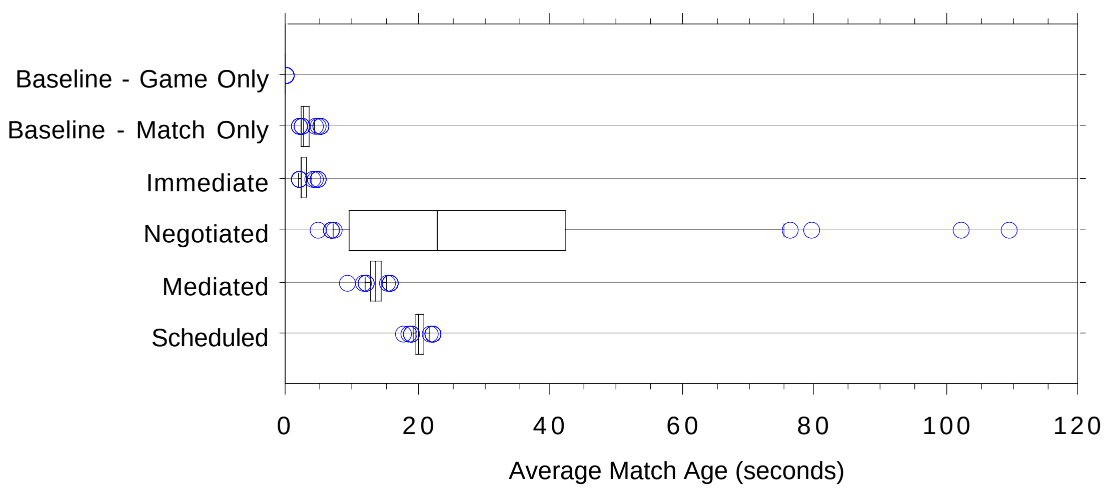

Figure 24. Box plot for average match age when completed, by treatment condition.

Table 40 — Simple Latin Squares ANOVA for Average Match Age When Completed

Source	SS	df	ms	F	p
Total	53626643381.055	215			
Between Subjects	4616515789.482	35			
Groups(IO_b)	299269386.409	5	59853877.282	0.416	.834
$Error_b$	4317246403.073	30	143908213.436		
Within Subjects	49010127591.573	180			
Interruption (I)	26719996732.005	5	5343999346.401	38.147	<.0001*
Order (O)	285933646.648	5	57186729.330	0.408	.843
Interrupt. X $Order_w$ (IO_w)	990658729.628	20	49532936.481	0.354	.996
$Error_w$	21013538483.292	150	140090256.555		

Note: a “*” indicates significance based on an alpha of .05.

Table 41 — Least Squares Means Table of Main Effect for Average Match Age When Completed

	Vs.	Diff.	Std. Error	t-Test	P-Value
Baseline - Match Only	Immediate	250.312	3054.299	.082	.9348
	Negotiated	-27952.965	3054.299	-9.152	.0001*
	Mediated	-10608.417	3054.299	-3.473	.0007*
	Scheduled	-17206.326	3054.299	-5.633	.0001*
Immediate	Negotiated	-28203.278	3054.299	-9.234	.0001*
	Mediated	-10858.729	3054.299	-3.555	.0005*
	Scheduled	-17456.639	3054.299	-5.715	.0001*
Negotiated	Mediated	17344.549	3054.299	5.679	.0001*
	Scheduled	10746.639	3054.299	3.519	.0006*
Mediated	Scheduled	-6597.910	3054.299	-2.160	.0327

Note: a “*” indicates significance based on a corrected alpha of .005 (.05 / 10).

The ANOVA found a significant main effect, therefore these data can be said to support the main hypothesis.

5.6.2 Support for Subhypothesis

The subhypotheses (Section 5.4.2 on pg 177) make predictions of differences between individual treatment conditions. Tests of these subhypotheses are made with the results of the post-hoc analyses of main effects reported in the previous sections for the five measures of performance.

The following table summarizes these results to facilitate testing the subhypotheses. The post-hoc analyses revealed the significant differences between pairs of treatments for all five measures of performance. A useful summary notation is introduced here to summarize these differences. The treatment conditions for each performance measure are sorted left to right by their corresponding value. (Note that not all treatment conditions are included in this summary because some are inappropriate for pair-wise comparison.) Greater-than, “>“, is used to denote all statistically significant separations of value between treatment conditions. Slash, “/ ”, is used to denote non-significant separations of value between treatment conditions. Brackets, “[“ and “]”, are used to denote non-significant grouping of treatments. Question mark, “?”, is used to denote pair-wise comparisons that are confounded by secondary effects.

Table 42 — Summary of Post-Hoc Pair-Wise Analyses for the Five Measures of Performance

performance measure	rank of treatments by post-hoc analysis
jumpers saved	base-game > negotiated > [mediated / immediate] > scheduled
task switches	immediate > [negotiated ? mediated] > scheduled
matched wrong	immediate > [negotiated / scheduled] > base-match; [immediate / mediated]; [mediated / negotiated / scheduled]
matches not done	[negotiated / scheduled] > [base-match / mediated / immediate]
age of matching task	negotiated > [scheduled / mediated] > [base-match / immediate]

Some performance measures are easily associated with success or failure on the experimental tasks. The “jumpers saved” measure, for example, can be easily viewed as a measure of successful performance of the game task. Therefore, the four different solutions for “Method of Coordination” of interruptions can be ranked from “best” to “worst” by looking at the post-hoc analysis results for the different performance measures. Only one of the five performance measures considered here is not easily associated with success or failure of the dualtask — “task switches.” Task switching is only indirectly related to saving jumpers and making matches.

The following table summarizes the "best" to "worst" ranking of the four different user-interface design solutions. All ranking is done from 1 (best) to 5 (worst). Ranking is based on the results of the post-hoc pair-wise analysis which is summarized in Table 42 above. Non-statistically significant ranking is represented with supplementary letter sequences. Table 43 ranks the competing solutions by which supported more success on the experimental task, therefore, those performance measures that are stated in a negative way ("matched wrong," "matches not done," and "age of matching task") reveal inverse rankings. Note that since this experiment emphasizes internal validity at the expense of external validity care must be taken in generalizing these findings to other user-interface design contexts.

Table 43 — Ranking of the Four Solutions for "Method of Coordination" of Interruptions

	base-game	base-match	immediate	negotiated	mediated	scheduled
jumpers saved	1	—	3.5(b)	2	3.5(a)	5
matched wrong	—	1	5(b)	3(b)	3(c)/5(a)	3(a)
matches not done	—	1(c)	1(a)	5(b)	1(b)	5(a)
age of matching task	—	1(b)	1(a)	5	3(a)	3(b)

5.6.2.1 SUBHYPOTHESIS 1

Subhypothesis 1 says "the total number of task switches encountered by a user is affected by which method of coordinating interruptions is implemented by the user interface" (pg. 177). The discovery of a significant main effect for the "switches between tasks" performance measure (see Section 5.6.1.2 on pg 210) is adequate evidence to confirm this prediction.

5.6.2.2 SUBHYPOTHESIS 2

Subhypothesis 2 says "there is an inverse relationship between the total number of times people switch tasks while performing a multitask and people's performance on those tasks" (pg. 177). This subhypothesis is based on the recognized phenomenon of automation deficit, i.e., the transition delays people experience when switching between tasks.

This subhypothesis predicts a directional correlation between two performance measures. The following table shows the results of the nonparametric Spearman rank correlation analysis for number of task switches and each of the other four performance measures.

Table 44 — Correlation between the number of task switches and each of the other four performance measures

	Rho	P-Value
jumpers saved	.340	<.0001*
matched wrong	.169	.0432*
matches not done	-.725	<.0001*
age of matching tasks	-.843	<.0001*

"*" indicates statistical significance at alpha = .05.

Spearman's Rho for "jumpers saved" is statistically significant but in the wrong direction. The results for these data do not support subhypothesis 2. Spearman's Rho for "matched wrong" is positive, but since the "matched wrong" is a negative measure of performance this statistically significant correlation is in the right direction. The results for these data do support subhypothesis 2.

Spearman's Rho for "matches not done" and "age of matching tasks" are both negative, but since these are negative measures of performance these statistically significant correlations are in the wrong direction. The results for these data do not support subhypothesis 2.

These results found conflicting support for the validity of subhypothesis 2. It is concluded that automation deficit is useful for predicting only some kinds of human performance.

Subhypothesis 2 is only relevant to this experiment because of its indirect statement about the four user interface solutions for coordinating user-interruption. The implication is that if such a negative relationship exists between task switching and performance, then the "best" user interface solution is whichever one results in the fewest number of task switches. Table 43 (pg 223) shows that the scheduled solution produces the fewest task switches. Is it then the "best" choice for a design solution? The results of the Spearman rank correlation suggests that the answer will be "only sometimes."

The scheduled solution is not the "best" solution in some cases. It is ranked the "worst" solution for the "jumpers saved" and "matches not done" data. The scheduled solution is ranked less successful than the immediate solution for the "age of matching tasks" data. However, the scheduled solution is ranked better than the immediate solution for the "matched wrong" data.

The correlation analysis and the performance rank comparison agree that the "matched wrong" performance measure is the only one that supports subhypothesis 2. For every matching task, the experimental design did not directly control whether subjects saw the game task or a matching task immediately prior. Instead, the preceding context of each matching task is a by-product of the interaction between how often the subject switches between tasks and the accumulation of randomly scheduled matching tasks.

Subjects made more errors on matching tasks when they switched tasks more frequently. The scheduled solution had the least task switching and the fewest wrong matches. In this solution, subjects performed long series of matching tasks that accumulated during the preceding 25 second scheduled interval. The other three user interface design solutions (immediate, negotiated, and mediated) had more frequent task switching and therefore subjects performed fewer matching tasks in sequence without switches.

The effects of automation deficit are only relevant to performance of the matching task in situations where subjects had to perform matching tasks after immediately switching from the game task. The scheduled solution had fewer of these situations than the other three solutions and also has the fewest matching errors.

Why is there no inverse relation between the number of task switches and the data for "jumpers saved," "matches not done," and "age of matching task" performance? The "matches not done" and "age of matching task" observations are influenced by other effects that outweigh the effects of automation deficit. These two performance measures deal with getting matching tasks done in a timely way. The results of this experiment show that user interface design solutions that support frequent task switching actually improve the timely handling interruptions. This results from the fact that the performance delays caused by automation deficit are small compared with the time required to perform matching tasks. The "jumpers saved" observations suggest that the time scale for playing the game task is much larger than the time scale of automation deficit. The game task is therefore relatively insensitive to the effects of automation deficit.

5.6.2.3 SUBHYPOTHESIS 3

Subhypothesis 3 says "the methods of coordinating interruptions that are most direct and immediate will create user interfaces that support higher user performance on multitasks than the methods of coordination that express interruptions with delayed timing or require interactive negotiation" (pg. 178). This subhypothesis is based on the assertion that people do not have sociolinguistic face-wants relative to their computers.

This view predicts that the immediate user interface design solution is the "best" solution. The "jumpers saved" data do not support subhypothesis 3. The immediate solution is ranked less successful than the negotiation solution (see Table 43 on pg 223). The "matched wrong" data do not support subhypothesis 3. The immediate solution is ranked "worst."

The "matches not done" data do support subhypothesis 3. The immediate solution is ranked "best" along with the mediated solution. The "age of matching tasks" data, also provides support subhypothesis 3. The immediate solution is ranked "best."

These results found conflicting support for the validity of subhypothesis 3. It is concluded that the theory of politeness is useful for predicting only some kinds of human performance.

The immediate method of coordinating user -interruptions is clearly the "best" solution for supporting good "matches not done" and "age of matching task" performance. The "Bald On-

Record" non-politeness implemented by the immediate solution puts people into the situation of performing well by necessity. Every matching tasks must be handled quickly as soon as it occurs.

The success of this non-politeness strategy, however, does not extend into performance for "jumpers saved" and "matched wrong." It seems that face-wants (see Section 2.14 on pg 80; and Brown and Levinson 1987) are not the main concern for all types of interaction. There are timing constraints related to the experimental multitask that have nothing to do with subjects' face-wants. Some parts of the multitask are sensitive to breaks in user-control. If a subject is unfortunately interrupted away from a task in a critical moment, then they can not help but make errors.

5.6.2.4 SUBHYPOTHESIS 4

Subhypothesis 4 says "the negotiation method for coordinating interruption will create user interfaces that support higher user performance on multitasks than the other methods of coordination" (pg. 178). This subhypothesis is based on the assertion that people have highly developed skills of negotiating for entrance into joint activities. Whenever people are interrupted, they automatically engage their negotiation skills to try to coordinate handling that interruption.

This view predicts that the negotiated user interface design solution is the "best" solution. The "jumpers saved" data do support subhypothesis 4. The negotiated solution is ranked second only to the base-game condition (see Table 43 on pg 223). The "matched wrong" data do provide some support for subhypothesis 4. The negotiated solution is ranked better than the immediate solution.

The "matched not done" data do not support subhypothesis 4. The negotiated solution is ranked "worst." The "age of matching tasks" data do not support subhypothesis 4. The negotiated solution is ranked "worst."

These results found conflicting support for the validity of subhypothesis 4. It is concluded that the theory of joint activities is useful for predicting only some kinds of human performance.

Subhypothesis 4 was based on the theory of joint activities. The negotiated interruptions solution can be viewed as an attempt to exploit people's natural ability to negotiate interruptions, i.e., for entrance into joint activities. The results from this experiment show that this strategy was partially successful. The negotiated solution produced the "best" performance on the "jumpers saved" and "good" performance on the "matched wrong." This combination would seem to indicate the negotiated solution as a clear winner.

The negotiated solution allows people to exercise their strength in dynamic negotiation. It, however, also allows people to exercise their weakness in handling interruptions in a timely way. The results for the "matches not done" and "age of matching task" support this conclusion. When people control of choosing convenient times to handle interruptions, they sometimes decide it is convenient to put them off indefinitely.

5.6.2.5 SUBHYPOTHESIS 5

Subhypothesis 5 says "the negotiation method for coordinating interruption will create user interfaces that result in users exhibiting more errors of omission on some parts of computer-based multitasks" (pg. 179). This subhypothesis is based on the assertion that people are unreliable in their ability to perform situational awareness during multitasks.

This view predicts that the negotiated user interface design solution is the "worst" solution for ensuring that all interruptions are completed. The "matches not done" data do support subhypothesis 5. The negotiated solution is ranked "worst" (see Table 43 on pg 223).

These results found support for the validity of subhypothesis 5. It is concluded that the theory of situational awareness is useful for predicting human performance in failing to perform some parts of multitasks.

5.6.2.6 SUBHYPOTHESIS 6

Subhypothesis 6 says "the mediated method for coordinating interruption will create user interfaces that support higher user performance on multitasks than the other methods of coordination" (pg. 179). This subhypothesis is based on the assertion that people experience transition points between stages of accomplishing activities, and that automatic mediation to present interruptions at those transition points will negate the problems of human interruption.

This view predicts that the mediated user interface design solution is the "best" solution. The "jumpers saved" data do not support subhypothesis 6. The mediated solution is ranked less successful than the negotiation solution (see Table 43 on pg 223). The "matched wrong" data do not support subhypothesis 6. The mediated solution is ranked "worst" along with the immediate solution.

The "matches not done" data do support subhypothesis 6. The mediated solution is ranked "best" along with the immediate solution. The "age of matching tasks" data do not support subhypothesis 6. The mediated solution is ranked less successful than the immediate solution.

These results found conflicting support for the validity of subhypothesis 6. It is concluded that the idea of mediation at natural task transition points is only useful for predicting people's ability to complete all interruption tasks in a timely way.

The mediated solution for coordinating user-interruption has produced mediocre results. The mediation solution promised to combine the benefits of the immediate solution with the benefits of the negotiated solution. The user would be freed from any direct responsibility or control for managing when to handle interruptions; and they would also only have to process interruptions at times that they found convenient. The mediated solution as implemented in this experiment did not deliver on this promise. Why not?

The answer to this question is probably the same answer to why the Pilot's Associate project (Hammer and Small 1995) was not fully successful in creating an intelligent decision aid that would mediate interrupting pilots of military single-seat aircraft. It was concluded that predicting people's interruptibility is a difficult problem. The complexity of the problem is illustrated by the large and diverse set of theoretical constructs present in Chapter 2 of this dissertation. The mediated solution is certainly interesting and the possible payoffs are great. However, without a practical theoretical model of human interruption no design guidelines can be formed.

The representative mediation solution implemented in this experiment is simplistic. Any good engineer could think of other, more "intelligent," algorithms for calculating people's interruptibility. However, an important goal for this experiment was to make progress toward a general solution to the problem of human interruption. The potential rewards of the mediated solution are attractive, however, the results of this experiment do not offer any general guidelines that could insure success.

5.6.2.7 SUBHYPOTHESIS 7

Subhypothesis 7 says "the scheduled method for coordinating interruption will create user interfaces that support higher user performance on multitasks than the other methods of coordination" (pg. 180). This subhypothesis is based on findings that people regard the predictability of response time an important influence on their behavior.

This view predicts that the scheduled user interface design solution is the "best" solution. The "jumpers saved" data do not support subhypothesis 7. The scheduled solution is ranked "worst" (see Table 43 on pg 223). The "matched wrong" data do provide some support for subhypothesis 7. The scheduled solution is ranked better than the immediate solution.

The "matches not done" data do not support subhypothesis 7. The scheduled solution is ranked "worst" along with the negotiated solution. The "age of matching tasks" data do not support subhypothesis 7. The scheduled solution is ranked less successful than the immediate solution.

These results found conflicting support for the validity of subhypothesis 7. It is concluded that the people's difficulty with uncertainty in timing is only useful for predicting people's error rate in complete interruption tasks.

The representative scheduled solution implement in this experiment for coordinating user-interruptions uses a fixed 25 second scheduling interval. This interval was chosen as appropriate for three reasons. First, it seemed appropriate to pilot test subjects. Second, it is a large enough interval to insure that subjects are given sufficient time to perform a significant amount of the game task without interruption. Each jumping character in the game task requires 13.7 seconds from its beginning jump until after its third (and last) bounce to safety. The 25 second interval insured that there was nearly enough time to successfully save two non-overlapping game subtasks. Third, it is a small enough interval to insure that several interval cycles could be completed within the 4.5 minute trials.

The observed level of success for this solution was "bad," with the exception of the "matched wrong" performance. It was, unfortunately, beyond the scope of this experiment to compare the relative success of a range of possible interval times. It may be true that another interval time, or perhaps one that was custom sized for each individual subject, would have be more successful.

The "matched wrong" performance may support the predictive power of the finding that people perform better on predictable tasks than not. However, for this experimental multitask, the most relevant measure of performance was "jumpers saved," and the scheduled solution was found to be the "worst" solution of the four possible.

5.6.2.8 SUBHYPOTHESIS 8

Subhypothesis 8 says "the total number of task switches encountered by a user will be less for user interfaces created with the scheduled method for coordinating interruptions than for user interfaces created with other interruption coordination methods" (pg. 180). This subhypothesis is based on the same assertions as subhypothesis 7. The predicted ranking is: "scheduled >

[negotiated mediated immediate].” The appropriate measure of performance for testing subhypothesis 8 is: task switches.

The “task switching” data do support subhypothesis 8. Its ranking (see Table 42 on pg 222), “immediate > [negotiated ? mediated] > scheduled,” found that the “scheduled” treatment condition produced the fewest task switches of all solutions.

These results found support for the validity of subhypothesis 8. It is concluded that the a scheduled solution for coordinating interruption can produce fewer task switches than the other solutions.

5.6.2.9 CONFLICTING RESULTS

Most subhypotheses were met with mixed support. The conflicting results from this experiment highlight the complexity of the interruption problem. Many of the theories and findings contained in the literature that have predictive power have only been tried in tightly constrained contexts. Their use for general interdisciplinary investigation is therefore difficult. There are many useful and relevant works that only investigate single task domains in isolation. These can be used to predict the relative appropriateness of alternative user interface design solutions for portions of user-interruption designs. However, good performance on one task does not insure good performance on all the tasks of a multitask.

5.6.3 The "Best" User Interface Solution

A naive strategy would be to just sum up the column totals for the design solution ranks and choose the one with the best total ranking. Table 43 (pg 223) is reproduced below with the column totals.

Table 45 — Ranking of the Four Solutions for "Method of Coordination" of Interruption (with Totals)

	base-game	base-match	immediate	negotiated	mediated	scheduled
jumpers saved	1	—	3.5(b)	2	3.5(a)	5
matched wrong	—	1	5(b)	3(b)	3(c)/5(a)	3(a)
matches not done	—	1(c)	1(a)	5(b)	1(b)	5(a)
age of matching task	—	1(b)	1(a)	5	3(a)	3(b)
Total	—	—	10.5	15	10.5/12.5	16

The columnar totals show that, everything being equal, the "immediate" and "mediated" solutions are equally the "best" solutions to the user-interruption problem. Problem solved, right? Well, no. Each system that interrupts its users will have its own user requirements. These requirements will emphasize which kinds of user performance are the most important for the overall success of the user interface design. The summary of rank solutions would have to be a weighted sum relative to each particular system's user requirements.

Also, even assuming that the five performance measures were equally important, the problem would not be solved. There is no "always best" solution for these four performance measures, and there are many other performance measures that have not been considered here. The results for those other kinds of performance could sway the straight sum totals in other directions.

The "best" solution may be some custom made hybrid that combines the different types of solutions in a meaningful way. It was beyond the scope of this experiment to examine interactions between the four kinds of solutions to the user-interruption problem. Perhaps general design guidelines can be discovered in the future that will guide the construction of such hybrid solutions.

5.6.4 Suggestions Toward the Creation of Design Guidelines

Tests of the subhypotheses of this experiment had mixed results. This supports a conclusion that there is no “best” interruption coordination solution for all measures of human performance. Instead, each solution has pros and cons for each different kind of human performance. Further investigation is necessary to uncover enough practical information to make reliable design guidelines.

These results alone do not have both the internal and external validity required for creating general design guidelines. They do, however, have some value and it may be useful to speculate about possible design guidelines. The following discussion of design guidelines is speculation and ignores concerns about validity.

The one result that seems clear from this experiment is that the appropriate method for coordinating interruption is relative to the kind of user performance that is most important. Therefore, the first step in choosing a user-interface design solution for a system that will interrupt its users is to perform a detailed user requirements specification. This specification must include careful analysis of how each kind of relevant user performance contributes to the overall successful function of the entire system.

General design guidelines must be made in terms of what is the best solution for a given set of important user performance types.

5.6.4.1 User interface Design Guidelines

These guidelines indicate the most appropriate method for coordinating user-interruptions in user interfaces for systems that support computer-based multitasks for people.

1. If accuracy on a continuous task is most important for success then choose the negotiated solution.
2. If accuracy on an intermittent task is most important for success then do not choose the immediate solution.
3. If completeness in performing every part of an intermittent task is most important for success then choose either the immediate solution or the mediated solution.
4. If promptness in performing every part of an intermittent task is most important for success then choose the immediate solution.
5. If completeness and promptness in performing every part of an intermittent task are both important for success then choose the immediate solution and do not choose the negotiated solution.
6. If accuracy on a continuous task and accuracy on an intermittent task are both important for success then choose the negotiated solution.
7. If accuracy on a continuous task and completeness in performing every part of an intermittent task are both important for success then do not choose the scheduled solution.

CHAPTER 6: CONCLUSIONS AND FUTURE WORK

6.1 AIMS AND OBJECTIVES

The usefulness of this dissertation is summarized and evaluated. After reading this chapter, the reader should be familiar with the contributions of this dissertation and their scope. The reader should also understand what kinds of future works are important for solving the problems associated with human interruption.

6.2 OVERVIEW

There are several parts of this dissertation. This work, however, pursues a single goal. This dissertation has been an attempt to make a first general investigation of the problem of human interruption during human-computer interaction (HCI). This recognized problem had not yet been addressed in a general way.

A review of the literature uncovered the fact that although many authors had identified this topic as a critical problem, none had made any general theoretical tools. Without such theoretical tools this dissertation had to take one of two different approaches: (1) conduct a series of empirical tests to discover non-generalizable information about the problem within one tightly limited context; or (2) create and partially validate the first theoretical tools for generalizable investigations of this problem. The second approach was chosen.

The different parts of the dissertation are reviewed to determine whether the main goal has been met. Contributions are identified, and the results of the experiment are discussed for indications of potential design guidelines. The future utility of the tools created here is discussed.

6.3 MOTIVATION

This dissertation has several parts, and a summary of its contributions makes these accomplishments more easily accessible.

6.4 CONTRIBUTIONS

The products of a successful dissertation must make original contributions to an important problem. The "Introduction" chapter (pg. 1) establishes the importance and timeliness of the problem of human interruption during HCI, and also elucidates the total lack of any general theoretical tools for addressing this problem. The major products of this dissertation are the General Definition of Human Interruption and the Taxonomy of Human Interruption. These are unique and significant.

It is useful to review the contributions of this dissertation. First, an identified but unsolved problem of human-computer interactions was highlighted. The problem was discussed with the perspective that interrupting people is not "bad" but only complicated and susceptible to causing human errors. Some illustrative real-world examples were provided. Background efforts to solve this problem were discussed and found to be of little general use. A useful theory of human multitasking was used to detail the different types of human cognition and the different possible types of human interruption. The importance, scope, and timeliness of this problem were thoroughly discussed. This description contributes to research by providing a detailed problem statement.

Second, 126 theoretical constructs of human interruption were identified from a broad survey of current literature. The theoretical constructs from a broad selection of relevant domains were identified and collected as a foundation for synthesizing generally useful theoretical tools. Representative literature from the following research domains, and their several sub-domains, were analyzed for relevant theoretical constructs: linguistics, psychology, social psychology, sociology, computer-supported cooperative work (CSCW), human-computer interaction (HCI), supervisory control, intelligent agents, and etymology. This analysis identified 126 theoretical constructs of human interruption. This set of theoretical constructs can be used as a theoretical foundation for conducting theory-based work. This rough theoretical foundation is a contribution in two ways: (1) as a foundation for creating useful research tools; and (2) as a stepping stone for the long-term effort of creating a general model of human interruption.

Third, the General Definition of Human Interruption and the Taxonomy of Human Interruption were synthesized as the first general theoretical tools for addressing the problems associated with interrupting people. The set of theoretical constructs identified in the preceding chapter was used as a foundation for building unique theoretical tools. Two tools were synthesized and presented: (1) a general Definition of Human Interruption with accompanying postulates and assertions; and (2) a practical Taxonomy of Human Interruption. Each part of the Definition of Human Interruption represents a useful concept for analysis and description of the problem. These parts were described and their foundations in theory discussed. The relevance and general utility of the definition was discussed within the context of each of the different research domains surveyed in the analysis chapter. The Taxonomy of Human Interruption identifies the most useful dimensions of the human interruption problem. Detailed examples identified in the current literature were provided for each dimension of the taxonomy. These tools are contributions to future efforts in analysis, description, prediction, and empirical investigation of this problem.

Fourth, the validation work presented in this dissertation contributed to the believability of the utility of the synthesized theoretical tools. These validation works, however, also have intrinsic worth. Fifth, the Taxonomy of Human Interruption was used to structure the first broad survey of published literature about human interruption by machine. The eight dimensions of the Taxonomy of Human Interruption were used to facilitate a general treatment of diverse works into useful focused discussions. This was the first time that these articles have been discussed in such a general and comprehensive way. The eight discussions brought out the general themes of existing research in each dimensions: (1) Source of Interruption; (2) Individual Characteristic of Person Receiving Interruption; (3) Method of Coordination; (4) Meaning of Interruption; (5) Method of Expression; (6) Channel of Conveyance; (7) Human Activity Changed by Interruption; and (8) Effect of Interruption. Extra survey coverage was given to the "method of coordination" dimension, because it was the focus of the experiment in the following chapter. This survey contributes to future investigations of the problem, by revealing commonalties between different previously unrelated articles.

Sixth, the Taxonomy of Human Interruption was used to guide the formation and operationalization of an important hypothesis about the effects of different methods of coordination on users' performance in HCI. The taxonomy presents the four recognized solutions for designing the coordination of human interruptions in user interfaces. A hypothesis was made that these four solutions will support different degrees of user performance on interrupt laden multitasks. An experiment with human subjects was designed and conducted to test this main hypothesis. Eight subhypotheses were proposed that address relationships between specific individual coordination methods. An experimental platform and multitask were created that conform to detailed requirements for controlling potential sources of confounding influence. Five measures of human performance were analyzed, and it was concluded that the main hypothesis is significantly supported. The results provided mixed support for the eight subhypotheses. This dissertation makes a contribution by conducting basic research on the problem of human interruption. The contribution of this experiment also is a first step toward discovering general design guidelines for solving this problem.

The multitask created for this experiment is itself a significant contribution. This unique multitask and its implementation as an experimental platform, may be useful for other investigations. This multitask, game task and matching task, represents a carefully designed testbed that supports study of human interruption in a well controlled way.

6.5 Future Work

This dissertation contributes to the problem of human interruption, but it does not provide a solution. The uncovered complexity of this problem implies that there is a vast amount of work that must be done before a final solution will be available. Some potentially fruitful additional works were considered during the creation of this dissertation, but which were outside the scope of this work. These additional efforts are outlined here and left as future works.

There are three categories of future works: (1) expansion of the theoretical foundation; (2) further validation of the theoretical tools presented in this dissertation; and (3) the creation of interdisciplinary user interface design guidelines.

6.5.1 A General Theory of Human Interruption

The set of theoretical constructs identified in Chapter 2 can serve as a useful foundation for some things. A finished general model of human interruption, however, would be much more powerful.

An analogy to automotive engineering illustrates the utility of high-fidelity human models. Modern cars are relatively heavy and fast, and put their occupants at risk of serious impact injury. Automotive engineers attempt to integrate structures into the design of new cars that protect their occupants during collisions. This engineering effort has a severe limitation. Many other engineering contexts have the luxury of being able to directly test prototypes under real conditions. However, real people can not be crashed on purpose to test the effectiveness of design prototypes of automobiles. The answer to this problem is a high-fidelity model of a person that can be crashed — the crash test dummy. Engineers can examine the damage inflicted on the dummy during a test and infer, with confidence, the damage that would have been incurred to a real person. The dummy supports inference because it is not just a mannequin, but a model of human anatomy. It's neck, for example, is affected by stress in much the same way that a human neck would be affected by stress.

Designing user interface solutions for the human interruption problem can have the same dilemma as that in the automotive engineering analogy. For some kinds of systems, human errors or failures are potentially so costly that they can not be tested directly under real conditions. What is needed is a high-fidelity model of human interruption that can act as a crash test dummy for testing user interface design prototypes.

The theoretical foundation created in Chapter 2 is not that final model of human interruption. Instead, this dissertation has identified and discussed a list of descriptors or what that crash test dummy should be like. The problem has been found to be very complex. A comprehensive and validated general model of human interruption is a very grandiose goal, and it is unlikely that any such thing will be built soon. This is because the human psychology is much more difficult to model than the human anatomy modeled by crash test dummies. There are, however, lots of possible milestones that could be pursued along the path toward such a grand model.

A first step in this work could be to expand the set of identified theoretical constructs. The analysis of the identified relevant domains of research could be deepened, and other relevant domains may be discovered. This expansion of the theoretical foundation could be used to refine and improve the Definition of Human Interruption and the Taxonomy of Human Interruption.

6.5.2 Further Validation of Theoretical Tools

The Definition of Human Interruption and The Taxonomy of Human Interruption presented in this dissertation need further validation through future work. The utility of these tools needs to be demonstrated extensively in order that their adoption can be more easily promoted. These theoretical tools have three uses: (1) generalization by structuring interdisciplinary discussions of findings; (2) guiding empirical research; and (3) analysis of existing problems. Partial validation for two of these uses was performed as part of this dissertation. Future work with these theoretical tools in any one of their three roles could potentially provide additional validation support and would also have intrinsic value.

6.5.2.1 GENERALIZATION

Chapter 4, “A Literary Framework” (pg 143), presents an interdisciplinary survey of the existing literature relevant to the coordination of human interruption during human-computer interaction. A useful future work would be to expand this survey. This would increase the usefulness of that survey for facilitating the generalization of diverse but relevant published works. Chapter 4’s only treats the “Method of Coordination” factor from The Taxonomy of Human Interruption in depth. Every one of the other seven factors also merit an in-depth interdisciplinary survey of existing work.

6.5.2.2 EMPIRICAL RESEARCH

The Definition of Human Interruption and the Taxonomy of Human Interruption could be used to suggest and guide many empirical investigations of the human interruption problem. These future works could produce important results, and further validate the theoretical tools presented by this dissertation.

The experiment reported in Chapter 5 had to fix several multitask variables to a single value in order to isolate the comparison of interest. An important question remains, “How sensitive are the results found to differences in those variables that were held constant for the experiment?” Some important examples of these fixed variables that could become independent variables to be varied in future experiments are: (1) duration of individual matching tasks; (2) predictability of duration of individual matching tasks; (3) degree of intelligence of mediating algorithm; (4) level of workload, i.e., number of jumping characters on the game task and/or number of matching tasks; (5) variability of workload over time; and (6) duration of scheduled interval. There were also some multitask variables that were used as dependent variables for the experiment that could be reworked as independent variables for future experiments. For example, an important future experiment could compare the four methods of coordinating interruption with a common fixed number of task switching between game tasks and matching tasks for each condition. This could be accomplished by designing the experimental software to not allow accumulated matching tasks to be performed in batches.

Only eight subhypotheses were tested for the experiment presented in this dissertation. Also, only five measures of performance were analyzed from the data recorded from the experiment. There are several other possibly important subhypotheses and different measures of performance that could be tested and analyzed in future works. For example, it can be hypothesized (or “subhypothesized”) that there is a relationship between people’s subjective preference of “Method of Coordination” design solution and their objective multitask performance. Do people like what they do best with? Another subhypothesis could state that there are correlations between individual differences of subjects and their multitask performance. Some possibly important individual differences that were recorded in the questionnaires were: sex, age, degree of previous computer experience, and degree of previous video game experience. Another future work subhypothesis could predict a temporal correlation between occurrence of matching errors and task switching events. Perhaps subjects’ matching errors are more likely to occur on the first matching task after a task switch. This might partially explain why there is a correlation between number of task switching and number of matching errors — the more times subjects switch tasks the more times matching tasks were attempted immediately after task switching.

An interesting topic that was beyond the scope of this dissertation is a more in-depth investigation of human interruption during human-computer dialogue. Dialogue, and not just interactive manipulation, is one of the recognized problems of interacting with intelligent semi-autonomous systems. The experiment conducted for this dissertation, however, focused on “Methods of Coordination.” An important future work would be to design an experiment that more directly looks at the issues of interruption and dialogue.

Large variances in performance were observed between individual subjects in this experiment. It should be an important vein of future work to try to find ways of exploiting that source of variability. Its existence means that it is unlikely that a single design solution will be “best” for all people. It may be true that there are subgroups of subjects that performed better under different conditions than those found by the means for all subjects considered together. It may be possible to find or construct a test of individual characteristics that would predict what kind of user interface solution would best fit each different kind of potential user. An intelligent user interface could then be constructed that would create custom user interfaces for each different person that used it.

A major influence on multitask performance observed in the experiment from this dissertation is the amount of time subjects spent away from the game task as it continued to run whenever backgrounded. One future work could be to remove this source of influence, and run the same experiment but have the game task pause whenever backgrounded. This would isolate the effects of interruption and eliminate the influence of timing of simultaneous multitask events.

6.5.2.3 ANALYSIS

Analysis is claimed to be a power of The Definition of Human Interruption and The Taxonomy of Human Interruption. However, this power of analysis was not tried for this dissertation. This effort is left as an important future work. Another useful future work would be to identify and document existing real-world systems that have the problem of user-interruption by machine. These could be potential proving grounds for theory-based analysis using the tools presented here.

6.5.3 User Interface Design Guidelines

General design guidelines for the "Method of Coordination" for user-interruption has not been solved here. Future work is needed to finish the job. The final solution may be some mixture of the four known individual design solutions. The data collected for this experiment support speculation that such a mixed solution may be more useful than any one solution in isolation. A few subjects remarked that their preferred solution for a user interface for this experimental multitask would be a "mediated interruption" solution that had a "manual override" capability. They liked the "mediated interruption" but wanted to be able to take control and switch to a "negotiated interruption" mode when they felt the mediator was failing to do a good job.

The creation of design guidelines that would be generally useful across different domains is an ambitious and complex future task. This is the final solution to the design question of the human interruption problem for human-computer interaction. Section 5.6.4 (p. 234) of Chapter 5 speculated about some possible design guidelines. Systems designers are not usually interested in speculation. In the absence of real design guidelines that are accepted as standards in the industry, systems designers will rely on their own abilities to engineer creative solutions to the user-interruption problem. Design failure will be both likely, because of the complexity and sensitivity of the problem, and costly.

BIBLIOGRAPHY

Abacus Concepts (1989), SuperANOVA: Accessible General Linear Modeling (Abacus Concepts, Inc., Berkeley, CA).

Abacus Concepts (1996), StatView Reference (Abacus Concepts, Inc., Berkeley, CA).

Adams, M.J. and R.W. Pew (1990), "Situational Awareness in the Commercial Aircraft Cockpit: a Cognitive Perspective," IEEE/AIAA/NASA 9th Digital Avionics Systems Conference, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 519-524.

Alford, R.C. (1992), in *Byte*. pp. 249-256.

Anand, T.S. and K.P. Lee (1989), "Extending CATS: Mixed-Initiative Inferencing," Artificial Intelligence for Applications, Fifth Conference on, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 107-114.

Anderson, R.C. and J.W. Pitchert (1978), "Recall of Previously Unrecallable Information Following a Shift in Perspective," *Journal of Verbal Learning and Verbal Behavior* **17**, 1-12.

Aronson, E. (1995), The Social Animal (W. H. Freeman and Company, New York).

Atkinson, J.W. (1953), "The Achievement Motive and Recall of Interrupted and Completed Tasks," *Journal of Experimental Psychology* **46**(6), 381-390.

Ballas, J.A., C.L. Heitmeyer, and M.A. Pérez (1992), "Direct Manipulation and Intermittent Automation in Advanced Cockpits," NRL Formal Report NRL/FR/5534—92-9375, The Naval Research Laboratory.

Ballas, J.A., C.L. Heitmeyer, and M.A. Pérez (1992), "Evaluating Two Aspects of Direct Manipulation in Advanced Cockpits," CHI '92 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 127-134.

Ballas, J.A., D.C. McFarlane, L.B. Achille, J.L. Stroup, C.H. Heithecker, and S.D. Kushiner (1996), "Interfaces for Intelligent Control of Data Fusion Processing," NRL/FR/5513--96-9806, Naval Research Laboratory, April 12.

Bannon, L.J. (1986), "Computer-Mediated Communication," in *User Centered System Design*, D.A. Norman and S.W. Draper, eds. (Lawrence Erlbaum Associates, Hillsdale, NJ) pp. 433-452.

Berger, T., A. Kamoun, and P. Millot (1988), "Real Time Measurement of Workload in Discrete Multitask Situations and Extensions to Continuous Tasks," International Conference on Human Machine Interaction and Artificial Intelligence in Aeronautics and Space, Toulouse-Blagnac, France, pp. 57-68.

Bergeron, H.P. (1968), "Pilot Response in Combined Control Tasks," *Human Factors* **10**(3), 277-282.

Bower, G.H. (1982), "Plans and Goals in Understanding Episodes," in *Discourse Processing*, A. Flammer and W. Kintsch, eds. (North-Holland, New York, NY).

Brennan, S.E. (1990), "Conversation as Direct Manipulation: An Iconoclastic View," in *The Art of Human-Computer Interface Design*, B. Laurel, ed. (Addison-Wesley Publishing Company, Reading, MA) pp. 393-404.

Brown, P. and S.C. Levinson (1987), *Politeness: Some Universals in Language Usage* (Cambridge University Press, New York, NY).

Bruning, J.L. and B.L. Kintz (1987), Computational Handbook of Statistics (HarperCollins-Publishers, Glenview, Illinois).

Burton, R.R. and J.S. Brown (1979), "An Investigation of Computer Coaching for Informal Learning Activities," *Int. J. Man-Machine Studies* **11**, 5-24.

Cabon, P., A. Coblentz, and R. Mollard (1990), "Interruption of a Monotonous Activity with Complex Tasks: Effects of Individual Differences," Proceedings of the Human Factors Society 34th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 2, pp. 912-916.

Callan, J.R., R.T. Kelly, J.W. Gwynne, III, and B. Feher (1990), "Communication interrupt effects on tactical decisions," International Conference on Systems, Man and Cybernetics Conference Proceedings, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 396-399.

Card, S.K. and G.G. Robertson (1996), "The WebBook & Web Forager," CHI '96 - Video Program, Association for Computing Machinery, New York, NY, Card, S.K., T.P. Moran, and A. Newell (1983), *The Psychology of Human-Computer Interaction* (Lawrence Erlbaum Associates, Inc., Hillsdale, NJ).

Cellier, J.M. and H. Eyrolle (1992), "Interference Between Switched Tasks," *Ergonomics* **35**(1), 25-36.

Chan, P. and R. Lee (1997), *The Java Class Libraries: An Annotated Reference* (Addison Wesley Longman, Inc., Reading, MA).

Chapanis, A. (1978), "Interactive Communication: A Few Research Answers for a Technological Explosion," ED168084/CS502432, ERIC (Educational Resources Information Center), Sewell, NJ.

Chapanis, A. and C.M. Overbey (1974), "Studies in Interactive Communication: III. Effects of Similar and Dissimilar Communication Channels and Two Interchange Options on Team Problem Solving," *Perceptual & Motor Skills* **38**(2), 343-374.

Cherry, E.C. (1953), "Some Experiments on the Recognition of Speech with One or Two Ears," *Journal of the Acoustical Society of America* **25**(5), 975-9.

Chignell, M.H. and P.A. Hancock (1988), "Intelligent Interface Design," in *Handbook of Human-Computer Interaction*, M. Helander, ed. (Elsevier Science Publishers B. V., New York, NY) pp. 969-995.

Chou, C.D. and K. Funk (1989), "Development and Validation of a Theory of Human Operator Behavior in a Multitask Environment," International Conference on Systems, Man and Cybernetics, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, Vol. 3, pp. 1154-1156.

Chu, Y.Y. and W.B. Rouse (1979), "Adaptive Allocation of Decisionmaking Responsibility Between Human and Computer in Multitask Situations," *IEEE Transactions on Systems, Man, & Cybernetics* **9**(12), 769-778.

Clark, H.H. (1996), *Using Language* (Cambridge University Press, New York, NY).

Clark, H.H. and E.F. Schaefer (1989), "Contributing to Discourse," *Cognitive Science* **13**(2), 259-294.

Cooper, R. and B. Franks (1993), "Interruptibility as a Constraint on Hybrid Systems," *Minds & Machines* **3**(1 (Feb)), 73-96.

Covey, S.R. (1989), *The 7 Habits of Highly Effective People* (Fireside: Simon & Schuster Inc., New York, NY).

Curry, M.B. and A.F. Monk (1995), "Dialogue Modeling of Graphical User Interfaces with a Production System," *Behaviour & Information Technology* **14**(1), 41-55.

Cypher, A. (1986), “The Structure of User's Activities,” in *User Centered System Design*, D.A. Norman and S.W. Draper, eds. (Lawrence Erlbaum Associates, Hillsdale, NJ) pp. 243-263.

Czerwinski, M.P., S.E. Chrisman, and M. Rudisill (1991), “Interruptions in Multitasking Situations: The Effects of Similarity and Warning,” Technical Report JSC-24757, National Aeronautics and Space Administration, Lyndon B. Johnson Space Center, Houston, Texas.

Davies, S.P., J.M. Findlay, and A.J. Lambert (1989), “The Perception and Tracking of State Changes in Complex Systems,” Proceedings of the Third International Conference on Human-Computer Interaction, Vol. 2, pp. 510-517.

Des Jardins, J. (1998), “Secrets of Professional Organizers,” *Working Mother* **December/ January**, 36-44.

Dillenbourg, P. and J. Self (1992), “A Computational Approach to Socially Distributed Cognition,” *European Journal of Psychology of Education* **7**, 353-372.

Duncan, S., Jr. (1972), “Some signals and rules for taking speaking turns in conversations,” *Journal of Personality and Social Psychology* **23**(2), 283-293.

Duncan, S., Jr. (1973), “Toward a Grammar for Dyadic Conversation,” *Semiotica* **9**, 29-46.

Edmondson, W.H. (1989), “Asynchronous Parallelism in Human Behaviour: A Cognitive Science Perspective on Human-Computer Interaction,” *Behavior and Information Technology* **8**(1), 3-12.

Endsley, M. (1995), “Measurement of Situation Awareness in Dynamic Systems,” *Human Factors* **37**(1), 65-84.

Enstrom, K.D. and W.B. Rouse (1977), “Real-Time Determination of How a Human Has Allocated His Attention Between Control and Monitoring Tasks,” *IEEE Transactions on Systems, Man, & Cybernetics* **7**(3), 153-161.

Field, G.E. (1987), “Experimentus Interruptus,” *ACM SIGCHI Bulletin* **19**(2), 42-46.

Forester, J.A. (1986), “An Assessment of Variable Format Information Presentation,” Proceedings of Information Management and Decision Making in Advanced Airborne Weapon Systems (AGARD-CP-414), North Atlantic Treaty Organization: Advisory Group for Aerospace Research and Development, Toronto, Canada, pp. 1-13.

Foushee, H.C. and R.L. Helmreich (1988), “Group Interaction and Flight Crew Performance,” in *Human Factors in Aviation*, E.L. Wiener and D.C. Nagel, eds. (Academic Press, Inc., San Diego, CA).

Galdes, D.K. and P.J. Smith (1990), “Building an Intelligent Tutoring System: Some Guidelines from a Study of Human Tutors,” Proceedings of the Human Factors Society 34th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 2, pp. 1407-1411.

Galdes, D.K., P.J. Smith, and J.W. Smith, Jr. (1991), “Factors Determining When to Interrupt and What to Say: An Empirical Study of the Case-Method Tutoring Approach,” International Conference on the Learning Sciences 1991, Association for the Advancement of Computing in Education, Evanston, IL, pp. 194-202.

Gaver, W.W. (1989), “The SonicFinder: An Interface that Uses Auditory Icons,” *Human-Computer Interaction* **4**(special issue on nonspeech audio), 67-94.

Gaver, W.W. and R.B. Smith (1990), “Auditory Icons in Large-Scale Collaborative Environments,” Proceedings of IFIP INTERACT'90: Human-Computer Interaction, International Federation for Information Processing, Laxenburg, Austria, pp. 735-740.

Gaver, W.W., R.B. Smith, and T. O'Shea (1991), “Effective Sounds in Complex Systems: The ARKola Simulation,” CHI '91 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 85-90.

Gifford, W.S. and D.L. Turock (1992), “The Electronic Receptionist: A Knowledge-Based Approach to Personal Communications,” CHI '92 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 113-114.

Gillie, T. and D.E. Broadbent (1989), “What Makes Interruptions Disruptive? A Study of Length, Similarity, and Complexity,” *Psychological Research* **50**(4), 243-250.

Gopher, D. and E. Donchin (1986), “Workload—an Examination of the Concept,” in Handbook of Perception and Human Performance: Volume 2, Cognitive Processes and Performance, K.R. Boff, L. Kaufman, and J.P. Thomas, eds. (John Wiley and Sons, Inc., pp. 41/1-41/49.

Grehan, R. (1990), “Multitasking for the Masses,” *Byte* **15**(2), 279.

Hall, B.L. and D.E. Hursch (1982), “An Evaluation of the Effects of a Time Management Training Program on Work Efficiency,” *Journal of Organizational Behavior Management* **3**(4), 73-96.

Hammer, J.M. and R.L. Small (1995), “An Intelligent Interface in an Associate System,” in *Human/Technology Interaction in Complex Systems*, W.B. Rouse, ed. (JAI Press, Greenwich, CT), Vol. 7,.

Hansen, W.J. (1971), “User Engineering Principles for Interactive Systems,” Proceedings of the Fall Joint Computer Conference, American Federation of Information Processing Societies (AFIPS Press), Vol. 39, pp. 523-532.

Harrison, B.L., H. Ishii, K.J. Vicente, and W.A.A. Buxton (1995), "Transparent Layered User Interfaces: An Evaluation of a Display Design to Enhance Focused and Divided Attention," CHI '95 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 317-324.

Hartson, H.R. and P.D. Gray (1992), "Temporal Aspects of Tasks in the User Action Notation," *Human-Computer Interaction* **7**(1), 1-45.

Hix, D. and H.R. Hartson (1993), "Developing User Interfaces: Ensuring Usability Through Product and Process," (John Wiley & Sons, Inc., New York, NY) pp. 197-198.

Holden, K.L. and M.R. O'Neal (1992), "The Utility of Various Windowing Capabilities for Single-Task and Multi-Task Environments," CHI '92 Conference Proceedings, ACM Press, pp. 680.

Hoogendoorn, E.L., K.B. Langton, N. Solntseff, and R.B. Haynes (1991), "A PC-Based Interface for an Expert System to Assist with Preoperative Assessments," Computer-Based Medical Systems, Fourth Annual IEEE Symposium on, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 251-257.

Hsu, S.H. and C. Shen (1992), "Multiple Monitors vs. Windowing Presentation Styles for Shop-Floor Controls," in *Proceedings of the Human Factors Society 36th Annual Meeting*,(Human Factors Society, Santa Monica, CA, , Vol. 1, pp. 351-355.

Husain, M.G. (1987), "Immediate and Delayed Recall of Completed-Interrupted Tasks by High and Low Anxious Subjects," *Manas* **34**(1-2), 67-71.

Ishihara, S. (1996), Ishihara's Tests for Colour-Deficiency, Concise Edition (Kanehara & Co., LTD., P.O. Box No.1 Hongo P.O., Tokyo, 113-91, Japan (Exclusive Wholesale distributor in U.S.A. & Canada: Graham Field, Inc., Hauppauge, New York 11788)).

Jessup, L.M. and T. Connolly (1993), "The Effects of Interaction Frequency on the Productivity and Satisfaction of Automated Problem-Solving Groups," Proceeding of the Twenty-Sixth Hawaii International Conference on System Sciences, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, Vol. 4, pp. 142-151.

John, B.E. and W.D. Gray (1995), "Tutorial 25 Notes: GOMS Analysis for Parallel Activities," CHI '92 — Tutorial 25 Notes, Association for Computing Machinery, New York, NY, Denver, Colorado.

Joslyn, S.L. (1995), "Individual Differences in Time Pressured Decision Making," Doctoral Dissertation, University of Washington.

Kanfer, F.H. and M.K. Stevenson (1985), "The Effects of Self-Regulation on Concurrent Cognitive Processing," *Cognitive Therapy & Research* **9**(6), 667-684.

Karis, D. (1991), "Evaluating Transmission Quality in Mobile Telecommunication Systems Using Conversation Tests," Proceedings of the Human Factors Society 35th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 217-221.

Katz, R. (1995), "Automatic Versus User-Controlled Methods of Briefly Interrupting Telephone Calls," *Human Factors* **37**(2), 321-334.

Kaye, A.R. and G.M. Karam (1987), "Cooperating Knowledge-Based Assistants for the Office," *ACM Transactions on Office Information Systems* **5**(4), 297-326.

Kendon, A. (1967), "Some Functions of Gaze-Direction in Social Interaction," *Acta Psychologica* **26**, 22-63.

Kendon, A. (1972), "Some Relationships Between Body Motion and Speech: An Analysis of an Example," in *Studies in Dyadic Communication*, A. Siegman and B. Pope, eds. (Pergamon Press, Elmsford, New York, pp. 177-210.

Keppel, G. (1991), *Design and Analysis: A Researcher's Handbook* (Prentice Hall, Upper Saddle River, NJ).

Kermis, M.D. (1977), "Intertask Correspondence in the 5 to 7 Shift: A Question of Competence Vs. Performance," ED136907, ERIC (Educational Resources Information Center), Sewell, NJ.

Kirlik, A. (1993), "Modeling Strategic Behavior in Human-Automation Interaction — Why an Aid Can (and Should) Go Unused," *Human Factors* **35**(2), 221-242.

Kreifeldt, J.G. and M.E. McCarthy (1981), "Interruption as a Test of the User-Computer Interface," Conference on Manual Control (17th Annual), pp. 655-667.

Kuperman, G.G. and W.A. Perez (1988), "A Frame-Based Mission Decomposition Model," in *Proceedings of the Human Factors Society 32nd Annual Meeting*,(Human Factors Society, Santa Monica, CA, , Vol. 1, pp. 135-139.

Kushnier, S.D., C.H. Heithecker, J.A. Ballas, and D.C. McFarlane (1996), "Situation Assessment Through Collaborative Human-Computer Interaction," *Naval Engineers Journal* **108**(4), 41-51.

Lane, D.M. and D.G. Jensen (1993), "Einstellung: Knowledge of the Phenomenon Facilitates Problem Solving," Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting, Human Factors and Ergonomics Society, Santa Monica, CA, Vol. 2, pp. 1277-1280.

Lashkari, Y., M. Metral, and P. Maes (1994), "Collaborative Interface Agents," National Conference on Artificial Intelligence, MIT Press, Cambridge, MA, Lee, W.O. (1992), “The Effects of Skill Development and Feedback on Action Slips,” Proceedings of the HCI'92 Conference on People and Computers VII, British Computer Society, Swindon, Wiltshire, England, pp. 73-86.

Li, J. and M.M. Mantei (1992), “Virtual Open Office: Supporting Effective "Open" Contact,” CHI '92 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 109-110.

Lieberman, H. (1997), “Autonomous Interface Agents,” CHI '97 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 67-74.

Liu, Y. (1993), “Visual Scanning, Memory Scanning, and Computational Human Performance Modeling,” Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting, Human Factors and Ergonomics Society, Santa Monica, CA, Vol. 1, pp. 142-146.

Lustig, M.W. (1980), “Computer Analysis of Talk-Silence Patterns in Triads,” *Communication Quarterly* **28**(4), 2-12.

Maes, P. (1994), “Agents that Reduce Work and Information Overload,” *Communications of the ACM* **37**(7), 31-40.

Maes, P. and A. Wexelblat (1996), “Interface Agents,” CHI '96 Tutorial Notes 22, Association for Computing Machinery, New York, NY.

Malin, J.T., D.L. Schreckenghost, D.D. Woods, S.S. Potter, L. Johannesen, M. Holloway, and K.D. Forbus (1991), “Making Intelligent Systems Team Players: Case Studies and Design Issues, Vol. 1, Human-Computer Interaction Design,” NASA Technical Memorandum 104738, National Aeronautics & Space Administration (NASA), Washington, DC.

Malone, T.W. and K. Crowston (1994), “The Interdisciplinary Study of Coordination,” *ACM Computing Surveys* **26**(1), 87-119.

Marcus, S. and S. Blau (1983), “Not Seeing Is Relieving: Invisible Writing with Computers,” *Educational Technology* **23**(4), 12-15.

McCarthy, J.C. and A.F. Monk (1994), “Channels, Conversation, Cooperation and Relevance: All You Wanted to Know About Communication but Were Afraid to Ask,” *Collaborative Computing* **1**, 35-60.

McFarlane, D.C. (1997), “Interruption of People in Human-Computer Interaction: A General Unifying Definition of Human Interruption and Taxonomy,” NRL Formal Report NRL/FR/5510—97-9870, Naval Research Laboratory.

McLeod, P. and J. Mierop (1979), “How to Reduce Manual Response Interference in the Multiple Task Environment,” *Ergonomics* **22**(4), 469-475.

Miller, G.A. (1956), “The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information,” *Psychological Review* **63**(2), 81-97.

Millot, P. and A. Kamoun (1988), “An Implicit Method for Dynamic Task Allocation Between Man and Computer in Supervision Posts of Automated Processes,” Analysis, Design and Evaluation of Man-Machine Systems 1988; Selected Papers from the Third IFAC/IFIP/IEA/IFORS Conference, International Federation for Information Processing, Laxenburg, Austria, pp. 77-82.

Minasi, M. (1993), in *Compute!*. pp. 60-62.

Miyata, Y. and D.A. Norman (1986), “Psychological Issues in Support of Multiple Activities,” in *User Centered System Design*, D.A. Norman and S.W. Draper, eds. (Lawrence Erlbaum Associates, Hillsdale, NJ) pp. 265-284.

Moray, N. (1986), “Monitoring Behavior and Supervisory Control,” in *Handbook of Perception and Human Performance: Volume 2, Cognitive Processes and Performance*, K.R. Boff, L. Kaufman, and J.P. Thomas, eds. (John Wiley and Sons, Inc., pp. 40/1-40/51.

Morrin, K.A., D.J. Law, and J.W. Pellegrino (1994), “Structural Modeling of Information Coordination Abilities: An evaluation and Extension of the Yee, Hunt, and Pellegrino Model,” *Intelligence* **19**, 117-144.

Morris, W.N. (1971), “Manipulated Amplitude and Interruption Outcomes,” *Journal of Personality & Social Psychology* **20**(3), 319-331.

Mouloua, M., R. Parasuraman, and R. Molloy (1993), “Monitoring Automation Failures: Effects of Single and Multi-Adaptive Function Allocation,” in *Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting*,(Human Factors and Ergonomics Society, Santa Monica, CA, Vol. 1, pp. 1-5.

Müller, H.J. and P.M.A. Rabbitt (1989), “Reflexive and Voluntary Orienting of Visual Attention: Time Course of Activation and Resistance to Interruption,” *Journal of Experimental Psychology: Human Perception & Performance* **15**(2), 315-330.

Nakagawa, M., K. Machii, N. Kato, and T. Souya (1993), “Lazy Recognition as a Principle of Pen Interfaces,” INTERCHI'93 Conference Proceedings — Adjunct Proceedings, Association for Computing Machinery, New York, NY, pp. 89-90.

National Transportation Safety Board (1988), “Aircraft Accident Report, Northwest Airlines Inc., McDonnell-Douglas DC-9-82, N312RC, Detroit Metropolitan Wayne County Airport, Romulus, Michigan,” NTSB-AAR-88-05, National Transportation Safety Board, Washington, DC, May 10, 1988.

National Transportation Safety Board (1979), “Aircraft Accident Report, United Airlines Inc., McDonnell-Douglas DC-8-61, N8082U, Portland, Oregon, Romulus, Michigan,” NTSB-AAR-79-7, National Transportation Safety Board, Washington, DC, December 28, 1978.

National Transportation Safety Board (1972), “Aircraft Accident Report, Eastern Air Lines, Inc. L-1011, N310EA, Miami, Florida,” NTSB-AAR-73-14, National Transportation Safety Board, Washington, DC, December 29, 1972.

Navon, D. and D. Gopher (1979), “On the Economy of the Human-Processing System,” *Psychological Review* **86**(3), 214-255.

Nintendo (1980 & 1981), Game & Watch: Fire (Nintendo of America Inc., P.O. Box 97032, Redmond, WA, 98073-9732).

Nintendo (1997), Game & Watch Gallery, for Game Boy (Nintendo of America Inc., P.O. Box 97032, Redmond, WA, 98073-9732).

Norman, D.A. (1992), *Turn Signals are the Facial Expressions of Automobiles* (Addison-Wesley Publishing Company, Reading, MA).

Norman, D.A. (1986), “Cognitive Engineering,” in *User Centered System Design*, D.A. Norman and S.W. Draper, eds. (Lawrence Erlbaum Associates, Hillsdale, NJ) pp. 31-61.

Norman, D.A. (1969), *Memory and Attention* (John Wiley & Sons, Inc., New York, NY).

Noy, Y.I. (1989), “Intelligent Route Guidance: Will the New Horse be as Good as the Old?,” Vehicle Navigation & Information Systems Conference (VNIS '89), Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 49-55.

Oberg, B. and D. Notkin (1992), “Error Reporting with Graduated Color,” *IEEE Software* **9**(6), 33-38.

Ochsman, R.B. and A. Chapanis (1974), “The Effects of 10 Communication Modes on the Behavior of Teams During Co-operative Problem-solving,” *International Journal of Man-Machine Studies* **6**, 579-619.

O’Donnell, R.D., Colonel and F.T. Eggemeier (1986), “Workload Assessment Methodology,” in Handbook of Perception and Human Performance: Volume 2, Cognitive Processes and Performance, K.R. Boff, L. Kaufman, and J.P. Thomas, eds. (John Wiley and Sons, Inc., pp. 42/1-42/49.

Osgood, S.S., K.R. Boff, and R.S. Donovan (1988), "Rapid Communication Display Technology Efficiency in a Multi-Task Environment," Proceedings of the Human Factors Society 32nd Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 1395-1399.

Pashler, H. (1993), "Doing Two Things at the Same Time," *American Scientist* **81**(1), 48-55.

Pattipati, K.R., D.L. Kleinman, and A.R. Ephrath (1983), "A Dynamic Decision Model of Human Task Selection Performance," *IEEE Transactions on Systems, Man, & Cybernetics* **SMC-13**(2), 145-166.

Pedersen, E.R. and T. Sokoler (1997), "AROMA: Abstract Representation of Presence Supporting Mutual Awareness," CHI '97 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 51-58.

Pérez-Quiñones, M.A. (1996), "Conversational Collaboration in User-Initiated Interruption and Cancellation Requests," Doctoral Dissertation, George Washington University.

Perse, R.M., K.P. Callahan, and T.B. Malone (1991), "Development of an AEGIS Combatant Integrated Survivability Management System (ISMS) Modeling Tool," Proceedings of the Human Factors Society 35th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 2, pp. 1201-1205.

Potjer, J., A. Russel, L. Boves, and E. den Os (1996), "Subjective and Objective Evaluation of Two Types of Dialogue in a Call Assistance Service," Interactive Voice Technology for Telecommunications Applications, Third IEEE Workshop on, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 89-92.

Preece, J., Y. Rogers, H. Sharp, D. Benyon, S. Holland, and T. Carey (1994), *Human-Computer Interaction* (Addison-Wesley Publishing Company, Reading, MA).

Random House (1989), *Webster's Encyclopedic Unabridged Dictionary of the English Language* (Portland House, New York, NY).

Rao, R., J.O. Pedersen, M.A. Hearst, J.D. Mackinlay, S.K. Card, L. Masinter, P.K. Halvorsen, and G.G. Robertson (1995), "Rich Interaction in the Digital Library," *Communications of the ACM* **38**(4 (April)).

Rayner, K. (1992), ed., *Eye Movements in Reading: Scene Perception and Reading* (Springer-Verlag, New York, NY).

Rayner, K. (1983), ed., *Eye Movements in Reading: Perceptual and Language Processes* (Academic Press, London).

Redding, R.E. (1992), “Analysis of Operational Errors and Workload in Air Traffic Control,” Proceedings of the Human Factors Society 36th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 2, pp. 1321-1325.

Rich, C. (1996), “Window Sharing with Collaborative Interface Agents,” *SIGCHI Bulletin* **28**(1), 70-78.

Rickel, J.W. (1989), “Intelligent Computer-Aided Instruction: A Survey Organized Around System Components,” *IEEE Transactions on Systems, Man and Cybernetics* **19**(1), 40-57.

Riley, P. (1976), “Discursive and Communicative Functions of Non-Verbal Communication,” ED143217/FL008785, ERIC (Educational Resources Information Center), Sewell, NJ.

Robertson, G.G., S.K. Card, and J.D. Mackinlay (1993), “Information Visualization Using 3D Interactive Animation,” *Communications of the ACM* **36**(4 (July)), 57-69.

Rouse, W.B. (1977), “Human-Computer Interaction in Multitask Situations,” *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-7**(5), 384-392.

Rubinstein, R. and H. Hersh (1984), *The Human Factor: Designing Computer Systems For People* (Digital Press.

Ryder, J.M. and W.W. Zachary (1991), “Experimental Validation of the Attention Switching Component of the COGNET Framework,” Proceedings of the Human Factors Society 35th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 72-76.

Sacks, H., E.A. Schegloff, and G. Jefferson (1978), “A Simplest Systematics for the Organization of Turn Taking for Conversation,” in *Studies in the Organization of Conversational Interaction*, J. Schenkein, ed. (Academic Press, Inc., San Diego, CA).

Salter, W.J. (1988), “Human Factors in Knowledge Acquisition,” in *Handbook of Human-Computer Interaction*, M. Helander, ed. (North-Holland, New York, NY) pp. 957-968.

Sanford, A.J. and S.C. Garrod (1981), *Understanding Written Language: Explorations of Comprehension Beyond the Sentence* (John Wiley & Sons, New York, NY).

Schank, R.C., G.C. Collins, E. Davis, P.N. Johnson, S. Lytinen, and B.J. Reiser (1982), “What's the Point?,” *Cognitive Science* **6**, 255-275.

Schegloff, E.A. (1968), “Sequencing in Conversational Openings,” *American Anthropologist* **70**, 1075-1095.

Schegloff, E.A. (1979), “Identification and Recognition in Telephone Conversation Openings,” in *Everyday Language: Studies in Ethnomethodology*, G. Psathas, ed. (Irvington Publishers, Inc., New York, NY) pp. 23-78.

Schegloff, E.A., G. Jefferson, and H. Sacks (1977), "The Preference for Self-Correction in the Organization of Repair in Conversation," *Language* **53**(2), 361-382.

Schneider, W. and M. Detweiler (1988), "The Role of Practice in Dual-Task Performance: Toward Workload Modeling in a Connectionist/Control Architecture. Special Issue: Human Information Processing: Theory and Applications," *Human Factors* **30**(5), 539-566.

Schweickert, R. and G.J. Boggs (1984), "Models of Central Capacity and Concurrency," *Journal of Mathematical Psychology* **28**, 223-281.

Senders, J.W. (1964), "The Human Operator as a Monitor and Controller of Multidegree of Freedom Systems," *IEEE Transactions on Human Factors in Electronics* **hfe-5**, 2-6.

Sheridan, T.B. (1987), "Supervisory Control," in *Handbook of Human Factors*, G. Salvendy, ed. (Wiley, New York, NY, pp. 1243-1268.

Sheridan, T.B. (1988), "Task Allocation and Supervisory Control," in *Handbook of Human-Computer Interaction,* M. Helander, ed. (Elsevier Science Publishers B. V., New York, NY) pp. 159-173.

Shneiderman, B. (1992), *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (Addison-Wesley Publishing Company, Inc., Reading, MA).

Smith, I. and S.E. Hudson (1995), "Low Disturbance Audio for Awareness and Privacy in Media Space Applications," MultiMedia '95, Association for Computing Machinery, New York, NY,

Soulsby, E.P. (1989), "Making Decisions Under Increasing Work Load Demands: the Need to Look Ahead," International Conference on Systems, Man and Cybernetics, Institute of Electrical and Electronics Engineers (IEEE), New York, NY, Vol. 3, pp. 1148-1153.

Spelke, E., W. Hirst, and U. Neisser (1976), "Skills of Divided Attention," *Cognition* **4**, 215-230.

Storch, N.A. (1992), "Does the User Interface Make Interruptions Disruptive? A Study of Interface Style and Form of Interruption," CHI '92 Conference Proceedings: Poster, Association for Computing Machinery, New York, NY, pp. 14.

Stroup, J.R. (1935), "Factors Affecting Speed in Serial Verbal Reactions," *Journal of Experimental Psychology* **18**, 643-662.

Stuart, R., H. Desurvire, and S. Dews (1991), "The Truncation of Prompts in Phone Based Interfaces: Using TOTT in Evaluations," Proceedings of the Human Factors Society 35th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 230-234.

Sullivan, N. (1993), "Ask Not for Whom the Phone Rings: It Rings for Thee," *Home-Office Computing* **11**(3), 104.

Taylor, M.M. and M.J. Hunt (1989), "Flexibility versus Formality," in *The Structure of Multimodal Dialogue*, M.M. Taylor, F. Neel, and D.G. Bouwhuis, eds. (North-Holland, Elsevier Science Publishers, Amsterdam) pp. 435-453.

Taylor, R.M. (1989), "Integrating Voice, Visual and Manual Transactions: Some Practical Issues from Aircrew Station Design," in *The Structure of Multimodal Dialogue*, M.M. Taylor, F. Néel, and D.G. Bouwhuis, eds. (Elsevier Science Publishers B. V., pp. 259-265.

Tsukada, K., K.I. Okada, and Y. Matsushita (1994), "The Multi-Project Support System Based on Multiplicity of Task," Eighteenth Annual International Computer Software and Applications Conference (COMPSAC 94), Institute of Electrical and Electronics Engineers (IEEE), New York, NY, pp. 358-363.

Tulga, M.K. and T.B. Sheridan (1980), "Dynamic Decisions and Work Load in Multitask Supervisory Control," *IEEE Transactions on Systems, Man, & Cybernetics* **SMC-10**(5), 217-232.

Van Bergen, A. (1968), *Task Interruption* (North-Holland Publishing Company, Amsterdam).

Vonnegut, K. (1950), "Harrison Bergeron," in Welcome to the Monkey House, (Dell Publishing, New York, New York).

Wærn, Y. (1989), *Cognitive Aspects of Computer Supported Tasks* (John Wiley & Sons, New York, NY).

Wagenaar, W.A. (1969), "Note on the Construction of Digram-Balanced Latin Squares," *Psychological Bulletin* **72**(6), 384-386.

Walden, R.S. and W.B. Rouse (1978), "A Queueing Model of Pilot Decisionmaking in a Multitask Flight Management Situation," *IEEE Transactions on Systems, Man, & Cybernetics* **8**(12 (Dec)), 867-875.

Ware, C., J. Bonner, W. Knight, and R. Cater (1992), "Moving Icons as a Human Interrupt," *International Journal of Human-Computer Interaction* **4**(4), 341-348.

Weiland, M.Z., B. Cooke, and B. Peterson (1992), "Designing and Implementing Decision Aids for a Complex Environment Using Goal Hierarchies," in *Proceedings of the Human Factors Society 36th Annual Meeting*,(Human Factors Society, Santa Monica, CA, Vol. 1, pp. 394-398.

Weiner, B. (1965), "Need Achievement and the Resumption of Incompleted Tasks," *Journal of Personality and Social Psychology* **1**(2), 165-168.

West, C. (1982), "Why Can't a Woman Be More Like a Man? An Interactional Note on Organizational Game-Playing for Managerial Women," *Work and Occupations* **9**(1), 5-29.

Wickens, C.D., I. Larish, and A. Contorer (1989), "Predictive Performance Models and Multiple Task Performance," Proceedings of the Human Factors Society 33rd Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 96-100.

Winograd, T. and F. Flores (1986), *Understanding Computers and Cognition: A New Foundation for Design* (Ablex, Norwood, NJ).

Zachary, W.W. and L. Ross (1991), "Enhancing Human-Computer Interaction through Use of Embedded COGNET Models," Proceedings of the Human Factors Society 35th Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 425-429.

Zachary, W.W., A.L. Zaklad, J.H. Hicinbothom, J.M. Ryder, and J.A. Purcell (1993), "COGNET Representation of Tactical Decision-Making in Anti-Air Warfare," Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting, Human Factors and Ergonomics Society, Santa Monica, CA, Vol. 2, pp. 1112-1116.

Zachary, W.W., M.C. Zubritzky, and F.A. Glenn (1988), "The Development of the Air Anti-Submarine Warfare Mission Testbed as a Tool for the Development of Operator Models," Proceedings of the Human Factors Society 32nd Annual Meeting, Human Factors Society, Santa Monica, CA, Vol. 1, pp. 1073-1077.

Zimmerman, D.H. and C. West (1975), "Sex Roles, Interruptions and Silences in Conversations," in *Language and Sex: Difference and Dominance*, B. Thorne and N. Henley, eds. (Newbury House Publishers Inc., Rowley, Mass.).

Zubritzky, M.C., W.W. Zachary, and J.M. Ryder (1989), "Constructing and Applying Cognitive Models to Mission Management Problems in Air Anti-Submarine Warfare," in *Proceedings of the Human Factors Society 33rd Annual Meeting,*(Human Factors Society, Santa Monica, CA, Vol. 1, pp. 129-133.

APPENDIX A: INDEX OF DOMAIN PERSPECTIVES: IDENTIFYING RELEVANT FIELDS OF RESEARCH

This index is a tool for discovering which domains and topics of research should be relevant for investigating interruption from a given perspective. Given a domain perspective or approach to a human-computer interaction (HCI) design problem with user-interruption, this index can be used to identify relevant domains and topics of published literature.

Interrupting people can cause critical and even life-threatening effects. The critical topic of user-interruption requires a comprehensive user interface design solution. However, comprehensive solutions require comprehensive identification and integration of all relevant published research results and theoretical tools from all sources of current literature. This index can be useful in identifying relevant domains within the current literature.

:Table A1 — Index of Domain Perspectives

Domain Perspective	**Sources of Useful Research Results and Theory**
Colloquial Meaning (Section 2.4 on pg 25)	etymology
Multitasking in HCI (Section 2.5 on pg 28)	multitasking; computer-supported cooperative work (CSCW); human-computer interaction (HCI); cognitive psychology; cognitive modeling
Multitasking in Linguistics (Section 2.6 on pg 36)	linguistics; cognitive psychology; HCI
Multitasking in Situational Awareness (Section 2.7 on pg 41)	situational awareness; attention; psychology; memory; dual-tasking
Management of Semiautonomous Agents (Section 2.8 on pg 54)	task off-load aids; supervisory control; intelligent software agents; multitasking
Human-Human Discourse (Section 2.9 on pg 61)	discourse analysis; communication acts
Human-Human Dialogue (Section 2.10 on pg 69)	dialogue
Psychology of Human Attention (Section 2.11 on pg 71)	attention; memory; cognitive psychology; perception
A Metaphor of Cognitive Momentum (Section 2.12 on pg 76)	informal literature on multitasking in office environments
Social Psychology of Conversation (Section 2.13 on pg 77)	sociology; social psychology; metacommunication; dialogue
Interactional Sociolinguistics of Politeness (Section 2.14 on pg 80)	interactional socio-linguistics; discourse analysis
Simultaneous Speech in Linguistics (Section 2.15 on pg 92)	discourse analysis; dual-tasking; dialogue
Language Use in Linguistics (Section 2.16 on pg 95)	linguistics; dialogue; signaling; signs; cognitive psychology; discourse analysis

Note: this index does not represent every possibility. Instead, this index cites only those topics and classes of literature that I found useful for writing the survey of theory for this report.

APPENDIX B: INDEX OF THEORETICAL CONSTRUCTS: IDENTIFYING COMMON CONCEPTS ACROSS DOMAINS

The interruption of people is a common phenomenon that occurs in several domains of research. I have identified a large set of theoretical constructs about this phenomenon from several diverse domains of published literature (126 identified theoretical constructs). However, it is difficult to use such a large set of theoretical constructs without an indexing aid. This Index of Theoretical Constructs provides such an index for finding all relevant theoretical constructs.

Some theoretical concepts cut across different fields of research in the published literature. Their ubiquitousness makes it difficult to locate all relevant research in the current literature. This index is a tool for easily indexing all the theoretical constructs of interruption identified in this report by the theoretical concepts to which they apply.

This index also categorizes theoretical constructs by whether they represent declarative or procedural information. Declarative theoretical constructs describe structures and how they function, i.e., “People have a single focus of consciousness; and this is how they switch processing steams into and out of it ....” Procedural theoretical constructs describe how the structures are employed to perform actions, i.e., “People accomplish dual-tasking by continuously switching their attention between the two tasks.”

Table B1 — Index of Theoretical Constructs

Theoretical Construct	declarative theoretical constructs	procedural theoretical constructs
PEOPLE	P6 (pg 28); P10 (pg 32); P14 (pg 34); P35 (pg 52)	P13 (pg 34); P16 (pg 37); P24 (pg 43)
Cognition	P9 (pg 29); P15 (pg 36); P74 (pg 77)	P11 (pg 32); P12 (pg 32); P75 (pg 77)
memory	P26 (pg 46); P27 (pg 47); P28 (pg 49); P29 (pg 49); P30 (pg 50); P31 (pg 50); P32 (pg 50); P33 (pg 50)	P34 (pg 51); P36 (pg 53)
attention	P37 (pg 53); P38 (pg 53); P39 (pg 54); P60 (pg 72); P61 (pg 73); P62 (pg 73); P63 (pg 73); P64 (pg 73); P65 (pg 73); P66 (pg 73); P67 (pg 73); P72 (pg 76); P73 (pg 76)	P68 (pg 75); P69 (pg 75); P70 (pg 75); P71 (pg 76)
Perception		P25 (pg 45)
Motor	P7 (pg 28); P8 (pg 29)	
Purposes	P1 (pg 27); P2 (pg 27); P3 (pg 27); P4 (pg 27)	P5 (pg 27)
social relationships	P80 (pg 82); P82 (pg 82); P83 (pg 82); P86 (pg 90)	P76 (pg 79); P77 (pg 79); P78 (pg 79); P79 (pg 79); P81 (pg 82); P84 (pg 82); P85 (pg 90); P87 (pg 91); P88 (pg 91)
joint activities	P94 (pg 97); P95 (pg 97); P96 (pg 99); P97 (pg 99); P98 (pg 100); P99 (pg 100); P100 (pg 107); P108 (pg 114); P109 (pg 114)	P93 (pg 97); P110 (pg 115); P111 (pg 115); P112 (pg 115); P113 (pg 115)
communi-cation	P17 (pg 40); P18 (pg 40); P19 (pg 40); P20 (pg 40); P21 (pg 40); P22 (pg 40); P52 (pg 66); P53 (pg 66); P101 (pg 107); P102 (pg 107); P103 (pg 108); P104 (pg 108); P105 (pg 108)	P23 (pg 40); P49 (pg 63); P50 (pg 63); P51 (pg 63); P54 (pg 66); P55 (pg 66); P56 (pg 66); P57 (pg 66); P58 (pg 68); P59 (pg 70); P89 (pg 94); P90 (pg 94); P91 (pg 94); P92 (pg 95); P106 (pg 108); P107 (pg 109)
work	P40 (pg 56); P41 (pg 56); P43 (pg 56)	P42 (pg 56); P44 (pg 56); P45 (pg 57); P46 (pg 60); P47 (pg 60); P48 (pg 60)
TASKS	T1 (pg 29); T2 (pg 36); T3 (pg 44); T4 (pg 52); T7 (pg 55); T6 (pg 55)	T5 (pg 52)
INTERRUPTIONS	In1 (pg 27); In3 (pg 91)	In2 (pg 36)
CONTEXTS	C1 (pg 60); C2 (pg 97); C3 (pg 100)	

Note that some theoretical constructs have a mix of declarative and procedural information. I have had to decide how to pigeon-hole these theoretical constructs into single categories. Readers will disagree with some of my decisions. Therefore, this index will be most useful if the reader pursues a broad range of related categories.

APPENDIX C: PROTOCOL DESCRIPTION FOR HUMAN SUBJECTS REVIEW COMMITTEE

PROTOCOL: HUMAN-COMPUTER INTERACTION IN MULTITASKING ENVIRONMENTS

Objective: User-interruption in human-computer interaction (HCI) is an increasingly important problem. Many of the useful advances in intelligent and multitasking computer systems have the significant side effect of greatly increasing user-interruption. However, no HCI design guidelines exist for solving this problem. The purpose of this experiment is to show whether methods for coordinating human-human interruption can be successfully applied to the design of user interfaces to mitigate the negative effects of human interruption by computer in HCI.

Method: Subjects will be asked to perform a dualtask composed of the following two concurrent tasks: (1) a continuous video game task in which keyboard key presses are made to play the game; and (2) an intermittent matching task in which decisions about matching are made with keyboard key presses. This dual task is a kind of routine task that normal people encounter ordinarily as part of their everyday work with computers. All stresses are within normal limits — there will be no stressors nor will subjects' behavior be manipulated (i.e., controlled through fraudulent influence). Subjects will use only one hand and perform all computer interaction with a keyboard. To perform the task, subjects will need to keep one hand in a comfortable position and make intermittent key presses with two fingers on closely spaced keys. Tasks will appear graphically on a color monitor -- no sound.

A single-factor within-subjects design will be used for this human subjects experiment. There are six conditions — three treatment conditions and three control conditions. Experimental trials are limited to a 10 minute maximum, and adequate rest periods will be provided between trials. A practice period will be used to introduce subjects to the experimental tasks during which subjects will be allowed to ask questions. All trials will be video and audio taped. A verbal protocol will be used for task instructions. Computer task and interaction events will be recorded through out the experiment. At the conclusion of the experiment, each subject will be asked to complete an exit questionnaire.

In order to insure that the subjects can interpret the visual information, they will be asked if their vision is corrected to 20/20, and given a screening test for color perception. The recruitment notice will mention that we need normal color perception vision, so if a subject does not do well on the color test, we will remind them that this is needed in the experiment.

Equipment: The experiment will be run on standard workstation computer equipment using a commercial CRT and keyboard. The subject will be seated in a comfortable chair and the keyboard input device will be adjusted for comfort and ease of use.

Risks: There are no risks of harm or discomfort anticipated in this experiment. The experimental tasks and environment are similar to those ordinarily encountered during daily routine work with computers. Experiments like this have been conducted previously at NRL and at GWU with no adverse effects on the participants.

Benefits to Common Knowledge: User-interruption is a previously innocuous HCI problem that has recently become critical to the successful function of many kinds of modern computer systems. This research will test a theory-based hypothesis about the possible application of human-human interruption methods to designing user interfaces that solve this critical HCI problem. The results of this experiment will be used for two additional purposes: (1) to synthesize user interface design principles about designing HCI that must include user interruption by machine and (2) to support the claimed utility of new theoretical tools -- a general definition and taxonomy of human interruption.

Benefits to Subject: participants will gain experience with a game-like multitask, and contribute to current research in human-computer interaction. No monitory compensation will be given to subjects at NRL; a $20 compensation will be provided to subjects at GWU.

Safety Measures: The experimenter will be present continuously and will monitor the safety of the procedure. In the unlikely event of a computer malfunction or a medical emergency or natural disaster, the experiment will be stopped immediately. At NRL on site Fire or Ambulance (767-3333), and at GWU local Fire or Ambulance (911) services will be called as appropriate

Subjects: No more than 40 total subjects will be recruited from NRL and GWU. This sample size has been sufficient in previous experiments using a similar paradigm. (For example, (Pérez-Quiñones and Sibert 1996)[17] successfully used 30 subjects in a similar experiment with three treatment conditions.) Each subject will participate in one experimental session for three hours or less. Participants will be free to withdraw at any time for any reason without prejudice. A copy of the results will be made available to all subjects at the conclusion of the experiment.

Safeguards: Instructions to the subjects do not mention performance levels or encourage extraordinary effort. Instructions to the subjects explain that the purpose of this experiment is to compare different user interface designs and not the subject's personal abilities. Subjects are instructed that there is no such thing as "good" performance or "bad" performance, but that they should try to maintain a consistent level of effort throughout the experiment. Recruitment of subjects will emphasize the voluntary nature of this study. There will be no coercion of people to volunteer as subjects — no NRL employees will be used as subjects who are subordinate to the principal investigator.

---

17. Pérez-Quiñones, M.A. and J.L. Sibert (1996), "A Collaborative Model of Feedback in Human-Computer Interaction," CHI '96 Conference Proceedings, Association for Computing Machinery, New York, NY, pp. 316-323.

Confidentiality: All data collected (including video recordings) will remain confidential and not be associated with the subject's name or identity outside the context of the experiment or in any published results. Each subject will be assigned a number and that number will be the only identification used to index the results files and the exit questionnaire.

Consent: Each subject will be asked to read and sign a consent form before the start of the experiment. Subjects will be given names, addresses, telephone numbers and email addresses of the experimenters so that they are able to voice concerns at anytime. The consent forms from subjects run at NRL will be stored at NRL in a locked cabinet only accessible to the principal investigator. The consent forms from subjects run at GWU will be stored at GWU in a locked file cabinet only accessible to Dr. John L. Sibert (Chairman of the GWU EE&CS Dept. Human Subjects Protection Committee).

Debriefing: At the completion of the data collection, each subject will be given the attached debriefing form and given an opportunity to ask questions, and provide comments about the study.

APPENDIX D: CONSENT FORM

CONSENT FORM — GWU

By signing this form you give your consent to be a subject in this experiment. Please read the following text and ask any questions you may have before signing.

Experiment: Human-Computer Interaction in Multitasking Environments

Duration of Participation: total time will be three hours or less.

Risks: There are no risks of harm or discomfort anticipated in this experiment. The experimental tasks and environment are similar to those ordinarily encountered during routine work with computers. If at any time you experience abnormal stress, please quit the experimental task and tell the Principal Investigator, Daniel McFarlane (see contact information below).

Benefits: participants will gain experience with a game-like multitask, and contribute to current research in human-computer interaction. A $20 compensation will be provided.

Confidentiality: The records of this study will be kept private. In any report we might publish, we will not include any information that might make it possible to identify you as a participant. Research records will be kept in a locked file accessible only to the principal investigator and the Chairman of the GWU EE&CS Dept. Human Subjects Protection Committee.

Contact Information:

Daniel C. McFarlane (principal investigator)
mcfarlan@seas.gwu.edu; 703-534-2723

Dr. John L. Sibert (GWU EE&CS Human Subject Protection Committee Chairman)
The George Washington University
Department of Electrical Engineering and Computer Science
801 22nd Street, N.W.
Washington, DC 20052
sibert@seas.gwu.edu; 202-994-4953

Voluntary Participation: Participation in the experiment is voluntary and the subject may discontinue participation at any time for any reason. A subject may discontinue participation at any time without penalty or loss of benefits to which the subject is otherwise entitled.

Purpose: To develop an understanding of how user interfaces for multitasking computer systems ought to be designed.

Procedures: Subjects will perform a dualtask similar to a computer video game. These two tasks will involve simple hand and arm movements to make keyboard presses with one hand. The dualtask is composed of the following two tasks: (1) a continuous video game task in which keyboard key presses are made to play the game; and (2) an intermittent matching task in which decisions about matching are made with keyboard key presses.

You will be given a copy of this form for your records.

As a voluntary participant, I have read the above information. Anything I did not understand was explained to my satisfaction. I agree to participate in this research.

____

(Participant)(date)

____

(Investigator)(date)

APPENDIX E: ENTRANCE QUESTIONNAIRE

To be filled out by the experimenter

Subject Number ___________ Code ___________

Location ____

Date ___________

Time ___________

[The answers you give will be kept strictly confidential. In any report we might publish, we will not include any information that might make it possible to identify you as a participant. (Please do not write your name on this questionnaire.)]

2. Your sex? (a) male (b) female

3. Your age? ___________

4. Your dominant hand? (a) right (b) left

5. High school graduate? (a) yes (b) no

   5a. If "yes," how many years of education since high school? 0 1 2 3 4 5 6 7 8 9 >9

[On the following questions please make a mark on the answer line to indicate where your answers fall between the two extremes of possible answers. These questions ask for your informal judgments of your own experiences, abilities, and preferences. Therefore, your answers can not be "correct" or "incorrect."

We realize that you may not have experience answering opinion questions on paper. Some people find it difficult because of this unnatural context. It may help to try to imagine how you would answer each question if it were asked in a more natural context. For example, suppose you are having lunch with some friends and acquaintances and one of them asks, ...]

6. How much computer experience do you have (i.e., amount of time spent working on computers)?

0 1 2 3 4 5 6 7 8 9 10

none considerable

7. How skilled are you with computers (i.e., proficiency with computer tasks)?

0 1 2 3 4 5 6 7 8 9 10

no skill expert

8. How much video game experience do you have (i.e., amount of time spent playing video games)?

0 1 2 3 4 5 6 7 8 9 10

none considerable

9. How skilled are you with video games (i.e., proficiency with video games)?

0 1 2 3 4 5 6 7 8 9 10

no skill expert

10. How skilled are you at juggling (i.e., proficiency juggling physical objects)?

0 1 2 3 4 5 6 7 8 9 10

no skill expert

11. How much typing experience do you have (i.e., amount of time spent typing)?

0 1 2 3 4 5 6 7 8 9 10

none considerable

12. How skilled are you at typing (i.e., proficiency typing)?

0 1 2 3 4 5 6 7 8 9 10

no skill expert

13. How skilled are you at touch-typing (i.e., proficiency typing without looking at keyboard)?

0 1 2 3 4 5 6 7 8 9 10

no skill expert

14. How much experience do you have performing more than one task at a time (i.e., amount of time spent performing multiple tasks at the same time by switching back and forth between different tasks)?

0 1 2 3 4 5 6 7 8 9 10

none considerable

15. How skilled are you at performing more than one task at a time (i.e., proficiency)?

0 1 2 3 4 5 6 7 8 9 10

no skill — expert

16. To what degree do interruptions affect you (i.e., to what degree do interruptions negatively affect your ability to perform tasks)?

0 1 2 3 4 5 6 7 8 9 10

none — considerable

17. To what degree do distractions affect you (i.e., to what degree do distractions negatively affect your ability to perform tasks)?

0 1 2 3 4 5 6 7 8 9 10

none — considerable

18. How much do you try to avoid distractions and interruptions when working (i.e., amount of effort and planning you normally expend to avoid distractions and interruptions when you must get things done)?

0 1 2 3 4 5 6 7 8 9 10

none — considerable

APPENDIX F: INSTRUCTIONS FOR SUBJECTS

INTRODUCTION

You will be presented a series of computer-based tasks. Each 4.5 minute session will be followed by a brief resting period.

You will be asked to perform two different kinds of activities: a game, and a matching choice. The game activity is similar to a Nintendo video game and runs non-stop during a session. The matching activity can happen at different times and each requires a brief choice. Sometimes you will be asked to perform the matching activity while you are playing the game.

All sessions of the computer-based tasks will involve the same two activities (game and matching). There are six different ways the computer is designed to provide these tasks. These different designs are different ways for controlling transitions between the game activity and the matching activity.

Overview
The experiment consists of a series of practice sessions followed by a series of experimental sessions. The six different computer designs will be given to you in a random non-meaningful order. After the experiment, there will be a brief paper-based questionnaire that will ask you about your opinions of the different computer designs.

Purpose
The purpose of this experiment is to compare different computer design methods and not your personal abilities. You can not perform "well" or "poorly" on these computer-based activities because your performance is not evaluated relative to other subjects' performance. We are not interested in how "well" you can do the activities.

The Two Activities

The game is to move U.S. Marine stretcher bearers to catch diplomats jumping from an overrun U.S. embassy and bounce them safely into a military truck. Each falling diplomat must be successfully bounced three separate times at three different locations. If a diplomat is missed at any of the three bounce points they are lost.

The game is trivial when diplomats jump one at a time. However, when more than one diplomats jump in quick succession it becomes a non-trivial game of juggling, and the stretcher bearers must be moved back and forth between the three jump points in quick irregular sequences to keep all the diplomats in the air at the same time. The diplomats jump at random times so the difficulty of the game varies over time. See Figure F1 below.

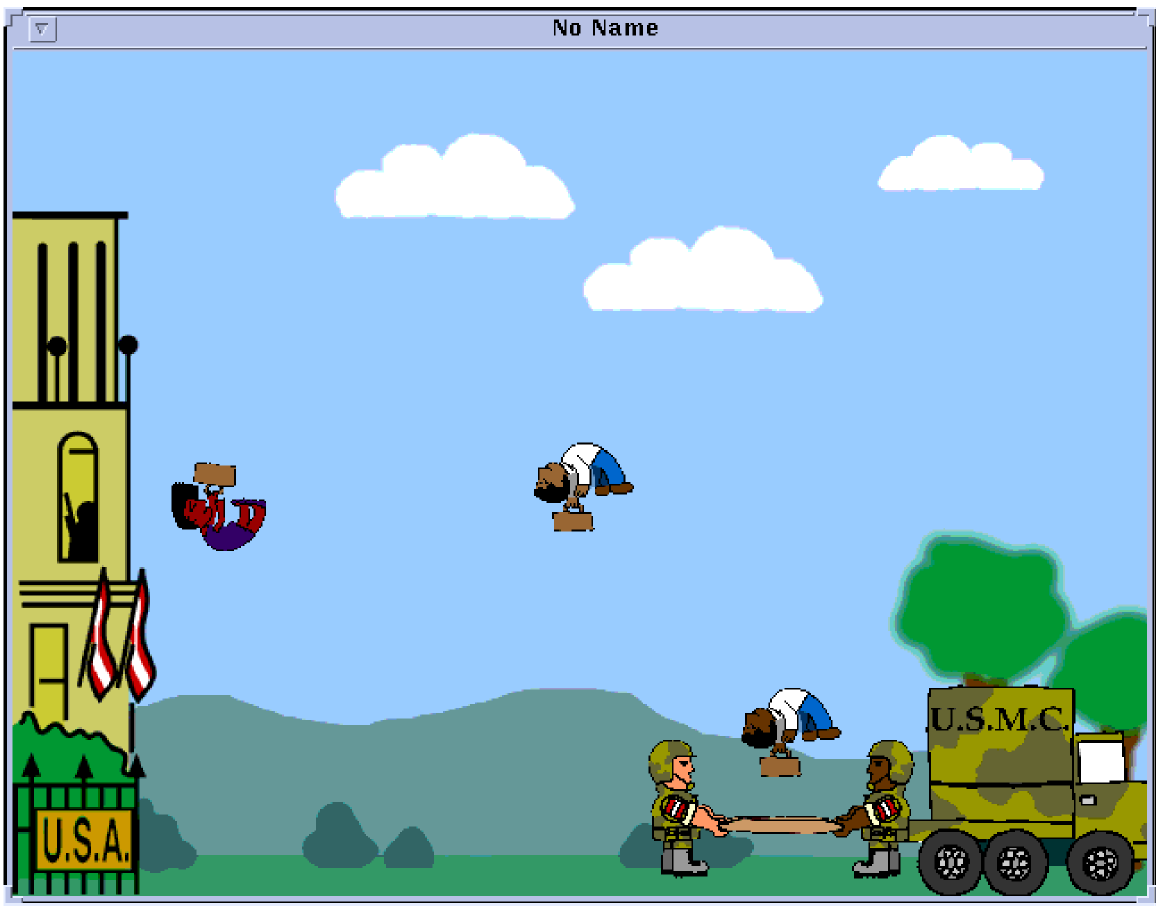

Figure F1.Game Activity: help the diplomats escape the overrun embassy by moving the Marine stretcher bearers to bounce them three times into the truck.

The second activity is a matching choice that requires you to make a matching decision either based on color or shape. A matching activity is presented on the window of a pager-like device. The matching activity totally obscures the view of the game which continues to run. (The game continues to run without possibility for pause regardless of whether it is visible or not.) A colored shape is presented at the top of the pager window, and the activity is to choose one of the bottom two colored shapes either by matching the color or shape of the top object. See Figure F2 below.

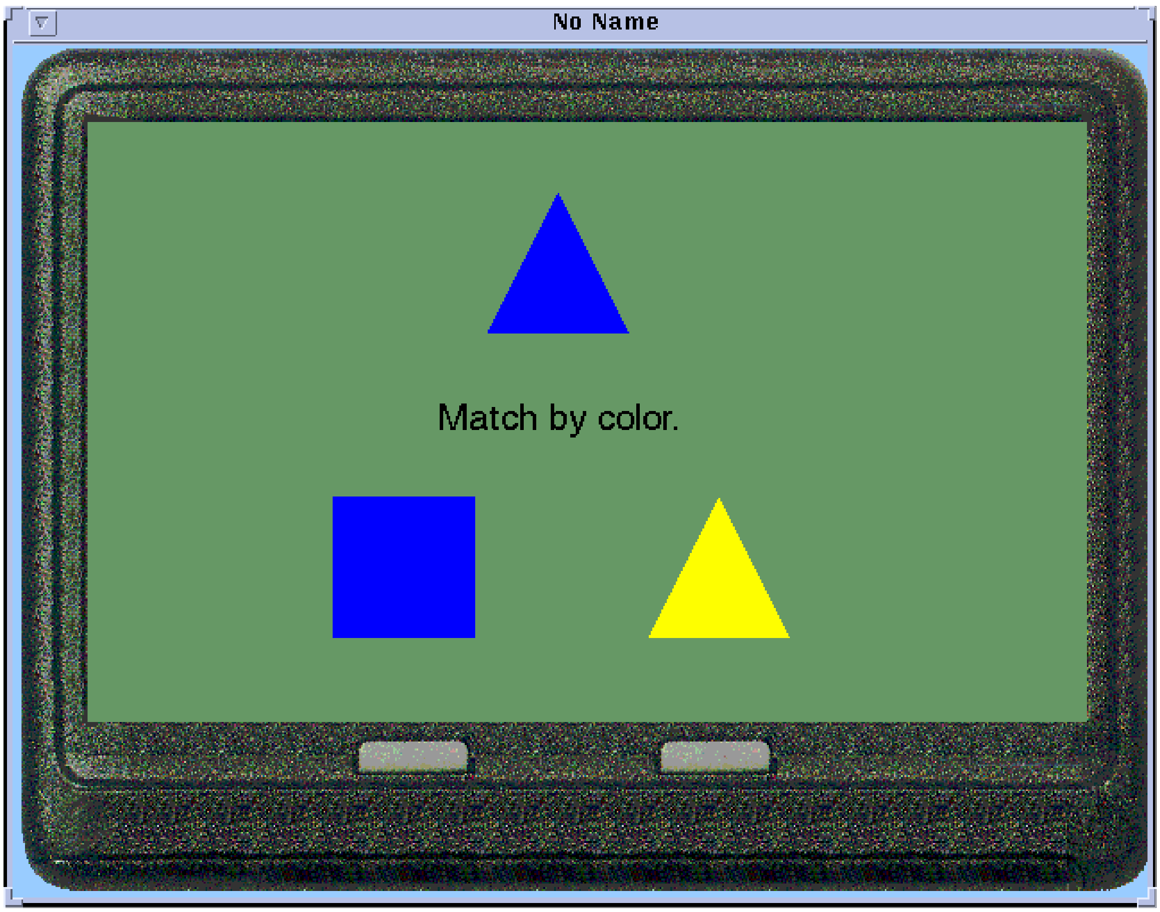

Figure F2.Matching Choice Activity: choose which of the bottom two colored shapes matches the top colored shape according to the displayed matching rule.

Computer-Based Task Instructions

There are six different versions of the computer design (four versions that include both activities, and two versions that include only one kind of activity). (Note, we identify these six different designs with Greek letters instead of numbers or English letters because we specifically do not want to imply any significance in their ordering.)

General instructions.
(1) Please use **only one hand** to perform all computer-based activities. (Please use whichever hand is your dominant hand.)
(2) All activities will be performed by pressing keyboard buttons.
(3) The **game** is performed by pressing the **"delete" and "page down" buttons** to control the back and forth movement of the stretcher bearers.
(4) The **matching choice** is performed by pressing the **"insert" and "page up" buttons** to choose either the left or right bottom shapes.
(5) Individual matching choice activities can **accumulate** and stack up.
(6) Completing the last waiting matching choice automatically returns you to the game in progress (except in the one case π where there is no game).

Specific instructions for the six different computer designs in no particular order.

ψ (psi)

Matching choices are not immediately presented as they arrive but instead accumulate; and all waiting matching choices are presented **every 25-seconds**.

ξ (xi)

Matching choices are presented **immediately** as they arrive.

τ (tau)

Game only.

δ (delta)

Arrival of matching choices are announced with a brief flash immediately as they arrive. **You are in control** of when to perform the waiting matching choice activity(ies). To show the waiting matching choice(s) press the "home" **key**. To hide the waiting matching choice(s) press the "end" **key**. Note that you can show and hide the waiting matching choices at will, so you do not necessarily have to perform all the waiting matching choices consecutively like you do in all other computer designs. (This is the only one of the six different designs in which the "home" and "end" keys are meaningful.)

λ (lambda)

When a matching choice arrives the computer tries to determine whether the game is currently difficult or not. **When the game becomes difficult,** the matching choices are not immediately presented but instead accumulate and are saved until a later time when the game becomes easier.

π (pi)

Matching activity only.

Performance Instructions

There are different performance instructions for the practice sessions and the experimental sessions.

The emphasis for the practice sessions is on learning. Your performance on these sessions is not important and will not be used in calculating the results of this experiment. Instead, what is important is to take the time to thoroughly learn the game and the matching choice activities and the different computer designs. Make an effort to get past the initial learning process so that your behavior is not still improving when you get to the experimental sessions. Please feel free to ask any questions during the rest periods between sessions and during the practice sessions themselves. It is important that you do not have any lingering uncertainty about what is expected of you once you reach the experimental sessions.

The emphasis for the experimental sessions is on consistency. Please do not expend extraordinary effort on the beginning sessions because if you become tired your performance will degrade on later sessions. Instead, concentrate on pacing your work so that you can maintain a consistently high level of effort. Also, please save any questions you may have for the rest periods between sessions.

The game activity and the matching choice activity are of equal importance. If one seems more compelling to you, please make an effort to treat them with equal importance.

APPENDIX G: EXIT QUESTIONNAIRE

To be filled out by the experimenter

Subject Number ____________

[The answers you give will be kept strictly confidential. In any report we might publish, we will not include any information that might make it possible to identify you as a participant. (Please do not write your name on this questionnaire.)]

[On the following questions please make a mark on the answer line to indicate where your answers fall between the two extremes of possible answers. These questions ask for your informal judgments of your own experiences, abilities, and preferences. Therefore, your answers can not be “correct” or “incorrect.”]

18. How much anxiety did you feel during this experiment?

0 1 2 3 4 5 6 7 8 9 10

no anxiety — considerable

19. How motivated did you feel while performing the experimental trials?

0 1 2 3 4 5 6 7 8 9 10

not motivated — extremely motivated

[The following questions ask about your perceptions and opinions of the different conditions of the experiment. Please refer to the written instructions as a reminder of the identities of the different user interface designs denoted with the Greek letters ψ ξ δ λ. In questions that ask for a ranking, no ties please.]

20. Please rank the conditions ψ ξ δ λ by how well you liked or preferred them as 1 2 3 and 4 (1 = most liked, 4 = least liked).

ψ ___

ξ ___

δ ___

λ ___

21. Rank the conditions ψ ξ δ λ by how easily they allowed you to perform the dualtask as 1 2 3 and 4 (1 = most easy, 4 = least easy).

    ψ ___
    ξ ___
    δ ___
    λ ___

22. Rank the conditions ψ ξ δ λ by how many errors you made on the **matching task** as 1 2 3 and 4 (1 = least errors, 4 = most errors).

    ψ ___
    ξ ___
    δ ___
    λ ___

23. Rank the conditions ψ ξ δ λ by how many errors you made on the **game task** as 1 2 3 and 4 (1 = least errors, 4 = most errors).

    ψ ___
    ξ ___
    δ ___
    λ ___

24. Rank the conditions ψ ξ δ λ by how much stress you felt while performing the computer dualtask as 1 2 3 and 4 (1 = least stress, 4 = most stress).

    ψ ___
    ξ ___
    δ ___
    λ ___

25. Rank the conditions ψ ξ δ λ by how interrupted you felt while performing the computer dualtask as 1 2 3 and 4 (1 = least interrupted, 4 = most interrupted).

    ψ ___
    ξ ___
    δ ___
    λ ___

26. Rank the conditions ψ ξ δ λ by how distracted you felt while performing the computer dualtask as 1 2 3 and 4 (1 = least distracted, 4 = most distracted).

    ψ ___
    ξ ___
    δ ___
    λ ___

27. Rank the conditions ψ ξ δ λ by how well you were able to predict the time interval between interruptions (i.e., how long it would be until you would stop performing the game task and begin performing a matching task) as 1 2 3 and 4 (1 = most predictable interruptions, 4 = least predictable interruptions).

    ψ ___
    ξ ___
    δ ___
    λ ___

28. Rank the conditions ψ ξ δ λ by how busy with the game task you were likely to be when interrupted (i.e., how busy with the game task you were likely to be when you had to stop performing the game task and begin performing a matching task) as 1 2 3 and 4 (1 = least busy, 4 = most busy).

    ψ ___
    ξ ___
    δ ___
    λ ___

29. Rank the conditions ψ ξ δ λ by how complex the game task was likely to be when you had to resume playing the game after being interrupted (i.e., how complex the game task was likely to be after you finished performing the matching task(s) and begin to perform the game task again) as 1 2 3 and 4 (1 = least complex, 4 = most complex).

    ψ ___
    ξ ___
    δ ___
    λ ___

30. In condition ψ, while performing the game task, how well were you able to anticipate the next 25 second cycle of interruptions, (i.e., the next switch to the queued matching tasks)?

0 1 2 3 4 5 6 7 8 9 10

no anticipation considerable

31. How much did you like the direct control over when to process interruptions provided by condition δ ?

0 1 2 3 4 5 6 7 8 9 10

none considerable

32. Was the direct control over when to process interruptions provided by condition δ useful for performing the computer dualtask?

0 1 2 3 4 5 6 7 8 9 10

not useful useful

33. In condition δ, how much extra work was it to have to deliberately switch the matching task on and off?

0 1 2 3 4 5 6 7 8 9 10

no extra work considerable

34. In condition δ, how distracting were the flashes of the pager that announced the occurrences of matching tasks?

0 1 2 3 4 5 6 7 8 9 10

not distracting extremely distracting

35. In condition δ, it was possible for a trial to end without you having attempted all of the announced matching tasks. How many of the total number of matching tasks did you complete before the trial ended?

0 1 2 3 4 5 6 7 8 9 10

none completed all completed

36. In condition λ, how well was the computer able to judge the difficulty of the game task (i.e., how well did the computer schedule the presentation of the matching tasks so that you performed the matching tasks only when the game task was less demanding)?

0 1 2 3 4 5 6 7 8 9 10

not well very well

37. Did you notice that the game task was less complex under any of the conditions ψ ξ δ λ (i.e., did some conditions have fewer total jumping diplomats)? (a) yes (b) no

    37a. If "yes," please describe.

38. Did you notice that the matching task was less complex under any of the conditions ψ ξ δ λ (i.e., did some conditions have fewer total matching tasks)? (a) yes (b) no

    38a. If "yes," please describe.

[Blank space is provided below for any comments you have. (Please refer to particular experimental conditions by their Greek letters. Please refer to particular questionnaire questions by their numbers.)]

APPENDIX H: DEBRIEFING

Purpose of This Experiment: This experiment is research about the design of user interfaces for human-computer interaction (HCI). HCI is a field of research that sits half way between computer science and human psychology, i.e., research about HCI attempts to find ways to make computers easier to use for people.

User-interruption by computer in HCI is an increasingly important problem. Many of the useful advances in intelligent and multitasking computer systems have the significant side effect of greatly increasing user-interruption. This previously unimportant HCI problem has recently become critical to the successful function of many kinds of modern computer systems. However, no user interface design guidelines exist for solving this problem. The purpose of this experiment is to examine the utility of applying methods for coordinating human-human interruption to the design of user interfaces. Four different methods for coordinating interruption are included in this experiment: (1) immediate interruption; (2) negotiated interruption; (3) mediated interruption; and (4) scheduled interruption (coordination by prearranged convention for interruption). Analysis of the recorded data will determine to what degree the different coordination methods do or do not mitigate the negative effects of human interruption by computer in HCI.

Possible Benefits from this Research: There are three possible benefits of this research. The first benefit is to test the specific hypothesis of this experiment -- whether known methods for coordinating interruption between people can be applied successfully to the context of computers interrupting their users in HCI. The second benefit would be that the results of data analysis could be used to synthesize some general user interface design principles for building systems that must interrupt their users. The third benefit would be to support the claimed utility of some new theoretical tools -- a general definition and taxonomy of human interruption.

The hypothesis of this experiment is based on these new theoretical tools, and positive results would support the claim that these new tools are generally useful. For more information about these new theoretical tools see the referenced paper (McFarlane 1997). The full reference is included at the end of this debriefing document. (Note, a copy of this paper can be obtained from the experimenter.)

Possible Results of this Experiment: An analysis of the recorded data from this experiment may reveal that user interfaces that include a method for users to coordinate interruption from the computer will allow them to perform tasks more successfully. We have no preconceptions about which method of coordination may be more useful for this particular task. The analysis may reveal differences in utility between different methods of coordinating interruption. If significant differences are discovered, they will motivate future studies.

Thank you for your participation in this study. Please do not discuss the experiment with others until we have run all the subjects, because foreknowledge of the purposes and details of an experiment will unnaturally influence people's behavior as subjects. If you have any questions, please contact Daniel McFarlane at 202-767-2116.

References for Additional Background Information:

Cypher, A. (1986), "The Structure of User's Activities," in *User Centered System Design*, D.A. Norman and S.W. Draper, eds. (Lawrence Erlbaum Associates, Hillsdale, NJ) pp. 243-263.

McFarlane, D.C. (1997), "Interruption of People in Human-Computer Interaction: A General Unifying Definition of Human Interruption and Taxonomy," NRL Formal Report NRL/FR/5510—97-9870, Naval Research Laboratory.

Miyata, Y. and D.A. Norman (1986), "Psychological Issues in Support of Multiple Activities," in *User Centered System Design*, D.A. Norman and S.W. Draper, eds. (Lawrence Erlbaum Associates, Hillsdale, NJ) pp. 265-284.