CIFE CENTER FOR INTEGRATED FACILITY ENGINEERING

INTERPRETIVE WORKFLOW MAPPING WITH THE LANGUAGE /ACTION PERSPECTIVE

by

Nancy (Tigger) Newman & Terry Winograd

TECHNICAL REPORT
Number 80

July, 1992

Stanford University

If you would like to contact the authors please write to:

c/o CIFE, Civil Engineering,
Stanford University,
Terman Engineering Center
Mail Code: 4020
Stanford, CA 94305-4020

Summary
CIFE Technical Report #80

ABSTRACT.
In this paper, we describe the general languge/action approach, lay out its analytic framework, describe a hypertext mapping tool that we developed for workflow mapping, and discuss the results of applying it in analyzing work in a coordination center of a medium size construction company.

SUBJECT.

This paper describes research into a method of mapping out the ways in which people interact to get their work done. The goal of the research was to develop a technique and a tool to identify the complex webs of requests, offers, and other languge actions that take place among a team of workers and the people with whom they interact in the process of getting their jobs done. Such a global view of work is necessary in creasting computer systems that support not just specified subtasks, but the entire process.

In previous work we have shown how a "langauge/action" appraoch can be used as the basis for designing workflow management technology. In this paper we report on a small-scale seed research project aimed at refining one aspect of this approach, the detailed mapping of workflow and the associated physical and communicative actions that support it.

OBJECTIVES/BENEFITS.
CIFE funded the research because coordination is a key element in all construction processes. Computer tools to improve coordination could have a broad payoff. Member companies could potentially benefit from new computer-based systems that use the theoretical elements produced in the research, and can also benefit from using the theoretical framework as a way of analyzing their own coordination activities. The research attempted to show that a systematic form of work mapping could be used to identify the ways in which comptuers could aid in the coordination of construction work.

METHODOLOGY.
This project involved creating a technique and a tool for determining the interactional context of the current activities in a workplace by gathering information about the interactional role that each language action plays. Gathering this information under the language/action perspective involved viewing the users' activities as a web of committments and separating the various strands of interactions. We identified a field study site that had multiple, interlinked parties carrying out activities requiring a high degree of coordination, with the possibility of frequent, identifiable coordination breakdowns and the potential for innovative computer use. Before choosing the experimental site, our research had already started with the development of the interpretive workflow mapping computer tool (IWM1) to store and analyze the experimental data and its contextual relationships. After subjects agreed to be studies, initial interviews and observations were made of approximately eight hours. During this time, notes were taken by hand and the entire session was audio-taped. After this first session, a rough mapping of the work structure was created using IWM1 and the questions for a second round of observations were generated from the information that was revealed as still missing during the mapping. These questions then served as the basis for several days of focused interviews and observation. The data gathered in this second round was not only used to fill in the blank field in the IWMI analysis, but more importantly, to shape the tool itself.

RESULTS:
When we set off to do this study, we anticipated being able to map a substantial part of the work structure, and to use that map as the basis for specific proposal for new possibilities for using computer-based tools to facilitate the work. We hoped that explicit mapping of individual instances of people handling breakdowns would be useful for revealing the structure of work that goes beyond the "standard Procedures."

As the work progressed, the scope of this kind of mapping became more apparent, and it was clear that it was not feasible to do broad mapping of the desired depth. In order to keep the data manageable, it would be necessary to have a much more focused set of specific design proposals (e.g., keep track on-line of equipment repair status) in order to find the right set of relevant areas to map. This is in accord with the observations that have been made more recently by proponents of traditional system design methodologies, that an overly thorough form of analysis can become a block to getting design done.

Since our primary goal in this project was to devlop the theoretical framework, we had not entered into a site with specific design needs. There was no clear guiding criterial for sorting out the interesting form the irrelevant details. We chose to continue the work in a broader, less-directed fasion, looking for those places where it led to new insights for us, and where we could make occassional suggestions about possibilities for changes in the work structure. The outputs included the prototype mapping tools, which could be used in further research but are not in a state for general distribution.

STATUS:
The research has led to better frameworks for analysis, but has not yet reached the point of applicable tools. The next step for making this kind of mapping a practical part of a design process is the integration of the kind of "bottom-up" structure that it implies with a more "top-down" focus on what the critical problems are that need solution and merit design efforts within the organization and its work. The mapping metaphor needs to be interpreted in a narrowed sense: not of making a complete map, but of using mapping techniques on selected details. This will be the subject of further studies, and will have an influence on the development of comercial systems (e.g. by Action Technologies) as well as our research prototypes. (This effort showed that a detail analysis of work flow without an overall focus for the study in anot an adequate method. You must start with a goal for systems analysis and then support it with a detailed method that reveals the transactions and logic necessary to support computer systems design).

Interpretive Workflow Mapping with the Language/Action Perspective

Nancy (Tigger) Newman & Terry Winograd
Computer Science Department
Center for Integrated Facilities Engineering
Stanford University
July, 1992

ABSTRACT

This paper describes research into a method of mapping out the ways in which people interact to get their work done. The goal of the research was to develop a technique and a tool to identify the complex webs of requests, offers, and other language actions that take place among a team of workers and the people with whom they interact in the process of getting their jobs done. Such a global view of work is necessary in creating computer systems that support not just specified subtasks, but the entire process of getting the work done.

In previous work we have shown how a "language/action" approach can be used as the basis for designing workflow management technology. In this paper we report on a small-scale seed research project aimed at refining one aspect of this approach, the detailed mapping of workflow and the associated physical and communicative actions that support it. In this paper, we describe the general language/action approach, lay out the analytic framework, describe a hypertext mapping tool that we developed for mapping, and discuss the results of applying it in analyzing work in a coordination center of a large construction company.

ACKNOWLEDGMENTS

This work was sponsored by a grant from the Center for Integrated Facilities Engineering (CIFE) at Stanford University, and was supported in part by a grant to the Center for the Study of Language and Information by the Ricoh Corporation. We want to thank Paul Teichholz and Bob Tatum at CIFE for helping make the connection for our field study and for insightful comments in the writing of this report. Foremost, we want to express our gratitude to Sam Free, Bob Jones, Terry McIntyre and Glen Young of the Raisch-Gradeway construction company for giving freely of their time and expertise to help us understand their work and see new possibilities.

SUMMARY

This CIFE seed project explored the use of detailed work mapping techniques within the larger context of the language/action approach to work redesign, which has been developed over the past few years in a number of projects at Action Technologies and Stanford (see Winograd and Flores, 1986, Winograd, 1987, Dunham, 1991, Kensing and Winograd, 1991, Medina-Mora et al., 1992). Our goal was to explore the possiblities for generating design possibilities for a computer-augmented work setting by explicitly relating two different dimensions of analysis: language/action, and physical activity. The main research effort was in developing a structured set of categories and questions (embodied in a computer-based hyptertext mapping tool) which could be used in the exploratory phase of design.

We began with a general framework from which to analyze work settings, which we wanted to test in examining an actual situation. The research was conducted by using on-site observations and interviews to test and then refine the analysis framework. We did not attempt to enter into the full design cycle, focussing instead on the problems and possibilities for the mapping tools. Therefore the results do not provide a justification of the technique, but an initial exploration of how it might be applied. Any actual design results are anecdotal and were not validated in terms of changes that could be made to the work structure.

Several primary things were learned from the experiment:

1) The basic ontological framework of the language/action approach is sound and can be used effectively in structuring the analysis of coordination work.

2) It is possible to effectively build and use computer-based tools to help deal with the complexities of analysis, and especially as a way of pointing out possible interactions that would not be identified with ordinary observation and interviewing techniques.

3) Even with computer-based analysis tools, it is overwhelming to do a comprehensive description of a real work setting. The amount of detail needs to be controlled by a more "top-down" approach in which only selected parts of the work organization and implementation are analyzed.

Our goal in future work will be to better integrate detailed work mapping with a more goal-directed approach, so that the advantages of possibility generation through the detailed observation of both the physical and linguistic dimensions can be combined with the effective focus provided by more highly directed design techniques.

I. BACKGROUND

Computers have become an integral part of people's work lives, and are being integrated into a wide variety of tasks. Recently there has been a growing recognition that individuals do not work in a vacuum: computers are not only to help individuals get tasks done but to help them work together more effectively. Creating ever better word processors and spreadsheets is a worthwhile endeavor but fails to reach the higher level at which these types of tasks are put into a context of interaction with other people.

There has been much recent focus on the ways that people interact *through* computers rather than just *with* computers. It has been characterized as "groupware," "computer-supported cooperative work," and "workflow design." (For overviews, see Greif, 1989; Johansen, 1988; Schrage, 1990; Greenberg, 1991; Grudin, 1991). The key idea is that by paying attention in the initial design to the way that people in an organizational setting will interact through computers, it is possible to generate new and more effective designs for their work, rather than simply automating the existing practices.

The language/action perspective on group work (see Winograd & Flores, 1986, Winograd 1987, Dunham, 1991, Kensing and Winograd, 1991) grew out of a theory of language action called "speech act theory." It focuses on the way people coordinate their work through language - or more broadly through communicative interactions which include both computers and other media. It is centered on facilitating the possibilities for people to request that someone do something for them, or offer to perform some service, and in other ways develop and manage the flow of work.

There are many approaches to designing computer systems and each one involves viewing the tasks to be done through a filter that highlights what is important from that particular perspective. Each possible perspective inevitably loses some of the flavor of the workplace activity, since the process of highlighting some aspects of the work de-emphasizes other aspects. Without some perspective, however, there is too much information to handle at once; some sort of guide is needed to help structure the collection of data so that the designer is not bogged down in overwhelming detail.

Some traditional design approaches focus on office layout, others on the objects manipulated, personality interactions, personal motivations, and so on. One of the most common perspectives used currently in system design for work places is information-oriented, with a focus on how people produce, relate, track, analyze, and have appropriate access to information stored in many forms, including

databases, documentation, display screens, etc. The design focus of this information approach is formulated in terms of the system's ability to accomplish or facilitate the generation or access of such forms of information. Thus, the design techniques tend to rely on formal job descriptions detailing each person's updating and use of this information. The goal is then to use such forms of information in creating solutions to regularly recurring points at which decisions must be made.

However, information as information is not interesting or useful; information is relevant because something needs to get done. It is used as part of a request, report, or the like, that has significance to people engaged in coordinated activity. For example, the task of typing a document is not of any value, except for the interpretation that it is a work schedule or a contract or memo. Information and tasks exist within a web of commitments that people make to each other and it is this network of requests, offers, promises, reports, and declarations that brings the information and tasks to life. The focus of the language/action approach is on revealing the interactional context in which information is produced. This facilitates design which goes beyond simply supporting the handling of the information, to capture and highlight its interactional purpose.

The first step of any design process is to understand what is being attempted—what breakdown or opportunity are we designing for and what constitutes a satisfactory intervention? With the language/action perspective, the analysis is in terms of effective interaction between people to accomplish tasks. This focus on purpose rather than form facilitates redesign since jobs and tasks are seen in terms of how groups of people interact to get things done. Since most new computer systems are used in pre-existing settings where manual or automated processes already exist for getting things done, designers needs to recognize that they are not working from scratch but are "redesigning" processes that already are in effect. This redesign may be simply to automate the current procedures to make the process faster, or may involve fundamental changes to the process itself. Redesigning a current work situation with the language/action perspective allows attention both to the current activities and to the underlying interactions that they are a part of. Thus, it can help avoid the possibility of slavishly designing a system to support the current ways of doing business or naively designing one that deals only with an idealized view of a task.

II. The Project: Developing interpretive workflow mapping

This project involved creating a technique and a tool for determining the interactional context of the current activities in a workplace by gathering information about the interactional role that each language action plays. The focus is on how people are communicating: what promises and offers they are making to each other, what declarations are being made and to whom, and so forth. Besides these patterns of interaction between people, the technique identifies recurrent sources of breakdowns—points at which someone declares that something is missing—and supports design for anticipating and coping with these breakdowns.

Gathering this information under the language/action perspective involves viewing the users' activities as a web of commitments and teasing apart the various strands of interactions. The process developed during this project disentangles the hows and whys of actions in a cyclical fashion with both observation and interviewing of the prospective users. This process is called "interpretive workflow mapping" because the workflow is interpreted by an observer who identifies the major interactive objects, called conversations, and maps the relationships between them and other entities identified as relevant in the workplace.

The current project has focused on two major goals: studying design-oriented analysis of work plus the development of the interpretive workflow mapping technique and tool. This process has resulted in the expansion of the language/action perspective framework as well, since the development of this technique has revealed those places where connections needed to be made between the basic structure of language acts and other observable work characteristics. It has been a cyclical process reaching towards the development of a useful set of questions to map out the current set of conversations. These questions are embedded into a tool that has been developed in the course of this project: IWM1, a computer program for exploring and recording information about workflows and the relationships among them that are revealed by language/action analysis.

II.A. The Experimental Site

We identified a field study site that had multiple, interlinked parties carrying out activities requiring a high degree of coordination, with the possibility of frequent, identifiable coordination breakdowns and the potential for innovative computer use. A San Francisco Bay area construction company's central dispatch office was chosen since it meets all of these requirements and also has a tradition of embracing

innovations in work methods. They currently have not computerized much of their operation, but have developed very effective non-computer methods of getting complex interactive jobs done. The choice of a coordination center allowed us to easily focus on the pattern of commitments between people since that was the very visible content of the work.

The dispatch office is staffed primarily by three dispatchers who communicate with a variety of other people each day, including laborers, equipment operators, foremen, area supervisors, equipment movers, equipment rental companies, supplies companies, shop workers, permit companies (who obtain equipment moving permits from the state government), union officials, trucking companies, sanitary companies, quarry companies, and dump facilities, as well as the various supervisory and support staff within the company headquarters. This communication is conducted in person, and through radios, phones, computer printouts, schedule boards, marked calendars, and forms of various sorts.

The three dispatchers are responsible for taking care of the construction employees by handling everything from dealing with the union when a person is hired, to getting them trained, telling them when and where to be for their next assignment, gradually increasing the complexity of the tasks that they are assigned, and balancing assignments to provide all employees with as much work as possible even when work orders are low. The dispatchers are responsible to the foremen and area managers for getting all of the employees, equipment, and trucks that are needed to the job site on time, changing any staffing assignments that won't work, and having any malfunctioning equipment repaired or replaced. They are responsible to the field engineer and the accounting staff for purchasing supplies and renting equipment at the lowest possible cost, consistent with quality and service requirements, and for keeping accurate records of all costs incurred by a job. They are responsible to the equipment movers for arranging a reasonable schedule and moving permits, to the shop staff for keeping accurate information on the current location of all equipment, and so forth. They are the central communication link for the company's daily activities, both for planning and for trouble-shooting.

II.B. The Research Activities

Before choosing the experimental site, our research had already started with the development of the interpretive workflow mapping computer tool (IWM1) to store and analyze the experimental data and its contextual relationships. After subjects agreed to be studied, initial

interviews and observations were made, comprising approximately eight hours. During this time notes were taken by hand and the entire session was audio-taped. After this first session, a rough mapping of the work structure was created using IWM1 and the questions for a second round of observations were generated from the information that was revealed as still missing during the mapping.

These questions then served as the basis for several days of focused interviews and observation. The data gathered in this second round was used not only to fill in the blank fields in the IWM1 analysis, but more importantly, to shape the tool itself. The set of structures represented in it was expanded and reshaped to handle additional real-world data that its original structure had not allowed for. As the structures expanded to fit the data, it challenged the original theory which was then expanded to include additional distinctions and types of information. The final round of interviews and observations were about four hours long and served to fill in specific gaps in the model of the work that had been created on the IWM1 tool. The following section on results contains a description of this IWM1 tool, the elements that it portrays, and the associated information that is appropriate to each.

II.C. The Technique of Interpretive Workflow Mapping

Designing computer systems with the language/action perspective involves first using a structured set of questions to identify current interactions within the web of commitments. The questions reveal both the "language/action domain" (the patterns of requests, offers, declarations and other conversational acts that show what work is getting done between people), and the "activity domain" that serves as its embodiment (the people, locations, objects, and physical actions that reveal how things are getting done). One of the key results of this research was clarifying the distinctions between elements of these two domains, and the relationships between them.

The data for the activity domain is the most easily accessible from direct observation while the data for the language/action domain tends to be generated in discussions with the people involved, asking questions such as "Who gets a copy of this report?", "Why did you start doing this when you did?", and "What happens if it doesn't get done on time?" The interplay between the two domains produces a continuing cycle of discovery during the analysis as the examination of something observable, such as an activity like filling out a form, prompts an examination of the interactions it participates in, while the identification of an interactional structure prompts an inquiry into the

observables that correspond to its embodiment in physical objects and actions.

Our work made it clear to us that at the current stage of development, this methodology served best as a relatively open-ended method for prompting discussion and discovery. It did not have the focus that would lead to structured ways of reaching conclusions and decisions about workflow enabling and the use of computer systems. A major focus of future work is the development of more directed methods that can take advantage of the kinds of information that come from work mapping.

II.D. The Structure of Distinctions Used in Mapping

The focus of the language/action perspective is the identification and use of workflow patterns, which generate the web of commitments that people make to each other: promising, offering, reporting, acknowledging, etc. The key elements are individual conversations, the acts that they are composed of, identification of the related breakdowns (or opportunities) that occur during, before, or after them, the roles of the agents involved in the conversations, and any routines that include acts within those conversations.

In order to clarify the differences between the language/action domain and the activity domain, it is necessary to make distinctions between elements of the two that have a strong correlation with each other, but are nevertheless different. There are three points at which the an element of the activity domain is roughly comparable to an element of the language/action domain: people with agents; activities with acts; and sessions with routines. These points and the internal dependencies of the various types of data within each domain are shown in Figure 1 and explained further below.

In this figure, the activity domain categories of data are on the left and the language/action categories are on the right, with an arrow between two categories indicating a relationship between them. The direction of each arrow indicates that the definition of a piece of data that is being pointed to, such as a gap, is dependent on the definition of one or more pieces of data of the type at the other end of the arrow, such as a role. A double-headed arrow indicates that the data on both sides of the arrow are mutually dependent. When a type of data is dependent on the definition of several other types of data, that is indicated by separate arrows for each or by an encompassing bracket sign with one arrow leading from it, such as the bracket enclosing the group of things, locations, and people with an arrow leading from it to activity. The dotted lines indicate the dependencies for the aggregate categories: session and routine.

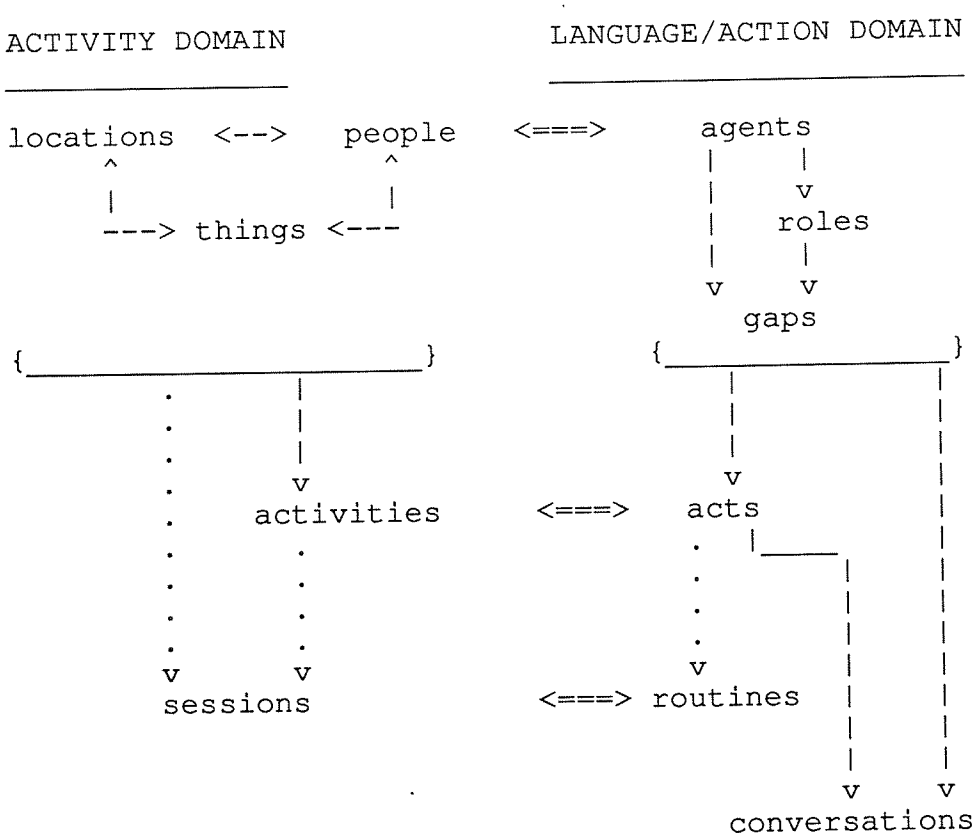

Figure 1: Categories in the activity and language/action domains

1. *The Language/action domain*

The categories of language/action domain data (on the right side of the figure) are:

Acts: Acts are the interactional speech acts that agents make, such as making declarations, promises, offers, reports, and requests. They are how agents create, change, and move within the web of commitments.

Gaps: Gaps are declarations by an agent that something is missing, a solution to a problem that needs to be handled or an opportunity that needs to be considered. These are sometimes referred to as "breakdowns" in language/action literature but the terms "gap" and "incompletion" were adopted here because of their more neutral connotation; even when work is running smoothly, there is always something that still needs to be resolved in an unfinished job, even if it is only the next standard step in the process.

Agents: Agents are any entities that can make language acts in their own name (as opposed to message-conveying). This includes people as well as organizations acting through people. (For example, it is the gravel company that is committed to the delivery, not the employee who receives the order.)We are using the term "agent" here in the sense that artificial intelligence and philosophy uses the term, as a being that can take action, not in the sociological or everyday sense of "agent" as a representative.

Roles: Roles determine the structure of possibilities for individuals to act in conversations. More than one agent may play the same role and more than one role may be played by one agent. Roles may correspond to official job titles and functions or may be unofficial, such as the role of "the person everyone relies on to keep a copy of every instruction manual."

Routines: Routines bring together a series of language actions and physical activities which recurrently are done together as part of carrying out a role. Typically they have a clear spatio-temporal continuity, as in the routine that pilots go through before takeoff, or the routines that people perform upon coming into their workplace in the morning.

Conversations: Conversations are coherent combinations of language actions by which parties coordinate their activity. A conversation is defined as a structure of language acts by which the participants move towards agreement and completion of some particular conditions of satisfaction.

The Activity Domain

The activity domain supports the language/action domain with a physical foundation. With any approach to analysis of activity, basic building blocks need to be selected and used to describe that activity. These building blocks include locations, people, types of information, times, objects, speech between people, etc. Developing a methodology involves selecting which categories are of interest and what level of detail is appropriate. For example, in analyzing a conversation, the researcher might note just the fact that it occurred, the general subjects, the words used, or the actual pronunciations, sounds and pauses in the utterances. Similarly, when recording the activities of a subject, the researcher might note each physical motion, the general point of each motion, the motions that have a work-related purpose, or the more interesting motions from some perspective. Without some bias, there is an infinite amount of data. Success in examining a situation relies

on selecting relevant details and structuring them in some useful fashion.

With the language/action perspective, the determination of *relevant* and *useful* are dictated by relationships to the underlying conversations and other constructs of the language/action perspective. It is impossible to reliably identify the offers and requests being made without noting who is making them, to whom, when, and through what medium. This project used the following categories of activity domain data:

Activities: Activities are any physical movements, including speech, picking up telephones and sticking a magnetic tag on a schedule board; the interesting ones are those which are part of creating and fulfilling commitments such as making a telephone call which involves a promise, etc. As noted above, we observe in terms of things that are relevant to the language/action perspective: e.g. "He said 'Okay'" rather than "He raised his head, vibrated his vocal chords and moved his lips."

People: People are physical entities, who in turn act as agents in the language/action domain (as described above). For example we might observe that a phone call is being made to a particular person, who in turn is acting in the identity of some supplier company.

Locations: Locations are any standard places in which the people are located when acting in the web of commitments. Equipment relevant to work activity is typically situated in standard locations, and coordination of movement to and from those locations is relevant to the flow of work.

Sessions: A session is a collocation of people, locations, things, and activities that can be usefully considered continuous in space/time, such as a staff meeting, a computer session, or a telephone call. Note that these are different from routines and from conversations, as their unity comes from a unity of place and time, not of underlying purpose.

Things: Things are physical objects which play a role in the language actions and in the physical activities involved in the web of commitments.

The Relationships Between the Basic Categories of Data

The different relationships among the basic elements of the two domains are described in detail in Appendix A. In general, the dependencies in the graph lend themselves to making five groupings, as follows:

Group 1: People, Locations, & Things—Activity Domain Basics
Group 2: Agents, Roles and Gaps—Language/Action Domain Basics
Group 3: Activities and Acts—Doing the work
Group 4: Sessions and Routines—Collections of Work
Group 5: Conversations—Apex of the Language/Action Domain

The first group of people, locations, and things is reminiscent of a children's game's basic categories of "people, places, and things." The three categories rely upon each other for their definitions, but can be described without reference to any other type of information.

The basic categories of the language/action domain are agents, roles and gaps. Agents have a strong correlation with the people category in the activity domain, since people can act as agents for themselves, as well as acting as agents for companies and other entities. The definition of roles depends on agents, since part of the definition of a role is the identification of those agents that can assume the role. Gaps depend on both agents and roles, since the recognition that a gap exists and other aspects of a gap need to be made by an agent who is filling a role.

The next group, activities and acts, are roughly analogous since both are the way in which work is done in their domain; activities are the physical motor and speech movements that perform the work in a physical sense and actions are those language actions that perform the work of moving through the web of commitments to fulfill promises, etc.

While activities and acts are how work gets done in both domains, sometimes the pattern of their occurrences is usefully viewed in clusters, such as the sessions of the activity domain (e.g. the session of a single telephone call between two people discussing several subjects) or the routines of the language/action domain (e.g. a pilot's checklist, including communication with the control tower).

Finally we have the conversations themselves, the focus of the language/action theory, tying everything together and depending on everything else except the collections of work: sessions and routines. Sessions and routines may be part of a single conversation or contain part or all of the actions or activities contained in several conversations. Although conversations are also a collection of actions

or activities in a sense, but they are structured around the goal of creating and meeting commitments and focus not only on what is being done, but also on the role it plays in completing gaps and reaching satisfaction of the commitment.

III. The Tool for Interpretive Workflow Mapping: IWM1

The interpretive workflow mapping tool (IWM1) was created on a Macintosh using SuperCard, which provides a flexible capability for creating display structures and making relationship links. This program includes the basic functionality of SuperCard and HyperCard programs, such as printing, etc. as well as providing some special functions for building a workflow map. For each major category of data (role, act, activity, etc.) in the language/action theory, we designed a card (a kind of description template) for prompting the entry, and then storing and displaying the pertinent information about that element and its relationships to other elements. Full information on the contents and intended use of each field is included in Appendix A. Most of this information is also accessible directly to the user as a help function during operation of the program.

III.A. Operation of the IWM1 tool

The program provides not only these templates for entering information but also methods of accessing and associating the data. It maintains a list of all data cards, as well as individual lists for each category of data and a list of all lists. In addition, users can design their own criteria (based on some combination of keywords, fields, and categories of cards) for constructing lists for a single use or continual maintenance. In addition, data cards or fields, or individual entries within fields can all be linked directly with any other card, allowing the user complete flexibility in setting and traversing any links in the tools that are useful for analyzing the data. Commonly these links are used to link a mention of a piece of data on some card to the card describing that data in detail.

Users can easily move from card to card, being able to traverse the links in several ways, and being able to move around within a category of data in several ways. In addition, users are able to view several cards simultaneously. Although the program separates different projects at the beginning of operation, it does allow the user to switch between projects or even view several simultaneously if desired for analysis. Special attention was paid to providing flexible methods of searching for cards by combinations of keywords, fields, etc. to provide maximum accessibility for all cards.

Full facilities are included for easy creation of cards, not only allowing creation of a card from scratch, but also copying (with renaming) of cards, and shortcut facilities for entering multiple cards. Cards can always be renamed, with these renamings showing up properly on updated lists and the data on the cards can be updated at any time.

Deletions of cards is also allowed and will affect the lists properly, but the user is prevented from deleting the templates. Comment cards can be added if desired and most of the fields permit scrolling to accommodate as much information as necessary.

IIIB. An example of workflow mapping using IWM1

A study of the work site began with interviews and observations of the people involved in the work, documenting the interviews with notes and audio tapes and sketching and photographing the layout of the office. This first interview is designed to give an idea of the scope of the work being done and to give the analyst enough information to begin blocking out the work patterns on the IWM1 tool.

Upon return to the analyst's own facilities, the first massive data input takes place. Using the notes, sketches, photographs, and audio tapes (preferably played on a dictation machine), a great number of data cards can be created at once. On this first pass, a card title and card type (conversation, action, location, gap, etc.) are input for each card, along with the "class" parameter.

Screen 1 (see Appendix B) shows a new conversation card being created.

After the first group of data cards are created, the cards are filled out and any new cards the analyst determines are needed can be created. Each newly created card already contains its name in the title section and the appropriate types of fields for its card type (conversation, action, location, etc.) As an aid to filling out the fields, the analyst is prompted by questions which are defined as the initial contents for each field. (Some fields must be scrolled to see the entire text of the question or all of the data input by the analyst.) As specific values are filled in, the questions are replaced (but can be accessed any time through the help function.) The questions are intended to suggest the ways that the different workflow elements can be linked. Appropriate sets of questions are provided for each type of card. For a complete set of the ones used in the initial tool, see Appendix A.

Screen 2 shows the newly created card for the conversation "Scheduling trucks" in its initial state with questions identifying the places to be filled in.

The filling in of the data cards reveals patterns in the work; questions that arise in filling in the fields can highlight areas that require more information. In the case of the "scheduling trucks" conversation, it was clear from the original interview that scheduling trucks was just part of the entire scheduling of job sites, which included arranging for

equipment and employees as well. Thus, at least one superclass, "scheduling jobs" is possible.

The next fields, "type" and "mood" refer to the kind of conversation is taking place. The most common answer of "action" for the conversation mood is the default for convenience, but other possible choices are possibilities, clarification, or orientation (see Winograd, 1987). The "type of workflow" field is used with workflows that have an action mood and the answers can consist of: offer/completion, request/completion, declaring, or informing. In this case, the type is request/completion, the initiator is the job foreman, with the dispatchers as participants. It's not clear on the first pass who the synchronous and asynchronous listeners to the conversation are, so these fields are left blank until further interviewing and observation.

Incompletions, also called "gaps", tend to become more apparent on subsequent iterations of the interview/observation/input cycle. In this example, two were mentioned in the first interview: the foremen giving insufficient information and the lack of a company truck to do the job. The justification is more clear: the truck is at the job site the next day, loaded with the right materials or empty , as required. The trigger cited in the first interview is the job foreman calling in but this expectation may not always occur -- further observation and interviewing may provide back-up trigger conditions. This conversation seems to be considered satisfactorily completed if the correctly loaded trucks are at the job site as requested, without any formal reporting being required as long as all goes well.

The last required field on the card is the "acts" field which will mostly have to wait for further interviewing and observation. The final field, "notes," is simply for adding notes and questions on the card for convenience. If more space is required, a "comments" card can be linked to extend the amount of total space.

Screen 3 shows the same card after filling in the card with the information available.

This process then continues, with the data input generating new questions for further research, the results of which can be used to finish the cards and to further define the network. There are various navigational commands within the program that can be used to find related cards to fill in or link. For example, the "find cards" command (provided in the menu under the "Actions" category in the menu bar) can be used to select and list cards based on the card type or on various aspects of the data stored on the card, including its type and the contents of its fields (using the string search capabilities provided by

SuperCard). This listing then allows direct access to the listed cards from its list, by simply clicking on the corresponding line in the list.

Screen 4 shows the "find cards" command being used to list and provide access to all conversation type cards.

For some fields (see Appendix A for details), simple text answers are enough, but the most central ones commonly involve linking to other cards. All fields are capable of supporting such links. There are items in the menus for navigating these links (jumping to a card that is referred to, returning to the previous one, etc.). To create a link, first the originating card and field are selected and then the linker called from the menu bar.

Screen 5 shows what appears if, when filling in the "renting trucks" card, the analyst knew that a possible incompletion would be if the rental truck company was busy and called the linker to make a corresponding link.

Any of the program's navigational commands (as well as the Make A New Card command) can be used to obtain the desired card for being linked to the "incompletions" field of the first card. In this case, the card "rental company trucks busy" already existed but had not had any of its fields filled in.

Screen 6 shows the situation after finding this card using the finder.

The linking is then completed by selecting either "link and go back" to return to the "renting trucks" card and continue working with it, or by selecting "link and stay" to keep the "rental company trucks busy" card as the most current card and further develop the analysis of what would happen in the face of that breakdown. The analyst is free to fill in the fields and create cards in whatever order makes sense, with the links linking the cards together in a easily navigable network and the cards and fields themselves generating appropriate questions. Thus the tool itself helps to prompt the analyst in uncovering the language/action interactions present in a workplace.

IV. Conclusions

This project ended up tackling several interrelated tasks: adapting the theoretical ontology on the basis of experience, adapting the methodology on the basis of experience, applying the framework systematically, and creating the IWM1 tool. The people whom we studied understood that we were exploring how to describe current work patterns, not yet new ways of designing those patterns and the tools to support them. However, it was still difficult to manage their expectations appropriately, especially since we had purposely chosen a company and department that had shown an eagerness to try innovations in their work patterns. We were able to make some minor suggestions based on our data but significant changes were inappropriate due to the early stage of our methodology and tool development.

IV.A. What Did We Learn in Adapting the Ontology on the Basis of Experience?

We began this project with the ontology of the language/action perspective, including acts of various sorts, gaps, and roles to describe how various people interacted to get work done through a network of commitments. Taking the theory and using it to shape observations of an actual work situation involved adapting the ontology to handle items that were directly observable. Although our intent was to map the underlying language/action structure of the work, our methods of determining that structure depended upon observations as well as interviews. Thus one of the major results of this research was the definition and inclusion of the activity domain into the ontology.

One of the areas that was brought into focus by this integration of the activity domain was the existence of routines and sessions, ways in which acts and actions were grouped, possibly across conversation boundaries. In addition, the distinction was made and sharpened between people, their roles, and agents (who could be people, companies, unions, etc.)

IV.B. What Did We Learn in Adapting the Methodology on the Basis of Experience?

Our methodology was based on the goal of comprehensive logging and mapping. We soon discovered that there was no hope of doing a "comprehensive" map, even for one part of the work (e.g. one person or task). Unlike other types of analysis that focus on effectively performing a particular function, such as typing a letter into a text editor, the very nature of the language/action perspective demands a

sensitivity to the interconnectedness of people and tasks. This is one of the strengths of the language/action perspective but it involves the recognition of an open-ended network where there is more potential detail everywhere.

One of the ongoing questions during the observations and interviewing was the level and breadth of detail to include. In the beginning, close to everything was treated as potentially relevant but as the methodology revealed the underlying language/action structure, the ability to focus on that structure increased. Questions could be prepared ahead of time to fill in missing details of workflow, as well as continuing to seek new information. However, the level of detail and the breadth of the study were constant issues. Our study focused on one office containing three people but those people were continually interacting with other offices in their company, leaders and workers in the field, the workers' union, the government, vendors, and many others. With each of these interactions, the question arose how far to carry the mapping.

We originally planned to include both classes of interactions and descriptions of specific instances (e.g., the scheduling of a particular truck to a particular work site, with the acts and gaps that actually occurred). The tool was intended as a way to record both instances and classes in a uniform way, and to link between them (showing what general classes applied to each particular instance, and which particular instances were justified or triggered by other ones.) But in actually carrying out the analysis, we found that the effort involved in creating instances was substantial and did not seem as much of a priority as a entering a broader set of general classes. The dilemma here is that it is often the details revealed by a particular instance that are the key to identifying new design opportunities, but the thorough cataloging of a set of instances (e.g., all those occurring during a period of observation) is so labor-demanding that it is hard to justify).

There is also the issue of multiple interpretations of the work patterns. We knew from the beginning that different people would have different interpretations of the work and that our methodology would have to live with that situation. Within the team that we focused on, we found that there were few differences in perspective between individuals. Because this team had interactions with a wide variety of other agents, though, many of the conversations incorporated the interests of these other agents. As more information was obtained about these other agents' work, the role that the conversations played in their work became easier to map but we never attempted to do more than sketch in the peripheral conversations of the other agents. In general, it was easy to add additional conversations and show their

relationships to the others. The language/action ontology supports and logically connects many different answers to the question of why something is done, accommodating different agents perspectives, even when focusing on one team of agents.

Another issue in applying the framework is that the observations and interviews inherently can affect the work, as the subjects explained what they were doing or had done and why, or simply were aware of the presence of researchers or recording devices. The framework also involves active interpretation since we don't yet have a structure that can be used in an observer-independent way. Thus, there is some subjectivity in the selection of detail and scope.

Interviews or work activities were often interspersed with the history behind the actions or projections about possible future improvements, neither of which fit neatly into the ontology of current work patterns. The ontology fit into the current snapshot of the work, but lacked a sense of the current direction and momentum of change. As the researchers, we personally were able to get a feel for this but there was no method in the ontology to encoding it. Perhaps this is not a lack in the ontology but rather is a part of the change management study that should accompany design. Either way, the question of how to couple attitudes towards conversations and change with the pattern of those conversations is still open.

IV.C. What Did We Learn in Creating the IWM1 Tool?

In creating the IWM1 tool, we needed both a rapid prototyping environment and a hypertext-like flexible linking system. We chose the SuperCard platform, defining different types of cards for each type of object in the ontology. SuperCard allowed us to easily add and change fields and cards, as well as allowing users to link fields with other cards, add comment cards, and search on a variety of criteria. The programming language was cumbersome and the platform limited us in some ways. The language/action perspective in not inherently stack-oriented, while the SuperCard platform works on a stack-based metaphor. In addition, although it was easy to create and follow links between cards, it was difficult to ever get an overall view of those links since the platform did not provide ready tools for graphically mapping links.

In the beginning, close to everything was treated as potentially relevant and after the original interview and observation, about 120 cards were created for various objects, people, gaps, etc. that had been observed or discussed. The accumulation of cards, fields, links, and tools was driven by the data from observations and interviews, thus from the ontology itself. We discovered searching tools to be very important and

developed many ways of accessing information. Because the platform was slow, lists of cards were kept updated for speed of access and easier analysis.

As the research continued and the ontology expanded, the tool was able to expand as well, incorporating new fields and card types as needed. With each new item from observation or interviewing, came the following questions:

> Is this item relevant to the language/action network?
> (Should it be put into the IWM1 tool?)
>
> Where does this data fit in the ontology?
> (Is there a type of card that fits this data?)
>
> What other type of data is this related to?
> (What other fields and links should be used?)
>
> What data can be seen now as missing?
> (What fields and links are currently empty?)
>
> What kind of tools (for searching, etc.) would be useful with this data?

Not only did the tool provide a data bank for the data we collected, but it also provided an opportunity for checking the consistency and completeness of the data and thus provided a proving ground for the ontology itself. The development of the tool was an integral part of the process of refining and expanding the ontology of the language/action theory.

IV.D. What does this imply for design?

When we set off to do this study, we anticipated being able to map a substantial part of the work structure, and to use that map as the basis for specific proposals for new possibilities for using computer-based tools to facilitate the work. We hoped that explicit mapping of individual instances of people handling breakdowns would be useful for revealing the structure of work that goes beyond the "standard procedures."

As the work progressed, the scope of this kind of mapping became more apparent, and it was clear that it was not feasible to do broad mapping of the desired depth. In order to keep the data manageable, it would be necessary to have a much more focused set of specific design proposals (e.g., keep track on-line of equipment repair status) in order to find the right set of relevant areas to map. This is in accord with the

observations that have been made more recently by proponents of traditional system design methodologies, that an overly thorough form of analysis can become a block to getting design done. (See, for example, Yourdon, 1989, Chapter 7 on "Changes in Systems Analysis.")

Since our primary goal in this project was to develop the theoretical framework, we had not entered into a site with specific design needs. There were no clear guiding criteria for sorting out the interesting from the irrelevant details. We chose to continue the work in a broader, less-directed fashion, looking for those places where it led to new insights for us, and where we could make occasional suggestions about possibilities for changes in the work structure.

The next step for making this kind of mapping a practical part of a design process is the integration of the kind of "bottom-up" structure that it implies with a more "top-down" focus on what the critical problems are that need solution and merit design efforts within the organization and its work. The mapping metaphor needs to be interpreted in a narrowed sense: not of making a complete map, but of using mapping techniques on selected details. This kind of analysis has been pursued in a number of studies done by Action Technologies and Business Design Associates, consulting companies which have developed methodologies for work redesign and workflow enabling, based on a language-action perspective. (See Dunham, 1991, Medina-Mora, Winograd, Flores, and Flores, 1992).

The detailed analysis of this study will serve as a basis for extending and enriching the language/action analytic methods to incorporate more of the constraints that have to do with the physical embodiment (both of people and things) that constitutes the workplace.

REFERENCES

Grudin, Jonathan (ed.) (1991), Special Section on Computer-Supported Cooperative Work, *Communications of the ACM* 34:12 (December, 1991).

Dunham, Robert (1991), Business design technology: Software development for customer satisfaction, IEEE Press, *Proceedings of the 24th Annual Hawaii International Conference on Systems Sciences,* 1991, 792-798.

Flores, Fernando, Michael Graves, Bradley Hartfield and Terry Winograd (1988), "Computer systems and the design of organizational interaction," *ACM Transactions on Office Information Systems* 6:2 (April, 1988), pp. 153-172.

Greenberg, Saul, (ed.) (1991), Computer-supported Cooperative Work and Groupware, New York: Academic Press.

Greif, Irene (Ed.). (1988). *Computer-supported cooperative work: A book of readings.* San Mateo, CA: Morgan Kaufmann.

Johansen, Robert (1988). *Groupware: Computer Support for Business Teams.* NY: The Free Press

Kensing, Finn, and Terry Winograd (1991), "The language-action approach to design of computer support for cooperative work," *Proceedings of the IFIP TC8 Conference on Collaborative Work, Social Communications and Information Systems,* Helsinki, Finland, 27-29 August, 1991.

Kukla, Charles, Anne Clements, Robert Morse, and Debra Cash (1992), Designing Effective Systems: A Tool Approach, in P. Adler and T. Winograd (eds.), *Usability: Turning Technology into Tools,* New York: Oxford University Press, 41-65.

Medina-Mora, Raúl, Terry Winograd, Rodrigo Flores, and Fernando Flores (1992), Workflow management technology: examples, implementations and new directions, *Proceedings of the 4th Conference on Computer-Supported Cooperative Work,* Toronto, November, 1992.

Schrage, Michael (1990). *Shared Minds: The New Technologies of Collaboration,* New York: Random House.

Winograd, Terry (1987), "A language/action perspective on the design of cooperative work," *Human Computer Interaction,* 1987/88:3, 3-30.

Winograd, Terry, and Fernando Flores, F. (1986). *Understanding computers and cognition—A new foundation for design*. Norwood, NJ: Ablex.

Winograd, Terry, Nancy Newman and Peter Yim (1991), "Including people in computer integrated manufacturing designs," *Proceedings of the International Conference on Computer Integrated Manufacturing*, Singapore, October 2-4, 1991.

Yourdon, Edward (1989), *Modern Structured Analysis*, Englewood Cliffs NJ: Prentice-Hall/Yourdon Press, 1989.

APPENDIX A: ANALYTIC CATEGORIES AND QUESTIONS USED IN WORK MAPPING

Group 1: People, Locations, & Things—Activity Domain Basics

PEOPLE

The following questions were asked for each person:

Location: What are the usual locations of the person while working? (Answer is of category location.) Common answers included various office locations but also places outside the office that were communicated with by phone or radio, as well as dispatcher's home since they sometimes called workers from home.

LOCATIONS

The following questions were asked for each location:

Constraints: What are the constraints on using the location? (The answer to this question is so far free-form and may be in the categories of agents, locations, things, or may involve time or, conceivably, may be related to the activities that can or can't take place there. Theoretical note: constraints concerning activities have not played a role in the current project but in later projects a dependency link from activities to locations might be called for; in some sense, a constraint may be a declaration by someone, which would also change the dependency graph.) One example of a constraint on a location in the current project was that the rental equipment area of the dispatch board was reserved solely for magnets representing rental equipment.

Access By: What agents can access the location, to view it or to touch objects within it? (Answer is of category agent plus a note about what type of access is allowed.) This is basically the inverse of the question "Controls location" under category agent. An example from the current project is that the dispatch office itself was allowed to be accessed for viewing by a variety of agents but only three could change anything without first obtaining permission.

Contains: What things does the location commonly contain? (Answer is of category thing.) Here is the question that focuses on the contents of locations. For example, one of the agents studied had on his desk purchase order forms, a speaker phone, three marker pens (red, black, and orange), a variety of standard office equipment like a

memo pad and ball-point pens, and a "Move Equipment" work sheet.

THINGS

The following questions were asked for each thing:

Description: What is an informal description of the thing? (Answer is free-form text.) The company equipment magnets, for example, are white and are labeled in English and with a computer bar code identifying each one individually.

Locations: Where is the thing commonly located? (Answer is of category "location".) This is basically the inverse of the question "what things does the location commonly contain?" in the category "location". Examples would include the locations of dispatch desk or dispatch board for the red and black markers and only the dispatch desk for the orange marker, which was only used for forms.

Manipulated: Which agents are allowed to manipulate the thing? (Answer is of category "agent".) This is basically the inverse of the question "What things does the agent control?" in category "agent". Of interest here was who could change the state of the thing—moving it, writing on it, etc. For example, only three people were allowed to transfer the contents of the dispatch board onto the computer with a light pen.

Observable: Which agents can observe the states of the thing? (Answer is of category "agent".) This is related to the previous question, except that who can see or hear or otherwise observe the state of an object is at issue instead of being able to actually change the state of the thing. In the above example, although three people were allowed to transfer the contents of the dispatch board onto the computer with a light pen, anyone else could watch.

Symbolic Map: How do the various states of the thing map onto meaning? (Answer is a free-form text description or left blank if there is no symbolic map for the thing. Theoretical note: In some ways, a symbolic map is or indicates the existence of a prior declaration.) For example, usually the symbolic map of a telephone would be blank but the placement of a green magnet on top of an employee magnet on the dispatch board meant that the employee needed to be notified.

General Use: What is the general, usual use of the object? (Answer is a free-form text description. Theoretical note: this too often seems to

imply a previous declaration of what the object should be used for.) For example, the telephone in the dispatch office is usually used to receive calls from job foremen and employees, and to call employees and equipment rental, equipment suppliers, and truck companies.

Group 2: Agents, Roles and Gaps—Language/Action Domain Basics

AGENTS

The following questions were asked for each agent:

Person: What people act as this agent (Answer is of category person). In the case of individuals the person and agent are equated. For groups or organizations, specific people can act as the agent.

Controls Location: What areas does the agent actually control? (Answer is of category location.) Since agents often have access to more places than they can actually control, this question focused attention on such things as area supervisors being able to come into the dispatch office but not to make any physical changes themselves, such as marking changes in the employee assignments on the dispatch board.

Controls Things: What things does the agent control? (Answer is of category thing.) Agents were counted as being able to control a thing if they had the right to affect its location or state. For example, only three people were allowed to write up purchase orders, no one else was supposed to touch the purchase order pads.

ROLES

The following questions were asked for each role:

Description: What is an informal description of the role that the agent(s) perform(s)? (Answer is free-form text.) This may be anything from a job title to a catch-all like "outside supplier" to a casual description such as "the person who handles the union leaders when they are upset."

Current Agents: Which agents currently perform this role? (Answer is of category agent.) For example, in the earlier example that only three people are allowed to change things on the dispatch board, this entry would contain their names.

Constraints: What constraints are there on agents playing this role? (Answer is free-form text.) Roles may be constrained to specific

people. Constraints might also consist of certain qualifications, such as the amount of experience in the field, or the simultaneous filling of another role, such as only agents who play the role of dispatcher can play the role of equipment renter.

GAP

The following questions were asked for each gap:

Gap Declarer: What role is being played by the agent who declares the existence of the gap? (Answer is of category role. Practical note: sometimes it is easier to specify the agent at first.) Gaps are declared by someone with the authority to do so; until this declaration, the gap does not exist, though there may be related gaps prior to this declaration. For example, the gap that a particular type of truck will be needed at the job site can only be declared by the foreman or supervisor for the job. Other people can point out the possibility or probability but their statements do not carry the weight of declaring that the situation exists.

Gap Handler: What role is being played by the agent who handles the gap? (Answer is of category role. Theoretical note: sometimes it is easier to specify the agent at first.) For example, the opportunity to assign an apprentice equipment operator on a job can be declared by a foreman and the dispatcher can then handle the gap by assigning an appropriate employee. Sometimes the same role involves both declaring and handling gaps, such as when dispatchers both declare the need for rental equipment and arrange its hire.

Completion Declarer: What role is being played by the person who says that a gap is completed, no longer of concern? (Answer is of category role. Practical note: sometimes it is easier to specify the agent first.) The person declaring a gap to be completed is saying that no further (if any) changes in how the gap is handled are to be taken. This often means that problems have been solved or opportunities taken advantage of, but can recognize a variety of final solutions, including simply declaring that the situation be ignored. For example, the gap that a truck is needed at a job site is declared complete by the foreman when the truck shows up as promised. This declaration of completion may be explicit, such as a confirmatory phone call, or implicit, such as not calling to complain.

Possible Causal Antecedents: What are the possible causal antecedents for this gap to occur? (Answer is of free-form text. Theoretical note: these often correspond to items of category act or activity.) The possible causal antecedents are those events that might encourage

the occurrence of this gap. For example, the increased experience and performance of an equipment operator would qualify that employee for more responsible assignments.

Possible Causal Anti-antecedents: What are the possible causal anti-antecedents for this gap to occur? (Answer is of free-form text. Theoretical note: these often correspond to items of category act or activity.) The possible causal anti-antecedents are those events that might discourage the occurrence of this gap. For example, the giving of precise directions and maps to the employees by the dispatchers discourages the possible gap of the employees getting lost trying to find the work site.

Grounding: What are the indications that cause the declarer to state that the gap exists? (Answer is free-form text. Theoretical note: this can be seen as a declaration that a certain situation exists, such as a particular state of things, people, and activities in the world.) For each declaration of a gap, there is a corresponding grounding, a set of indications that cause the declarer to state that the gap exists. The grounding may be as simple as the observation that a green dot is on the employee's magnet on the dispatch board indicating that the employee needs to be called, or may be a complex web of interactions.

Completion Declaration: What is the declaration that causes this gap to be completed? (Answer is of type act.) The completion declaration is that declaration made that causes the gap to no longer be handled in any way. This does not necessarily mean that a problem has been resolved or an opportunity has been taken advantage of, nor even that anything has been done about the gap at all other than it has been both declared and discussion about it closed. This declaration may take the explicit form of an acknowledgment, such as an acknowledgment that a payment was received, or simply an implicit acknowledgment in that no protest is raised, such as the simple lack of additional bills generated for that same balance.

Completion Can Cause: What further gaps can the completion of this gap cause? (Answer is of type gap.) The completion of a gap can trigger other acts or gaps; which ones result will depend on the reports and declarations of the situation and the method of completing the gap (whether a problem or opportunity was solved or taken advantage of completely, partially, or not at all.) For example, the completion of the gap of a bulldozer being needed at a job site might bring into a variety of acts or gaps depending on whether the bulldozer was already at the site (no special effort concerning the bulldozer is required), whether the bulldozer was at

another site (in which case the bulldozer needs to be moved—a new gap), or the bulldozer must be rented from an outside vendor (in which case the necessary arrangements need to be made with the vendor—a new gap.)

Typical Acts: What acts are typically involved in handling this gap? (Answer is of type act.) These acts may be simply the declaration that the gap exists, such as the Vice President in charge of Operations telling the dispatchers that they missed an opportunity to use the company's own equipment instead of renting—the mistake is now irreparable but the declaration is intended to discourage the repetition of the mistake. There may be alternate acts that are taken, such when a bulldozer is needed this may be a list of the acts involved in assigning a company bulldozer and the acts involved in the alternate, arranging a rental bulldozer.

Group 3: Activities and Acts—Doing the work

ACTIVITY

The following questions were asked for each activity:

Meaning: If the activity is communicatory, what meanings are expressed in the activity? (Answer is free-form text.) This was at the level of the illocutionary content of the message. For example, when a dispatcher talked to a employee, the fact that directions to the employee's new job site were given was recorded rather than the actual words used.

Purpose: If the activity is not communicatory, what was the reason that the activity was performed? (Answer is free-form text.) The determination of reasons was based both on interviews and observation, thus involving interpretation by the observer. An example of the purpose of an activity, such as rolling the computer table to the other side of the room, would be so that the dispatch board's bar coded schedule can be input into the computer.

Trigger: What is the activity that triggers the performing of this activity? (Answer is an activity.) For example, the printing out of the daily schedule triggers the copying of the schedule which in turn triggers the distribution of the schedules.

Preparatory: What is needed before doing this activity? (Answer is an activity or thing.) This involves what is required to be done or available before the activity is possible. For example, copying the daily schedules is not possible until a computer printout of the schedules is available. Similarly, printing out the schedules is not

possible until the activity of reading in the scheduling board's bar codes (or equivalent hand-entry activity) has been finished.

Competence: What skills or general knowledge does performing the activity require? (Answer is free-form text.) This involves a variety of abilities or types of knowledge, such as being able to type or knowing where the best prices for traffic cones can be found.

Actor: What agent performs the activity? (Answer is of type agent.) For example, entering in data from the scheduling board into the computer is usually performed by two of the three dispatchers, though it is occasionally performed by the third as well.

Attendees: What agents can observe the activity? (Answer is of type agent.) This does not include the observation of the results of the activity, unless those results are immediate; for example, if reading in the bar codes off the dispatch board makes printing the daily schedule possible and an agent observes the printed schedule, that fact is not recorded here but rather in the attendees section of the activity of printing the schedule.

Things: What things are directly involved in the activity? (Answer is of type thing.) For example, phones are involved in phone calls, while the moving schedule and orange pen are involved in scheduling equipment moves.

Locations: What locations are involved in performing the activity? (Answer is of type location.) For example, speaking on the Gradeway radio requires standing in a particular part of the room, while speaking on the radio requires standing in a different part of the room.

Time: When are the customary and possible times for this activity to be performed? (Answers are times of the day, week, month, etc.) For example, planning meetings for the following week are held on Thursday mornings, while foremen call in their next day's requirements between 10 am and 2 PM each day.

ACTS

For each type of act, similar questions were asked, as shown below:

Actor: What is the role of the agent doing the act? (Answer is of type role.) For example, foremen request work crews, while dispatchers offer or promise the crews and then report on the completed crew and declare the new schedule.

Recipient: What is the identity of the agent to whom the act is addressed? (Answer is of type role.) For example, dispatchers receive the work crew requests, and foremen and field engineers receive the declaration of the new daily schedule.

Attendees: What roles are played by the persons who are witnesses to the act? (Answer is of type role.) For example, when a dispatcher and a foreman discuss a particular employee's qualifications over the radio, the other foremen, other dispatchers, and the equipment movers overhear the declarations being made.

Authority: Who has the authority to make the act and where does this authority come from? (Answer is free-form text related to roles.) Requests, offers, promises, reports, and declarations tend to be effective only if backed by the proper authority. This tendency is especially pronounced with declarations. There are many examples in which a group of people all are aware of a situation and talking about it but lack the authority to declare the situation as existing and thus trigger the appropriate acts and activities. For example, until the dispatcher assigns a piece of equipment or an employee to a job, it is not enough that the foreman has said that the employee should work the next day on that job. Similarly, only a dispatcher has the authority to purchase traffic cones and other supplies or request new state permits for moving equipment.

Produced: What activities are involved in producing this act? (Answer is of type activity.) Declaring the work crews is done by printing and distributing daily schedule reports; another example of producing an act may simply be calling a supplier on the telephone and asking for traffic cones to be delivered.

Interpreted: What activities are involved in interpreting this act? (Answer is of type activity.) For example, the interpretive counterpart of a spoken act is listening or turning on a device and listening; written acts involve reading to be interpreted, and so forth. The main point here has been to recognize that producing and interpreting an act are different and may not always mesh well without careful planning.

Group 4: Sessions and Routines—Collections of Work

SESSIONS

The following questions were asked for each session:

Actor: What persons are involved in the session? (Answer is of type person.) For example, determining crew, equipment, and truck

needs for the next day is usually done by one of the dispatchers with the job's foreman, superintendent, or area manager.

Attendees: What persons can observe the session? (Answer is of type person.) This does not include the observation of the results of the session, unless those results are immediate; for example, although all foremen can see the result of assigning crews, equipment, and trucks to each job when they receive the daily schedule report, they are only participants of the session itself if they are listening to the radio while the dispatcher and job representative discuss the needs on the radio, etc.

Things: What things are directly involved in the session? (Answer is of type thing.) For example, phones are involved in phone conversations, while the black and red marker pens are involved in scheduling each job. *Locations:* What locations are involved in performing the session? (Answer is of type location.) For example, radio conversation sessions require standing in a particular parts of the room, while inputting the dispatch board into the computer involves standing in a different part of the room.

Time: When are the customary and possible times for this session to be performed? (Answers are times of the day, week, month, etc.) For example, once a week on Thursday mornings, planning for all jobs for the next week is done. Similarly, on a daily basis, the next day's crew, equipment and truck needs for the next day are determined between 10 am and 2 PM.

Activities: What activities are being performed within the session? (Answer is of type activities.) These activities may be part of the same goal or of several different goals, as long as they are performed together in some sense. For example, questioning about various types of crew personnel needed for the next day are activities that set up the crew schedules, questioning at the same time about equipment and truck needs sets up those other different schedules, while adding into the conversation discussions of any new employee's performance adds a different type of activity entirely.

Continuity: What is the major source of continuity for the session? (Answer is free-form text.) The answers to this question may be as simple as "all activities between 9 am and 10 am in the dispatch office" or "all activities during that telephone call" or "all activities involving the computer's light pen", etc.

ROUTINES

The following questions were asked for each routine:

Actors: What are the roles being played by the agents involved in the routine? (Answer is of type role.) For example, requesting and promising crew, equipment, and truck assignments for the next day is done by a dispatcher and a foreman, superintendent, or area manager.

Attendees: What are the roles being played by the agents who can observe the routine while it is happening? (Answer is of type role.) For example, during the routine of a dispatcher requesting employees to work at various job sites the next day and telling each employee the exact location of the job, the attendees would be the other dispatchers and any other people such as foremen, job superintendents, and area managers that walk into the office.

Acts: What acts are being performed within the session? (Answer is of type activity.) These acts may be part of the same conversation or of several different conversations, as long as they can be perceived as being performed together. For example, the acts of requesting crew, requesting equipment, requesting trucks, and reporting on the performance of a new employee may all be part of the same telephone call.

Continuity: What is the major source of continuity for the procedure? (Answer is free-form text.) The answers to this question will often parallel the answers within the sessions of the activity domain, when a session is an example of a typical routine. It may also have nothing to do with time-space continuity and instead focus on something else, such as a particular object, such as the dispatch scheduling board.

Trigger: What is the act or time that causes an instance of this routine to be started? (Answer is of type act or a time.) This may be a daily event, such as the routine of crew/equipment/truck requests being made between 10 am and 2 PM, or a weekly event such as Thursday morning weekly planning meeting. It may also be an act that triggers it, such as a request by a foreman for supplies to be bought or a declaration that an employee did not show up to work.

Incompletion: What are the gaps that may occur during the routine? (Answer is of type gap.) This includes any problems or opportunities that may occur, in fact, anything identifiable as a gap that can interrupt the flow of the routine. For example, the routine of entering information off the dispatch scheduling board into the computer can be interrupted by phone calls from the field or the water truck company representative coming in to look at the board.

Justification: What are the set of gaps that are identified as potential (anticipated) or actual (to be resolved) that justify the invocation of the routine? (Answer is of type gap.) An example of a potential gap would be that the work at the job sites might not get done as planned if the employees don't know when and where to work, thus the routine of calling the employees and letting them know where to work the next day is justified. An example of an actual gap would be equipment that is broken or needs maintenance, thus the routine of letting the shop know of problems and the location of all equipment is justified.

Group 5: Conversations—Apex of the Language/Action Domain

CONVERSATIONS

The following questions were asked for each conversation:

Mood: What is the mood of the conversation? (Answer is one of the following: action, possibilities, clarification, or orientation.) Conversations for action involve an push for something to be accomplished, such as getting employees to the work site, equipment delivered, and other goals. When the goals are less well defined and instead there is a feeling of exploration of goals, these types of conversations are called conversations for possibilities. Conversations for clarification involve determining the conditions of completion of another conversation or the nature of a request, offer, report, etc. These range from a simple "what?" requesting repetition to "how does she spell her name?" to asking if any additional crew members are needed, such as another grade/checker. Conversations for orientation involve all the conversation and other activities that do not directly relate to getting some specific goal accomplished, exploring possible goals, or clarifying some part of those processes, instead allowing people to build a shared understanding, to integrate through speech acts; examples of this include discussions of an employee's skills, arguments about the relative service provided by various companies, or even gossip about the power structures within an organization.

Type: What is the type of conversation? (Answer is needed only when the mood is "action"; then the answer must be one of the following: offer/completion, request/completion, declaring, or informing.) Conversations for action can take several different types. Conversations of completing are conversations in which there is a network of requests, offers, promises, and reports. Within conversations of completion type, offer/completion conversations begin with an offer and request/completion conversations begin

with a request. Conversations of declaring and informing consist of people making statements that either establish a new state, such as a dispatcher declaring an union member to be a company- sponsored applicant for apprenticeship as an equipment operator, or assert something about the world, such as a foreman informing a dispatcher that a piece of equipment is broken.

Initiator: Who starts an instance of the conversation? (Answer is of type role.) For example, in a construction company dispatch office, if a foreman requests over the radio that a particular type of equipment to be delivered the next day, then the initiator would be the foreman.

Participants: Who also actively participates in the conversation? (Answer is of type role.) In the previous example of a foreman requesting over the radio that a particular type of equipment to be delivered the next day, then the other participant would be the dispatcher who handles the call.

Simultaneous Observers ("Synchronous Listeners"): Who also observes the conversation while it is happening? (Answer is of type role.) In the previous example of a foreman requesting over the radio that a particular type of equipment to be delivered the next day, then the simultaneous observers would be all people listening to the radio, which includes the foremen, job superintendents, area managers, field secretary, and Vice President in charge of Operations.

Non-simultaneous Observers ("Asynchronous Listeners"): Who also observes the direct results of the conversation? (Answer is of type role.) In the previous example of a foreman requesting over the radio that a particular type of equipment to be delivered the next day, then the simultaneous observers would be all people listening to the radio, which includes the foremen, job superintendents, area managers, field secretary, and the Vice President in charge of Operations, as well as the accountants and shop manager.

Trigger: What act or time causes an instance of the work flow to begin or reoccur? (Answers are either acts or times.) Like triggers for routines, these triggers may be specific acts such as receiving a phone call from a particular agent, particular times of the day such as early afternoon, and so forth.

Justification: What are the set of gaps that are identified as potential (anticipated) or actual (to be resolved) that justify the invocation of the routine? (Answer is of type gap.) In the example of the foreman requesting equipment, one justification for the conversation is that

the work can progress faster if the equipment arrives. This might allow the company to avoid being late or even to free up equipment earlier than planned for a later project, thus cutting down the possible rental cost.

Incompletion: What are the gaps that may occur during the routine? (Answer is of type gap.) This includes any problems or opportunities that may occur, in fact, anything identifiable as a gap that can interrupt the flow of the routine. Incompletions include any events (including both problems and opportunities) that might occur within the work flow that would require some sort of resolution, even if only a decision to avoid making any other decision. In the prior example of a foreman requesting a piece of equipment, an incompletion may occur if the particular piece of equipment is already being used elsewhere. At that point, a decision would need to be made whether to move the equipment, rent equipment, refuse to send the equipment, etc.

Conditions of Satisfaction: What are the conditions of satisfaction—what will it take for the conversation to be declared successfully completed? (Answer is free-form text.) The conditions for satisfaction are the events that must occur for the work flow to be considered completed by all parties. In conversations of informing or declaring, the informing or declaring itself may be sufficient for the conversation to conclude. In conversations of completing, there is an implicit or explicit report that must be made and accepted for the conversation to be completed. In the case of requesting equipment to be sent to a construction job site, the report and acceptance may be implicit in the satisfactory arrival of the equipment, or explicit in the verbal acceptance of an employee of a new job assignment.

Acts: What are the acts that take place as part of the conversation? (Answer is of type acts.) This list may contain only a single declaration for a conversation of informing or declaring but will typically contain multiple acts in conversations of completing. These acts can include offers, requests, promises, reports, and declarations.

APPENDIX B: SAMPLE SCREENS FROM THE IWM1 MAPPING TOOL

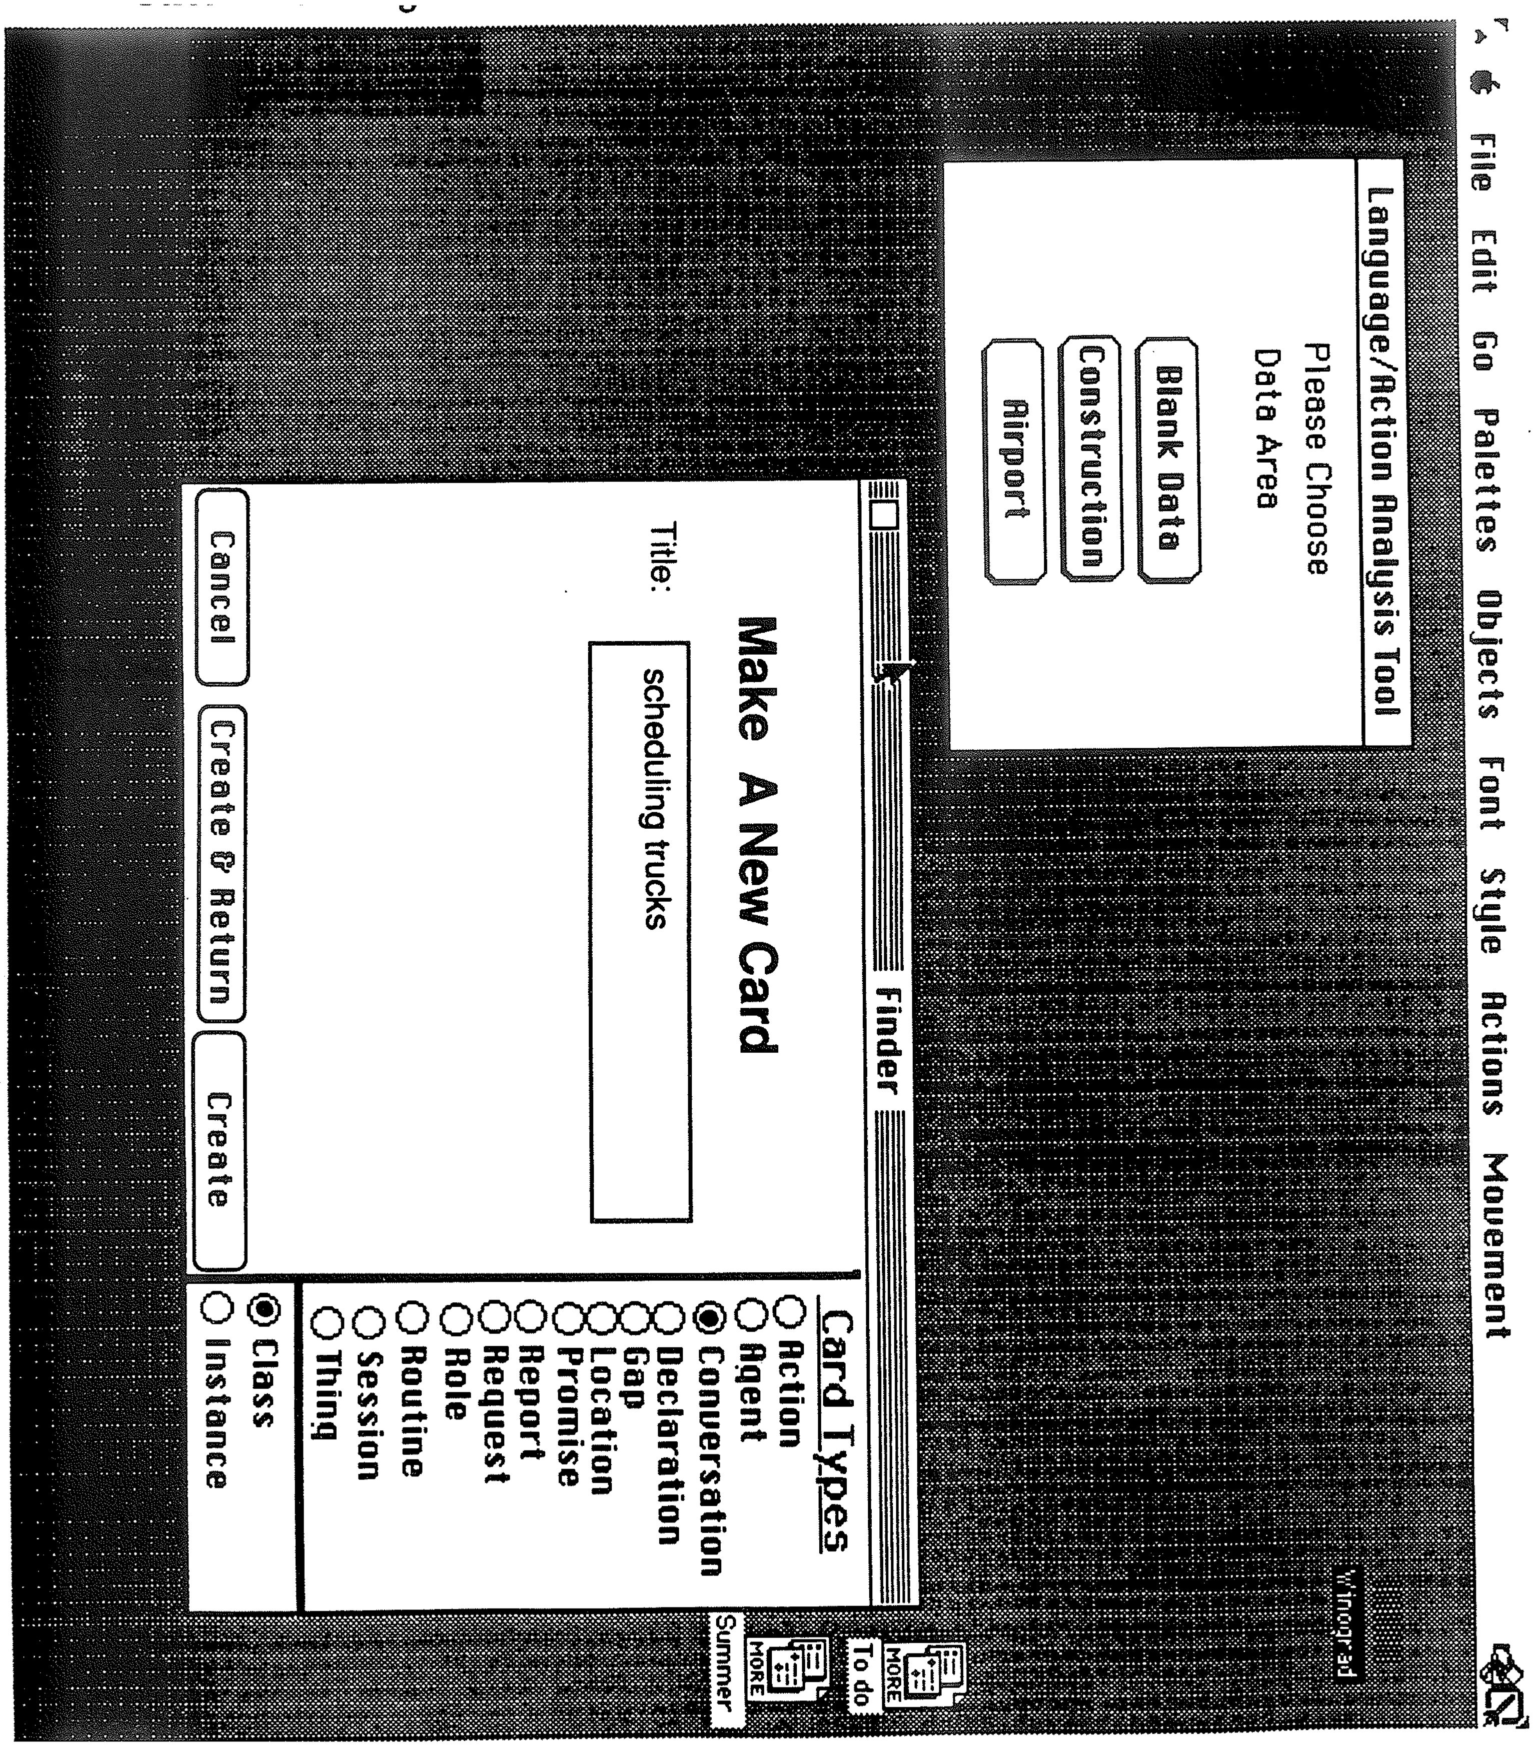

Screen 2 : Newly created conversation card with prompting questions

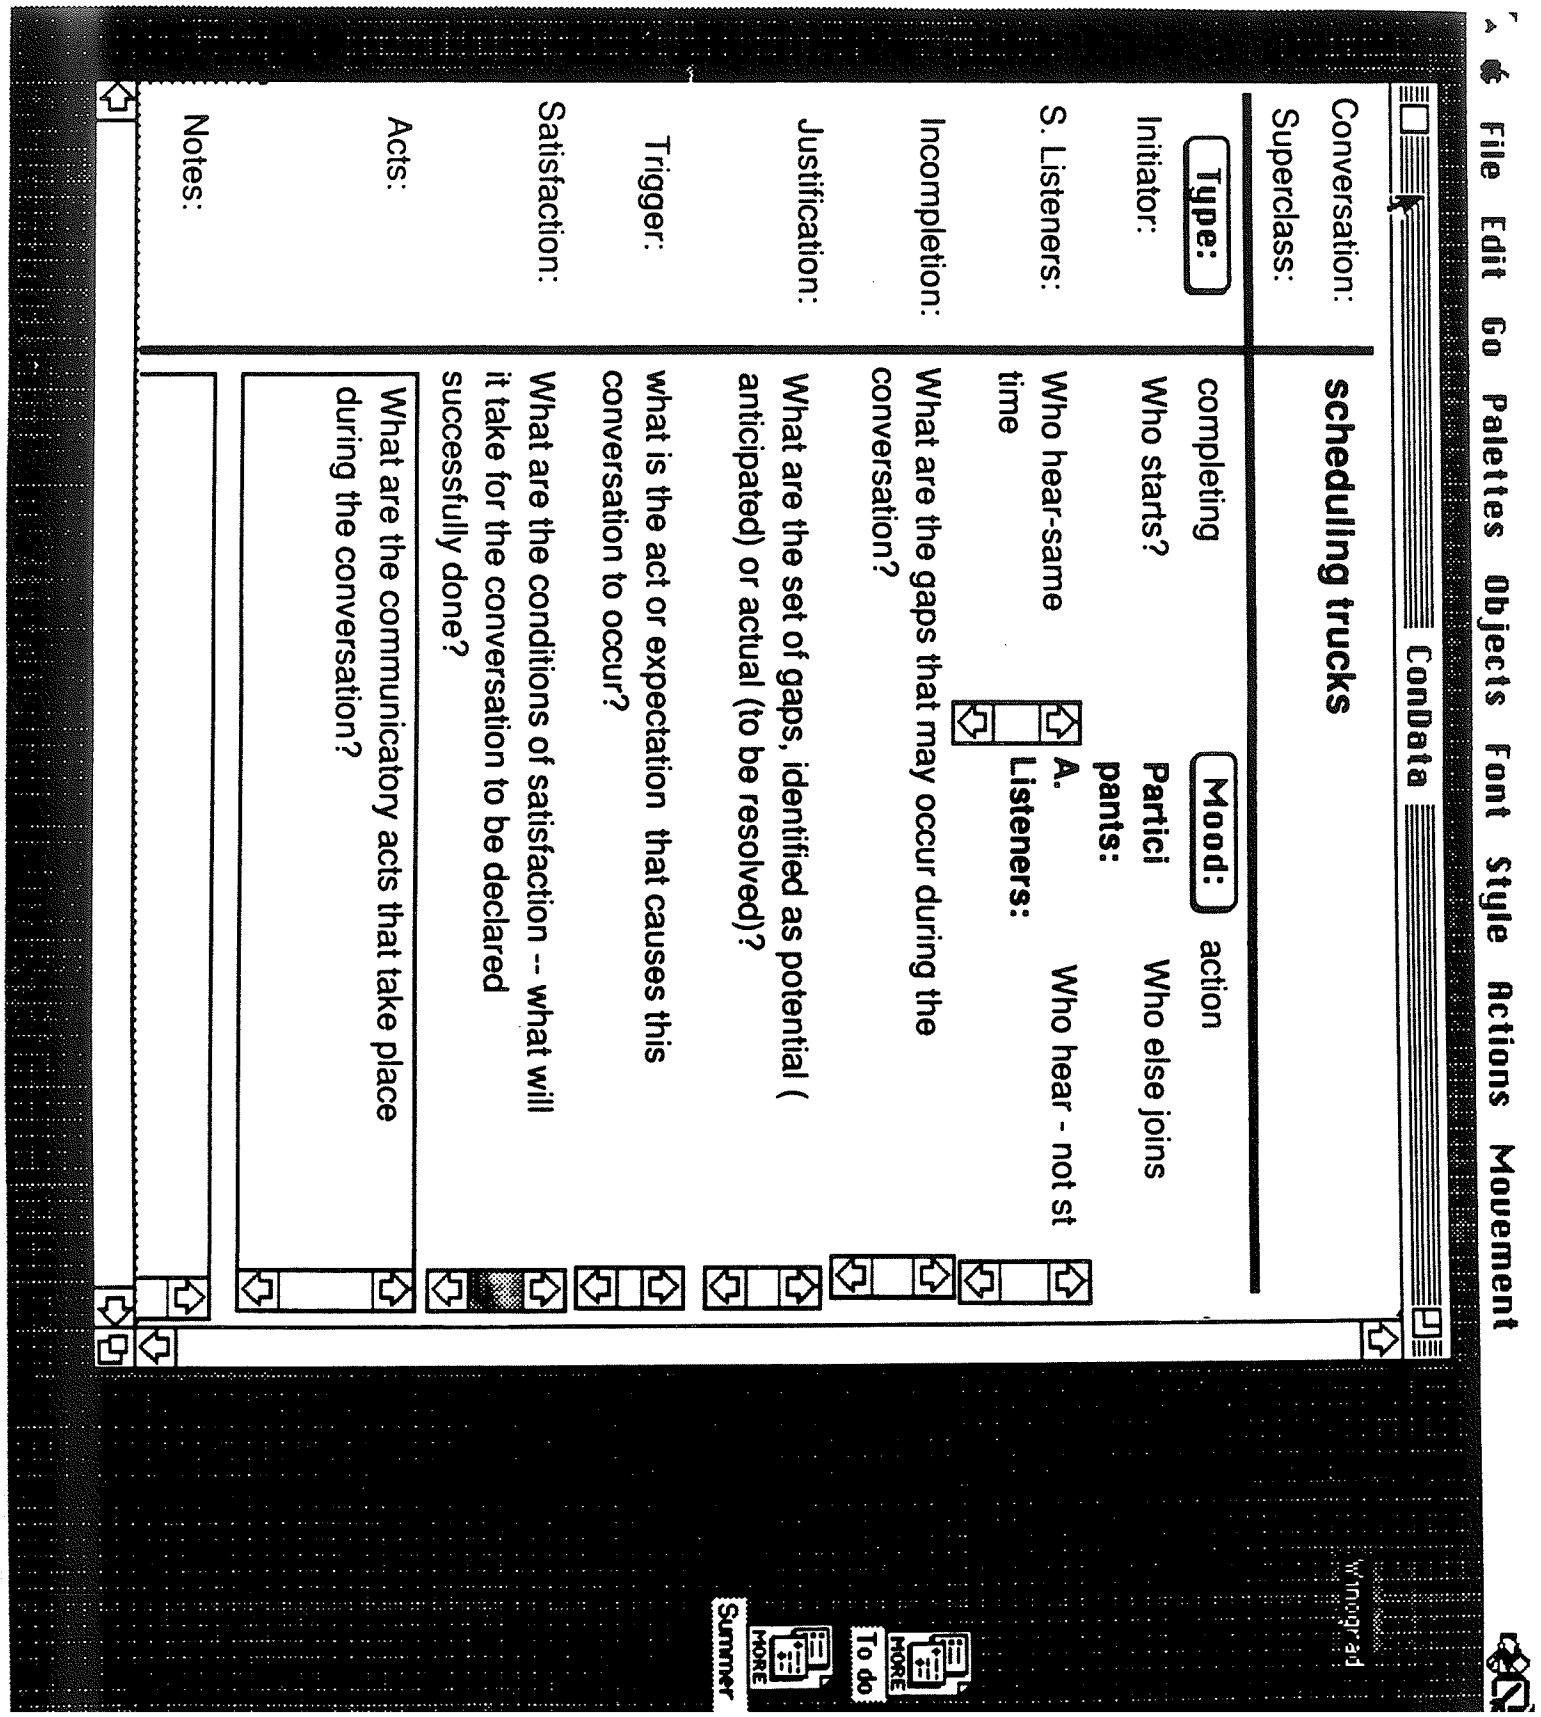

Screen 3 : Information filled in to the conversation card "Scheduling trucks"

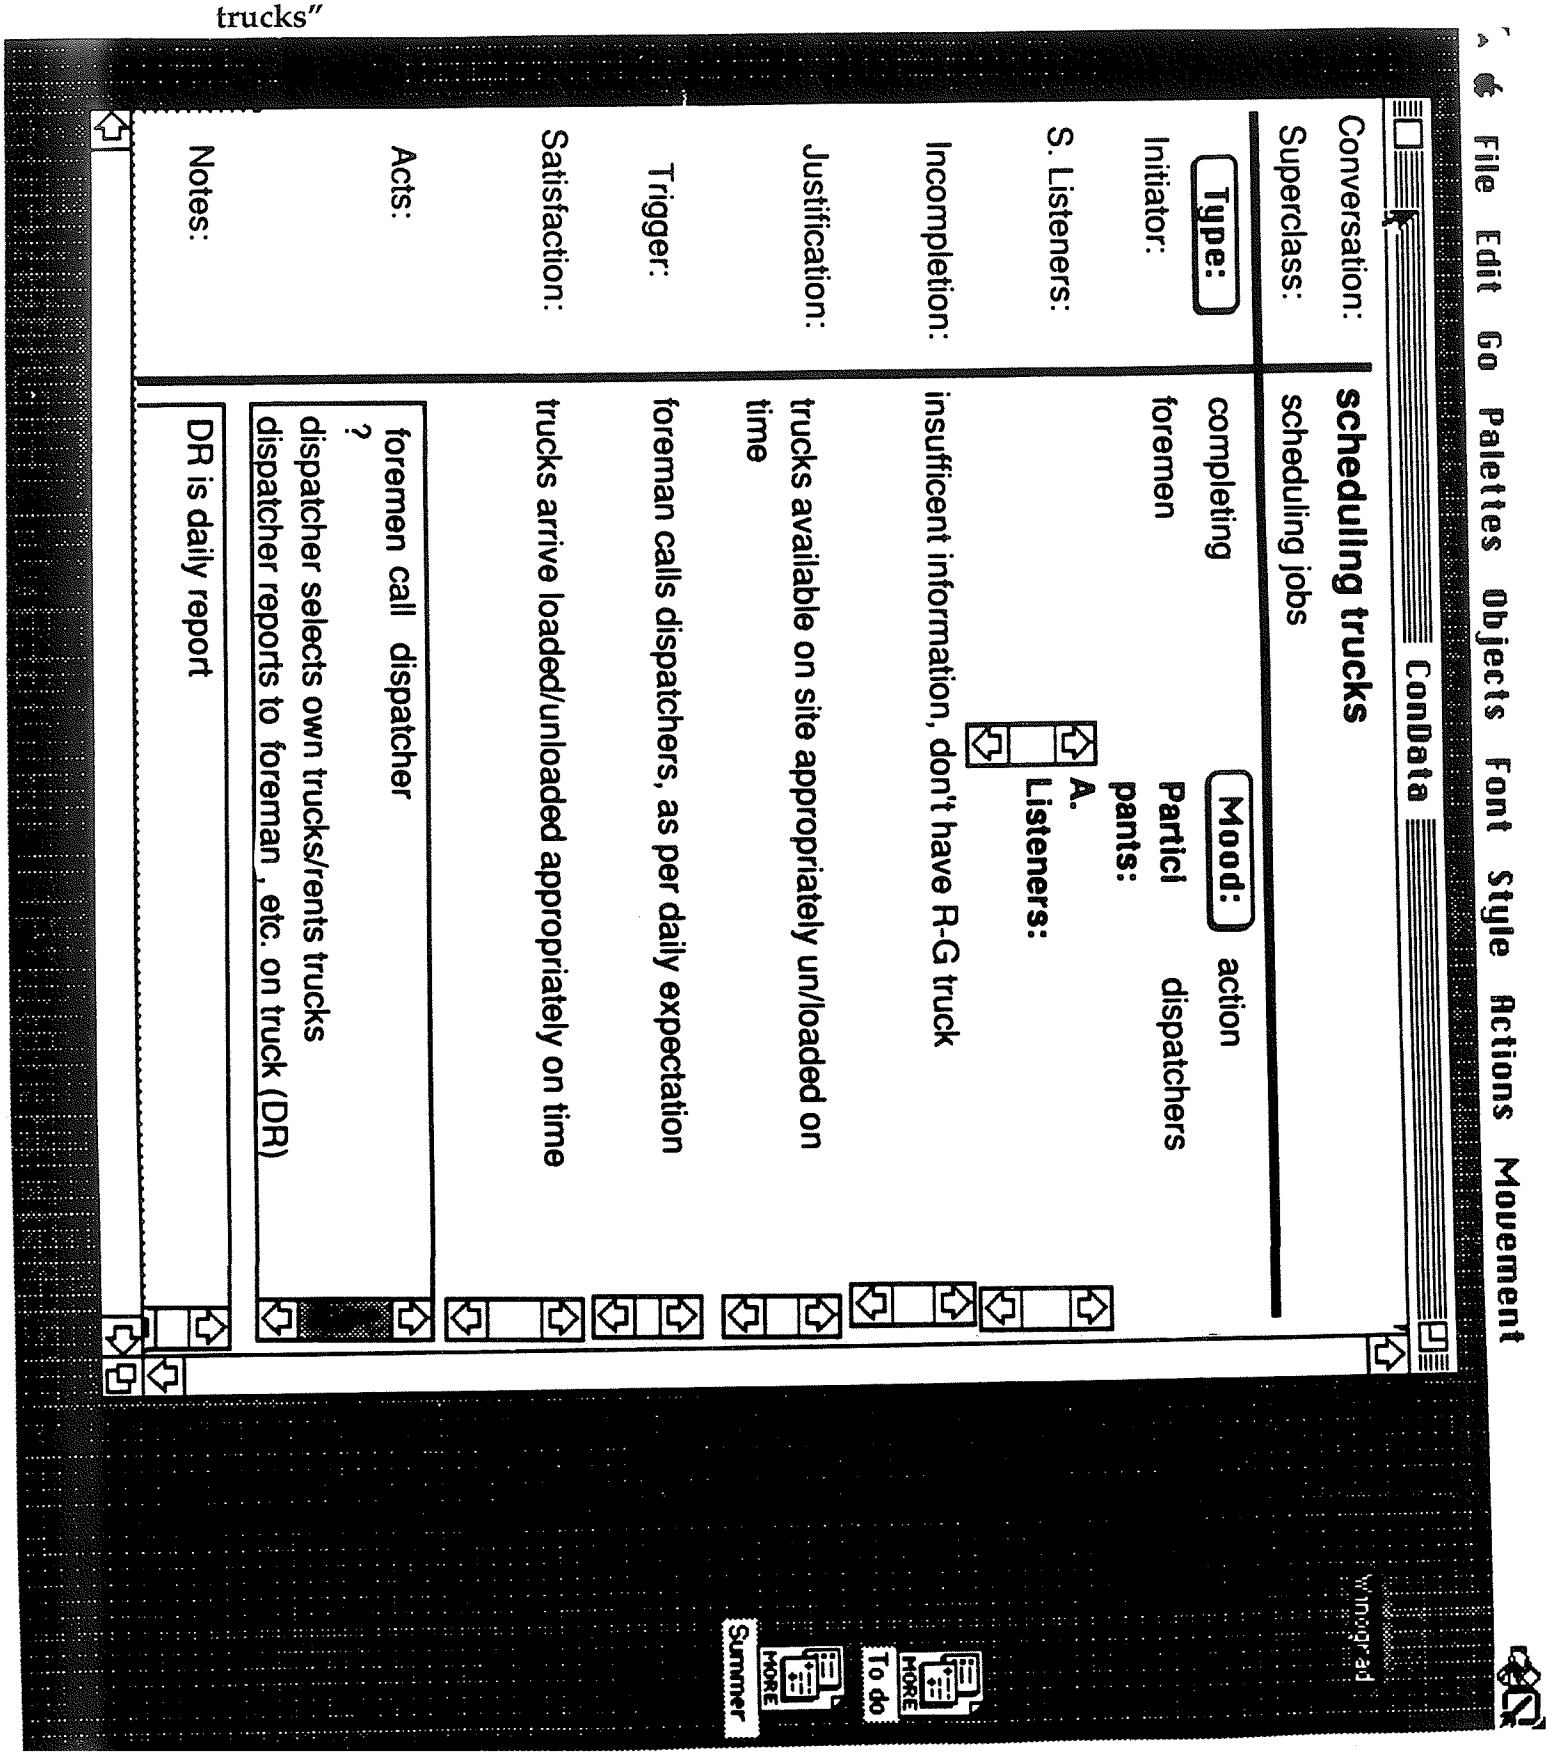

Screen 4 : Making a list of cards satisfying a specified criterion

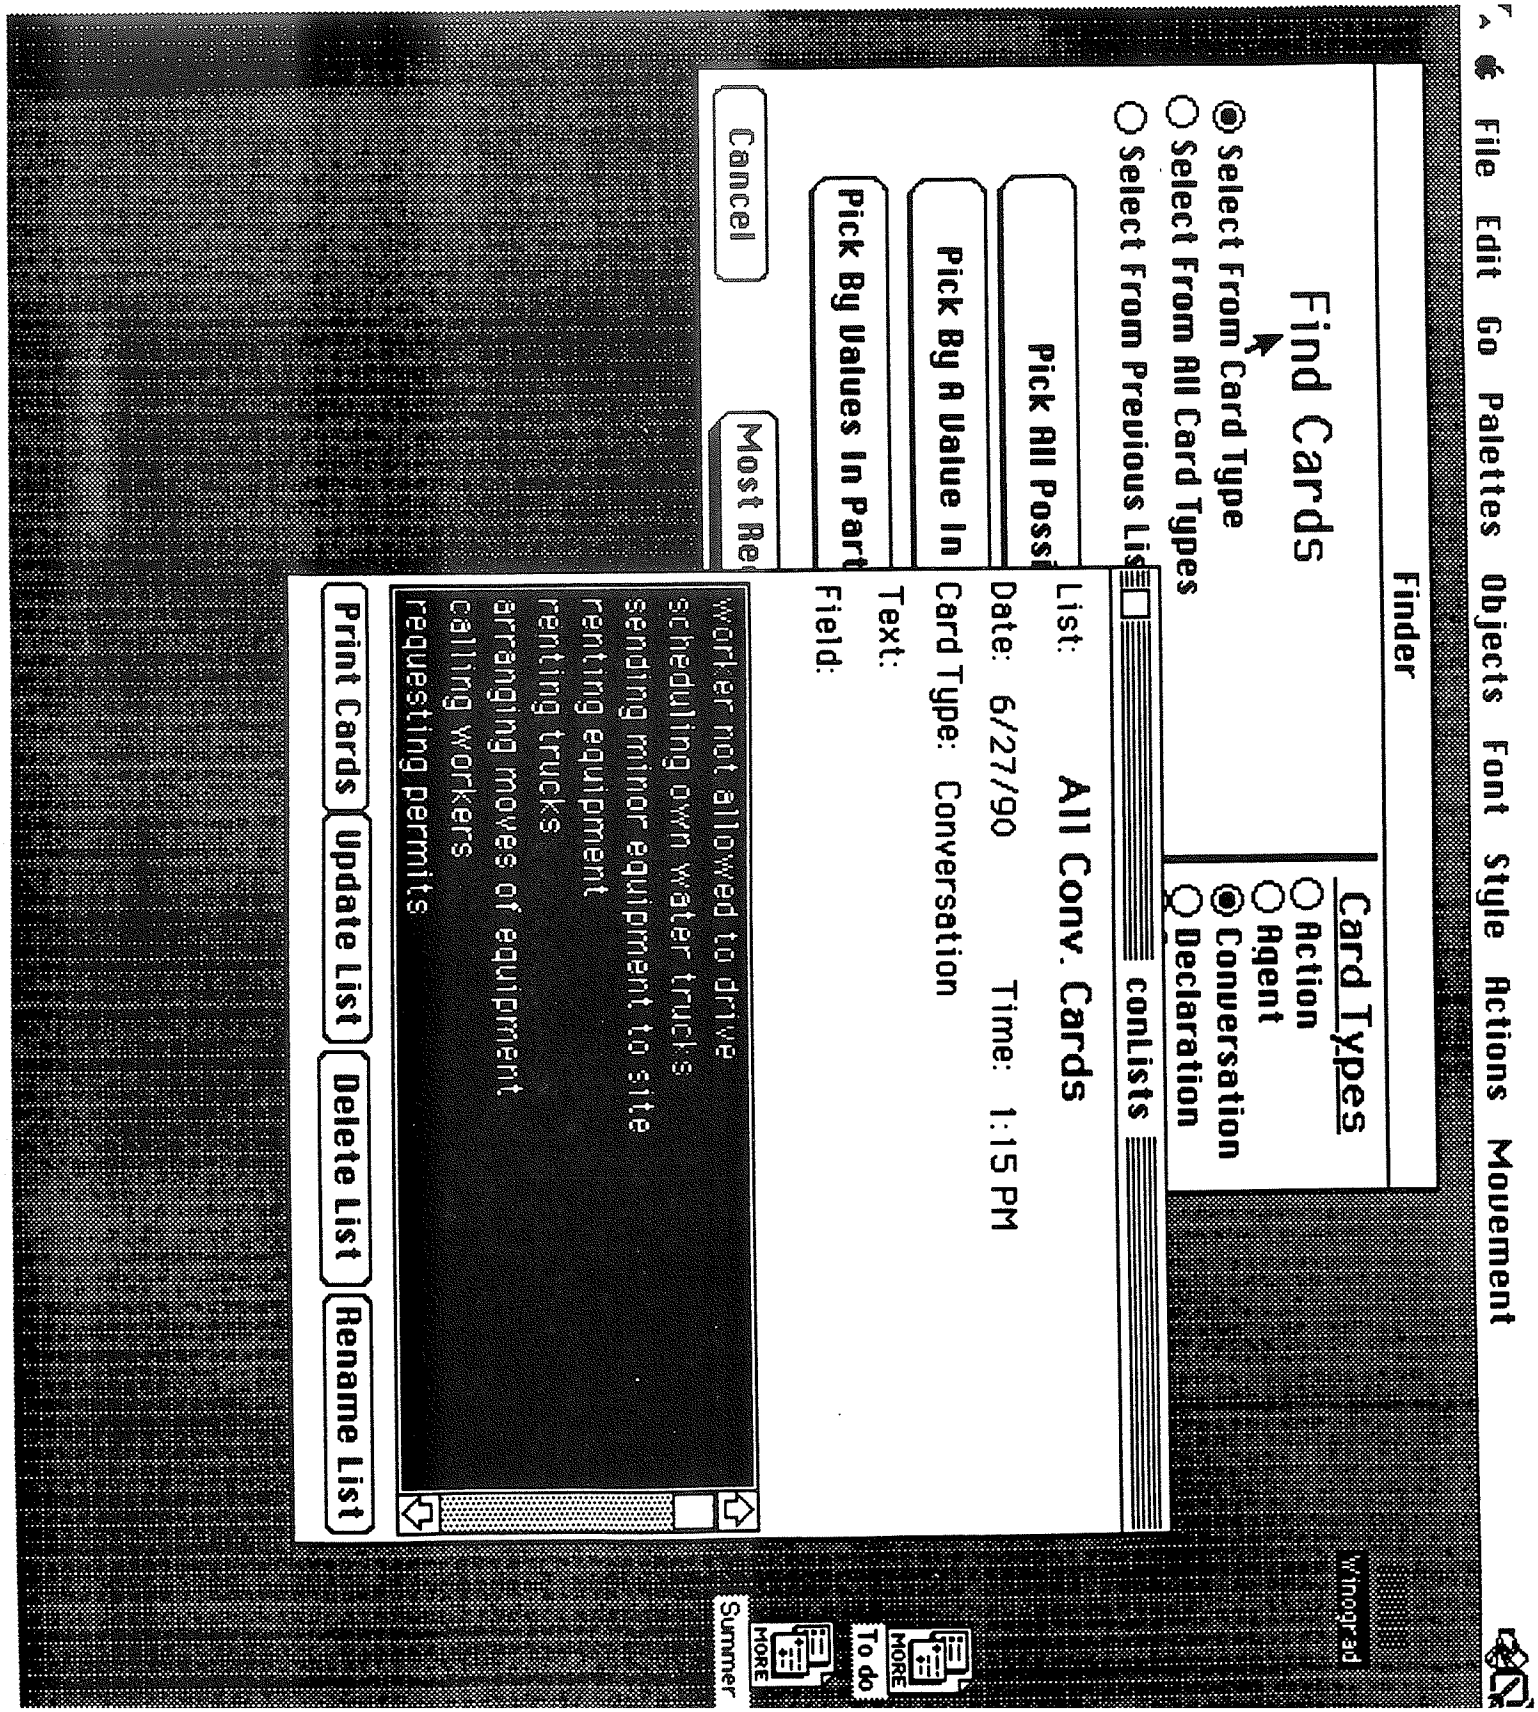

Screen 5 : Creating a link to a related card

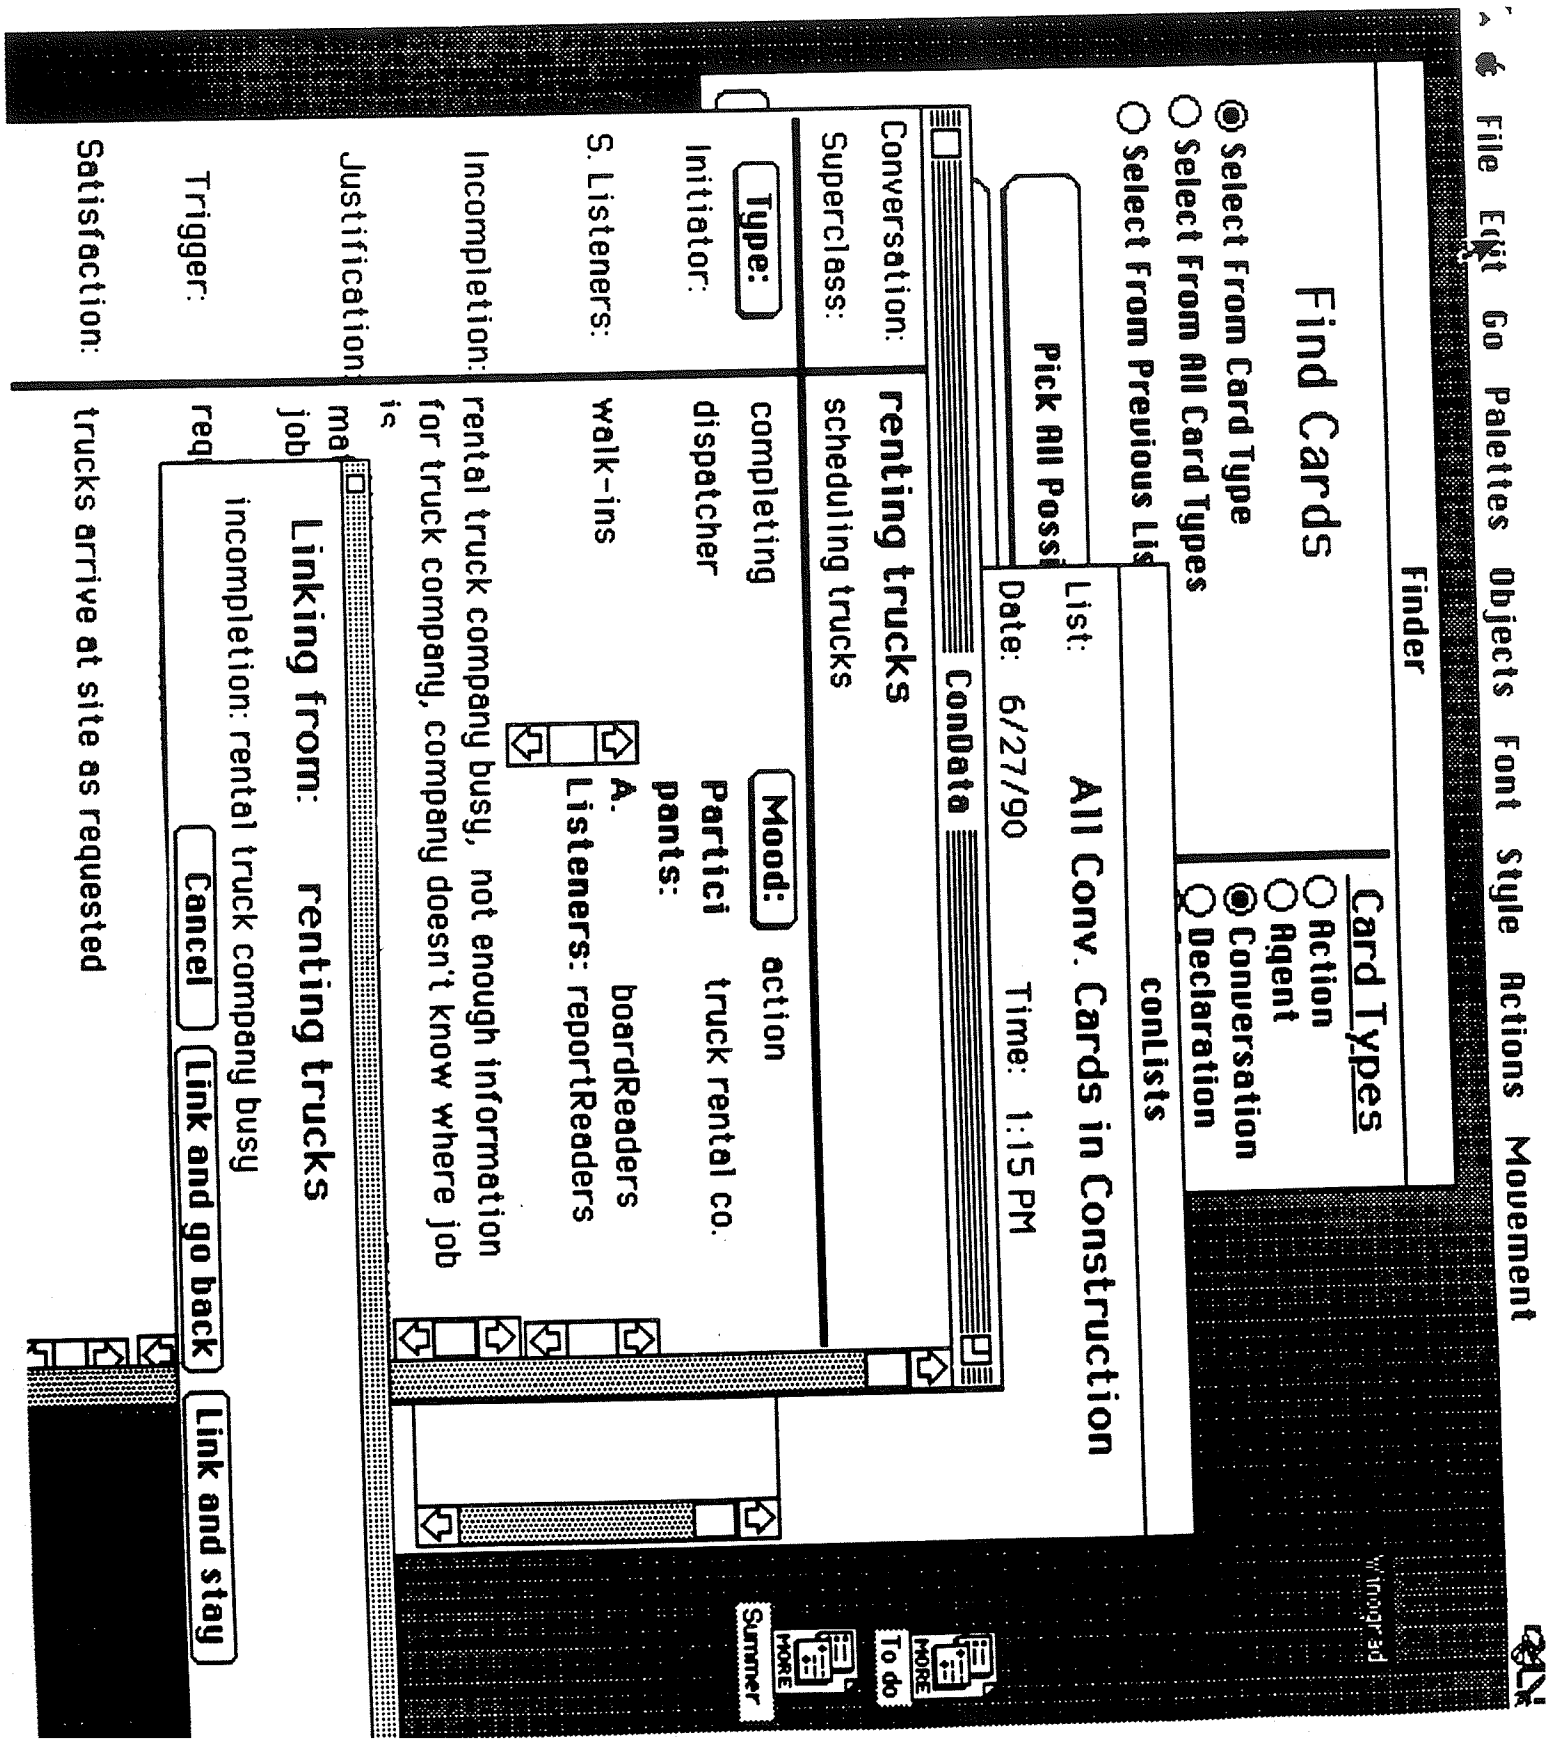

Screen 6 : Establishing the link

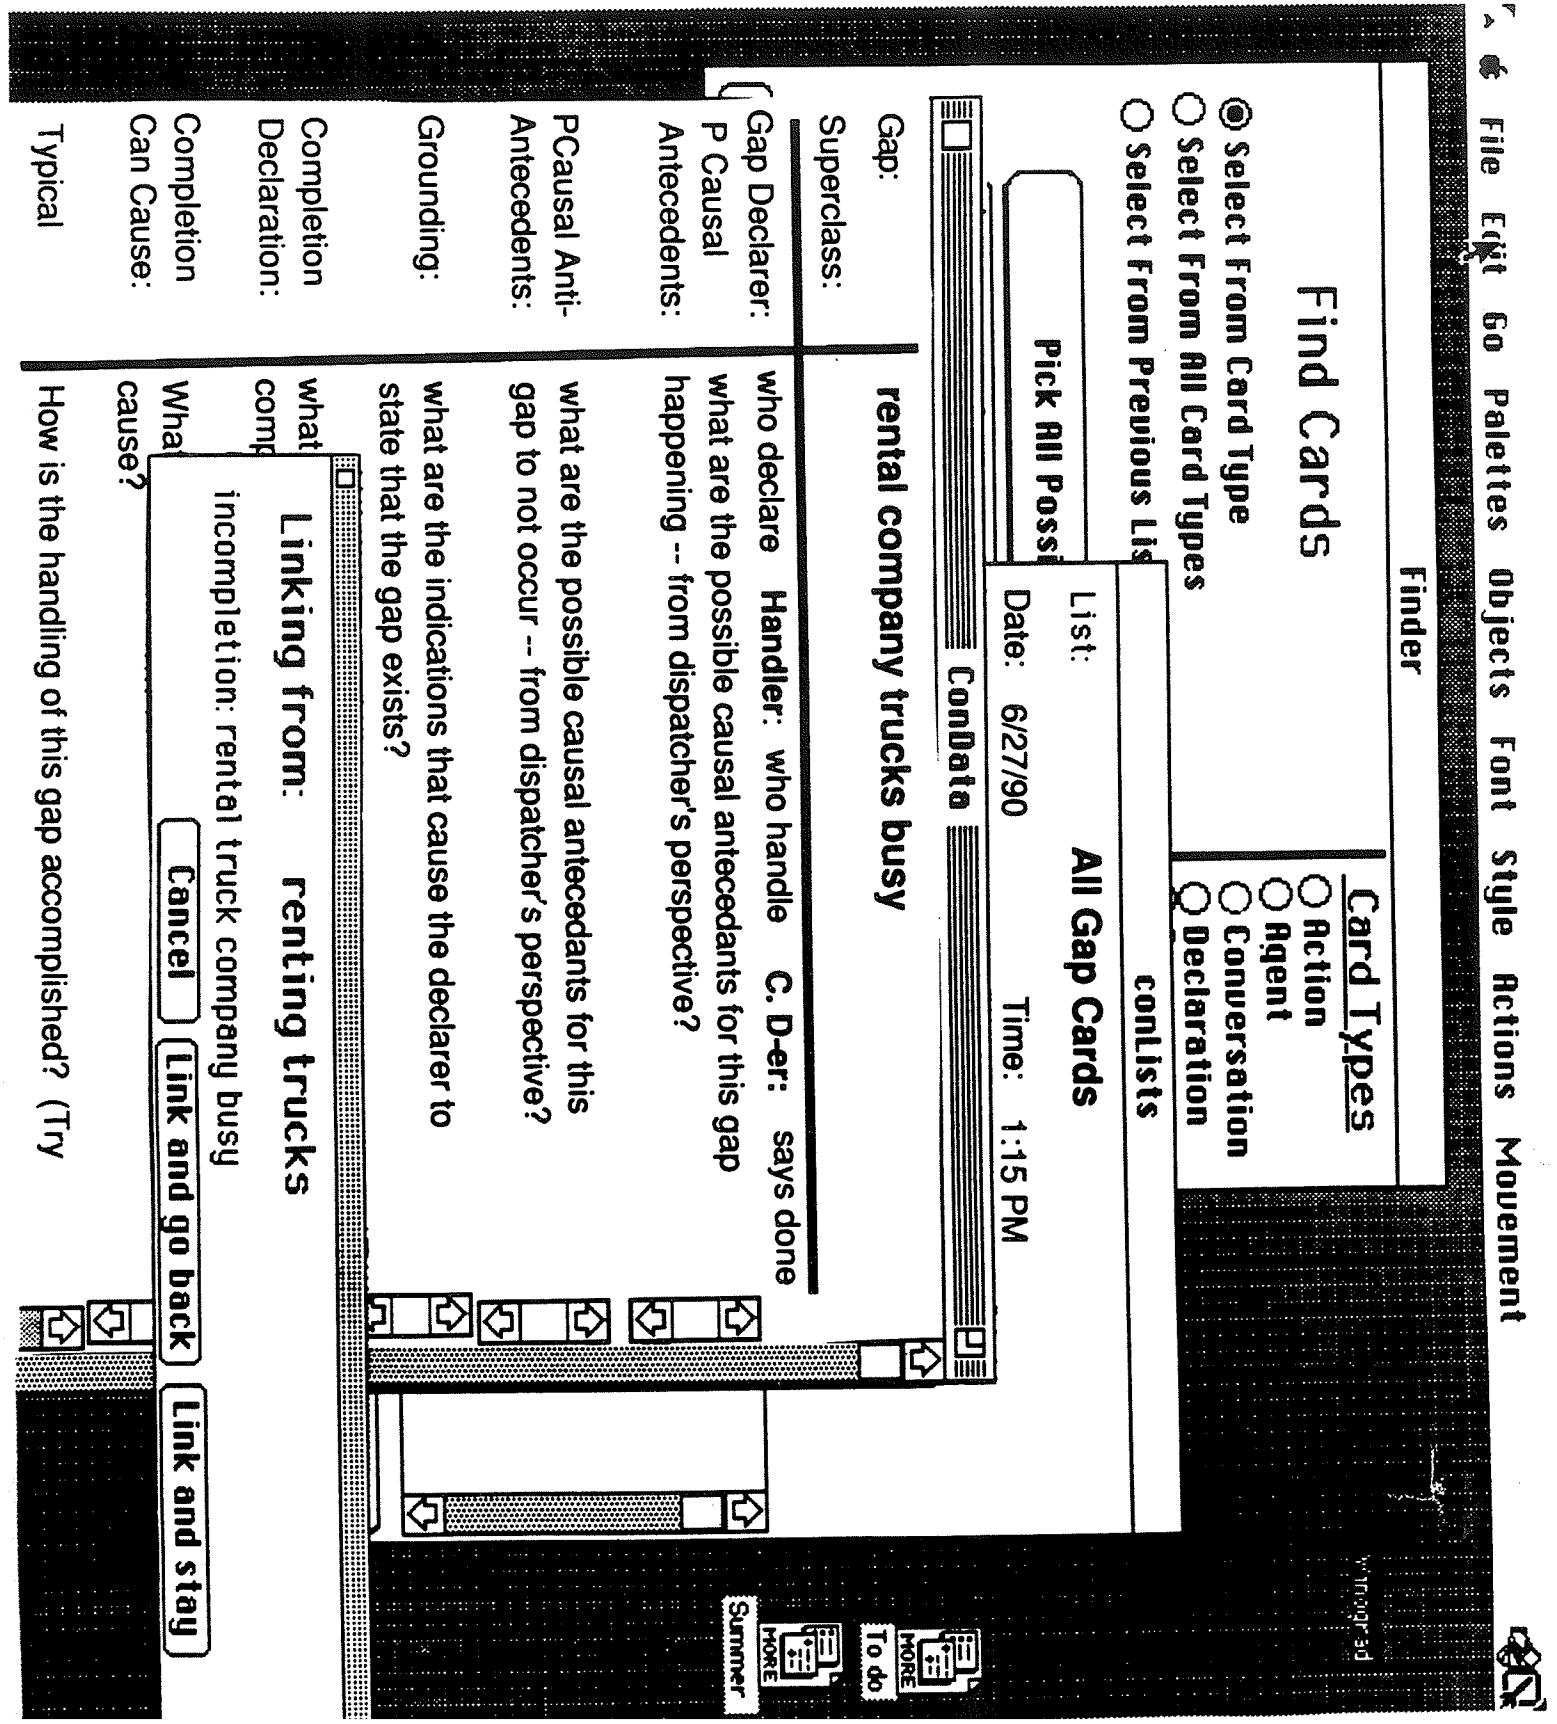

THE ACTION WORKFLOW APPROACH TO WORKFLOW MANAGEMENT TECHNOLOGY

Raúl Medina-Mora, Terry Winograd, Rodrigo Flores, Fernando Flores

Action Technologies, Inc.
1301 Marina Village Parkway
Alameda, CA 94501
(510)521-6190 rmm@ati.mhs.compuserve.com

ABSTRACT

This paper describes ActionWorkflow™ approach to workflow management technology: a design methodology and associated computer software for the support of work in organizations. The approach is based on theories of communicative activity as language/action and has been developed in a series of systems for coordination among users of networked computers. This paper describes the approach, gives an example of its application, and shows the architecture of a workflow management system based on it.

KEYWORDS

Workflow, ActionWorkflow, Coordination, Coordinator, Business process.

INTRODUCTION

In introducing new technologies into a workplace we are not simply augmenting the work, but are in effect reorganizing it. Technological innovation offers an opportunity for organizational innovation. In providing computer support for cooperative work, we are directly concerned with its potential for business process redesign.

For the past ten years we and our colleagues at Action Technologies have been developing computer software for organizational communication and action, based on a theory of work structure as language action. Previous publications [2][3][10][12] have described the basic elements of the theory and explained its application to computer-supported cooperative work:

- *Language acts*, classified according to a speech-act taxonomy.
- *Conversations*, which are coherent sequences of language acts with a regular structure of expectations and completions.

- *Time tokens* for *completions* in conversations.
- *Explicit mutually-visible* representations of acts, conversations, and times, as a way of facilitating communication in an organization.

This has led us to a new way of characterizing workflow, based on the identification and construction of atomic "loops" of action in which a performer completes an action to the satisfaction of a customer (internal or external). The overall workflow in any organization is an interweaving of these action workflow loops, some of which are highly recurrent (done in a structured way time after time) and others are *ad hoc* (unique to a situation). Our experience with workflow management technology has demonstrated the effectiveness of action workflow analysis in redesigning the action structure in an organization to improve the workflow, along with providing computer support.

BUSINESS PROCESSES

We distinguish three different domains in which to describe activities of an organization:

Materiel processes.

Human activities are rooted in the physical world. Nothing happens without physical things moving and changing state. If we ask "What is happening?" the obvious answer is a description of physical activity.

In the tradition of factory automation this was the relevant domain, in which physical components were transformed and assembled into product unities. Materiel process redesign and technologies have been used to move and process objects more efficiently, from the early analyses of Taylor and the production innovations of Ford, through the sophisticated techniques of modern industrial engineering.

Information processes

With the twentieth-century shift to "information work," the materiel process domain fails to capture what is important about everyday activity. With computer workstations, all of the physical work becomes indistinguishable—talking to people and tapping keys in front of display screens. What is relevant is the nature of what the talk and tapping is about.

Theorists and information technology providers have developed sophisticated ways to analyze and facilitate the flow of information. Current techniques of data flow analysis, database storage and retrieval, transaction processing, network communication, and many more have provided a structure of effective information processing. This is the heart of the applications offered by the computer industry today.

Business processes

What is lost in the information perspective is the recognition that information in itself is uninteresting. Information is only useful because someone can do something with it, and we can't define "do something" circularly as just the handling of more information. What do people do that matters?

Here we find the domain of business processes, in which people enter into language actions that have consequences for their future activities. When a customer hands a supplier an order form, there is a physical activity (transferring a piece of paper) and an information dimension (communicating a form with information about a particular set of goods, delivery instructions, etc.). But the true significance is in the business process dimension: It is a request for the supplier to perform some particular actions, in return for which the customer is committed to perform other actions (e.g., payment).

Our theoretical work has been identifying the basic structure of the business process dimension: workflows, roles, acts and the incompletions they lead to, which constitute expectations for further behavior by the participants. It is important to note that business processes are implemented in information processes, just as information processes are implemented in materiel processes. In moving to a focus on the language/action structure of workflow, rather than on the forms or database transactions used when acting, we are revealing a higher existing level of organization.

WORKFLOW

Most current approaches to workflow management are structured around the domain of information processes [9]. They begin with a class of information objects, such as forms or stored images, and define workflow as a sequence of actions to be done on those objects. The primary organizing structure is the "routing" of information objects among users, and the specification of automatic actions to be taken in that routing. In a way, this is very much like the materiel process view, in which parts are passed along from one "station" to another in a factory for processing, and some of the component tasks are taken over by automated machinery.

Traditional work management is well suited to highly structured "heads-down" paper processing, but is not adequate for supporting the realities of work in the 90s, with its emphasis on better educated workers who combine structured work with opportunity-based initiative and individual responsibility for quality and customer satisfaction.

Although our approach also includes capacities for generating and managing forms, these are grounded in the dimension of business process structure, which is constituted of action workflow loops. This provides the basis for allowing individuals to deal directly with the consequences of their work for completion and satisfaction.

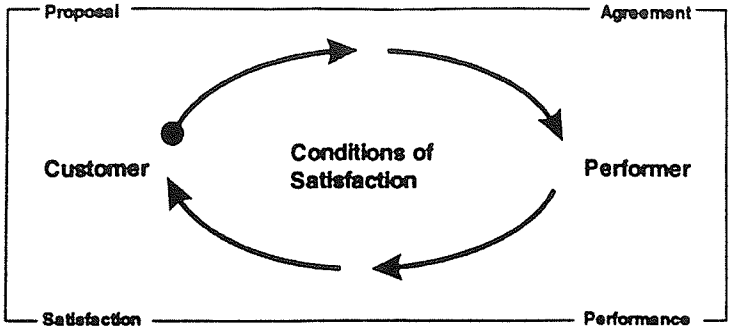

Figure 1. ActionWorkflow Loop

Figure 1 shows the basic sequence of actions in the action workflow loop. There is always an identified customer and a performer, and the loop deals with a particular action that the performer agrees to complete to the satisfaction of the customer.

The loop proceeds in four phases:

1) Proposal

The customer requests (or the performer offers) completion of a particular action according to some stated conditions of satisfaction.

2) Agreement

The two parties come to mutual agreement on the conditions of satisfaction, including the times by which further steps will be taken. This agreement is only partially explicit in the negotiations, resting on a shared background of assumptions and standard practices.

3) Performance

The performer declares to the customer that the action is complete.

4) Satisfaction

The customer declares to the performer that the completion is satisfactory.

At any phase there may be additional actions, such as clarifications, further negotiations about the conditions, and changes of commitments by the participants. (For a more detailed analysis of these possibilities, see [12], p. 65). The structure is defined by the language acts through which people coordinate, not the actions done by individuals to meet the conditions of satisfaction. The key difference in our approach is this shift from the task structure to the coordination structure. In a more traditional workflow approach, actions of coordination are seen as one kind of task or as a flow of information between tasks. In our perspective, tasks are defined by the requests and commitments expressed in the loops. This shift is analogous to moving

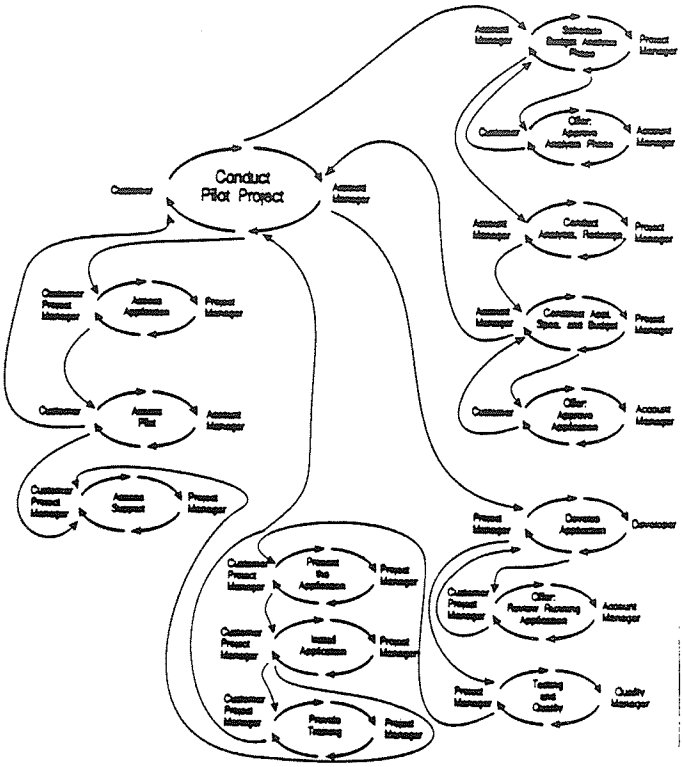

Figure 2. Business process map for pilot projects

from a view of a network as a collection of nodes (with links between them) to seeing it as a collection of links (with shared nodes). Although all the elements are still there, the different starting point leads to different potentials for representing and supporting the activities.

The simple workflow loop structure is both general and universal. It is general in that it occurs whenever there is coordination among people, regardless of what they are doing. The words "customer" and "performer" apply to people within a single organization as well as across boundaries. The loop structure is universal in that it is independent of any culture, language, or communication medium in which it is conducted. There are endless variations in the specifics of how the steps are taken, what other loops are triggered, and how people respond to breakdowns within them, but the basic structure is the same. The action workflow loop is like an atomic element of the chemistry of interactions. By combining these loops, all the complex phenomena of organizations are generated.

Our initial designs, such as The Coordinator™ [1][11][12], based their utility on the universality of this basic structure. They provided tools for creating and managing records of conversations (which correspond to workflow loops) based on the universal vocabulary of speech acts. The research described here follows later developments [2], which expand on this elementary structure as the basis for doing business process design. In place of the sequential tracking of forms found in other approaches to workflow support, we design (and help redesign) a business process as a collection of interrelated loops, each with its own completions and possibilities for breakdown.

Figure 2. shows an example of a business process map that was created to manage the conduct of pilot projects in the Action Technologies development group. The lines connecting loops show dependencies between them, with each connected to the appropriate quadrants of the loop, according to which aspects of the workflow structure they complete. We will examine a smaller example in some detail below.

We approach the task of designing a workflow management system by first analyzing the workflow structure and its possibilities for improvement and new functionality and for new or improved conditions of satisfaction that can be offered to customers. This analysis process, or "work mapping," uses theory-guided observations and interviews to generate explicit representations of the acts, roles, and incompletions that make up the flow of work. We have experimented with more detailed forms of mapping, in which we represent material and information structures in their relationship to the language/action structure [7], but the primary focus in our applied work has been on tools for revealing and highlighting the key elements of workflows and their relationship to completions and incompletions that are vital to the organization.

New opportunities to improve performance come from the ability to identify, observe, and anticipate potential "breakdowns," or failures to reach satisfactory completion. From the maps and associated discussions it is possible to identify places where breakdowns may occur on a recurrent basis and to see what additional steps or workflows can be put into place to anticipate and/or cope with them. The explicit articulation of the structure of customers, performers, and conditions of satisfaction leads to identifying new kinds of offers or requests that can be made. On the basis of these, new workflow structures can be instituted. While "breakdown" (by other names) is a standard concept in other forms of workflow analysis, the loops with their associated completions are unique to our approach.

Finally, we can identify those places where technological support can be valuable:

- Notifying users about actions that need completion.
- Providing users with the specific tools and information to complete a task, in a ready-to-hand way associated with identifying it.
- Managing reminders, alerts, follow-ups, etc. to keep processes moving along.
- Giving users an overview of where their tasks fit into the overall processes, both dynamically and through maintaining records of workflow history and providing structured access to them.

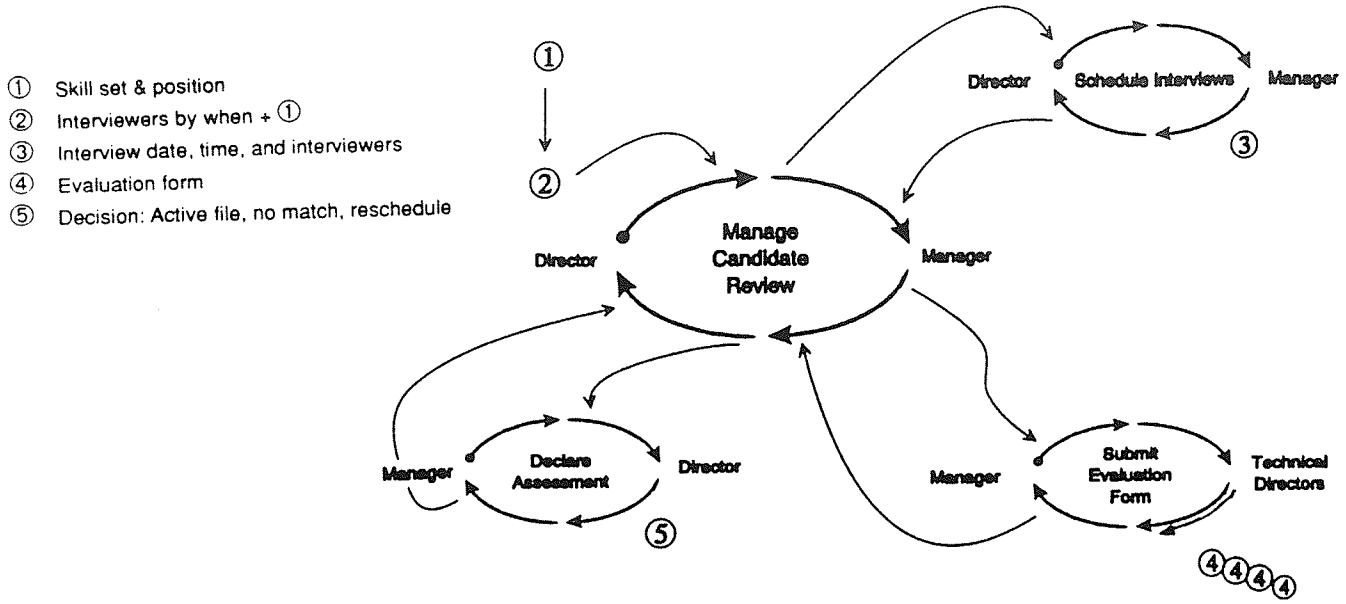

Figure 3. Core structure of the candidate review process

- Giving managers an overview of the status of workflow in the organization, both on demand and through generating regular reports and measures based on workflow structure.
- Automating standard procedures and individualized responses, on the basis of the action workflow structure.

Our methodology for providing workflow support is based on creating a unified conceptual structure and data representation that ties these functions into a coherent whole based on the explicit representation of workflow loops and their interconnections.

AN EXAMPLE

We will illustrate business process analysis and support with an application that was developed for managing the review of job candidates. This process is part of a larger business process for staffing, which is based on several dozen interconnected workflows, including advertising for positions, receiving and evaluating resumes, etc.

The process centers on four central loops, as shown in Figure 3. Each loop stands for a recurrent workflow, with the customer identified on the left and the performer on the right. Lines connecting workflow loops indicate triggering and dependency relationships between them. Numbered circles indicate forms and other external representations that play a role in the process.

The candidate review process starts when the director of personnel makes a request to a personnel manager to manage the review of a particular candidate. The manager starts the process by filling in an on-line form with information such as the interviewers, positions sought for the candidate, required skills, etc., as shown in Figure 4.

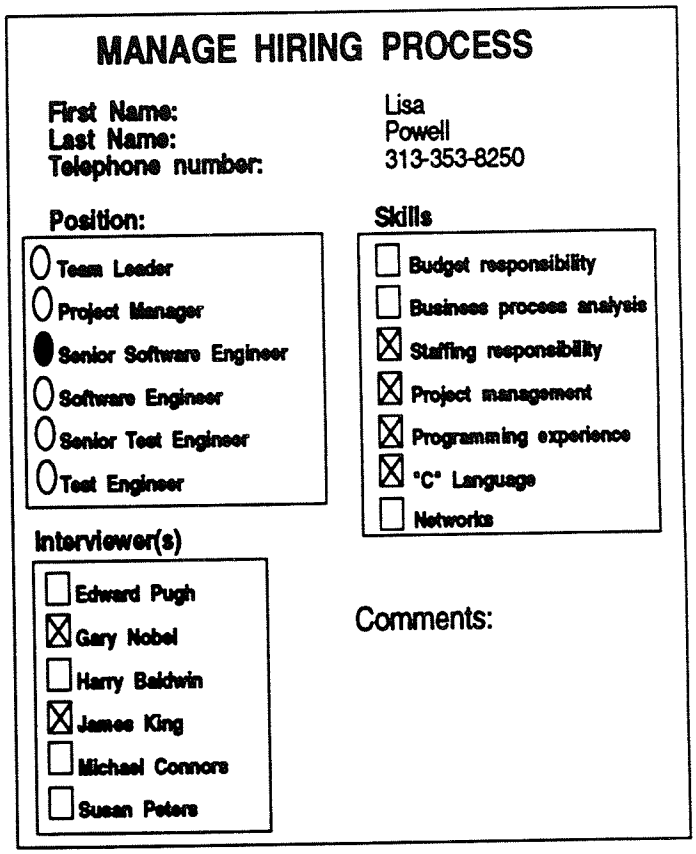

Figure 4. Form for initiating candidate review

This application was developed using the Lotus Notes version of the workflow management system, so the form was defined using the standard facilities for designing Notes forms. Other implementations differ, as described below. The structure of the review process has been defined by the analyst, working with the participants, and is stored in the definitions database maintained by the Workflow Management Server (see below). The server instantiates instances of all the workflows of the process and starts the "Schedule interviews" workflow automatically.

The "Schedule interviews" workflow corresponds to the second phase (agreement) of the main workflow: the manager agrees to do the work as requested by the director once the interviews have been scheduled. By including this scheduling in the agreement phase, a specific completion time can be promised.

Once the review process reaches agreement, the "Performance" phase starts and the "Submit evaluation forms" workflows are automatically started, one for each of the selected interviewers. Again, forms are defined for each of the participants and used in making actions in the workflow.

Once an interview has been scheduled for a particular date, all the workflows for submitting evaluation reports are initiated and directed to the selected interviewers to be completed on the specified date.

Each interviewer can use the workflow database to identify the set of workflows in progress. Figure 5 shows the status of interviews organized by interviewer. The lines showing next actions and times are generated from the action workflow database, using names defined specifically for this workflow.

Candidate	Recommended Action:	By When:
Edward Pugh		
Medina, Raul	Schedule an interview date	03/13/92
Bush, George	Check status of evaluations	0
James, Henry	Check status of evaluations	0
Samson, Dick	Check status of evaluations	0
Frank, Teddy	Thank you for submitting evaluatio	02/29/92
Gary Noble		
James, Henry	Recommit to evaluate	0
Henry Baldwin		
James, Henry	Recommit to evaluate	0
Wilson, Peter	Recommit to evaluate	02/29/92
James King		
Wilson, Peter	Recommit to evaluate	02/29/92
Bush, George	Recommit to evaluate	0
James, Henry	Recommit to evaluate	0
Michael Connors		
Wilson, Peter	Recommit to evaluate	02/29/92

Figure 5. Status display of interviews

By selecting one item, the interviewer brings up the on-line evaluation form for the candidate, which can be filled in incrementally and submitted when completed (this submission of a completed form constitutes a "declaration of completion" action in the workflow action structure). If the interviewer does not submit the evaluation report by a day after the agreed-upon completion date, the definition has been structured to cause the system to send a "follow-up" reminder to submit the report.

The definition of a workflow structure includes definitions of the forms that are used by customers, performers and observers of each workflow at each phase. When an interviewer accesses the document for the interview, it shows up as an evaluation form to be completed, since the interviewers are the performers of the workflow "Submit evaluation form." Other participants would see the forms relevant to the actions they are able to take, with fields available or protected from editing as suited to their roles.

Once all of the interview workflows have been completed, the system automatically declares the main workflow complete and moves to the fourth phase, where the personnel director declares (or not) satisfaction with the process. The system sends a mail message to the personnel director, as a prompt to act on the workflow for final assessment of the candidate.

At any time the manager can get an overview of the status by examining the workflow database through an appropriate view, as illustrated in Figure 6.

Recommended Action:	By Whom:	By When:
Not in process		
Harris, Mike		
Commence interview process	Manager	
Schedule Interviews		
Medina, Raul		
Schedule an interview date	Manager	03/13/92
Complete Evaluations		
Bush, George		
Check status of evaluations	Manager	0
James, Henry		
Check status of evaluations	Manager	0
[redacted]		
Check status of evaluations	Manager	0
Decision Pending		
Jones, Tom		
Decide on candidate	Director	02/29/92

Figure 6. Status overview of workflows

ARCHITECTURE

We have defined a general Workflow Management System architecture for interoperability among different applications and across diverse platforms, integrating the coordination of specific applications along with system enhancements and utilities from users and third-party developers. This architecture has been the basis for several implementations, including a DOS based "Business Process Management" system (BPM1) [2], an extended version of The Coordinator in the Windows environment, and a workflow application development environment in Lotus Notes (from which our example was drawn).

The overall architecture consists of one or more client applications (called *workflow-enabled applications*), and the structures and components that enable them to interact with the workflow management server and receive services from

it. Figure 7 shows the major components of a Workflow Management System.

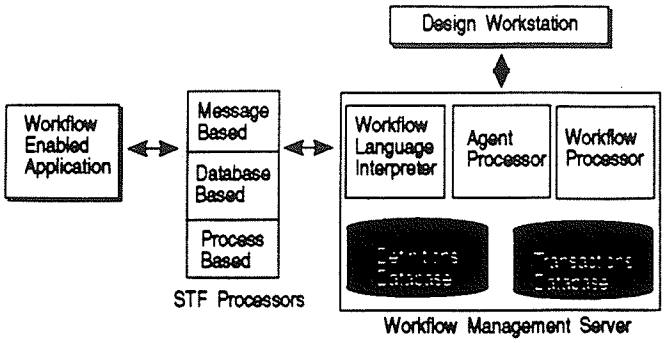

Figure 7. Workflow Management System Architecture

Workflow enabled applications
The goal of the Workflow Management System is to provide workflow capabilities to new and existing computer applications. Adding or integrating an existing or new application is referred to as "workflow-enabling."

Workflow-enabled applications are of three types:

1. Workflow-initiating applications.
An example of a workflow-initiating application would be an existing order-entry application that has been modified to initiate a fulfillment workflow. The task of the order-entry application is complete once the fulfillment workflow takes over and starts a sequence of actions to verify the new order, define customized requirements, alert manufacturing, etc. This level of integration can be done with little or no modification to the existing system.

2. Workflow-participating applications.
In the order-entry example discussed above, the participating applications are those that perform the details of the fulfillment process. The order-entry application first initiates a workflow to verify credit, for example, by sending an e-mail form to a credit manager to which she or he can respond by checking Yes or No. The addition of those buttons to an existing e-mail form, plus the work of defining Yes or No as they are to be understood in this case by the workflow processor are the only steps required to workflow-enable this aspect of the application.

3. Workflow management applications.
Workflow management applications provide managerial views and actions in addition to the operational ones needed to conduct the work. In the above order-entry example, an application that had workflow management built in could be used to keep track of fulfillment cycle times, sources of breakdown, etc. The candidate review example includes workflow management.

STF processors
STF Processors translate between an application's native data format and the Standard Transaction Format of the Workflow Language Interpreter. STF Processors isolate the Workflow Management Server from the interface used by the application and provide a layer for integrating different protocols and technologies. By providing an appropriate STF Processor, any existing database, messaging, or networking system can be incorporated into a workflow management network. If an application communicates by writing to a database, for example, the STF Processor will read the database and look for the records that hold STF transactions.

This architecture makes it possible for existing line-of-business applications, databases, networks, and protocols to be orchestrated by the ActionWorkflow system. Organizations already have tools in place to manage parts of tasks, and parts of workflows. It is an important requirement of a workflow system to integrate with the existing infrastructure, or the benefits will not outweigh the costs of moving to it.

There are three types of STF processors:

1. Message-based.
Message-based applications interact with the Workflow Management System by sending and receiving messages. The STF Processor receives the messages from applications and interacts directly with the Workflow Management Server. Similarly, it constructs messages to be sent back to the application. Message-based STFs are independent of the message transport. Our current implementations use MHS as the messaging system.

2. Database-based.
The client application writes and modifies records in an external database that is concurrently accessed by an STF Processor that has been built for the particular database platform. Applications initiate and participate in workflows by modifying records in this shared database. The STF Processor monitors changes to the database and interacts with the Workflow Management Server for recording and updating transactions. Applications can manage workflows and business processes by querying this shared database to obtain reports about the status of the workflows. We have implemented transaction databases in Lotus Notes and on SQL servers.

3. Process-based.
In the inter-process communication STF interface, a client application receives services from a server by making a process-to-process service request (a remote procedure call, for example). In this case, the STF structures are embedded in the parameter blocks of the service request and service result calls.

Workflow Management Server
The Workflow Management Server uses stored definitions of the workflow structure and of the history of transactions

to interpret and initiate acts. It comprises a number of interacting components:

a) Definitions Database
This database describes the workflow of the organization. The definitions include several basic structures. The core is the set of loop types and act names, with associated forms For example, the loop type "Manage candidate review" would have an associated form as shown in Figure 4, and an "accept candidate" act as one of its ways of reaching completion. The definitions database also specifies the linking relationships connecting the different loops, and the actions to be taken automatically by the agent processor.

The linking relationships are used to generate the appropriate sets of "next actions" for each participant as the workflow proceeds, and for automation. They can be of several kinds:

1. Subordinate workflow loops:
In order to complete a part of one workflow it is necessary to initiate and complete a subsidiary one. For example, in order to do the review it is necessary to schedule interviews.

2. Independent triggered workflow loops:
An action in one workflow triggers the initiation of another, which proceeds independently. For example, in a sales workflow the selling of an item from stock may trigger reordering, but the reordering is not a part of completing the sale that triggered it.

3. Resolving workflow loops:
The decision as to which action to take in one workflow requires the initiation and completion of another workflow. For example, a credit approval must be received before accepting or rejecting an order.

In each of these cases, there may be several triggered loops of a given kind instead of just one, with concurrency relationships among them. In the candidate review example all workflow loops for interviews are started in parallel at the moment the agreement is reached in the main loop. The definition of the process indicates that the performance phase of the main loop is completed once all the interview loops are complete.

b) Transactions Database.
This database contains the history of completed workflow loops and workflows-in-progress. It is accessed both for carrying out transactions and for providing status reports and overviews.

c) Workflow Language Interpreter
The Workflow Language Interpreter receives service requests from STF Processors in the form of workflow language constructs: workflow declarations, workflow actions, and requests for workflow management services. It instructs the workflow processor to calculate workflow states and next actions based on specified criteria (such as the current state of the workflow and the role of the person taking an action). It takes actions and makes reports based on the calculations of the workflow processor and the logic of the workflow definitions.

d) Workflow Processor
The workflow processor generates and manages transaction records in the transactions database, which keep track of the current state and history of the workflow, organized according to the component loops and associated completion times.

e) Agent Processor
The agent processor maintains a queue of events and times to trigger workflow actions that have been specified in the definition. We have taken the approach of *incremental automation*, initially assuming human action at each point, and then introducing a program-determined action at any point where rules can be effectively specified. Agent code is written in the workflow definition language and initiated on the basis of the workflow type and act that triggers it. It can take actions both within the workflow structure (making acts and initiating new workflows) and in other functions (printing reports, sending email messages, running other applications, etc.).

There are three ways in which agents are triggered:

1. Triggering act.
For example, a cancellation in a particular workflow initiates a request to a manager to deal with problems caused by cancellation.

2. Status changes in a workflow.
For example, a workflow moving to the state "completed" may trigger actions to cancel all of the subsidiary workflows in progress, whether or not the termination resulted from a cancellation, success, failure, etc.

3. Incompletion times.
For example, a follow-up request to a performer may be initiated when the time for completion of a loop has been reached without a declaration of completion.

Design Workstation
The design workstation is a separate application that is used to generate, modify, and maintain the definitions. We have developed a graphical notation for high-level workflow maps, and have implemented interactive structured drawing tools for creating and manipulating those maps, which can used for business process redesign, both with and without workflow management system development.

CONCLUSIONS

The approach and architecture described here have been developed in a number of prototypes and products. In addition to the development of computer support systems, the theory and analysis methodology has been used as the basis for consulting about redesign of business processes in a number of organizations (For a general discussion of business process redesign, see [5][6]. Kukla [8] describes a case study in a chemical plant, using earlier versions of our approach).

Our experience has demonstrated the effectiveness of business process redesign and computer support based on an ActionWorkflow analysis. The theory provides a starting point that is very different from conventional approaches to workflow. When an analyst first asks people in an organization "What is the work here?," the natural response is to start looking at the forms and procedures. We explicitly reject this, ignoring the forms and asking "What are you actually *doing*?" Without the action workflow structure, this question might seem meaningless, but with it there is a specific direction to move. Who are the customers and performers? What are the conditions of satisfaction in each loop? How is each of the four stages is carried out? How are the loops related to one another?

This questioning leads to identifying those places where gaps and confusions lead to incomplete workflows, misunderstanding of results, and ineffective information flow. This can then lead to new forms and procedures, rather than simply automating the old ones. Traditional methods have been *production-centered*, focusing on efficiency (as measured in standard output for input) and control. Our approach is *satisfaction-centered*, with a central focus on commitments, conditions of satisfaction, and timely completion.

In a significant way, this new methodology corresponds to the shift of concerns in business as we move into the 90s. Guiding concerns of productivity and efficiency have been replaced with others, such as quality (how are conditions of satisfaction set, met, and declared by customers; responsiveness (how are cycle times related to the completion of the structure of loops and how can they be systematically reduced); and customization (how can secondary loops be designed and managed to effectively tailor conditions of satisfaction in the main loops).

Our current efforts are to provide a general platform for action workflow management, which can be incorporated into existing information systems in an incremental way, providing the basis for new understanding of the business processes, and facilitating business process design on a larger enterprise-wide scope. Our goal is to open up the potential to radically improve the functioning of the workflow-enabled organization.

ACKNOWLEDGMENTS

We wish to thank Chauncey Bell, Tom Buchler, Bob Dunham, Pablo Flores, Michael Graves, Manuel Jasso, Juan Ludlow, John McAfee, Russell Redenbaugh, Tom White, and others at Action Technologies who have participated in the development of the concepts and systems described in this paper.

REFERENCES

1. Bullen, Chris, and John Bennett, Groupware in Practice: An Interpretation of Work Experiences, in C. Dunlop and R. Kling (eds.), *Computerization and Controversy*, Academic Press, 1991, 257-287.

2. Dunham, Robert, Business design technology: Software development for customer satisfaction, IEEE Press, *Proceedings of the 24th Annual Hawaii International Conference on Systems Sciences*, 1991, 792-798

3. Flores, Fernando, Michael Graves, Bradley Hartfield and Terry Winograd, Computer systems and the design of organizational interaction, *ACM Transactions on Office Information Systems* 6:2 (April, 1988), 153-172.

4. Flores, Fernando, *Management and Communication in the Office of the Future*, Dissertation, University of California, Berkeley, 1982.

5. Hammer, Michael, Reengineering Work: Don't Automate, Obliterate, *Harvard Business Review*, July-August 1990, 104-111

6. Keen, Peter, *Shaping the Future: Business Design through Information Technology*, Boston: Harvard Business School Press, 1991.

7. Kensing, Finn, and Terry Winograd, The language-action approach to design of computer support for cooperative work, *Proceedings of the IFIP TC8 Conference on Collaborative Work, Social Communications and Information Systems*, Helsinki, Finland, 27-29 August, 1991.

8. Kukla, Charles, Anne Clements, Robert Morse, and Debra Cash, Designing Effective Systems: A Tool Approach, in P. Adler and T. Winograd (eds.), *Usability: Turning Technology into Tools*, New York: Oxford University Press, 1992, 41-65.

9. Moad, Jeff, ViewStar faces many Goliaths, *Datamation* 38:11 (May 15, 1992), 30-31.

10. Winograd, Terry, A language/action perspective on the design of cooperative work *Human-Computer Interaction* 3:1 (1987/88), 3-30.

11. Winograd, Terry, Where the action is, *Byte*, December, 1988, 256-260.

12. Winograd, Terry and Fernando Flores, *Understanding Computers and Cognition: A New Foundation for Design*, Norwood, NJ: Ablex, 1986. Paperback issued by Addison-Wesley, 1987.