Considering Costs of Interruption and Deferral in Routing Interpersonal Communications

Eric Horvitz Johnson Apacible Paul Koch Raman Sarin Muru Subramani
Microsoft Research
Redmond, Washington USA 98052
{horvitz, johnsona, paulkoch, ramans, murus}@microsoft.com

***Abstract*— We describe a prototype service that performs cost-benefit analysis about whether to route a phone call to a user versus reschedule the communication for later. The system balances the expected cost of interruption with the cost of deferral of a communication. We first present a decision-theoretic perspective on the handling of interpersonal communications. Then, we describe a prototype system named Bestcom-X that makes decisions about the routing of telephone calls coming into a corporate PBX, based on models of the cost of interruption that were built via machine learning. We then describe a derivative of Bestcom-X, named Bestcom-ET which is now being used actively used by approximately one thousand employees at the Microsoft Corporation.**

Keywords—Statistics, economics, call routing

I. INTRODUCTION

Today, people seeking to communicate with others make personal decisions about the best timing and channel of the communication. They select and execute a communication modality or set of modalities based on their own needs and preferences—as well as on their knowledge and intuitions about the preferences and context of the person being contacted. Attempts to communicate are often suboptimal for both the initiator and recipient of a communication attempt.

Attempts by a *contactor* to establish real-time telephony may interrupt the recipient or *contactee* of the communications attempt during a poor time for conversing, or frustrate the contactor with a voice message capture that may lead to costly delays for both the contactor and contactee. Contactees employ multiple methods to filter incoming communications selectively. Some people may employ well-trained assistants, while others rely on the manual screening of incoming telephone calls, voice messages, and batches of email messages.

Limiting or deferring real-time communications so as to minimize disruptions and maximize privacy is only one piece of the challenge. Depending on the caller and the situation, contactees may often desire to be reached in real time rather than be missed by a caller. With current, ad hoc communications, it not uncommon for two people trying to speak to one another to note frustration about a nonconvergent volley of attempted communications, referred to as "playing phone tag."

The Bestcom project at Microsoft Research centers on call handling via an automated consideration of communication preferences, and of the current situations of both contactees and contactors. Other efforts in the spirit of Bestcom on context-sensitive communication include the work on the Nomadic Radio project [9]. Bestcom differs from the earlier work in the centrality of employing decision-theoretic analysis to identify the utility maximizing communication actions—actions we refer to as *best means* communication. Communication actions include the routing and rescheduling of real-time communications. The project also focuses on infrastructure and integration with legacy systems, such as the use of centralized context servers and decision making modules with existing corporate PBX systems or voice over IP infrastructure.

In best-means prototypes, we seek to build methods that make communication decisions that maximize the overall expected utility of participants, considering a directly assessed or computed expected *cost of interruption* for the contactee and the *cost of delayed communication.* One key decision involves automated rescheduling of communications based on considerations of the identity of the caller, and the current and future status of a user's interruptability.

We shall first present a decision-theoretic model for best means decision making. The model encodes key uncertainties and preferences about the context-sensitive handling of communications. We shall pause to consider key design considerations with regard to *principal agency* and *privacy* for a decision-theoretic call handling service. Then, we briefly review the use of Bayesian models of the cost of interruption, and the use of statistical forecasting methods that provide presence and availability predictions via consideration of a user's calendar, time of day, day of week, and the monitoring of users' activities on multiple devices. Next, we present an initial Bestcom implementation for handling real-time telephone calls, named Bestcom-X (for *best-means communication—experimental*), which has been integrated with the PBX at the Redmond campus of the Microsoft Corporation. Finally, we describe a larger deployment of a less-sophisticated descendant of Bestcom-X, named Bestcom-ET, in use by 933 users at the time this manuscript was composed.

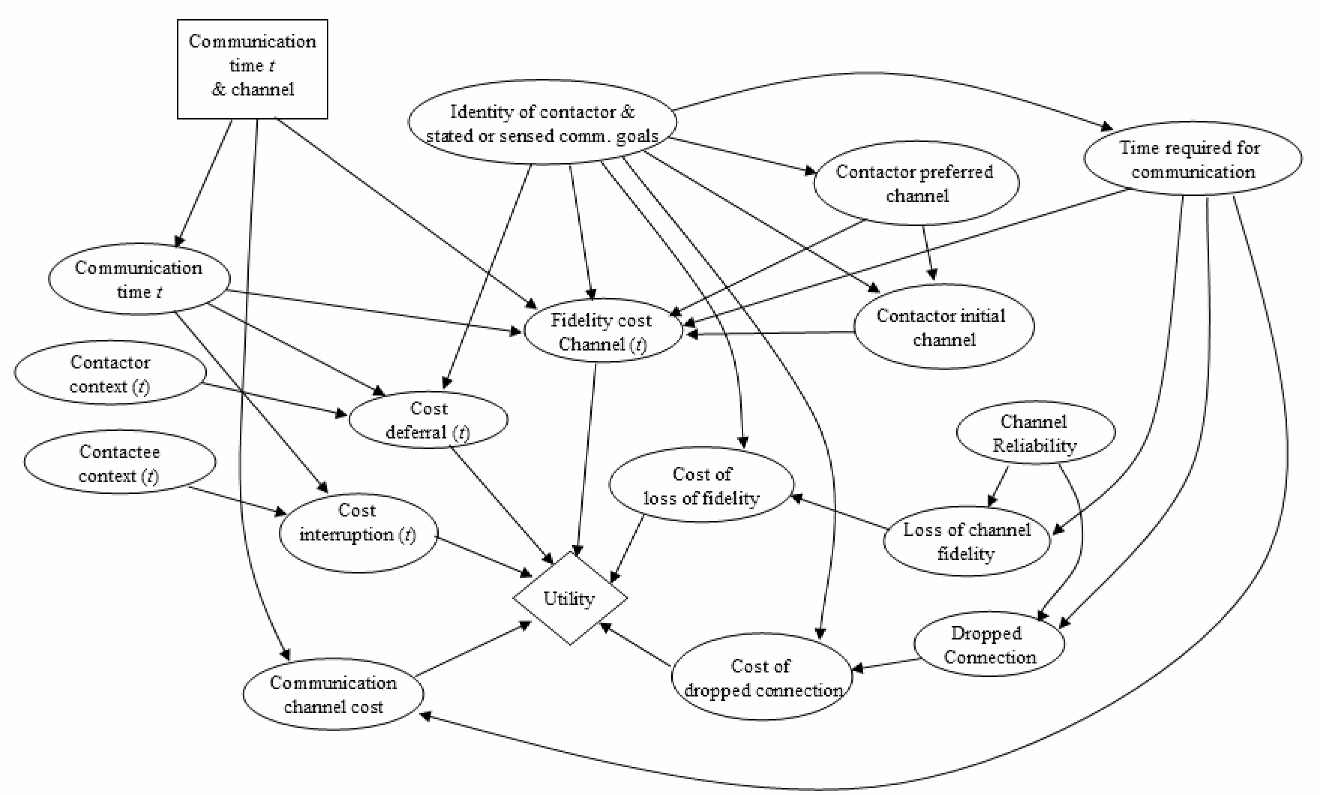

Figure 1. Decision model for best-means communication actions, represented as an influence diagram. Utility maximization algorithms identify the communication action with maximum expected utility given the expected cost of interruption, the cost of communication deferral, the cost of channel usage, and the losses in fidelity associated with the use of alternate modalities.

II. Decision-Theoretic Approach to Communications Handling

We take a decision-theoretic perspective as a foundation for best-means interpersonal communications. A decision-theoretic communication routing system employs components that infer the probability of key states of the world in real-time and attempts to take actions that maximize the expected value of the communication to the people they serve by harnessing an encoding of *communication preferences*.

A. Representing and Reasoning about Incoming Calls

In our detailed models, we reason about the current and future cost of interruption of taking a call and the cost of deferring communications. Probabilistic reasoning typically leads to probability distributions over these values. Beyond the detailed models, we have worked to formulate simpler methods that are approximations of the more sophisticated decision analyses. We shall start with a brief discussion of a detailed decision model as background before describing simpler approximations that operate in the real-world implementations we have constructed.

Figure 1 displays an influence diagram for the best-means communication problem, representing dependencies among critical variables, including random variables (ovals), a decision node (square), and a value node (diamond) capturing important dimensions of the interpersonal communication decision problem. At the heart of the model is a consideration of the tradeoff, captured in the contactee's value function, of the cost of deferring a call until another time *t*, and the cost of interruption (COI) of relaying the communication until another time, when the cost of disruption will be lower Other factors considered include the loss of fidelity with use of different channels (e.g., moving from a high-bandwidth video channel to a video channel with low frame rate), the cost of channel usage and the reliability of the different channels.

When key uncertainties and preferences, represented as the random variables and utility node of this graph, are assessed from users, such detailed models perform well. However, the burden of assessment can be severe. We have designed tools for minimizing the assessment burden, and for learning about costs from observations with data collection and machine learning. We have also developed methods that allow users to encode preferences directly about the costs and benefits of communication actions. By developing an expressive set of conditioning contexts, we can allow users to encode directly preferences and then perform simple cost-benefit analysis at run time. Such context specification includes the time of day and day of week, current device activity, and properties of a contactee's appointments that are accessible from electronic representations of meetings.

B. Criticality of Principle Agency and Privacy

In our work to build decision-theoretic communication routing systems, we discovered that we must make key design decisions about the critical issues of *principal agency* and *privacy*. Assumptions about agency and privacy have important implications for design guidelines, methods, and usage of best-means communication services.

1) Principal Agency

In decision theory, the *principle agent* of a decision is the actor who is charged with responsibility for a decision; we seek to optimize the expected utility of actions for the principal agents. Issues of agency arise in many real-world applications of decision analysis. For example, when a physician works with a patient on a medical decision problem, the default principal agent is the patient. Although the physician may do her best to advise a patient on the best course of action, it is the patient's preferences about outcomes and uncertainties that should be considered. In cases where a patient is incapacitated, others, such as family members may take on the role of principle agent of the medical decision making.

Communication interactions involve two or more agents. One can imagine different approaches to agency for a communication-handling agent. As a guiding design principle, we treat the *recipient* of a communication as the principal agent as it is the contactee whose attention is being sought by the contactor. Taking this perspective, automated actions about the *if, when, and how* of communications are based on the contactee's preferences. Thus, we consider the contactee's preferences to guide decisions about the acceptance, rejection, rescheduling, and shifting modalities of a communication. This doesn't mean that Bestcom services overlook the preferences of the contactor. Indeed, the recipient's preferences may typically take into consideration the preferences and situation of the contactor. At times, a contactee may assign agency to the contactor for portions of the decision making, but we seat this decision in the contactee's hands.

Beyond a recipient-centric version of best-means services, other formulations are feasible. For example, a system could select actions that would have the greatest value to both contactees and contactors, per a utility model treating both as equals. In another formulation, Bestcom decisions are guided by communication guidelines or a specific objective function specified for an enterprise. Nevertheless, we seek a recipient-centric approach as the dominating Bestcom paradigm because it is the contactor who seeks, typically without prior arrangement, the attention of the contactee.

2) Context and Preference Privacy

There has been growing interest in providing technologies that provide colleagues with rich presence information, so as to assist contactors with making communication decisions [4,8]. With work on Bestcom, we take as a design guideline that, by default, the rationale behind any communication decision made by a Bestcom agent is kept confidential. Only the user's communication agent has access to rich preferences and context information. Keeping rationale of decisions, and, more generally, the context of contactees confidential by default resonates with seating the agency of decisions with the contactee. A user must explicitly grant individuals privileges to review real-time or forecasts of presence or availability. As we shall see, new context sensing and availability forecasting tools reveal a great many details about a user, and keeping such information private by default resonates with the wishes of users who become familiar with the richness of the sensing and inference.

III. Modeling the Cost of Interruption

We shall now review our work on models of the expected cost of interruption based evidence about desktop activity, from analysis of a user's calendar and presence information, and considerations of time of day and day of week. Our work comes in the context of several studies of interruption [6,7].

We have taken a machine learning approach to the problem. We build statistical models that provide probability distributions over the cost of interruption for a user given real-time observations. The modeling methodology involves gathering and labeling data with relevance to a user's interruptability, including calendar information, device activity, and, when available, information from other sensors, such as an analysis of the visual scene (*e.g.,* to identify head presence and pose) and ambient sound in a user's office (*e.g.,* to detect conversation).

In earlier work, we pursued the construction of comprehensive models that apply to any situation. However, we have found it useful to collect data and build models that are tailored for different situations. In particular, we have constructed models with machine learning procedures for the cases where a user is sensed to be in their office versus situations where a user is away from their office, where we often have to rely more centrally on appointment data, sometimes extended with location information. We have also explored building distinct models for the case where activity is sensed on a computing device and where there is no activity sensed.

To assess the cost of interruption, we start by providing users with assessment tools that allow them to assess default costs of interruption as a function of the time of day and day of week. Figure 2 displays a week-long *time-pattern palette* provided in Bestcom-ET for assessing the default costs of interruption for different

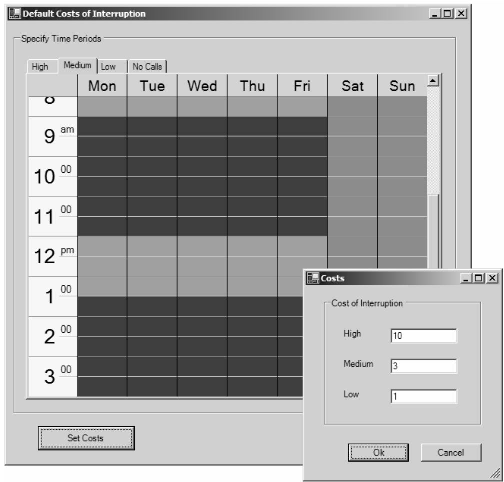

Figure 2. Time-pattern palette for assessing default costs of interruption by time of day and day of week.

times of day and days of week. These default costs are used when appointments are not indicated on a user's calendar and when device activity is not sensed. With the time-pattern palette, users indicate costs over time by clicking and dragging over regions of time on different layers representing low, medium, and high costs of interruption.

Next we gather data about the cost of interruptability associated with different kinds of meetings. We provide users with an appointment interruptability assessment tool. This application extracts events from a user's MS Outlook calendar and provides a means for them to assess the events as being associated with low, medium and high cost of interruption [4]. Following the labeling of events, a data analysis system harvests multiple attributes of each event by making calls with MAPI and the Outlook Object Model to build a set of cases that includes nearly thirty properties of each calendar item (*e.g.,* location and duration of meeting, number of invitees at a meeting, relationships with attendees, duration, recurrent meeting versus standalone, etc.). A library of labeled cases is created and we employ Bayesian machine learning to learn a Bayesian network model. At run time, the Bayesian model supplies a probability distribution over states of interruptability based on appointments and location (employing 802.11 for coarse location information [1]). Fig. 4 displays a learned Bayesian network that predicts the likelihood that a meeting will be attended and the cost of interruption for the meeting. These likelihoods are combined into an expected cost of interruption based on time of day and calendar. Details of this combination are described in [4].

We developed a separate set of tools for building models that predict the costs of interruptability in office situations, including consideration of the relevance of desktop activity. The methods leverage logs of several hours of users' activities in their offices. We create a videotape of activities in a user's office, including a view of the user's computer display. The videotape is synchronized with a time-stamped stream of events from the client devices and other sensors (if present) provided by the Eve event-monitoring system [3]. With regard to computing events, we capture such low-level states as the application being used, whether the user is typing, clicking with the mouse, as well as a set of higher-level events such as the pattern of switching among applications (*e.g.,* single application focus versus switching among applications) and indications of task completion (*e.g*., a message being sent, a file being closed, an application being closed, etc.). We also note whether a conversation or other acoustical signal (fingerprint of shuffling activity) is identified, and whether a user's head is being tracked, and if so, if the user is gazing at or away from the computer. The visual analysis is performed by a Bayesian head tracking system developed at Microsoft Research [10]. Finally, we consider, as additional evidence the user's appointment information for the periods of time logged.

Following the capture of computer data and video, we use a tool, called the Interruption Workbench [5] to display the data for labeling. The Interruption Workbench displays video and audio of a training session and that tracks the status of events from the log of events gathered during the training session. The workbench allows users to define high, medium, and low costs of interruptability states. A screenshot of the tagging tool is displayed in Figure 3.

When a user has completed tagging one or more sessions of office activity, the system creates a data file containing a vector of event states for each two-second period, and associates these periods with different interruptability labels. The system then performs a Bayesian learning procedure and outputs a Bayesian network model that predicts the cost of interruption in

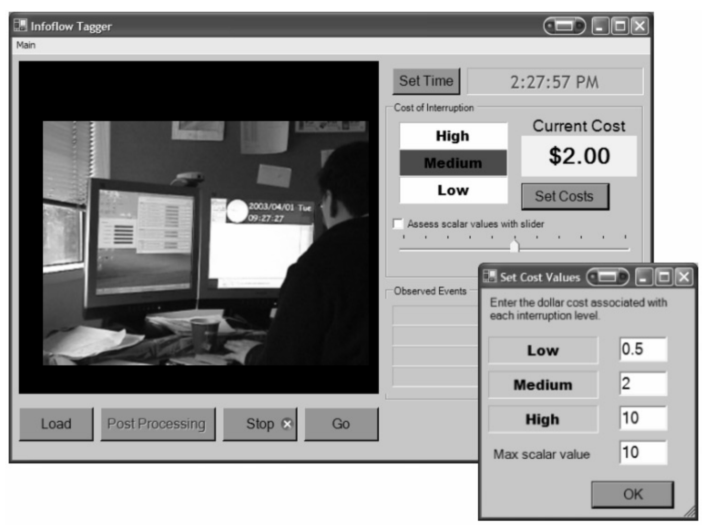

Figure 3. Interruption-cost workbench, employed to capture and synchronize sensed events for use in tagging periods of time with cost of interruption with an incoming call.

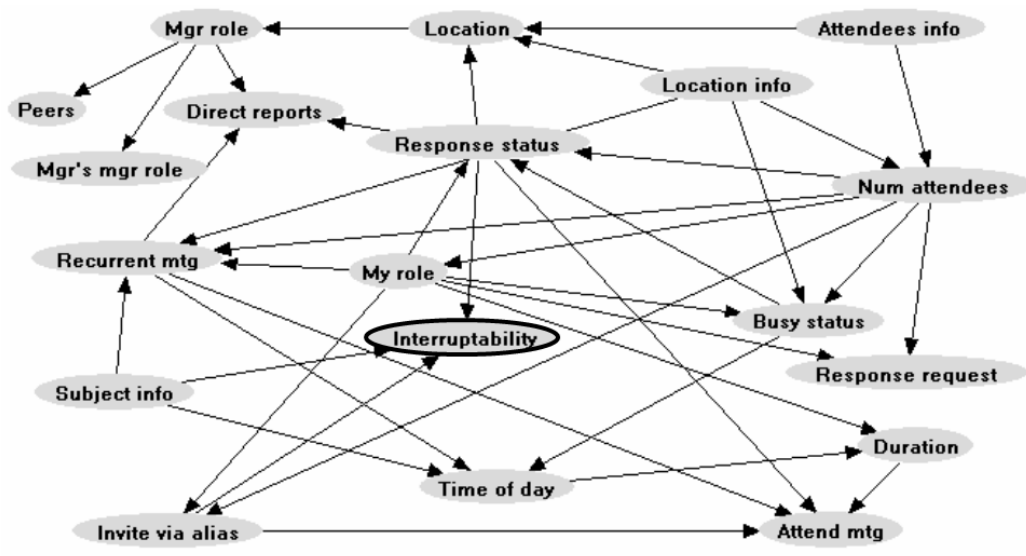

Figure 4. A Bayesian network learned from tagged appointment data provides inferences about the cost of interruption with a phone call given multiple properties of calendar items encoded in MS Outlook appointments.

real time, based on calendar information, computer interactions, and sensor data.

We note that beyond inferring the current cost of interruption, variables are automatically created from the data set that make predictions about the future states of a user. These include inferring the probability distributions over times until states of low, medium, or high costs of interruption will be reached, and predictions about the times until low, medium, and high interruptability will be achieved for different amounts of time, *e.g.*, the time until a user will remain in a state of low cost of interruption for at least 15 minutes.

Finally, we assess from users the dollar values that the user would pay to avoid taking phone calls at each of the cost levels. This enables us to compute the *expected cost of interruption* (*ECI*), by taking the expected value of the cost associated with different states of interruptability under uncertainty as follows,

$$ECI = \sum_i p(I_i \mid E)c(I_i)$$

where $p(I|E)$ refers to the probability distribution over the state of interruptability of the user given vector of evidence E and $c(I_i)$ is the cost assigned to interruptability state i. Fig. 5 shows inferences about the cost of interruption over time based on a model built from labeled calendar information.

The methods yield relatively accurate predictions. In a set of experiments, predictions about the interruptability of meetings had a classification accuracy of 0.81 [4]. Predictions of interruptability, from observations of device usage in an office computer use setting showed a classification accuracy of 0.73 in correctly assigning the state of interruptability of a user for the holdout cases. The system could predict the time that a user would next be in a state of high cost of interruption with a 0.78 accuracy [5].

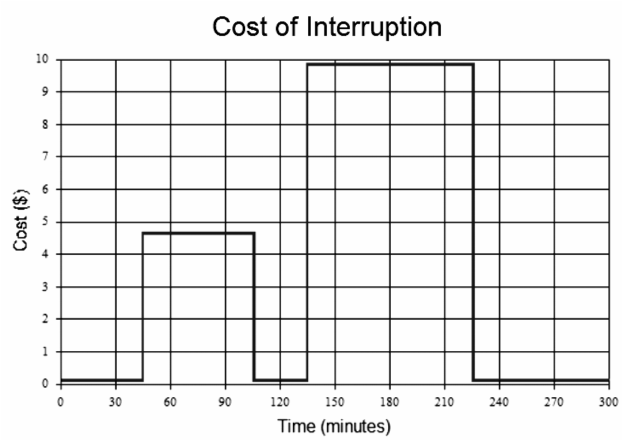

Figure 5. A forecast of the expected cost of interruption over time, based on inferences about meeting attendance, meeting interruptability, and default costs of interruption by time of day and day of week.

IV. Research Prototype: Bestcom-X

We shall now describe Bestcom-X, an implementation that makes use of dynamically computed costs of interruption in controlling the routing and rescheduling of incoming phone calls. Bestcom-X has centered on the handling of real-time telephone communications coming into the Microsoft PBX system. Bestcom-X monitors a user's appointments and activities and considers the identity of people calling the user. The system decides whether to route the call immediately to voicemail, to ring a users desk phone, to route the call to another phone or a mobile phone, or to seek to reschedule the call until later. It also logs all of its interactions and provides, in email a report on what it did with each call for auditing purposes.

The initial challenge of integrating Bestcom-X with the PBX were substantial; elucidating available programming interfaces and developing communications with the PBX were not straightforward and required detailed work with the PBX manufacturer and third-party integrators. We will not cover here the details of this low-level integration, and they are not important for understanding the critical functionality of Bestcom-X.

Bestcom operates on personal computers in conjunction with a centralized Bestcom-X server. The Bestcom server maintains accounts for all contactees, and stores preferences about the handling of communications as well as the current state of the user. The Bestcom server maintains a whiteboard of contextual information, similar in design to earlier work on context services [2]. Such information includes whether a user is currently on the telephone, the current meeting status of the user, settings of presence made with instant messaging, the user's proximal schedule, and key events sensed about a user's activity on registered client devices. The system makes calls to a subsystem that provides the current and future costs of interruption over time, employing the models described in Section III.

A. Abstracting Users into Static and Dynamic Groups

Bestcom-X has access to the current cost of interruption and to forecasts of the cost of interruption over time for a user. To perform cost-benefit analysis, we allow users to assess the cost of deferring a cal until later. Rather than assessing a complex function of the cost of deferral over time, we assess single numbers that capture the amount of dollars that a user is willing to pay to avoid deferring a call from different callers. It would be burdensome to assign costs to every caller. Thus, Bestcom-X provides abstraction tools that allow users to assess the cost of deferral for users in terms of groups, including the default group, *Others*, representing people not in a special, called out group. Predefined groups are provided based on several different properties that characterize relationships and activities. Such groups provide an essential abstraction for reducing the burden of preference assessment. Figure 6 displays the group manager of Bestcom-ET.

As displayed in the figure, users can create ad hoc *static* groups such as *Critical colleagues, Close friends,* and *Bestcom core* with an editor that allows users to create groups and add people from inside or outside the organization to the groups.

Beyond ad hoc groups, the system allows users to define or choose predefined classes of *relationship* and *dynamic* groups, capturing different relationships and classes of activity, respectively. The members of such groups are populated automatically via an examination of relationships and activities. Available classes of groups include people associated with meetings within different time frames gleaned from a user's online calendar, organizational relationships gleaned from an online directory, a tracking of communication history, and projects, as captured by the authoring of registered documents, software development tasks, and contributors to project-specific servers. Meeting-centric groups are populated by an ongoing analysis of appointments encoded in a user's Microsoft Outlook calendar. These dynamically assembled groups include such potential contactors as people a user is scheduled to meet with in the next hour and the rest of the day.

Dynamic groups also include sets of people based on the history of communications via different modalities. These groups include people who have either contacted the user or had been contacted by the user within different time horizons. Such communication groups include people whom the user telephoned within a day or week and people who had successfully reached the user by telephone within a day or week. Groupings of people by relationships are constructed by making calls to Microsoft's Active Directory Services. Such groups include organizational peers, direct reports, manager, all people above the user, and all people below the user in the organizational hierarchy. People falling outside of static and dynamic groups are considered to be in a special group called *Others*. This group is important for handling the default case of people who fall outside of the explicit group specifications.

Once users define new groups or add groups of potential contactors from a *group chooser*, they can assign privileges and properties to the groups. As displayed by the foreground pop-up form at the lower right-hand side of Figure 6, users can assign special ring tones, forwarding privileges, and rescheduling privileges to different groups. Users can also invoke cost—benefit analysis for incoming communications from contactors from a group and assess a cost of deferral, capturing the value of allowing a contactor to breakthrough to a user in real time.

B. Call Handling in Bestcom-X

Bestcom-X allows users to define best telephone numbers based on time of day and day of week, sensed location, and context. Users employ a time-pattern palette similar to the palette used for assessing costs to indicate phones by time. For example users can specify that their office phone is best when they are sensed as being their office. If they are not sensed as being in their office, then a primary cell phone should be used, except on weekends when their home phone would be best.

The system also allows users to define a cost of deferral for each group and to grant reschedule and forwarding privileges to people by group. For example,

Figure 6. Group Manager in Bestcom-ET. Groups of individuals can be created that capture relationships and activities, including meetings, communication history, and projects. Foreground displays a form used for assigning privileges and properties, including an assessment of the value of allowing a call to breakthrough to a user.

a user can tell the system that callers in the group, "People with whom I am meeting later today," have a \$10 cost of deferral, and have forwarding and call reschedule privileges.

When a call comes into the PBX, the Bestcom-X resolves the group of the caller, and contacts the Bestcom-X server, which provides the cost of interrupting the user at the current time and over the next several hours. The system also accesses the cost of deferral and privileges that the owner of the account has granted to the individual.

Because contactors may be members of multiple groups (*e.g.,* a peer and a member of a meeting with the contactee in the next hour), contactors inherit a cost of deferral associated with the group with the greatest cost of deferral. Contactors who are not unified with any specially defined groups inherit properties assigned to the *Others* category.

When a contactor calls a contactee the Bestcom service seeks to identify all groups that contain the contactor and to access the cost of interruption associated with a user's current situation. If no activity is reported by the Eve event systems operating on client devices and no appointments appear as currently active on the user's calendar, the system accesses the default costs for the time of day and day of week. If user activity on a client device is registered on the Bestcom server, the cost associated with the activity is noted. The system also notes the cost of interruption associated with a meeting appearing as currently active on a user's calendar. Users can specify whether activity or appointments take precedence or whether the highest cost of the two sources of contextual information should be taken as the cost of interruption associated with a context.

The basic cost—benefit analysis employed by Bestcom is represented by the graphic in Figure 7. In summary, if the cost of deferral assigned to a caller exceeds the current cost of interruption, the call is relayed to the user at the best number, established by time of day, day of week, and user presence. On the other hand, if the cost of interruption is larger than the value assigned to taking a communication from a contactor, the system either takes a message or attempts to reschedule the call, depending on whether the contactor is in a group that has *seek reschedule* property. If the contactor has a reschedule privilege, the contactor is engaged in selecting from a list of times, composed automatically by checking both the contactor's calendar and the contactee's cost of interruption over time.

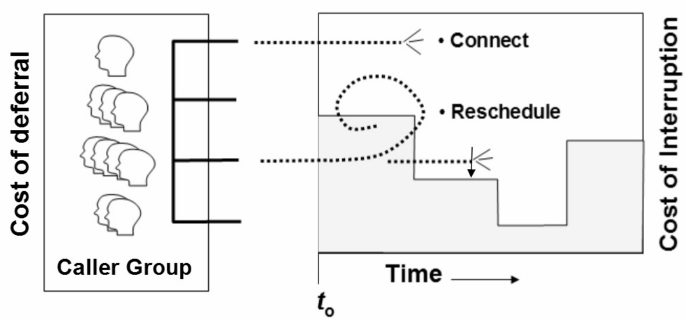

Figure 7. Graphic demonstrating call breakthrough versus reschedule (swirl) for a better time depending on the cost of interruption and cost of deferral for a communication.

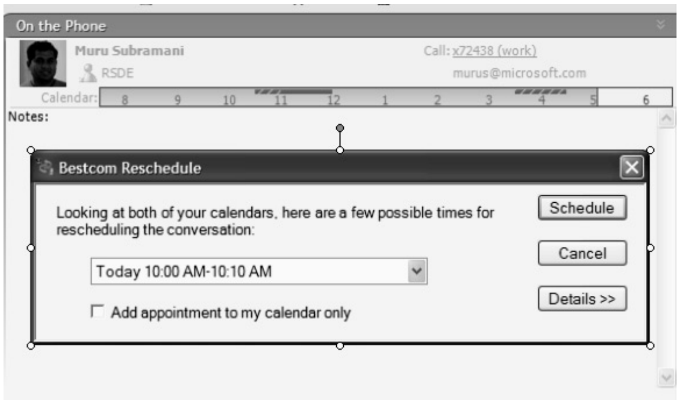

Figure 8. Contactor's experience with Bestcom rescheduling. Automated rescheduling in Bestcom-X, consider the cost interruption over time to finds one or more candidate times and also allows the contactor to provide background and documents for the call.

User's can specify at configuration the minimal amount of time to delay until making an appointment for a real-time call, so as to be sure to be notified about the forthcoming coming communication. In a variant of the system, the expected time until a user will see an alert is inferred from activity and users can specify a preferred likelihood that they will be notified about a rescheduled call. This threshold is used to push the meeting options later in a dynamic manner.

Figure 8 displays a contactor's experience with rescheduling a real-time call. A rescheduling attempt is made by the system and, given a go ahead, when available, the schedule of the contactor is also considered. Candidate times for the real-time conversation are offered. As displayed in the figure, the contactor can provide details about the call, and is provided with an opportunity to add links to documents to be reviewed before or during the conversation. The contactor can also request that the call be set up so as to share screens during the scheduled conversation.

Overall, Bestcom-X is solving a portion of the more comprehensive best-means decision problem represented by the best-means decision model displayed in Figure 1. Figure 9 shows the variant of the best-means decision subproblem considered by Bestcom-X, centering on decisions about the best time and channel to use for taking calls. The contactee's context is expanded to show the multiple classes of observations considered in inferences about the cost of interruption of the user.

V. LARGER-SCALE FIELDING: BESTCOM-ET

Bestcom-X was deployed with our research team and tested by a group of six users. The system served as an exploration of core principles and multiple interacting components for routing and rescheduling phone calls. In demonstrating the system and the process required to configure models and preferences, we discovered that the effort required to learn personalized models of the cost of interruption would limit typical users from using the system. Thus, we set out to create a more basic platform for research on real-world incarnations of Bestcom-X.

We set out to initially field on a larger scale a simpler system that provides a set of basic policies for routing and rescheduling based on logical statements about a user's state and about priorities assigned to callers. The system, named Bestcom-ET (for Bestcom-*Enhanced Telephony*), has been informed by insights about desirable functionality learned on the Bestcom-X effort. Bestcom-ET also leverages the work on integration of the Microsoft PBX with the Bestcom server, and links to the software client.

Bestcom-ET provides similar access to the static and dynamic caller groups. However, it does not employ the Bayesian models of the cost of interruption. Instead Bestcom-ET allows users to set forwarding to mobile devices based on assertions about the time of day and basic situations, such as "Forward to my cell phone when my computer is locked," "When my screen saver is running, When my Messenger status is set to Busy," etc. We also provide the time-pattern palette to assess busy states and to identify the best device to route call, given routing to a mobile device is indicated by the logical rules.

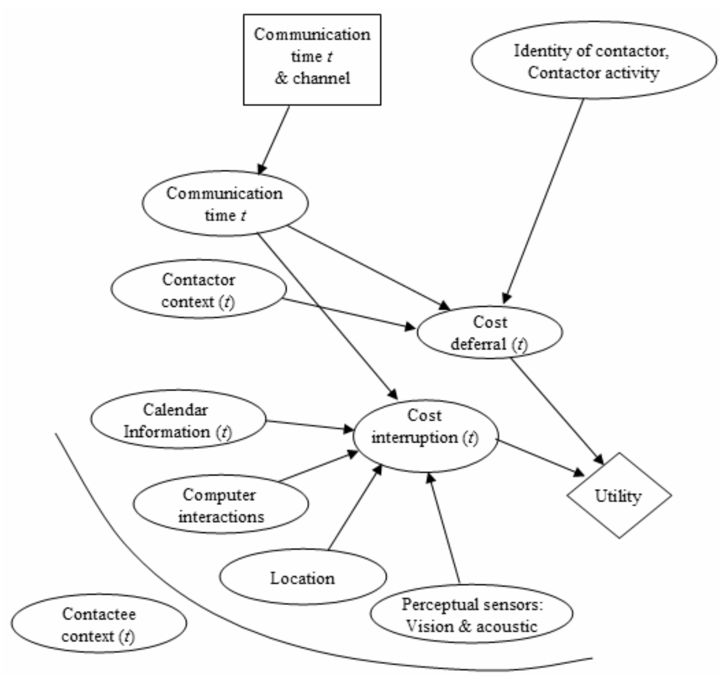

Figure 9. Inflence diagram representing key random variables and relationships in the best-means communication decision problem addressed by the Bestcom-X prototype. The model can be viewed as a focused variant of the more comprehensive best-means communication problem displayed in Figure 1.

Bestcom-ET also provides rescheduling services. However, instead of a providing a formal cost-benefit analysis as is performed in Bestcom-X, Bestcom-ET rescheduling is invoked based on logical rules about availability (e.g., if a caller has rescheduling privileges, invoke rescheduling when I am on the phone, etc).

As part of our research, we have been logging multiple aspects of the use of Bestcom-ET and have been monitoring and responding to an email feedback list. As of writing this article, 933 users are actively using Bestcom-ET. 524 users have set up custom-tailored forwarding rules. The most popular logical states are forwarding a phone call to a mobile device when the user's computer is locked, or when the Messenger instant messaging system is set to *away*, and forwarding a call directly to voicemail when the user is in a meeting or when Messenger is set to busy. 430 users have set up custom static groups and 338 users have added dynamic groups to their policy specifications.

We have continued to monitor user questions, suggestions, and feedback. At the initial fielding of the system, this message was typical of many enthusiastic new users:

> **Friday, March 07, 2003 2:35 PM**
>
> **To:** Bestcom Team
>
> This is the COOLEST application I have seen in years, also it is going to increase my productivity!!! (Despite today that I have being playing with it a little bit too much, I think). GREAT JOB!!!

We are continuing to push more sophisticated functionality into Bestcom-ET. For example, we will soon be fielding, via an upgrade to Bestcom-ET, a simple cost—benefit analysis that allows users to build, via direct assessment, models of the cost of interruption. The models are constructed via the time palette for costs of interruption by time of day and day of week, and with tools that allow users to directly assign costs to appointment "types" (e.g., "I am in a scheduled meeting with a colleague higher up on the organizational chart," "I am in a meeting with the people I manage, etc.) and to basic system states (e.g., "I am typing in an Office Application," "I am using Powerpoint is in presentation mode," etc.). The simple cost—benefit analysis feature also allows people to assign "caller priorities" to caller groups and define policies that weight the cost of interruptability with the caller priorities.

We plan to monitor closely the usage patterns of the basic cost—benefit analysis, and to review ongoing feedback and perform surveys on the functionalities to guide our continuing effort to make available the more sophisticated features to users. We intend to eventually make available tools that allow users to employ machine learning methods used in Bestcom-X to build more powerful models of interruptability.

VI. Bestcom Research Directions

We are currently exploring several challenges with extending core methods on the basic research versions of our system as well as continuing to evaluate the overall experience with Bestcom-ET. As sample challenges, we are exploring means for making best-means communication services easier to configure and manage. Ease of configuration, training, and overall set up would bring the power of the advanced technologies to general users at minimal cost. Some of the work on ease of initial configuration centers on the design of intuitive controls and transparent policies. In one approach, we found that users appreciated our adding to Bestcom-ET explanations about how each call was handled and why. In a next step, we will be overlaying user-interface conventions on top of such statements about the reasoning of the system to allow immediate, in place editing of policies.

In the basic research arena, we are pursuing the generalizability of models of interruptability. Demonstrations that models trained in the lab can provide useful inferences to multiple users—or useful inferences with just a small amount of personalization—would allow us to provide rich models of interruptability

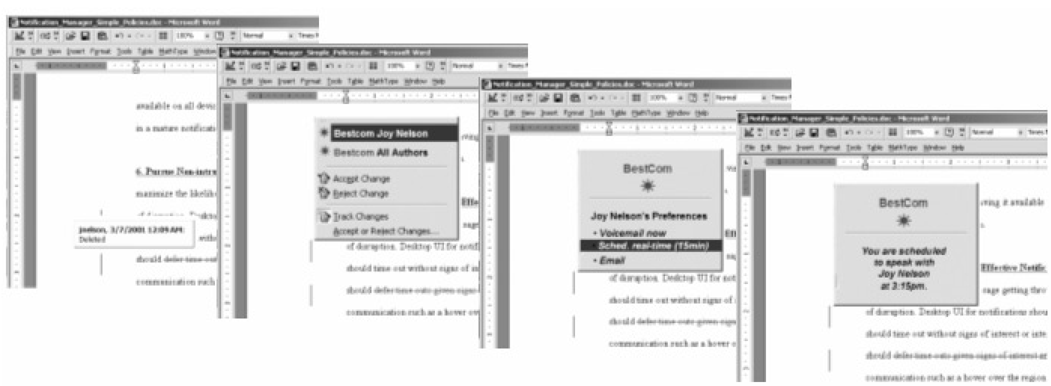

Figure 10. Bestcom service as embedded in MS Word. Here, best-means communication is invoked via a menu item. In this case, the Bestcom server reports back a list of recommended communications actions, sorted by the contactee's preferences.

to users without requiring imposing the burden of training such models. Also, we are interested in the loss of fidelity that comes with using the more approximate cost—benefit analyses. We shall be exploring the comparative performance of the different approaches.

We are continuing to explore the development of richer models that can support about the expected utility associated with the use of *different channels* for different kinds of communications and contexts. Such decision making is captured in decisions and random variables in the influence diagram in Fig. 1. This work includes study of the tradeoffs between delays, cost, and considerations of information loss with suboptimal channels that may be available earlier or at lower costs than more ideal channels.

Also, a significant thrust of our current work centers on the integration of Bestcom more deeply into desktop software applications. When meshed with computer applications, best-means communications can serve to enable users to initiate communications in the context of the usage of software applications, dragging key pointers and bits into a conversation. Automated decisions about the best timing and channel for a communication can be critical in such integrated communications.

An example of an early prototype of integrated best-means communication is displayed in Figure 10. Here a user reviewing a document in Microsoft Word has executed a right-click of their mouse on an edit recently entered by a co-author of a document. A menu has appeared that includes an option to "to Bestcom" the co-author to discuss the changes. By invoking the service, a best-means communication schema is transmitted to the Bestcom server, indicating the identity of the user, the current work context, and the place in the document that is at the user's focus of attention. Information is also transmitted about the user's device, relaying that the user currently has access to a full-display client.

The contactee's computed or accessed preferences are relayed to the user. As highlighted in the figure, the suggested options include *voicemail now*, *scheduling a real-time conversation in 15 minutes*, or *sending email*. The contactor chooses to schedule to speak with the contactee in 15 minutes and invokes the Bestcom rescheduling service. We believe that such scenarios will be commonplace one day and that context- and preference-centric best-means communications decision making will be a critical enabler.

VII. Summary

We introduced research on best-means communication decision making. We first described the challenge and promise of employing decision-theoretic reasoning to make optimal decisions about the timing and routing of communications. After reviewing the issues of principle agency and privacy of best-means communications decision making, we focused on Bestcom-X a research project exploring the use of machine learning to support cost—benefit analysis of the tradeoff between the cost of interruption and the cost of communication deferral. We described methods for building statistical models of the cost of interruption based on office activity and on a user's calendar information. We then described a descendant of Bestcom-X named Bestcom-ET that we view as a platform for continuing field work on best-means communication. Finally, we touched on active areas of research, including our efforts to integrate best-means communications services with software applications.

We are excited about the potential for enhancing interpersonal communication with automated and semi-automated decisions that work to maximize the value of communication to people. We believe that best-means communications services, implemented on a wide scale could fundamentally enhance the way people communicate with one another.

References

1. Bahl, P. and V.N. Padmanabhan. *RADAR: An In-Building RF-Based Location and Tracking System*. in *IEEE INFOCOM 2000*. 2000. Tel-Aviv, Israel.

2. Dey, A.K., Mankoff, J., Abowd, G.D., Carter, S.A. *Distributed mediation of ambiguous context in aware environments*, UIST 2002, Fifteenth Annual Symposium on User Interface Software and Technology, CHI Letters 4(2), pp. 121-130, October 27-30, 2002.

3. Horvitz, E., Breese, J., Heckerman, D., Hovel, D., Rommelse, D. The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users, *Fourteenth Conference on Uncertainty in Artificial Intelligence* (Madison WI, July 1998). Morgan Kaufmann Publishers, 256-265.

4. Horvitz, E. Koch, P., Kadie, C.M. Jacobs, A. Coordinate: Probabilistic Forecasting of Presence and Availability. *Proceedings of UAI '02, Proceedings of the Eighteenth Conference on Uncertainty and Artificial Intelligence*, Edmonton, Canada, July 2002, pp. 224-233.

5. Horvitz, E. and Apacible, J. Learning and Reasoning about Interruption, Proceedings of ICMI 2003, IEEE International Conference on Multimodal Interfaces, Vancouver, Canada, November 2003.

6. Hudson, J.M., Christensen, J., Kellogg, W.A., and Erikson, T.*' "I'd Be Overwhelmed, But It's Just One More Thing to Do*:" Availability and Interruption in Research Management.' Proceedings of CHI 2002, pp. 97-104.

7. McFarlane, D. Coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction - INTERACT'99*, IOS Press, Inc., The Netherlands, pp. 295-303, 1999.

8. Tang, J., Yankelovich, N., Begole,J., van Kleek, M., Li, F., Bhalodia, J. ConNexus to Awarenex: Extending awareness to mobile users, CHI 2001, Seattle, 2001.

9. Sawhney, N. and Schmandt, C. *Nomadic Radio: Scaleable and Contextual Notification for Wearable Audio Messaging*. In Proceedings of the 1999 Conference on Human Factors in Computing Systems (CHI '99), Pittsburgh, PA, May 15-20, 1999.

10. Toyama, E. and Horvitz, E.. Bayesian Modality Fusion: Probabilistic Integration of Multiple Vision Algorithms for Head Tracking. Proceedings of ACCV 2000, *Fourth Asian Conference on Computer Vision*, January 2000.