Rhythm Modeling, Visualizations and Applications

James "Bo" Begole, John C. Tang
Sun Microsystems Laboratories
2600 Casey Ave
Mountain View, CA 94043 USA
<Firstname>.<Lastname>@Sun.COM

Rosco Hill
University of Waterloo
200 University Ave. W.
Waterloo, Ontario, Canada N2L 3G1
rhill@engmail.uwaterloo.ca

ABSTRACT
People use their awareness of others' temporal patterns to plan work activities and communication. This paper presents algorithms for programatically detecting and modeling temporal patterns from a record of online presence data. We describe analytic and end-user visualizations of rhythmic patterns and the tradeoffs between them. We conducted a design study that explored the accuracy of the derived rhythm models compared to user perceptions, user preference among the visualization alternatives, and users' privacy preferences. We also present a prototype application based on the rhythm model that detects when a person is "away" for an extended period and predicts their return. We discuss the implications of this technology on the design of computer-mediated communication.

Keywords
Awareness, context-aware computing, rhythms, CSCW, user modeling, instant messaging, visualization, CMC.

INTRODUCTION
People exhibit temporal patterns, or rhythms, in their daily behavior such as when they typically arrive at the office, take breaks, attend recurring meetings, and commute to and from work sites. Researchers of cooperative work in co-located settings have found that coworkers share a sense of these patterns and they use their awareness of rhythms to coordinate their work activity and form expectations of availability [12, 18]. When coworkers are geographically distributed, however, it becomes difficult to form and maintain awareness of rhythmic patterns. Diminished awareness of coworkers' rhythms increases the cost of coordinating communication and work activities among geographically remote coworkers.

For example, when a physical office neighbor is away, we may have an idea of when she will likely return based on past behavior. However, when a remote coworker is away, it is difficult to form an idea of when he will likely return because we have had less information about his comings and goings over time which is needed to form an awareness of his rhythms. Currently, we might solve this by contacting another person who has close physical proximity to our remote coworker and ask if they know when he'll be back. This involves interrupting another person which some may not be comfortable doing. In addition, this may not succeed because our coworker's neighbors may also not be present, or our coworker may work from home or a satellite office among people who do not track his rhythms.

A partial solution can be found using awareness systems [5, 11, 14] which provide realtime information about remote coworkers' online presence. Such systems are increasingly popular in the form of commercial Instant Messaging (IM) such as AIM [2], Sun™ ONE Instant Messaging [13] and others. Over time, the information provided by such systems can help coworkers form a sense of each other's temporal patterns.

Another solution is based on observations of rhythmic patterns in the records of use of IM and other computer-mediated communication (CMC) technologies, as we reported in previous work [3]. These observations suggest a number of applications using computer inferencing of the rhythmic patterns. The applications are potentially useful for distributed coworkers who do not have a strong sense of each other's rhythms and may also be useful to coworkers who are newly introduced and have not yet had time to form awareness of each other's rhythms.

In this paper, we describe algorithms for detecting and modeling rhythmic patterns, visualizations of rhythms, user perceptions of their mental models of rhythms, and prototype applications. The next section describes the human-observable patterns in the data, desired applications, and requirements of a computer model to support those applications. We discuss related work in modeling user behavior and the extent to which previous models meet the application requirements. We then describe our modeling technique. A number of end-user visualizations of rhythmic patterns are presented and compared. We describe a design study that examined the accuracy of the computer model against user perceptions of their rhythms, user preferences among the end-user visualizations and user concerns about privacy. We also describe initial prototype applications of the model based on information gathered in the design study. The paper ends with a discussion of the implications for other CMC technologies and privacy considerations of rhythm inferencing technologies.

BACKGROUND
In past work [3], we described human-observable patterns seen in the record of individuals' computer, email, instant messaging and phone activity collected using a research

awareness and communication system called **Awarenex** [14]. We observed a number of patterns that had implications for predicting when someone would likely be reachable on their computer. The following were the key predictive features we observed in the data.

- The shape of the distribution of arrival and departure times may have predictive power.
- There are recurring breaks in activity that do not appear in a person's calendar. Most people exhibit a lunch break and some people have other recurring, unscheduled breaks.
- Activity may occur throughout recurring scheduled appointments, implying the person tends to be reachable throughout that recurring appointment.
- Activity may differ from the scheduled start and stop times of a recurring appointment. For example, the pattern may show that the person tends to be active, and therefore reachable, earlier/later than the scheduled end/beginning of an appointment.
- Patterns of transitions between locations are sometimes evident. For example, a person may start their day by logging in at home and later traveling to the office.
- Patterns may differ according to location (e.g., office, home).
- Patterns may differ according to day of the week.

Other researchers have also observed a temporal component to an individual's reachability and availability. J. Hudson *et al.* [9] found that research managers' receptiveness to interruption varied regularly with the time of day. Horvitz *et al.* [8] also report variations in patterns of computer activity at different times of day. Tyler and Tang [17] observed that email correspondents maintain a rhythmic pace in their email exchanges and that an awareness of typical pacing is among the factors used to judge when a response breakdown has occurred.

Rhythm Model Requirements

From the observations of rhythmic patterns previously described, it is evident that humans can detect patterns in the record of activity data, suggesting some degree of predictability to a person's online reachability. We were interested to design an algorithm by which a *computer* could detect patterns and make comparably meaningful inferences about a person's future state of reachability.

A computer model of rhythms could support a number of applications. First, a representation of remote coworkers' overall rhythms may help form expectations of when and where they can reach each other. Another application is to automatically set the “away” status of a person in their instant messaging client during recurring periods of inactivity such as lunch, commute, regular breaks and when they leave for the day. Rhythm data can also be used to predict when they are likely to return from such an absence. Rhythm information can supplement calendar information by providing accurate information about reachability during scheduled appointments. Rhythm information along with scheduled appointments can be used to find periods of overlapping availability among a number of people. Another application is to predict an individual's likely location at the time of a package delivery. For example, if a package is scheduled to arrive on a Wednesday and the recipient usually works from home on Wednesdays, the package could be routed to her home.

These applications impose a number of requirements on the computational rhythm model. The model should

- Predict probability of reachability throughout the day
- Describe the temporal variations in reachability so that end-user representations can be constructed
- Detect patterns within specific days of the week and locations as well as patterns across multiple or unspecified days and locations
- Identify significant recurring periods of inactivity such as lunch, regular breaks, end of the day
- Identify recurring transitions between locations
- Describe the range and distribution of start and end times of recurring features

Related Modeling Approaches

A number of approaches to modeling temporal patterns have been investigated by past researchers and are described in this section. The past approaches do not support all of the application requirements identified above.

One approach to modeling temporal variation is using time-series analysis techniques such as spectral analysis and auto-regressive integrated moving average (ARIMA) models [4]. Generally, this kind of analysis is applicable for problems where the values over time appear to exhibit a pattern of non-random behavior, but the underlying causative processes are unknown or are too complex to model directly. A common example is stock price fluctuations. We explored using an ARIMA model with limited success. In the end, we found ARIMA modeling to be unsuitable for our purposes because, although it satisfies the first three requirements, it does not meet the last three.

Other researchers have previously explored computer models for related applications. The **Priorities** system described by Horvitz *et al.* [8] contains a presence-forecasting subsystem to infer the likelihood that someone is away now and, if so, route alerts to a mobile device. The probability distributions used in the prediction are based on the record of past presence categorized by special regions of the day: morning, lunch, afternoon, evening and night.

Coordinate [8] and **Augur** [16] used Bayesian networks and Decision trees that include time of day along with other factors to construct models that predict the likelihood of attending a meeting. S. Hudson *et al.* [10] also employed these and other machine-inferencing techniques to detect a person's level of interruptibility from simulated sensor input and time of day.

Our applications require a model that is both *predictive* and *descriptive* of the temporal patterns. Bayesian networks and Decision trees are predictive, being able to answer queries based on the state of input parameters, and they are descriptive in the sense that they describe the network of dependencies among factors that contribute to a prediction. However, they do not describe the changes in a factor's influence over time, such as the timing of a recurring appointment, commutes, lunch, and recurring breaks. Our model differs in that it describes temporal patterns in a way that supports the applications we identified previously. In addition, this temporally descriptive model allows humans to assess how much credence to attribute to the model by comparing it to their own perception of a person's rhythm. It can also augment a person's mental model of rhythms, enabling her to make inferences of her own, rather than relying solely on the result of an opaque machine decision.

Another distinction is that while the models employed by Horivtz *et al*. [8] and S. Hudson *et al*. [10] are referred to as “predictive,” they are not predicting future state, rather they are detecting current state. Inferencing about current state does not meet the requirements of some rhythm applications, such as predicting when someone will return from an absence or searching for the best times to make contact.

Another difference is that our model is a custom modeling technique, specifically designed for detecting and modeling temporal patterns. Bayesian networks and Decision trees are general models that can be applied to a variety of tasks. Such models can be constructed automatically with machine-learning techniques by feeding in a record of inputs and known correct responses. While our model is more specialized, it is generalized in three senses: it is user-independent, it is constructed (“learns”) from a record of values, and it can analyze records of a variety of sources: computer, email or telephone activity, presence sensors, online calendar, and other sources.

RHYTHM DETECTION AND MODELING

The data structure of a rhythm model is a container of *features,* which are regions of time that identify significant occurrences in a pattern. Examples of features are start- and end-of-day, lunch, recurring breaks, recurring appointments, and transitions between locations. A feature consists of a label, the frequency of occurrence, and the probability distributions of the start and end time and duration.

Many rhythmic patterns did not conform to common parametric statistical distributions, such as Gaussian, Poisson, etc. For example, the distribution in Figure 1 has peaks at approximately 20 minute intervals as a result of being constrained by a mass transit schedule [3]. Therefore, our model does not assume parametric distributions.

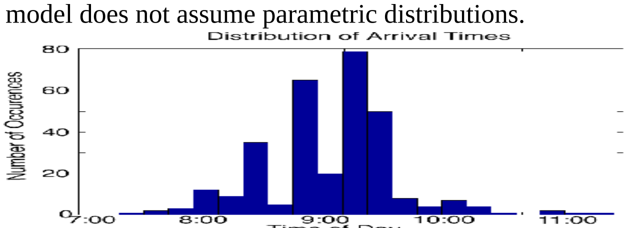

Figure 1. An example start-of-day distribution that does not conform to a parametric distribution.

Another goal of our modeling technique is to minimize *a priori* knowledge of the structure of a person's day. Beyond the coarsest generality that people tend to start activity in the mornings and end in the evenings, peoples' patterns vary widely. Even a common feature like “lunch” does not show up in all people's data. Because it is impossible to enumerate all of the rhythmic patterns we should look for, we discover features dynamically.

Our model building algorithm has three main steps: feature discovery, classification and refinement. The first step is to detect candidate features by discovering significant dips in the activity level. The next step is to classify instances of inactivity according to which feature they belong to, if any. The final step is to refine each potential feature by calculating the start, end and duration of the feature from the instances in that feature's class. Details of the model building steps are described next.

Feature Discovery

The first step in building the model is to detect points in time that potentially designate a significant rhythmic feature. We are looking for two types of pattern.

1. Recurring transitions between locations. These are easily detected by noting cases where an inactivity period begins in one location and ends in a different location.
2. Recurring periods of inactivity. These occur due to recurring meetings, lunch and other regular, scheduled or unscheduled breaks. Detecting this type of feature involves identifying significant dips in the person's aggregate activity level.

To discover recurring periods of inactivity, we first filter the historical data according to factors that significantly influence rhythm: day of week, and location [3]. We next aggregate the record by calculating the percentage of time the person was active at each minute of the day, weighted by recency of the activity. Figure 2 shows an example of one individual’s percent-active levels throughout the day on his office computer on Mondays, along with segmentation and threshold values described next.

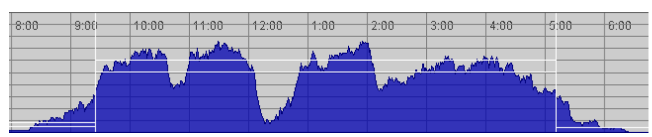

Figure 2. Graph of percent-active levels for Mondays in the office. Candidate features occur where the percent-active level crosses the thresholds.

To detect significant periods of inactivity, we calculate thresholds which identify candidate features as the points where the percent-active level crosses the threshold values. Upper and lower bounds are computed for threshold values to minimize the mis-identification of spurious features from small changes in activity due to natural variance (noise). When activity is declining, it must cross the lower threshold and when activity is rising, it must cross the upper threshold. We determine the optimal upper and lower thresholds by checking all possible positions and applying a weighting function which favors longer spans between threshold crossings to short, potentially spurious spans.

Before calculating threshold values, we segment the day into regions of sustained (more than four hours) “high” and “low” activity. Different threshold values are calculated within each segment. This alleviates cases where long spans of sustained “low” activity confounded the threshold optimization. We use a heuristic based on the distribution of percent-active values to determine “low” versus “high” values. Optimal threshold values are calculated for each segment and candidate features that are detected in a low region are ignored. Figure 2 shows an example of segmentation (vertical white lines at approx. 9:20 and 5:10) and thresholds within the segments (horizontal white lines).

Classification

The next step is to associate instances of inactivity with the feature classes. This classification is determined by how “close” an instance of inactivity is to a feature. We use the l_2-norm, or Euclidean distance, function shown in Equation 1 to compare time periods according to their start, end and duration. A lower value indicates greater similarity. Research has shown that Euclidean distance is both normative and descriptive of human cognitive processes involving non-obvious objects of comparison [6].

$$d(p_1, p_2) = \sqrt{\left(\frac{\Delta\, start}{\hat{\sigma}_{start}}\right)^2 + \left(\frac{\Delta\, end}{\hat{\sigma}_{end}}\right)^2 + \left(\frac{\Delta\, duration}{\hat{\sigma}_{duration}}\right)^2}$$

Equation 1. Metric to determine the distance, or similarity, between two time periods.

The numerators, $\Delta start$, Δend and $\Delta duration$, are the absolute value of the difference between the properties of the two periods being compared. In building the model, p_1 is defined by the estimated values of the feature based on the threshold crossing and is compared to each instance of an inactivity period in the data set, which are treated as p_2.

Each dimension in the space is normalized by weighting the parameter so that large differences significantly affect similarity while small differences have little affect. The distance function uses $\hat{\sigma}_{start}$, $\hat{\sigma}_{end}$ and $\hat{\sigma}_{duration}$ which are the estimated standard deviations of the start and end times and duration of the rhythm feature. Although we do not assume a normal distribution, standard deviation provides a useful measure of variance with which to normalize the numerators of each term. Estimates of standard deviation are bootstrapped using an initial value of one third the duration as detected in the threshold-crossing step. The refinement pass is iterated until the proportional difference between the estimate and the refined value converges or a maximum number of iterations is exceeded.

We use a maximum distance of three to determine whether a time period is considered an instance of a feature. Three allows all of a period's properties to lie within one normalized unit from the rhythm feature, or any two, but not all three, to lie between one and two normalized units. For example, a period that starts within one standard deviation of the expected start time and is within two standard deviations of the end time and duration, would have a distance less than three and be considered an instance of the feature. In contrast, another period which has start, end and duration greater than one standard deviation away from the feature would not be considered an instance of the feature.

Feature Refinement

The next step is to refine the estimates of the start, end and durations of the potential features. From the set of instances in the feature's class, we calculate statistics and determine the probability distribution for each property of the feature: start, end and duration. We also calculate the *occurrence frequency* of the feature, which is the percent of days in which we detect an instance of the feature. We use a simple algorithm to name features. The feature closest to a canonical lunch period is named “Lunch.” A feature that corresponds to a recurring appointment is named after that appointment. A location-transition feature is named after

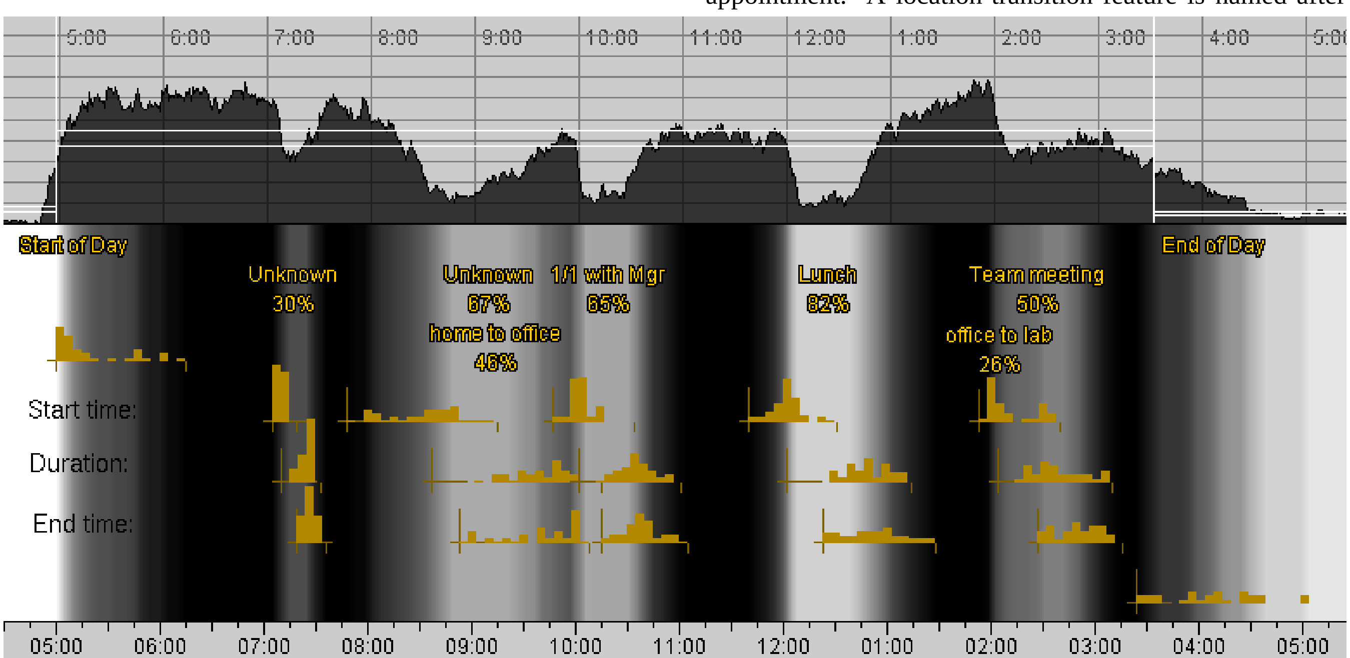

Figure 3. Example rhythm model. Percent-active levels in the upper portion of the graph identify rhythmic features shown in the lower portion along with the occurrence frequency and probability distributions for start time, duration and end time of each feature.

the locations in which it begins and ends, such as "Office to Home." Other features are labeled "Unknown."

The lower portion of Figure 3 shows an example of the the rhythm model extracted for one individual's Mondays across all locations. The probability distributions for the start, duration and end are drawn below the label of each feature and on different rows to avoid overlap. The distributions use a bin size of 5 minutes. Midday features are referred to as *two-sided* because they have a start and end time, while start-of-day and end-of-day are *one-sided,* having either a start or end time but not both.

In the example shown in Figure 3, the daily activity regularly begins around 5:00am. There is an unscheduled recurring break that is detected between 7:00 and 7:30, generally lasting less than 20 minutes, on 30% of the days. There is another unscheduled break that is detected 67% of the time between 8:00 and 10:00 which lasts between 45 and 90 minutes. It corresponds to a location transition feature "home to office." There is a recurring meeting "1/1 with Mgr" that is detected 65% of the time which starts promptly at 10:00, lasting around 30 minutes. A lunch break is detected 82% of the time between about 11:40 and 1:15, generally lasting between 30 and 75 minutes. The "Team meeting" is detected 50% of the time, promptly starting at 2:00 lasting between 20 and 75 minutes. This feature corresponds to a location transition from "office to lab." The day ends with equal probability anywhere between 3:30 and 4:30 and occasionally later than that.

VISUALIZATIONS OF RHYTHMS

We were concerned that the analytic visualizations we had used to in past work [3] to examine rhythm information might be too complex for end-users to interpret. To address this, we designed a number of visualizations, each having strengths and weaknesses described in this section.

Percent-active graph

The first visualization we considered was the aggregation of multiple days of activity seen in Figure 4. For each minute of the day along the horizontal axis, the height of the graph represents the percentage of days that this person was active at that minute. This visualization is useful because significant dips in the activity level suggest typical periods when the person is not likely to be available. By presenting the information to users, they can interpret what is a "significant dip" for themselves.

Figure 4. Percent-active graph. The vertical height represents the percent of times this person was active at each minute of the day.

A disadvantage is that the 'V' shape of dips conveys misleading impressions of the variability and correlations of a feature's start, end and duration. For example, consider the dip between 8:00 and 10:00. At the top of the dip, there is a span of approximately 2 hours and at the bottom the span is as narrow as 20 minutes. The slope between these end points is nearly linear, suggesting that the duration of this break varies with equal probability between 20 minutes and 2 hours. Looking at the same break in the model of Figure 3, we can see that the duration is actually between 45 and 90 minutes, most often in the later range. Furthermore, the 'V' shape may misleadingly suggest that when the feature starts early, it lasts longer and when it starts late, it takes a shorter time.

Color Saturation Gradient

An alternate visualization of the same data is shown in Figure 5. Here, the percentage of activity is represented by varying the level of color saturation. Because they represent the same data, the gradient visualization has many of the same disadvantages as the percent-active graph. However, the gradient is vertically more compact and arguably more aesthetically pleasing.

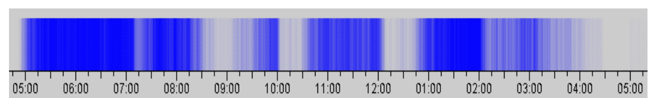

Figure 5. Gradient shading of activity levels. This displays the same data as in Figure 4 varying the color saturation according to the activity level.

The major downsides of the gradient visualization are that users cannot perceive as many levels in the color shades as they can in the activity graph and that interpreting the shades is prone to inaccuracy resulting from perceptual effects of edge fluting and simultaneous contrast [15]. In the activity graph, however, it is easy to distinguish even small differences in level from one minute to the next by the height of the curve. For example, the upward trend in activity between 1:00 and 2:00 is obvious in Figure 4 but hidden in Figure 5.

This difference has advantages and disadvantages. The gradient serves to perceptually smooth the data without in fact manipulating the values. This helps de-emphasize the variation due to noise. It also hides unnecessary detail which may alleviate privacy concerns. However, in some cases, such details have predictive power which is lost if not detected by a user. Furthermore, while users may discern that there is a difference in color saturation at different levels of activity, the saturation does not convey the *amount* of difference between the levels as effectively as the graph. To partially compensate for this, we exaggerate the saturation levels in the visualization by exponentially weighting the activity values such that low values are pulled lower and high values higher.

Compressed Actogram

Another visualization is a compressed actogram as seen in Figure 6. Each row represents an actual day of activity, stacked chronologically with the most recent day at the bottom. In this example, we see activity in three locations: home (white), office (black) and lab (dark gray). Significant rhythmic features emerge where the background shows through in gaps of activity.

This visualization has a number of advantages over the others. First, it addresses the misleading nature of the 'V'-shaped dips in the activity graph. Also, by presenting actual instances of daily activity, users can see the actual variation in start times, end times and durations for themselves. Another advantage is that we can interleave location information rather than having to plot each location separately. Because the days are sorted by recency, this visualization inherently portrays recent changes in rhythm. For example, this person has recently started working in a different location (lab) in the afternoons. A disadvantage is that stepped shapes of a feature are not as evident in a compressed actogram as in the percent-active graph. For example, the stepped departure seen in Figure 2, would not be evident in a compressed actogram. Another disadvantage of this visualization is that it raises greater concern about privacy as it exposes details about activity on specific days. To alleviate the privacy concern, we provide an option to include random noise to obscure individual data points while maintaining the same aggregate values [1]. The main disadvantage is that this visualization is more complex and visually cluttered than the previous ones.

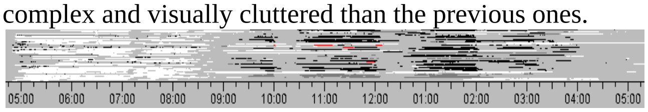

Figure 6. Compressed actogram. Each row represents a day of activity colored by location: home (white), office (black) and lab (dark gray).

Model gradient

Another visualization is a variation of the color-saturation gradient visualization but using the data of the rhythm model. An example of the gradient visualization of the rhythm model has previously been displayed in the lower portion of Figure 3. The color saturation varies according to the cumulative probability distribution of the start and end times of each feature. This portrays the shape of the probability distributions of start, end time and duration of a rhythmic features more precisely than the previous gradient. Note, for example, the steps in the shading of the 'End of Day' feature in Figure 3 which correspond to the peaks in the probability distribution of that feature. For features other than 'Start of Day' or 'End of Day', the shading plateaus at the occurrence frequency value of the feature, such that higher frequency features (e.g., “Lunch”) are whiter than lower frequency features (e.g., “1/1 with Mgr”).

Note that the model in Figure 3 is derived from the activity across all locations for Mondays. Figure 7 shows separate models derived from activity in the home (top) and office (bottom) locations for Mondays.

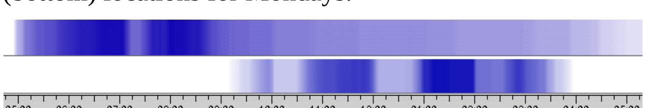

Figure 7. Rhythm model gradient for activity in home (top) and office (bottom) locations.

Shaped Ribbon

Another visualization of the model data can be seen in Figure 8. The edge of each feature is shaped by the cumulative probability of the start or end time of that feature and gives a sense of the range as well as the shape of the variance. For example, in Figure 8, the day typically begins between 5:00 and 5:45, and usually (~75% of the time) by 5:15. The typical duration of a feature is apparent from the gap between the start and end edges of the feature (e.g., the feature that starts near 7:00 has a typical duration of 15 minutes). The gap is shaded to convey the feature's occurrence frequency. Because the edges are parallel, this visualization has the advantage of avoiding overlaps in the start- and end-time distributions within a feature, which sometimes overlap in the gradient visualization However, when the times of two different features overlap, there is a collision in the visualization, such as the one around 10:00.

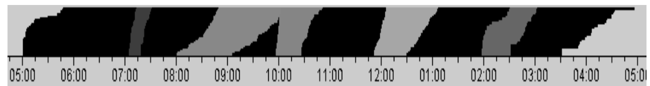

Figure 8. Shaped model ribbon. The edges of each feature are shaped by the cumulative probability distributions of the feature's start and end times.

While this visualization is appealing in its precision, it's complexity has disadvantages for end-user applications. First, interpreting the graph of a cumulative probability is not natural to many people. Second, the tails of the distribution are perceived more prominently than in the other visualizations, because of the narrowing shapes at the top and bottom edges. To get a sense of a person's typical pattern, one should look at the horizontal midrange within the ribbon, but one's eye is drawn more to the top and bottom which represent the extreme ends of the distribution. Even after truncating the tails of the probability distributions as much as 15% on each side and fading the top and bottom edge of the gradient, the prominence of the tails was distracting. Although we found this to be a useful analytic visualization, it requires training to interpret it correctly, which we wanted to avoid for end-user visualizations.

DESIGN STUDY

Having developed a model and visualizations of users' rhythmic patterns in their computer activity, we wanted to get user input on how they could interpret and use this information. We conducted a study to get end user feedback to guide the design of applications. The study explored

- Reactions to the visualization approaches
- Reflections on how well the model aligned with people's perceptions of their own rhythms
- Reactions to the privacy implications of sharing this information on rhythmic patterns with others

We could only interview a limited number of people, so the results are preliminary and reported as qualitative observations. Still, they are useful in guiding the design at this early stage.

Method

We interviewed current **Awarenex** users who are not closely involved in this rhythm modeling research and for whom we had at least three months of data. We conducted a structured interview with nine subjects. The interviews lasted approximately an hour and were audiotaped and analyzed to identify recurring issues. Anecdotally, we found it helpful to use the subjects' rhythm models in scheduling these interviews, as they helped suggest when interview slots would fit within the subjects' days.

During the interviews, subjects were presented four of the visualizations of an overall rhythm (Figures 4-7) together on one web page. The data in these visualizations depicted the activity of a person who was in a time zone that was three hours later, to suggest coordinating with a remote person in a different time zone. The subjects were prompted to compare and contrast among the visualizations to elicit what information they could get from each. They were also asked to depict their mental model of their "daily features," which they were then asked to compare with the computationally-generated rhythm models (e.g., Figures 3 and 7). Finally, having seen what information can be depicted in the rhythm models and how accurately the models portray their daily activity, the subjects were asked how comfortable they would feel sharing these rhythm visualizations with others.

Visualization Reactions

While there were quite a variety of responses to which visualization people preferred, the percent-active graph (Figure 4) and the model gradient (Figure 7) emerged as favorites. The percent-active graph afforded precisely comparing levels at different times of the day, and some said that they could most clearly see features in this representation. When asked to draw their own model of their daily rhythm, six of the nine subjects drew a representation like the percent-active graph. It should be noted that since almost all the subjects were members of a computer research lab, there might be a bias toward being familiar with this kind of graph. Others who do not share a scientific background may find this display to be overly complex.

Those who liked the modeled gradient liked that it distinguished which locale the activity occurred in. Knowing which locale the user was in provided important contextual information, such as whether it would be appropriate to call the user at 5:00 am at home. A few people would have preferred the percent-active graph if it had also shown activity according to locale. While some commented that the gradient did not allow making the fine distinctions in activity level that the percent-active graph offered, others mentioned that they did not need such detail for how they would imagine using these visualizations. Thus, they preferred the model gradient because it identified the important daily features without unnecessarily exposing more details and the associated privacy concerns.

Some found that the compressed actogram (Figure 6) conveyed the most information, especially since it could indicate changes in patterns over time. Yet most found the compressed actogram too complex to easily interpret. Since there was quite a range in preference for visualizations, perhaps this should be a setting that users could choose to match their personal preference and what kinds of information they need from the visualization.

Accuracy of Computational Model

As one way to compare the rhythm models with people's perceptions of their daily activity, we counted the number of rhythmic features (e.g., lunch, recurring meetings) that subjects depicted and compared it with the number of features that the model found. On average, subjects depicted 2.67 features per person per day, which means that in addition to lunch, they typically indicated 1-2 features per day. We found a matching feature in the model 79% of the time. Sometimes the model even prompted subjects to remember a feature that they had forgotten to depict. The median duration of these correctly identified features was 41 minutes and the median occurrence frequency was 45% (note that to be identified as a feature, it must have at least a 10% occurrence frequency).

On the other hand, the model detected several features that the subjects did not perceive, nor could they provide an explanation for them. On average, the model detected 2.77 features per person per day that the subjects did not depict. We expect the majority of these are spurious while a few may be features of which the person is unaware. This number of “false hits” varied depending on the person, showing that the model was more accurate for some people/job roles than others. The lowest average of false hits per day for a person was 0.8, and the highest average was 5.0. The median duration of these false hit features was 19 minutes, and the median occurrence frequency was 26%.

One of the main reasons for inaccuracies in the model is that it does not adequately account for changes in people's daily routines. Seven of the nine subjects had substantial changes in their routine recently (e.g., taking a class during the day, working from home for a couple of months, changing the day of a weekly group meeting). Given how likely these kinds of routine changes occur, the rhythm model would need to do a better job of accounting for these changes. Although it currently does weigh recent data more heavily, the lag before a new trend overwhelms an old one is too great. It will improve the model to detect a change in routine more quickly.

A somewhat common pattern that also introduced inaccuracy is regularly occurring meetings that occur less often than weekly. While the model did indicate a percentage of how often a feature occurs, the subjects did not readily interpret that value as corresponding to bi-weekly or monthly meetings. The effect of such non-weekly patterns could be more accurately represented if the model kept track of the periodic pattern of those series of meetings.

The subject's job role also affected the accuracy of the rhythm model. The rhythm models for those whose jobs involved a lot of computer keyboard work (e.g., programmer, administrative assistant) were more accurate than those with more interrupt-driven work (e.g., managers).

In addition to asking subjects to depict their own daily patterns, we also asked them to indicate if there were regions when they were especially open to interruptions or would rather not be disturbed. Seven of the nine subjects were able to identify such regions. According to this self-reported data, subjects indicated that they would rather not be disturbed at the beginning or end of the day or right after lunch, which is consistent with what J. Hudson *et al.* [9] found. Two subjects indicated they would rather not be disturbed during high productivity times (late afternoon, in their cases). Three subjects also indicated regions right before recurring meetings when they would rather not be disturbed, as they felt they were often needing to make last-minute preparations for the meeting.

Privacy

When asked about sharing their rhythm models with others, three had no reservations about doing so. Most, however, were concerned about wanting to provide more context to help explain and interpret the rhythm models, or would only want to share it with select people. Two subjects were uncomfortable sharing their rhythm models with anyone, although one of these subjects could think of another person for whom they would like to be able to see a rhythm model (because this person is involved in providing time-sensitive approvals).

APPLICATION: INFERRING "AWAY" STATUS

Several comments directed our attention to applications other than visualizing an overall daily rhythm. One subject mentioned that he was not interested in information beyond a couple hours into the future. If he needed to coordinate contacting someone further in the future than that, he would simply send email and coordinate asynchronously. Others also mentioned being more interested in when to expect someone to return from a period of inactivity than in an overall view of the person's day. Focusing more narrowly also addresses some privacy concerns, as providing a local prediction of when someone will return from being inactive would reveal much less information about a person's typical daily patterns.

These comments lead us in two directions. First, the rhythm visualizations can provide a more temporally localized view of a person's rhythm, rather than an entire day's view. We plan to redesign the rhythm visualizations in future work. Second, applications of rhythms should focus on reachability at the current time and in the *near* future.

Inferring Away Status

Our first application of the rhythm model is to infer situations when a currently inactive person will likely continue to be inactive, or "away," corresponding to a rhythmic feature such as lunch. The general approach for this inference is that when the person is inactive during the range of a feature, we calculate the likelihood that the current inactivity period will become an instance of the feature, based on how long they've been inactive so far.

The algorithm for this application uses both the static model of the person's rhythm, described previously, and also dynamically constructs a model of the specific feature using instances in the past that are similar to the person's current state. This dynamic modeling is guided by suggestions from the design study that users' questions should be answered from data that is local to the current time, as opposed to the entire day.

To explain the specific steps of the algorithm, consider the example illustrated in Figure 9. Here, the individual has been inactive within the range of their "lunch" feature, which was identified in their rhythm model. The current period of inactivity began at 12:15 and has continued for 10 minutes. We search the data set for past instances of inactivity that are at least as long as the current one and that started "near" the same time, which is the start of the current one plus or minus two standard deviations of the feature's start time. Next, we examine each prior instance to determine whether it is considered an instance of the lunch feature using the same criterion used in the feature refinement pass of the model's construction (i.e., the period has a Euclidean distance less than three units from the feature). The likelihood that the current inactive period will become an instance of lunch is the percentage of these past periods that were considered to be instances of lunch. In this example, 8 of the 12 periods that started around the same time and are at least as long as the current inactivity period were instances of lunch. This implies a 66% chance that the current inactivity period will be an instance of lunch. Once the probability exceeds half the feature's occurrence frequency, we label this inactivity as a detected feature.

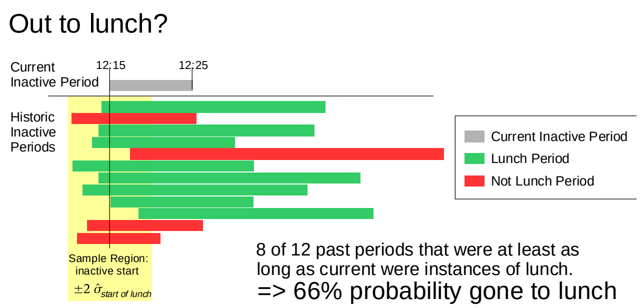

Figure 9. Example of determining the probability that current inactivity period is an instance of lunch.

Note that as the length of the current inactivity period grows, the total number of periods in the calculation will decrease, as periods shorter than the current one are excluded. This will initially increase the probability that the current inactivity is an instance of lunch, as the proportion of "true" lunch instances in the set grows. The likelihood will begin to decrease, however, when the current inactivity becomes longer than the "true" instances of lunch because the shorter "true" instances will be dropped from the set, decreasing the proportion of "true" lunches in the set. For example, after the next minute, the current inactivity period will be 11 minutes long and the bottom-most period which is 10 minutes long will be dropped. This raises the probability to 8 out of 11, or 73%. When there are fewer than a minimum number (5) of periods in the comparison set, we do not attempt a prediction.

Figure 10 shows a screenshot of the integration of status inferencing with **Awarenex**. In this example, the current time is 12:14 on a Thursday in the U.S. Pacific time zone. The first entry is a normal **Awarenex** entry indicating that Bo has been inactive for 50 minutes. This inactivity does not correspond to a rhythm feature – it may be due to the appointment he has scheduled for the current time, which is indicated by the clock icon. The second person on the list, John, has been inactive for 11 minutes. The system infers that he is at lunch (75% probability) and predicts that his return time (ETA) will be on or before 12:50. The third person in the list, Rosco, works in the U.S. Eastern time

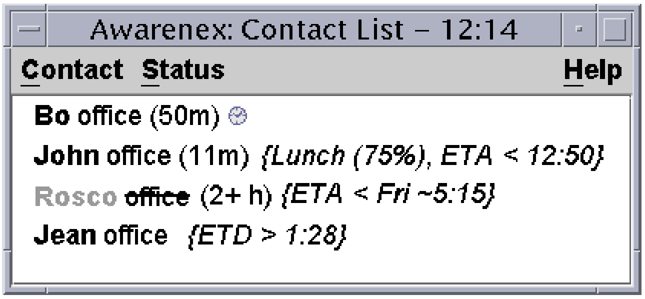

Figure 10. Screenshot illustrating integration of inferred status in the **Awarenex** contact list.

zone and has been logged out for more than two hours. The system predicts he will return on Friday around 5:15 (8:15 eastern). In the last entry, Jean, who also works in the U.S. Eastern time zone, is currently active and the system predicts that her departure (ETD) may be as early as 1:28 (4:28 eastern) which provides an indication of how long Jean is likely to remain available for those who may want to reach her before she leaves for the day.

DISCUSSION

Leveraging the *Computer* in CMC Technologies

Rhythm information complements real-time awareness by providing information about future availability when someone is not currently present. Such inferencing illustrates the potential for computer-mediated communication (CMC) technologies to provide functionality beyond traditional communication technologies. Many existing CMC technologies primarily re-create the functionality of traditional forms of communication. One example is that of Voice over Internet Protocol (VoIP) telephony products which primarily mimic the functionality of conventional telephony. Email also largely emulates the send and delivery model of conventional mail, though with a speed that approaches a synchronous medium. These CMC technologies often have conveniences (e.g. integration with address books or web access to voice mail) that provide an incremental enhancement over conventional technologies, but they continue to miss the opportunity to address the fundamental problem of finding a good time to contact someone.

The rhythm awareness research described here is one example that takes advantage of the *computer* in CMC technologies to gather and process data. Since much of our work activity involves using a computer, it can naturally capture some of the context of our work activity. This contextual information can help work colleagues plan and coordinate making contact. While conventional communication technologies and much CMC to date focuses primarily on providing a contact channel among people, rhythm inferencing and similar computation goes beyond that to provide the contextual information that people use to find good times to make contact.

Privacy

As with all systems that provide awareness information, this work raises the concern that making this kind of information available may encroach on individuals' privacy. We explored this question with participants of the design study. Several said they were comfortable sharing rhythm information with anyone, though most had some reservation. Those who were concerned primarily expressed that they would like to be able to interpret the pattern for people seeing it, to avoid projecting a negative impression. This mirrors how we deal with rhythm awareness among physically co-located colleagues where we are aware of who is aware of us and we can manage their interpretations by offering explanations for deviations from our typical rhythm. For example, if you were to start coming in later than usual, you might drop hints that you are also working later than usual, or some other appropriate explanation. Managing others' perceptions is more difficult when information is distributed electronically where, as Grudin notes [7], it does not necessarily come along with the context needed to interpret it appropriately.

Because privacy is socially negotiated in a dynamic and ongoing process, we don't believe a solution that relies solely on technology (e.g., strict access control) can completely solve the problem. However, there are steps that technologies can take to mitigate the concern, and we describe a few here as they specifically relate to rhythm information. One step is to store and present only the information necessary to answer a user's question. That is, the activity on individual days is not needed when users are interested in the aggregate pattern. An example is that the inferred status application shows only the result of the inference, rather than the data that leads to it. When daily details must be shown, as in the compressed actogram visualization, they can be obscured by introducing random noise that maintains the aggregate values.

Another step is to gradually expose details of a person's rhythm over time. It takes time to develop a sense of rhythmic awareness of neighboring workers, and that time also allows trust to build as the relationship forms. Similarly, users should control how much rhythmic information is being electronically conveyed to pace the amount of disclosure to be appropriate for the social relationship.

SUMMARY AND FUTURE WORK

We have presented an algorithm to build a model of the temporal patterns in a record of a person's computer activity and online presence. The model is designed to support a number of applications of rhythm information aimed at helping distributed teams coordinate communication and work activity. The model is both predictive and descriptive of temporal patterns, allowing users to augment their own mental model for making their own inferences, as opposed to relying solely on a computer inference. We described and compared a number of visualizations based on values from aggregate activity and the derived model.

Results of a design study found that subjects preferred the activity-level graph and model gradient to the other visualizations and preferred the visualization to show location information. Those subjects who were most concerned about sharing rhythm information expressed concerns about the information being taken out of context and that they wanted to be able to interpret the patterns for others.

Regarding accuracy, the preliminary results from the design study indicate that the model's accuracy is promising along with a number of ways it can improve. While we do plan to enhance the model in several ways, what is not clear is what level of accuracy is needed to be useful. For example, “false hits” in the model may be interpreted as bad times to reach someone. However, the “false hits” have low occurrence-frequency values relative to correct hits which suggests we may be able to exclude many “false hits” by raising the occurrence-frequency threshold for features that are considered part of the final model. Furthermore, as long as users do find times to make contact, incorrectly blocked off periods may not be perceived as greatly harmful. In addition, the way the model is used by the inferred-status application, the existence of a feature only identifies a region of *potential* unreachability which, in the case of a “false hit,” may not be realized. In future work, we will deploy the inferred-status prototype to **Awarenex** users to better understand what level and type of accuracy is required.

The subjects' responses also helped determine how the model could best be applied. Subjects were mainly interested in determining the person's availability in the near future, as opposed to over the entire day. Such local visibility into a person's rhythm also serves to minimize users' privacy exposure.

We described a prototype application that automatically infers when a person will be away for an extended period and provides a prediction of their expected return. In addition to using the static rhythm model, the inferencing module constructs a model of the specific rhythm feature of interest. The model is refined dynamically and the inference increases in confidence as the length of inactivity grows. This approach is partly guided by the design study subjects who emphasized interest in rhythms that are local to the current time.

The rhythm modeling and applications presented here complement real-time awareness systems, such as IM, by providing information about coworkers' future availability. The techniques may also be applied to other CMC technologies to help people find good times to make contact and support the overall aim of helping restore rhythm awareness among members of a distributed team.

ACKNOWLEDGEMENTS

We thank all of the participants of the design study for agreeing to let us collect and analyze their activity information and for providing insights on the use of this information. We also thank Randy Smith for consultations in the analysis of the rhythm information.

REFERENCES

1. R. Agrawal, R. Srikant: “Privacy-Preserving Data Mining,” *Proceedings of ACM-SIGMOD 2000*, Dallas, May 2000, pp. 439-450
2. AOL Instant Messenger (AIM), available at <http://www.aim.com/>
3. J. Begole, J. Tang, R. Smith and N. Yankelovich, “Work rhythms: Analyzing visualizations of awareness histories of distributed groups”, *Proceedings of CSCW 2002*, ACM Press, pp. 334-343.
4. C. Chatfield, *The Analysis of Time Series*, fifth edition, Chapman & Hall/CRC, Boca Raton, FL, 1996.
5. P. Dourish and S. Bly, “Portholes: Supporting Awareness in a Distributed Work Group,” *Proceedings of CHI 92*, Monterey, CA, May 1992, pp. 541-547.
6. W. Garner, *The processing of information and structure*, Wiley, New York, 1974.
7. J. Grudin, “Desituating Action: Digital representation of Context,” *Human-Computer Interaction*, Vol. 16, Nos. 2-4, 2001, pp. 269-286.
8. E. Horvitz, P. Koch, C. Kadie, and A. Jacobs, “Coordinate: Probabilistic Forecasting of Presence and Availability”, *Proceedings of the 2002 Conference on Uncertainty and Artificial Intelligence*, July 2002, AAAI Press, pp. 224-233.
9. J. Hudson, J. Christensen, W. Kellogg, T. Erickson, “‘I'd be overwhelmed, but it's just one more thing to do’: Availability and interruption in research management”, *Proceedings of CHI 2002,* ACM Press, pp. 97-104.
10. S. Hudson, J. Fogarty, C. Atkeson, J. Forlizzi, S. Kiesler, J. Lee and J. Yang, “Predicting Human Interruptibility with Sensors: A Wizard of Oz Feasibility Study”, *Proceedings of CHI 2003,* ACM Press, to appear.
11. E. Isaacs, J. Tang, and T. Morris, “Piazza: A desktop environment supporting impromptu and planned interactions,” *Proceedings of CSCW 96*, ACM Press, pp. 315-324.
12. M. Reddy and P. Dourish, “A Finger on the Pulse: Temporal Rhythms and Information Seeking in Medical Work,” *Proceedings of CSCW 2002,* ACM Press, pp. 344-353.
13. Sun™ ONE Instant Messaging, available at <http://wwws.sun.com/software/products/portal_icp/home_portal_icp.html>
14. J. Tang, N. Yankelovich, J. Begole, M. Van Kleek, F. Li and J. Bhalodia, “ConNexus to Awarenex: Extending awareness to mobile users”, *Proceedings of CHI 2001*, ACM Press, pp. 221- 228.
15. E. Tufte, *Envisioning Information*, Graphics Press, Cheshire, CT, 1990
16. J. Tullio, J. Goecks, E. Mynatt and D. Nguyen, “Augmenting Shared Personal Calendars,” *Proceedings of UIST 2002*, pp.11-20.
17. J. Tyler and J. Tang, “When Can I Expect an Email Response? A Study of Rhythms in Email Usage,” *Proceedings of ECSCW 2003*, in press.
18. E. Zerubavel, *Hidden Rhythms: Schedules and Calendars in Social Life*, Chicago: The University of Chicago Press, 1981.