

Rhythm Modeling, Visualizations and Applications

James “Bo” Begole, John C. Tang
Sun Microsystems Laboratories
2600 Casey Ave
Mountain View, CA 94043 USA
Bo.Begole@Sun.com, John.Tang@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 [14, 20]. 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. A partial solution can be found using awareness systems [6, 13, 16] which provide realtime information about remote coworkers' online presence. Such systems are increasingly popular in the form of In-

stant Messaging (IM) such as AIM [2], Sun™ ONE Instant Messaging [15] 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 temporal 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** [16]. We observed a number of temporal patterns in a person's activity.

- Most people exhibit regular arrival and departure times.
- Patterns may differ according to day of the week.
- Patterns may differ according to location.

- The shape of the distribution of arrival and departure times has predictive power in some cases (e.g., when arrival/departure is constrained by a transit schedule).
- There are recurring breaks in activity that do not appear in a person's calendar. Most exhibit a lunch break and some 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 regularly differ from the scheduled start and stop times of a recurring appointment (e.g., the person may be regularly active earlier/later than the scheduled end/beginning of an appointment).
- Some people exhibit patterns of transitions between locations. For example, a person may start their day by logging in at home and later traveling to the office.

These activity patterns have implications for predicting when someone would likely be *reachable* for communication: if they are active on a computer, they can be reached by IM or email; and if they are near a phone and not using it, they can be reached by phone. However, the use of activity data in determining someone's availability is limited in at least two ways. Firstly, one may be reachable while not interacting with a device. For example, one could be sitting near a computer while reading a printed document and therefore reachable although not active. Secondly, being physically reachable does not necessarily equate to being mentally receptive, or “available”, to communication.

Nevertheless, activity does imply reachability and may be useful enough for a substantial proportion of communications, as evidenced by the popularity of instant messaging systems where computer activity indicates “presence.” The users studied by S. Hudson *et al.* [12] considered themselves to be in some degree “interruptible” 68% of the time. This suggests that the recipient of a communication is receptive to interruption more often than not.

Other researchers have also observed a temporal component to an individual's reachability and availability. J. Hudson *et al.* [11] found that research managers' receptiveness to interruption varied regularly with the time of day. Horvitz *et al.* [10] also report variations in patterns of computer activity at different times of day. Tyler and Tang [19] 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

The observations of rhythmic patterns previously described suggest a number of applications. First, a representation of remote coworkers' overall rhythms may help form expectations of when and where coworkers 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 predicting reachability for

scheduled appointments during which the user is regularly active. 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 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 suggest a number of requirements for a computational model of rhythms. The model should

- Predict probability of reachability throughout the day
- Describe the temporal variations in reachability so that end-user representations can be constructed
- Model 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 transitions

Related Modeling Approaches

A number of approaches to modeling temporal patterns have been investigated by other researchers.

A common 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. An ARIMA model can characterize different levels of activity from one moment to the next, but it does not isolate or identify recurring regions of low activity.

Other researchers have previously explored computer models for similar applications. The **Priorities** system described by Horvitz *et al.* [10] contains a presence-forecasting subsystem to infer the likelihood that someone is away now and for how long. If they are likely to be away for a substantial amount of time, and the message is determined to be of a high-enough priority, the message is routed 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. Priorities does not provide an end-user representation of the computational model.

Coordinate [10] and **Augur** [18] 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. Bayesian networks and Decision trees are general models that can be applied to a variety of

inferencing tasks. Such models can be constructed automatically with machine-learning techniques by feeding in a record of inputs and known correct responses. S. Hudson *et al.* [12] 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 model is a custom technique, specifically designed for detecting and modeling temporal patterns. Although it is more specialized, our model is generalized in three senses: it is user-independent, it is constructed (“learns”) from a record of values, and it can analyze data from a variety of sources: computer, email or telephone activity, presence sensors, online calendar, and other sources.

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 depict the network of factors that contribute to a prediction, one of which may be time. However, temporal patterns are not apparent when examining such models. Our model differs in that the structure represents the temporal patterns of activity. A useful aspect of such a temporally descriptive model is to allow human observers to augment their own mental model of rhythms, enabling them to make inferences of their own, perhaps based on information not known to the computer, rather than relying solely on an opaque machine inference.

RHYTHM DETECTION AND MODELING

The structure of a rhythm model is a container of *transitions*, which are regions of time that identify significant changes in the pattern. There are three types of transition.

1. Recurring transitions between locations.
2. Start- and end-of-day transitions.
3. Recurring periods of inactivity due to meetings, lunch and other recurring, scheduled or unscheduled breaks.

The first two transition types are easily detected. Location-change transitions occur when an inactivity period begins in one location and ends in another. Start- and end-of-day transitions are detected as the first and last activity in the day. Start- and end-of-day are *one-sided* transitions and have a start or end time, respectively. Location-change transitions and inactivity transitions that occur between the start- and end-of-day (e.g., lunch) are considered *two-sided*, having a start time, an end time and a duration.

The transition data structure consists of a label, the frequency of occurrence, and the probability distributions of the start and/or end time and, for two-sided transitions, the duration. The properties of a transition may not conform to common parametric 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, the transition data structure does not assume parametric distributions and records the probability distributions minute-by-minute.

A 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




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

mornings and end in the evenings, peoples' patterns vary widely. Even a common transition 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 transitions dynamically by detecting points in time that potentially designate a significant change from one level of activity/reachability to another.

To detect and model two-sided transitions, we use an Expectation Maximization (EM) algorithm, first introduced by Dempster, Laird and Rubin [5] and summarized in Data Mining texts such as Dunham's [7]. Our implementation of EM has three main steps: transition discovery, clustering of similar inactivity periods and estimate refinement. The first step is to detect candidate transitions by discovering significant changes in the activity level. This step provides an initial estimate of the start, end and duration of the transition, seeding the clustering in the next step. The clustering step finds instances of inactivity periods in the record that are similar to each candidate transition. The final step is to refine the estimates for the start, end and duration of the transition from the instances found in the clustering step. We repeat the last two steps until the property estimates converge or a maximum number of iterations is exceeded. Details of each step are described next.

Two-sided Transition Discovery

To discover the two-sided 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.

Thresholds identify candidate two-sided transitions as the points where the percent-active level crosses the thresholds. Upper and lower thresholds are used to minimize the mis-identification of spurious transitions from small changes in activity due to natural variance. When activity is declining, it must cross the lower threshold and when activity is rising, it must cross the upper threshold. We determine the upper and lower thresholds by comparing all possible positions using a “penalty” function which favors longer spans between threshold crossings to shorter, potentially spurious spans. The upper and lower threshold positions with the lowest final penalty are used.

Before calculating threshold values, we segment the day into regions of sustained (more than two hours) “high” and “low” activity. Different threshold values are calculated within each segment. Segmentation alleviates cases such as that in Figure 2 where a long span of sustained “low” activity would otherwise confound the threshold level determination. Such cases arise when the person usually di-



Figure 2. Graph of percent-active levels for Mondays at home showing segment boundaries (vertical white lines) and thresholds (horizontal white lines). Candidate two-sided transitions occur where the percent-active level crosses the thresholds.

vides their work day between locations but occasionally stays in one of them.

We use a heuristic based on the distribution of percent-active values to determine “low” versus “high” segments. A low segment is a span of greater than two hours in which the aggregate activity never exceeds half of the 80th percentile value for activity levels.

Clustering

The next step is to associate instances of inactivity from the data set with the candidate transitions. This clustering is determined by how “close” an instance of inactivity is to a transition. 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. Although duration is a redundant measure, we include it to account for periods that have similar duration but are offset in start and end times. A lower value indicates greater similarity. Garner [8] found that Euclidean distance is both normative and descriptive of human cognitive processes involving non-obvious objects of comparison.

The numerators, $\Delta start$, Δend and $\Delta duration$, are the absolute value of the difference between the properties of the two periods being compared. p_1 is a two-sided transition whose property estimates are initially based on the threshold-crossing and refined later. It is compared to each instance of an inactivity period in the data set, treated as p_2 .

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

Equation 1. Distance metric to determine the similarity between two time periods.

Each term is normalized by the corresponding estimated standard deviation for that property of the transition, $\hat{\sigma}_{start}$, $\hat{\sigma}_{end}$ and $\hat{\sigma}_{duration}$. Although we do not assume a normal distribution, standard deviation provides a useful measure of variance with which to normalize each term. Estimates of standard deviation are bootstrapped using an initial value of one third the duration found in the previous step; the estimates are refined in the next step.

A maximum distance of three determines whether a period of inactivity is considered an instance of a transition. Three allows all three of the period's properties to lie within one standard deviation from the transition's estimated properties, or any two, but not all three, to be as much as two standard deviations away. 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 transition. In contrast, another period which has start, end and duration of two or greater standard deviations away from the transition would not be considered an instance of the transition. If any one property is more than three standard deviations away, then the distance will be greater than three and the period will not be associated with the transition.

Transition Property Refinement

The next step is to refine the estimates of the start, end and durations of the transitions. From the set of instances associated with the transition, we determine the probability distribution and recalculate the estimates for mean and median for each property of the transition. The model uses the median as the new estimated value of start, end and duration properties. The clustering and refinement steps are repeated until the initial and the refined estimates converge or a maximum number of iterations (ten) is exceeded.

We also calculate the *occurrence frequency* of the transition, which is the percent of days in which we detect an instance of the transition. In addition to indicating how often a transition occurs, the occurrence frequency is a measure of how significant a potential transition is. If we did not find many instances that match the potential transition, it will have a low occurrence frequency.

We use a simple algorithm to name transitions. The transition closest to a canonical lunch period (12:00-1:00), without being too far away, is named “Lunch.” Too far away is defined as more than twice the normal maximum distance (i.e., six). A transition that corresponds to a recurring appointment is named after that appointment, such as “Staff Meeting.” A location-change transition is named after the locations in which it begins and ends, such as “Office to Home.” Other transitions 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 transition and on different rows to avoid overlap. The distributions use a bin size of 5 minutes. In this example, the daily activity regularly begins around 5:00am. There is an unscheduled recurring break 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-change transition “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 transition 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.




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

VISUALIZATIONS OF RHYTHMS

While the visualizations we had developed so far were designed to help us analyze the data for rhythmic patterns [3], we were concerned that they might be too complex for end-users to interpret. We designed a number of visualizations to explore how to best convey rhythmic patterns to end users. The first three visualizations below are based on the raw data set and the last two are derived from the model. We were interested to explore whether users would prefer representations of the raw data including the irregularities due to variance or the modeled data which abstracts the salient patterns.

Percent-active graph

The first visualization we considered was the graph of percent-active levels used in the detection of two-sided transitions, seen previously in Figures 2 and 3 and repeated below in Figure 4 to simplify comparison against subsequent visualizations. 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.

One drawback to this visualization is that the 'V' shape of dips conveys a misleading impression of the variability and correlations of a transition's start, end and duration. For example, in the dip between 8:00 and 10:00 there is a span of approximately 2 hours at the top of the dip, but a span as narrow as 20 minutes at the bottom. The slope between



Figure 4. Percent-active graph.

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 actually varies between 45 and 90 minutes, most often in the later range. Furthermore, the 'V' shape may mistakenly suggest that when the transition 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. The gradient perceptually smoothes the data, which helps de-emphasize the variation due to noise. It also hides unnecessary detail which may alleviate privacy concerns. In some cases, however, such details have predictive power which is lost if not detected by a user. Furthermore, while users may discern 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 percent-active 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.

Although the gradient visualization is vertically more compact and perhaps more aesthetically pleasing than the percent-active graph, users cannot perceive as many levels in the color shades as they can in the percent-active graph,




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

and interpreting the shades is prone to inaccuracy due to perceptual effects of edge fluting and simultaneous contrast [17]. The percent-active graph affords distinguishing even small differences in level from one minute to the next. For example, the upward trend between 1:00 and 2:00 is obvious in Figure 4 but hidden in Figure 5.

Compressed Actogram

Another visualization is a compressed actogram of a limited span of data, shown in Figure 6. Each row represents an actual day of activity, stacked chronologically with the most recent day at the bottom. While similar to the actograms we presented in past work [3], here the activity is one pixel high and colored by location. In this example, we see activity in three locations: home (white), office (black) and lab (dark gray). Significant rhythmic transitions emerge where the background shows through in gaps of activity and where the location color changes.




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

This visualization has a number of advantages over the previous ones. First, it addresses the misleading nature of the 'V'-shaped dips in the percent-active 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 this visualization raises greater concern about privacy as it exposes details about activity on specific days. To alleviate this 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.

Model gradient

In contrast to visualizing the raw data, we explored two visualizations based on the model which abstract the salient patterns of a person's rhythm. Since the rhythm model identifies specific transitions of interest to the end user (i.e., when is a colleague likely to be active), we used it to create a variation of the color-saturation gradient. An example of the gradient visualization of the rhythm model was shown earlier 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 transition. For two-sided transitions, the shading plateaus at the occurrence frequency value of the transition, such that higher frequency transitions (e.g., "Lunch") are whiter than lower frequency transitions (e.g., "1/1 with Mgr").




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

While the model in Figure 3 combines activity across all locations for Mondays, Figure 7 separates activity in the home (top) and office (bottom) locations for Mondays.

Shaped Ribbon

Another visualization of the model data can be seen in Figure 8. The edge of each transition is shaped by the cumulative probability of the start or end time of that transition, giving a sense of the range and 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 transition is apparent from the gap between the start and end edges of the transition (e.g., the transition that starts near 7:00 has a typical duration of 15 minutes). The gap is shaded to convey the transition's occurrence frequency. Because the edges are parallel, this visualization avoids overlaps in the start- and end-time distributions within a transition, which sometimes overlap in the gradient visualizations. However, when the times of two different transitions overlap, there is a collision in the visualization, such as the one around 10:00.




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

While this visualization is appealing in its precision, its complexity has disadvantages for end-user applications. First, interpreting the graph of a cumulative probability is not intuitive to many people. Second, while it is the horizontal band near the middle that gives a sense of a person's typical pattern, the extreme ends of the distribution create tails at the top and bottom edges that can distort accurate perception of the patterns. In Figure 8, the tails were truncated by 15% on each side, but they still have a distracting perceptual effect. Although we found this to be a useful analytic visualization, it required too much training to interpret correctly for an end-user visualization.

DESIGN STUDY

We wanted to get user input on how they could interpret and use the model and visualizations of rhythmic activity patterns. We conducted a study to get end user feedback on the following and to guide the design of applications.

- 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 for 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, using the subjects' rhythm models helped suggest good times to propose for scheduling these interviews.

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. We also asked them to depict their mental model of their own "daily transitions," and to compare it with the computationally-generated rhythm models (e.g., Figure 7). Finally, having seen what information can be depicted in the rhythm models and how accurately the models portrayed 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 transitions 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. We had not expected this visualization to be favored, due to the noise and misleading characteristics of dips as described previously. It should be noted that since almost all the subjects were members of a computer research lab, there may be a bias toward familiarity 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 the locale where the activity occurred. 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 transitions 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

To compare the rhythm models with people's perceptions of their daily activity, we counted the number of two-sided transitions (e.g., lunch, recurring meetings) that subjects depicted and compared it with the number of transitions in the model. On average, subjects depicted 2.67 two-sided transitions per person per day, which means that in addition to lunch, they typically indicated 1-2 transitions per day. We found a matching transition in the model 79% of the time. Sometimes the model even prompted subjects to remember a transition that they had forgotten to depict. The median duration of the correctly identified transitions was 41 minutes and the median occurrence frequency was 45%.

On the other hand, the model detected transitions that the subjects did not perceive, nor could they provide an explanation for them. On average, the model detected 2.77 transitions per person per day that the subjects did not depict. The majority are likely spurious, while a few may be transitions of which the person is unaware. The 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 transitions 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 months, changing the day of a weekly group meeting). Although the model does weigh recent data more heavily, the lag before a new trend overtakes an old one is too great. Given how likely these kinds of routine changes occur, the rhythm model needs to more quickly detect these changes.

A somewhat common pattern that also introduced inaccuracy is regularly occurring meetings that occur less often than weekly. While the model does indicate the percentage of how often such a transition occurs (the occurrence frequency), 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 model. As expected, the 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).

These preliminary results indicate that the model's accuracy is promising along with a number of ways it can improve. "False hits" have low occurrence-frequency values relative to correct hits which suggests that a higher thresh-

old may exclude many “false hits.” While we do plan to enhance the model in this and other ways, it is not clear what level of accuracy is needed to be useful. For example, although “false hits” may be interpreted as bad times to reach someone, 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, described later, the existence of a transition only identifies a region when the person has the *potential* to be inactive. In the case of a “false hit,” that potential would not be realized.

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.* [11] 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 rhythms with anyone, although one of these could think of another person for whom they would like to be able to see a rhythm model (because that person is needed for time-sensitive approvals).

APPLICATIONS OF RHYTHM MODELING

Providing an end-user visualization is one application of the rhythm modeling work that allows viewers to make inferences about good times to coordinate with that person. In addition to that and the applications described previously in the discussion of “Rhythm Modeling Requirements,” several interview comments suggested other applications. One subject mentioned that he was only interested in information within a couple hours into the future. If he needed to coordinate contacting someone beyond that, he would simply send email and coordinate asynchronously. Others 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 addresses some privacy concerns, as providing a local prediction of when someone will return from being inactive reveals less information about a person's typical daily patterns. These comments suggested applications that focus on reachability at the current time and in the *near* future.

Inferring Away Status

One 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 two-sided transi-

tion such as lunch. The general approach for this inference is that when the person is inactive during the range of a two-sided transition, we calculate the likelihood that the current inactivity period will become an instance of the transition, based on how long they have 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 transition 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” transition, as identified in their rhythm model. The current period of inactivity began at 12:15 and has lasted 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. “Near” is considered as the start of the current inactive period plus or minus two standard deviations of the transition's start time. Next, we examine each prior instance to determine whether it is considered an instance of the lunch transition using the same criterion used in the transition clustering pass of the model's construction (i.e., the period has a Euclidean distance less than three units from the transition). The likelihood that the current inactive period will become an instance of lunch (i.e., the probability that the person is currently “at 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 transition's occurrence frequency, we consider the person to be in the transition (e.g., “at lunch”).

Out to lunch?




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. 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%. 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. When there are fewer than a minimum number (5) of periods in the comparison set, we do not attempt a prediction.

Predicting Times Around Rhythmic Transitions

Once the model identifies a person as being in a rhythmic transition, it can use the statistical descriptions of the transition's start, end, and duration to also predict the end time of that transition. For example, if an inactivity period is identified as lunch, the model can predict the return from lunch by taking the 80th percentile of lunch durations and adding that to when the lunch period started.

Figure 10 shows a screenshot of the integration of status inferencing and return time prediction 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 modeled transition but may be due to the appointment he has scheduled for the current time (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 he will return (ETA) on or before 12:50. The third person in the list, Rosco, works in the U.S. Eastern time 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). This provides information for those who may want to reach her before she leaves for the day.

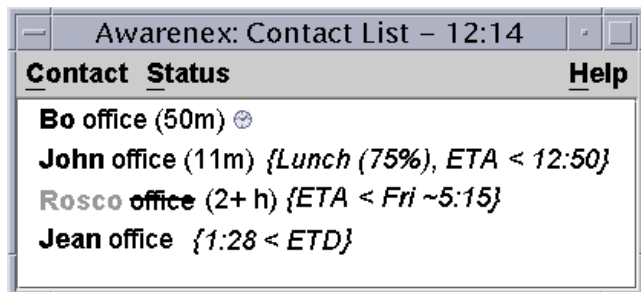


Figure 10. Screenshot illustrating integration of inferred status in the **Awarenex** contact list.

The functionality described here is similar to that of **Time-Wave** and **SmartOOF** described by Horvitz *et al.* [10], with the addition of predicting and presenting the approach to the end of a work day. In addition, our system integrates the functions in an awareness and multi-channel communication system, rather than calendar and email. Status inferencing fits naturally in an awareness system where the information is available for coordinating a communication before a contact is attempted rather than as an asynchronous response after an email message has been sent.

Coordination Applications

Rhythmic modeling could also help the coordination of activities among people. Comparing a model of your daily activities with someone else's model (along with any scheduled appointments) could help identify overlapping times of availability to try to make contact with that person. This would be especially useful between differing time zones where you may not have intuitions about the overall rhythms in that time zone.

Furthermore, models of a group of people could be compared to help identify candidate times for scheduling a group meeting. Especially if the participants span different time zones or temporal contexts (e.g., early bird, night owl), the models of rhythmic patterns could help identify overlap times for scheduling a meeting.

DISCUSSION

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 know who is aware of us and can manage their interpretations by offering explanations when we deviate from our expected rhythm. For example, if you started coming in later than usual, you might drop hints that you were also staying later, or another appropriate explanation. Managing others' perceptions is more difficult when information is distributed electronically where, as Grudin notes [9], 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 that 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 [1].

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.

Leveraging the Computer in CMC Technologies

The rhythm modeling and applications presented here complement real-time awareness systems, such as IM, by providing information about coworkers' future reachability. Such inferencing illustrates the potential for computer-mediated communication (CMC) technologies to go beyond traditional communication technologies. For example, typical use-case scenarios for Voice over Internet Protocol (VoIP) primarily mimic the functionality of conventional telephony, using computers for media codecs, routing and address directories. Recreating telephony on a data network provides few capabilities beyond what telephony already provides and misses the opportunity to take advantage of other ways that computing resources can address the fundamental problem of finding a good time to contact someone.

The rhythm awareness research leverages the computer in CMC technologies to gather, process and present context data. Since much of our work activity involves using a computer, it can naturally capture and analyze some of the context of our work activity. As demonstrated in previous awareness systems [6, 13, 16], this contextual information can help work colleagues plan and coordinate making contact. Rhythm inferencing extends context awareness to information about one's temporal context.

We have presented an algorithm for building a model of the temporal patterns in people's computer activities and online presence. The model is both predictive and descriptive of temporal patterns, allowing users to augment their own mental model for making their own inferences, rather than relying solely on a computer inference. Exploring what information can be gleaned from the temporal context enables new applications in supporting the coordination of distributed teams. One area we are currently exploring is using sensors to capture additional sources of context information (e.g., presence, audio activity) to augment the rhythm model. 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 and Nicholas Matsakis for analysis assistance.

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), <<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. A.P. Dempster, N.M. Laird, and D.B. Rubin, "Maximum Likelihood from Incomplete Data via the EM algorithm". *Journal of the Royal Statistical Society, Series B*, 39(1), 1977, pp. 1-38.
6. P. Dourish and S. Bly, "Portholes: Supporting Awareness in a Distributed Work Group," *Proceedings of CHI 92*, Monterey, CA, May 1992, pp. 541-547.
7. M.H. Dunham, *Data Mining: Introductory and Advanced Topics*, Prentice Hall, 2002.
8. W. Garner, *The Processing of Information and Structure*, Wiley, New York, 1974.
9. J. Grudin, "Desituating Action: Digital representation of Context," *Human-Computer Interaction*, Vol. 16, Nos. 2-4, 2001, pp. 269-286.
10. E. Horvitz, P. Koch, C. Kadie, and A. Jacobs, "Coordinate: Probabilistic Forecasting of Presence and Availability", *Proceedings of the Eighteenth Conference on Uncertainty and Artificial Intelligence*, Edmonton, Alberta, July 2002, Morgan Kaufmann, pp. 224-233.
11. 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.
12. 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, pp. 257-264.
13. 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.
14. 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.
15. Sun™ ONE Instant Messaging, <http://www.sun.com/software/products/instant_messaging/>
16. 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.
17. E. Tufte, *Envisioning Information*, Graphics Press, Cheshire, CT, 1990
18. J. Tullio, J. Goecks, E. Mynatt and D. Nguyen, "Augmenting Shared Personal Calendars," *Proceedings of UIST 2002*, pp.11-20.
19. 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.
20. E. Zerubavel, *Hidden Rhythms: Schedules and Calendars in Social Life*, Chicago: The University of Chicago Press, 1981.