MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

SocialMotion: Measuring the Hidden Social Life of a Building

Christopher Wren, Yuri Ivanov, Ishwinder Kaur, Darren Leigh, Jonathan Westhues

TR2007-034 September 2007

Abstract

In this paper we present an approach to analyzing the social behaviors that occur in a typical office space. We describe a system consisting of over 200 motion sensors connected in a wireless network observing a medium-sized office space populated with almost 100 people for a period of almost a year. We use a tracklet graph representation of the data in the sensor network, which allows us to efficiently evaluate gross patterns of office-wide social behavior of its occupants during expected seasonal changes in the workforce as well as unexpected social events that affect the entire population of the space. We present our experiments with a method based on Kullback-Leibler metric applied to the office activity modelled as a Markov process. Using this approach we detect gross deviations of short term office-wide behavior patterns from previous long-term patterns spanning various time intervals. We compare detected deviations to the company calendar and find and provide some quantitative analysis of the relative impact of those disruptions across a range of temporal scales. We also present a favorable comparison to results achieved by applying the same analysis to email logs.

LoCA 2007

MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

SocialMotion: Measuring the Hidden Social Life of a Building

Christopher R. Wren Yuri A. Ivanov
Ishwinder Kaur Darren Leigh Jonathan Westhues

TR2007-034 July 2007

Abstract

In this paper we present an approach to analyzing the social behaviors that occur in a typical office space. We describe a system consisting of over 200 motion sensors connected in a wireless network observing a medium-sized office space populated with almost 100 people for a period of almost a year. We use a *tracklet graph* representation of the data in the sensor network, which allows us to efficiently evaluate gross patterns of office-wide social behavior of its occupants during expected seasonal changes in the workforce as well as unexpected social events that affect the entire population of the space. We present our experiments with a method based on Kullback-Leibler metric applied to the office activity modelled as a Markov process. Using this approach we detect gross deviations of short term office-wide behavior patterns from previous long-term patterns spanning various time intervals. We compare detected deviations to the company calendar and find and provide some quantitative analysis of the relative impact of those disruptions across a range of temporal scales. We also present a favorable comparison to results achieved by applying the same analysis to email logs.

To Apeear in LoCA 2007

September 20, 2007: to appear in the Third International Symposium on Location- and Context-Awareness in Oberpfaffenhofen, Germany

1 Introduction

The social fabric of an organization is largely hidden from direct observation. It is not the same as the organizational structure. It is not completely defined by email or phone communication. The social fabric is the interpersonal connectivity in a group that is largely created and maintained by physical interactions in the space [3]. Measuring the structure and dynamic evolution of these connections is essential to understanding the health and productivity of an organization. At the same time, the details of these interactions are very sensitive from a personal privacy standpoint, so it is important to treat them with respect [16].

This social fabric is embedded within the architectural structure of a building. There will be social spaces such as lounges or kitchens. There may be specific associations between individuals and particular places in a building, such as offices. There are likely to be organizational functions that are concentrated in particular locations, such as departments and work groups. We are motivated by these facts to take a building-centered approach to measuring the social fabric. By measuring the way the social fabric drapes over the physical structure of the space we indirectly measure the structure of the social fabric itself.

This space-centric perspective has several advantages. By trying to understand the way the building is being used we can avoid the need to instrument or catalog each person in the space. This is a significant advantage in cost, reliability, and the impact on privacy. One of the insights that we gained during the experiments with the data that we collected is that in order to find events that disrupt the established patterns of social interactions, we do not need to look at every individual participant of the daily routine of the office space, but rather at how all together the inhabitants behave differently in response to a socially disruptive event. This implies that all analysis can be done *en masse*, without resorting to finding patterns in behavior of each individual, nor establishing his/her identity. This approach also makes it possible to fluidly handle visitors and other transient populations.

We have installed a 200 node sensor network in a corporate facility and have used it to collect the data for the period of over 10 consecutive months. Each sensor node has the ability to sense motion when a moving object appears in its field of view. We consciously tried to avoid using video cameras in order to accommodate the sense of comfort among the people populating the office. This on one hand made the problem of data collection easier, as the occupants felt that their privacy was respected, but on the other hand, made the task of data analysis harder, as no ground truth can be derived from the sensor data.

However, it should be pointed out that even with the use of the video cameras, it would be nearly impossible to obtain ground truth for the data on such scale. Video storage issues aside, a simple review of the recorded videos could take a very long time. In view of this we opted for using other means for analysis - company calendar and email logs. This approach inevitably leads to questions of quantitative validity of the analysis. In the light of this argument we would like to emphasize that it is not a classification approach that we present, but rather a ranking system. We attempt to explain found irregularities in the behavior patterns the best we can, but it is nearly impossible to frame the experiments in a purely quantitative fashion for lack of exact transcription of a year worth of movements of 100 people.

Using a tracklet graph representation that we have developed in our earlier work [10] this paper focuses on the statistics calculated over a large population of these graphs. The graphs record the possible destinations for a person leaving his/her office. Aggregating these graphs over extended periods of time allows us to estimate Markov models of the connectivity of the space. The Kullback-Leibler Divergence (KLD) is then used to compare models from different time periods. This technique enables analysis of the time-varying patterns of social behavior of the entire organization.

During the ten months of observation the organization experienced several disruptive events: the arrival and departure of a large transient worker population, a change in the senior management structure within the organization, a change in management in the parent company, and other events such as holidays, etc. We will present evidence that the sensitivity of this building-centered social analysis to disruptions compares favorably to a classical email-based approach.

2 Background

Social relationships in organizations have been studied for quite some time [18]. The methods developed for the analysis of these relationships have ranged from conducting individual surveys and handing out questionnaires to mining data in online communities and using wearable sensors to monitor the social activities of the people. The analysis of the social dynamics and understanding roles of individuals in social networks [20] may help organizations to improve their productivity by leveraging the groups naturally forming among their staff.

In the recent years, data gathered from the company email communications has been used to identify social networks in an organization [21, 22]. The success of this approach is largely due to a directed and unambiguous nature of point-to-point email communications, as well as to the convenience of the centralized storage of this data. When compared to traditional surveying techniques this approach is both cheaper and less prone to human reporting error. Though most such analyses use the `To` and `From` data fields in the email to form directed graphs of social relationships, some scholars have gone further and included the content and topic of message to allow determination of expertise and social roles of the correspondents [12]. More detailed analyses are made possible by fusing of informationfrom multiple data sources, such as web pages [8].

Recently swarms of simple sensors have been used as perceptual tools for living and office spaces. Much of this work has its focus on prediction of motion [2] and monitoring activity [23, 14, 1]. That body of work so far has avoided the question of social network analysis.

In another very different approach, wearable electronic sensors have been used to gather information about social networks in organizations [9, 6]. In order to gather information about verbal interactions and the nature of the conversational dynamics in social networks, people in the organization wore sensors on their bodies over extended periods of time. The goal of this analysis was to understand the social network in an organization and identify the people informally entrusted with leadership roles in various groups and subgroups. The authors found that conversational cues such as turn-taking pattern strongly correlates with the speakers' roles in the organization.

Up until recently the research in this area has been focusing on the roles of individuals, viewing social networks from the perspective of the individual participants. By necessity such analysis can be privacy invasive. All measurements that carry identity or cognitive information such as emails or voice recordings can be viewed as compromising individual's privacy and are subject to strict regulations and limitations on the scale of deployment. In contrast, measuring less descriptive information in a process that allows a person to preserve his or her anonymity enables massive deployments, that can be rich sources of information about social dynamics without having to sacrifice significant individual privacy. This data can be freely distributed, analyzed and reported on. An extensive discussion of privacy is beyond the scope of this paper hoeever. We invite th interested ereader to see our prior work in this area [17].

The social network of any organization is not a fixed immutable feature to be discovered through a one-time effort. The network, like the community that it describes is alive and changes over a period in time. It is affected by the influx of workers, by top down management changes or directives and in response to physical factors such as building space reallocations. This paper is a step towards developing methods that can be used to identify and quantify the effects of such interruptions and activities on social relationships between people. Our hope is that this information could help planning such moves more effectively to minimize their disruptive impact.

3 Models

A gross estimate of the level of activity in a space can be cheaply and easily obtained from a collection of motion sensors. Even the most basic such measurements will reveal the patterns of usage in a space. This point is clearly illustrated in Figure 1, which shows the total activity observed in a space over many months. It can be seen that there is a dominant weekly pattern that governs the behavior of this group of people. It is also possible to see the added activity contributed by a pulse of seasonal workers from approximately June 1 through August 31. Finally, there is a clear depression of activity around the winter holidays.

While these are interesting results, they are very coarse. There are a number of more subtle, though no less significant events that have been missed. For example, the plot in Figure 2 shows the number of people who are out of the office, as recorded in the company calendar. It is easy to see that there are

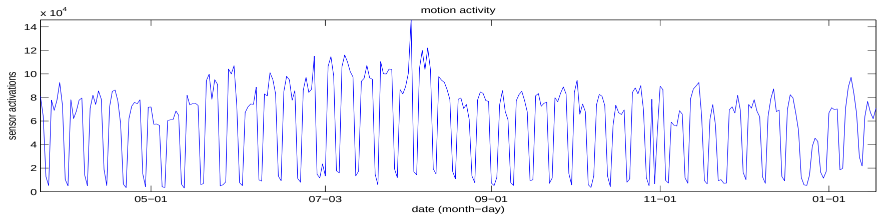

Figure 1: Gross activity in the experimental space, as measured by the number of raw motion sensor activations.

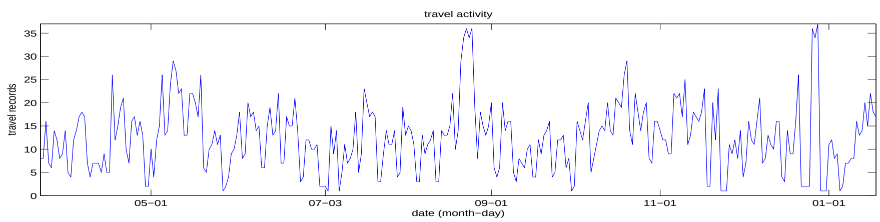

Figure 2: A plot showing the number of people away from the facility on a particular day.

periods of time in late August and late December when many people are away from the office (the two primary peaks in that plot). The August travel pulse, with a third of the people away, is not represented in the simple motion analysis of Figure 1 . The absences are masked by some other activity in the space.

Capturing more subtle events in the space requires a more sensitive model. In this section we focus on modeling the structure of the interconnections between parts of the space. To do this we extract the trips that people take in the space and model how those trips connect the various parts of the space.

It is not possible, with the very little information provided by motion detectors, to flawlessly extract the tracks of individuals in the space. To compensate for this perceptual shortcoming we instead use a representation that encodes the inherent ambiguities in the data efficiently. We briefly review a method introduced in [10] that allows us to extract a structure called a *tracklet graph*.

One of the main contributions of this paper is in embedding these tracklet graphs into a new probabilistic framework. This embedding results in a Markov chain model that describes the probability of seeing a trip between any two points in the space, even if we never see a single, unambiguous example of such trip. This model captures the social fabric of the space.

Finally we end the section with a description of the Kullback-Leibler Divergence (KLD) formulation for Markov chains that we use to detect changes in that social fabric over time.

3.1 Tracking

Classical tracking accommodates ambiguity by maintaining a set of hypotheses that describe every consistent explanation of the data. This is referred to as multiple hypothesis testing [4]. The problem with approaches similar to that is that with each new ambiguity the hypothesis set may grow exponentially. The goal of the tracking algorithm is to use new observations to prune away the ambiguities and arrive at a single, true explanation for the data.

Most of the tracking literature focuses on the use of high-quality sensors such as cameras [13]. Since we only use motion detectors we cannot prune our hypotheses: all motion activations look the same. Unlike a camera-based system, we cannot prune hypotheses by noticing a blue shirt, or a particularly tall person, or recognizing a face. All the data collected by the system is therefore inherently ambiguous about the person's identity. A person may emerge from a particular office, but there is no way to tell who that person is, or even that it is in fact a single person. And when that person or group of people passes other people in the hallway there is no way to tell with any certainly who went where.

Conversely, we *do* need to know where people go in order to model the social fabric. Maintaining an exponentially growing hypothesis set is not feasible. To circumvent this we utilize a representation that folds the hypotheses set into a graph structure that does not grow exponentially.

Rather than requiring perfect tracking from high-quality sensors, or settling for low-quality inference from low-quality sensors, we show one possible method that allows us to draw inferences from these collections of imperfect graphs.

3.1.1 Tracklets

The idea of a *tracklet* [10] was developed as a way to represent ambiguity in tracking for forensic applications. It allowed a system to efficiently query a human operator for input to refine the imperfect tracking results.

The basic concept of a tracklet is that it aggregates observations that are unambiguously related to each other. This is illustrated in Figure 3 where at the top we see a collection of observations strung out through space and time. If the observations are sequential in time, come from sensors that are close to each other is space, and are isolated from other distracting observations, then a simple model of movement in the space will allow us to aggregate those observations into a single chain (the second line of Figure 3). Further, if the collection of observations is itself isolated in space and time, then we say that the tracklet begins or terminates (third line of Figure 3). In subsequent figures we will drop the observations themselves, only showing the high-level tracklet abstractions, such as the bottom line of Figure 3.

We shall refer to the simple directed graph in Figure 3 as γ_0. The graph γ_0 shows that the probability of the terminating event Z being generated by the same individual (or group of individuals) that generated the originating event

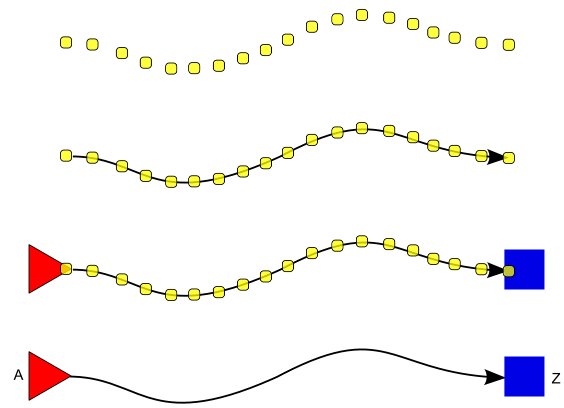

Figure 3: The Tracklet derived from a sequence of observations.

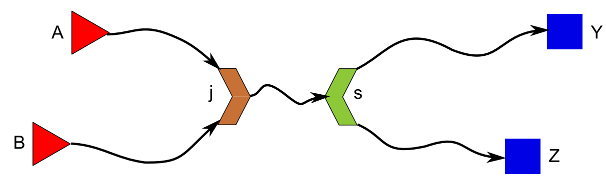

Figure 4: A simple tracklet graph with a single split s.

A is high, without loss of generality we will say

$$P(A \prec Z|\gamma_0) = 1$$

the probability that Z flows from the same cause as A, given graph γ_0 is unity.

3.1.2 The Tracklet Graph

Ambiguities in the environment will generate more complicated graphs. All observations that share a possible common cause will be linked together in the same connected graph. In Figure 4 we see a more complicated graph representing two or more individuals crossing paths in the space.

The new symbols represent a split and a join. The brown chevron represents a join. The green chevron represents a split. That is, several observations follow the split, s, within a small enough temporal and spatial neighborhood to generate ambiguity. In the case of a split, there must be several individuals moving together or very near each other to be considered as joint actors in a tracklet. After the split, the fate of each individual that traversed the tracklet from j to s is ambiguous: did they move along the tracklet from s to Z or from s to Y.

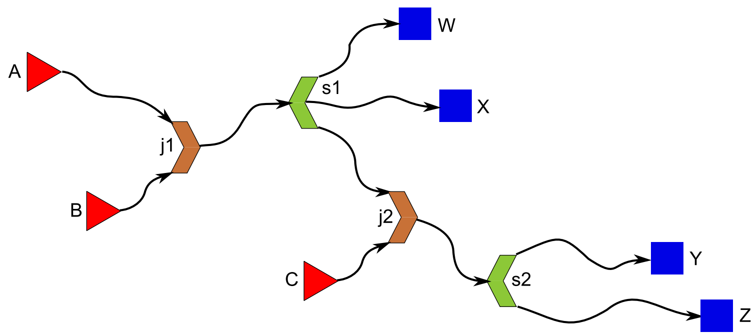

Figure 5: An example tracklet graph, called γ_2.

Given this graph, which we shall call γ_1, we do not know the fate of an individual originating at node A, this ambiguity is expressed as

$$P(A \prec Z|\gamma_1) = p \tag{1}$$
$$P(A \prec Y|\gamma_1) = q \tag{2}$$

It is worth noting that $p+q$ only equals unity if we are told that individuals are moving in isolation. Otherwise it is possible that a small group (indistinguishable from individuals by these sensors) could have originated at A and split, some to Y and some to Z.

Observed graphs may become arbitrarily complex. The graph γ_2 in Figure 5 is a much more complicated graph with several splits an joins.

By tracing the directed links in the graph we can see that it is possible to connect the origin A with any destination $\{W, X, Y, Z\}$.

The tracklet graph in Figure 5 folds onto a single representation all the ambiguity in the situation that would have required several dozen separate hypothetical outcomes to cover all the possible combinations of destination and possible group sizes.

3.2 Populations of Graphs

Imagine that several individuals are moving about a space at a particular instant in time. If they happen to cross paths they will generate ambiguities, and thereby will contribute to a particular tracklet graph instance, γ_i. The framework above allows us to model the fate of those individuals at that moment in time. We now extend that framework by considering collections of graphs, Γ that describes the behavior of a population of individuals over a longer span of time.

We can ask questions of the form, "What does this collection of graphs tell us about the probability that individuals travel from one place to another within a building." We can answer this question by retrieving all the graphs that contain an origin at location A and then counting the number of graphs that contain a possible connection between A and Z:

$$P(A \prec Z|\Gamma) = \frac{\sum_{i=0}^{N} P(A \prec Z|\gamma_i)}{M}$$

where N in the number of graphs in Γ, and M is the number of graphs in Γ that contain the node A as an origin.

It is possible to accumulate evidence for a repeated behavior even if there is not one unambiguous tracklet showing the behavior in its entirety. It is easy to see that if traffic originating at A *always* terminates at Z then all the graphs in Γ that contain A as an origin will contain a plausible path from A to Z and therefore a significant quantity of evidence. The $P(A \prec Z|\Gamma)$ will end up being the average over the evidence $P(A \prec Z|\gamma_i)$ from the graphs in the collection. That will be a smaller number than if every example of that behavior were completely unambiguous, then $P(A \prec Z|\Gamma)$ would be identically unity.

The maps in Figure 8 demonstrates the output of this form of analysis. The illustrations show $P(A \prec Z|\Gamma)$ as a line between A and Z which are the sensors arranged on the circumference of the circle. Darker links represent higher probabilities. These maps illustrate the social fabric of the building.

3.3 Trip Model

We want to extract models from graph populations captured over different periods of time and then compare those models to find changes in the social fabric. If we interpret these models as Markov chains [19], then we can use known formulations of the Kullback-Leibler Divergence to measure the differences between two models [15]

The probabilities $P(A \prec Z|\Gamma)$ between all pairs A and Z in the space represent an estimate of the probability that, given an origin at the point in space A then there will be a trip that ends at point Z. We interpret this as the transition probability between states in a Markov chain where each element is:

$$T_{A,Z} = P(A \prec Z|\Gamma)$$

A complete Markov model also requires a stationary probability that captures the likelihood of seeing the beginning of a chain in a particular state. The elements of this vector can be trivially estimated from data:

$$S_A = \frac{N_A}{N\omega}$$

where N_A is the number of tracklets that originate at A and $N\omega$ is the number of all tracklet originations. Taken together S_A and $T_{A,Z}$ describe a Markov chain that is a generative models of the social fabric during a window of time.

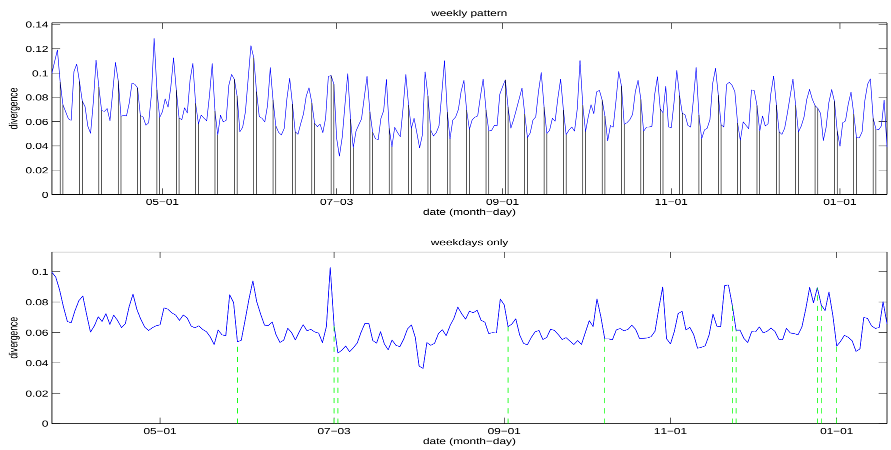

Figure 6: Human patterns are cyclo-stationary. The top plot shows a comparison of all days, including weekends (marked). The bottom plot shows a comparison of weekdays only (holidays marked).

3.4 Comparing Models

The Kullback-Leibler Divergence is a way to measure the "distance" between two probability distributions. Rached, et. al. showed that the calculation of the divergence when the probability distributions take the form of Markov processes [15] can be done as follows:

$$D(i,j) = \sum_r S_i(r) \sum_c T_i(r,c) \log \frac{T_i(r,c)}{T_j(r,c)}$$

where r are the rows and c are the columns of the matrices.

The plots in Figure 6 are examples of the KLD applied to the analysis of the social fabric of an organization. For each instant in time two models are extracted: one describing the social structure prior to that time, and one describing the structure after that time. The value of the plot is the divergence between these two models.

3.5 The Cyclostationarity of Humans

Figure 6 illustrates a problem with our model as stated. Human social behavior is periodic. In particular the top plot in Figure 6 shows that the model is overwhelmed by the weekly rhythms of the social system. It is not fruitful to blindly compare weekdays to weekends. There has been significant work on cyclostationary processes, for example the work of Kuhl [11]. In our analysis we will generally exclude weekends and assume that weekdays are relatively homogeneous. It is possible to see in the lower plot of Figure 6 that this simplification greatly improves the analytic results. It is now easy to find significant events such as holidays that were largely obscured by the variations induced by the weekly rhythm.

4 Data Collection

We present results on two datasets: motion data from a sensor network, and a collection of email extracted from user archives. Both datasets cover the 10 months from March 22, 2006 to January 22, 2007.

4.1 Motion

The motion dataset is a record of 10 months of activity at a corporate facility. The facility hosts 80 full-time employees and at least as many transient workers. The site hosts a large number of visitors who may be at the site only a few days a year. The facility also hosts janitorial staff, maintenance staff, electricians, plumbers, carpenters, couriers, caterers and other workers who are not under the administrative control of the organization. Each and every one of these individuals contributed data to this experiment because there was no need to outfit individuals with any kind of tag or hardware.

The data was collected with a motion-based sensor network. The network is built of over 200 nodes that communicate over a wireless radio link to 10 base stations that copy data to the wired local area network. Each node is comprised of an off-the-shelf passive infrared motion detector module and a combination microprocessor and radio board designed at the Massachusetts Institute of Technology[14]. When the nodes observe motion, they broadcast their serial number. The base stations timestamp the data and then inserts a record into a datlabase. The database used to generate our results contains nearly 21 million motion events.

The output of the tracklet building process is a list of tracklets, a list of joins and splits, and a membership map that associates the raw motion events with the tracklets. The list of tracklets in the database contains 3.4 million entries. The average length of a tracklet in the database is 8.5 seconds. This is a meansure of the crowdedness of the building, and implies that even short walks through the space are likely to encompass several ambiguities.

Extracting the tracklet graphs from the raw data takes a few hours to process a database covering over 7000 hours of data. All of this processing is causal, so it can occur during collection. The sensor network is designed such that it can easily be scaled up. In contrast, storage and computation requirements of a network of video cameras can be extremely complicated and costly.

4.2 Email

The email database was built through voluntary participation. Building occupants were asked to visit a web page and provide authentication credentials to a script. The script used those credentials to access the participants' email

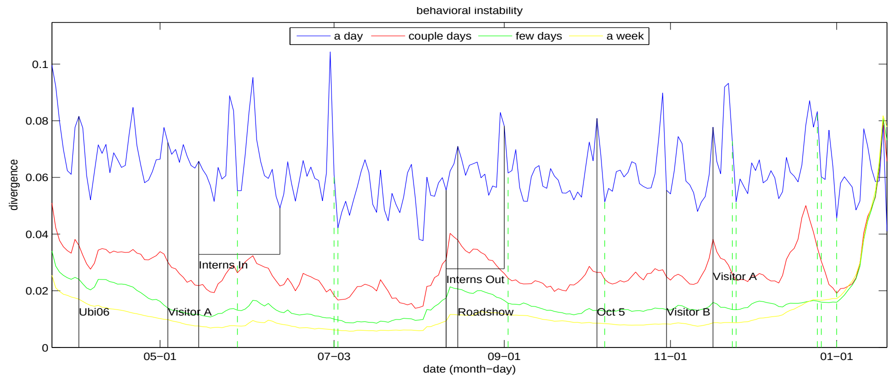

Figure 7: Responses over a range of temporal scales from a day to a week.

archives. From each email in the desired time window the script extracted sender and recipient information from the headers. Each unique address was assigned a random integer as a token. Sender and recipient token pairs were then recorded in a database, along with the time of the message. In this way a single email may generate more than one pair, since emails may have multiple recipients. Email lists were treated as distinct entities. Email addresses outside the organization were assigned to a single special token, and so are indistinguishable to the analysis. Once data collection was complete the mapping between tokens and addresses was destroyed. This creates a sanitized view onto the structure of the email habits of the organization. The database contains over 37,000 address pairs. There is no noise in this data, the origin and destination of each communication is explicit, so it is trivial to extract Markov models from this data[19].

4.3 Ground Truth

Ubiquitous computing has struggled with the concept of ground truth. Even the best methods are intrusive, inturrupting the user with intrusive questionnaires many times per day [7]. It is difficult to see how to scale these methods to year-long trials. There is also a long-standing concern over the reliability of self-reporting, particularly on the subject of interpersonal communication [5]. In this study we rely on the corporate calendar as a source of lables for disruptions.

5 Results

The plot in Figure 7 shows the result of the Kullback-Leibler divergence analysis described in Section 3. It is possible to process at different time scales by estimating models from larger temporal spans of data. The four plots in that figure show models ranging from roughly a day to roughly a week. We use an exponential window to create differentiable plots. However sinc ethe windows have theoretically infinite extent, it is difficult to precisely specify the time windows. The “a week” window mixes 10% of the current day distribution with 90% of the estimated distribution to form the new estimate. For comparison, the “couple days” window mixes 40% of the current day with 60% of the estimated distribution. Hard windows are equivalent to convolving with a boxcar filter, and induce similar aliasing artifacts into the results. The Black (dark) vertical lines indicate disruptive events. Green (light) vertical lines indicate holidays.

The events chosen on the plots as disruptive events include things we have taken from the company calendar. It’s not possible to record the daily activity of each individual every day for a year. We are instead looking for group events: events that likely affect the whole group. Those can be externally driven: such as holidays. They can also be internally driven, such as staff changes or large corporate events. By insuring that all the major events are represented in the data, with a statistically significant degree of confidence, and further that all the statistically significant peaks in the data are explained by these events, we demonstrate that our tool is sensitive to these macroscopic patterns of behavior.

The time window indicates the amount of data used to build the models that will be compared. So on the daily time scale we compare a model of today with a model of yesterday. On the weekly timescale we compare the model of this week with a model of last week, in a sliding fashion. The weekly model is not simply a filtered version of the daily model. This can be seen in practice by comparing the holiday behavior. For example, July 4th is the single most disruptive event at a daily timescale (top plot). However it completely disappears with only a slight expansion of the time window. In contrast the winter holidays are associated with extended leaves from the office by many people, as can be seen in Figure 2. This holiday has less of an impact at the daily level because people depart in a staggered fashion: no two days taken in isolation are very different from each other. However the flow of people has a significant impact on the longer time scales since the week before the winter holidays is very different from a week during the winter holidays.

The largest long-term disruption is in August when summer vacations, a major pulse of business travel, and the reduction in seasonal staff typically occur. This creates a massive disruption that has a significant impact even at the weekly time scale. It is interesting to note that this major disruption is not even noticeable in the gross activity plot in Figure 1. This indicates that it is not a simple matter of there being less people in the space: it is a sustained, *structural* change in the social fabric of the organization.

The final peaks to notice are the ones labeled “Visitor A”. In addition to the senior management termination in October, another disruptive development was a change in management above the organization. It is interesting to note that the May visit is associated with a daily spike but not a lasting disruption. However the late November visit is associated with one of the largest spikes in the entire “couple days” plot. While the visit is proximate to the traditional US holiday, it is clear form the travel records in Figure 2 that this holiday does not induce a pulse of travel similar to the August or December pulses. We hypothesize that this visit was itself massively disruptive due to lingering sensitization after the

October 5th event. "Visitor A" is a highly placed corporate executive who is visiting a little over a month after the sudden termination of a senior manager at the site on October 5th.

5.1 Details

We visually explore some of the Markov models in Figure 8. The three rows in the figure illustrate three different points in time. The top row shows July 17th, a day in the middle of the summer when the organization is enjoying relative stability. The second row illustrate October 5th, the day of a disruptive event: the termination of a senior manager. The third row centers on a major holiday.

In each case two models are presented. The divergence plots, such as Figure 7, compare two models. In Figure 8 the left column is the model of past history. The parameters are set so that these models are estimated from approximately one week of data. The right column is a model of the near future. These models are estimated from approximately a day of data.

The Markov model describes the probability of seeing a trip starting at one sensor and ending at another. This is represented graphically by arranging the sensors on a circle and connecting them with straight lines. Line brightness corresponds to probability, with darker lines showing more probable transitions.

5.1.1 Midsummer

represents a period of stability. The left and right models at the top of Figure 8 are estimated from non-overlapping data. There are small differences in the plots. However it is visually apparent that the structure of the social connections in the space are stable. This is reflected in the relatively low divergence rates from this period of time on the plots in Figure 7.

There are some features that are worth noticing. The markers around the outside of the plot indicate important locations in the building. The kitchen is a key resource and is marked with a circle on the right side of the plot. The restrooms are marked with squares at the top of the plot. The elevator and stairways are marked with triangles at the bottom of the plot. On the left are stars that mark a range of sensors that are near cubicles typically occupied by seasonal staff. Finally the diamonds in the upper right mark a range of sensor associated with the administrative group in general, and the senior managers in particular. Staff offices are mostly near the lower right and left side of the plot.

5.1.2 Early October

is a period of organizational disruption for the group. A senior manager who leads about a third of the organization is unexpectedly terminated. This causes a wave of fear and gossip to sweep through the organization. This disruption shows up in Figure 7 both as a spike in the daily model and as an extended ridge in the longer-scale plots. In the model view on the middle right we can see significantly less connectivity between the administrative wing in the upper

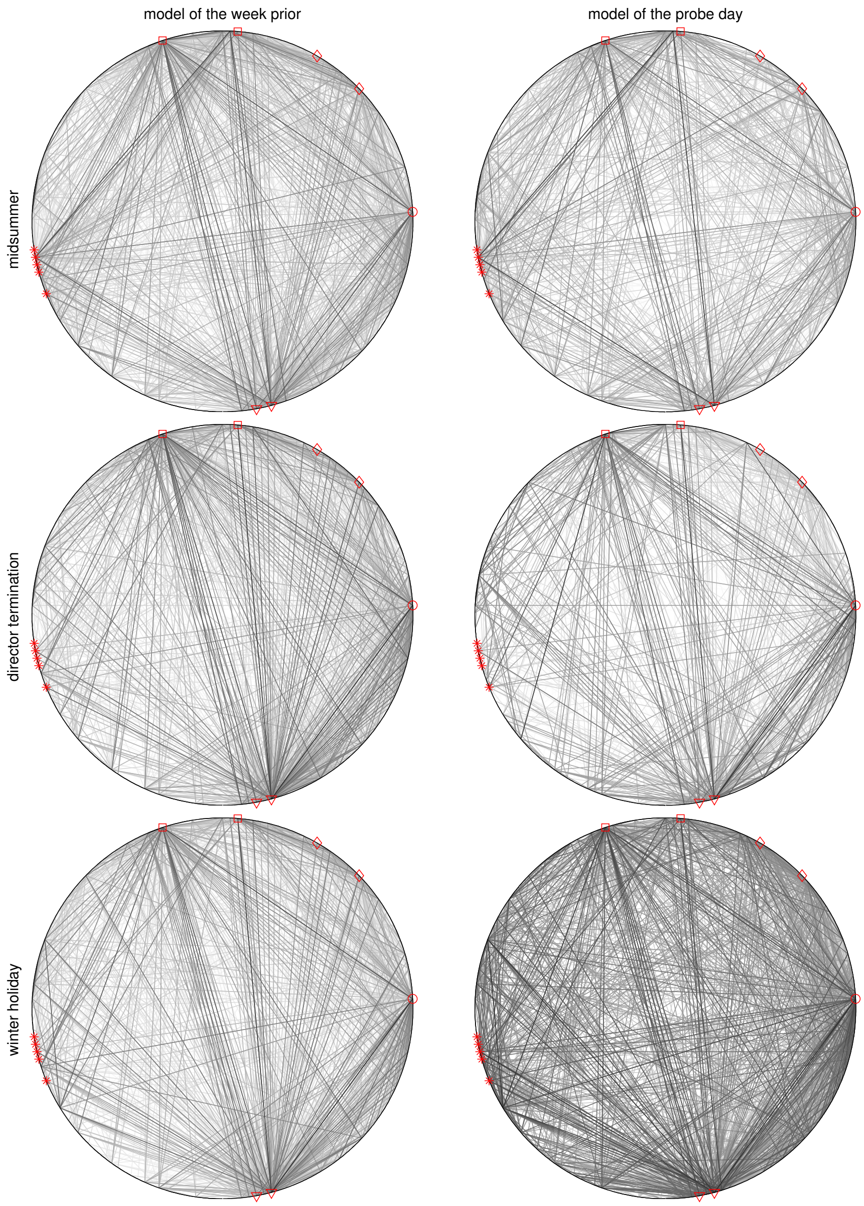

Figure 8: Example transition probabilities (darker means higher probability). Squares mark bathrooms. Stars mark seasonal employees. Triangles mark elevators and stairs. The circle marks the kitchen. Diamonds mark executive offices.

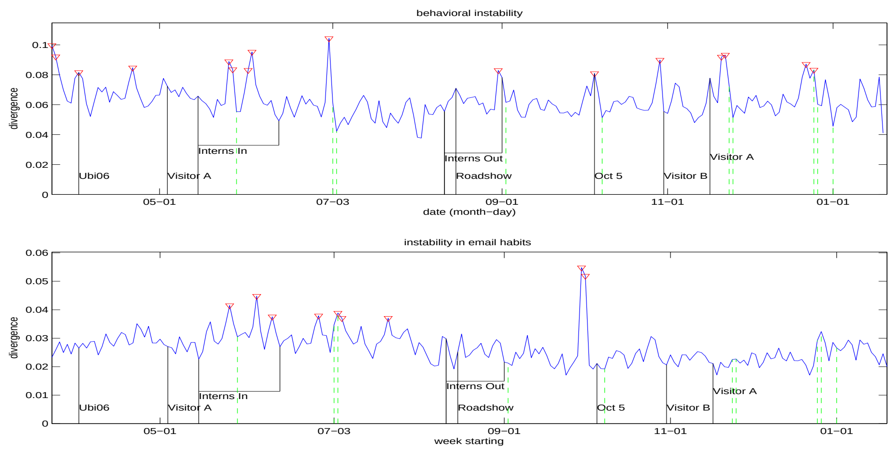

Figure 9: Physical behaviors (top) show more sensitivity to disruptive events than email behaviors (bottom).

right and the rest of the building. At the same time we see many more strong links among the staff offices on the left and lower right.

It is also interesting to compare the structure of the top left plot to the structure of the middle left plot. These are both models of relatively stable periods. However the top plot includes a much larger fraction of seasonal staff. The top model shows a dominant quadrilateral marking the interactions of the seasonal employees with the central resources of the bathrooms, the kitchen, and the elevators. In the middle plot the left side of the quadrilateral has largely disappeared as those areas of the building have become depopulated. This shift is responsible for the long timescale disruptions in June (addition of seasonal staff) and August (subsequent reduction).

5.1.3 Holidays

marked with light (green) vertical lines in Figure 7, are some of the most disruptive events. The plot on the bottom right of Figure 8 appears much darker than the rest. This is because of the normalization of the transition matrix. Where there is no dominant path then there are many more paths that appear nearly dominant. As a result a much higher proportion of lines are drawn dark. It is also interesting to note that the characteristic quadrilateral reasserts itself in the lower right. Even though there are far fewer temporary employees in the winter than there are in the summer, they do not receive vacation time, so the few who are present during the winter comprise a disproportionate fraction of the population during holidays.

5.2 Email Comparison

A more classical approach to social network analysis is to instrument communication media such as email systems or telephones. We have collected an email database and performed the same analysis on that data. Once the transition and prior probability distributions are estimated from the data, the exact same code computes the Kullback-Leibler divergence for the plots. In Figure 9 we present a side-buy-side comparison of the stability of email behaviors and the more physical behaviors observed by the sensor network. The red triangles on the peaks indicate that a data point has a high confidence of being an outlier, with a p-value of $p < 0.05$ according to a χ^2-test. The significance testing helps us objectively differentiate meaningful peaks from the noise in the plot that our eye may find interesting.

The most stunning thing about the email plot is how insensitive it is to disruption. The major winter holidays and the presence of the seasonal employees generate significant peaks. The rest of the disruptive events that were detected by the motion system, including most holidays, are completely buried in the noise. It could be that people do not, in fact change their email habits while traveling. It could be that there are simply too many distracting emails in the system that are hiding the informative structures in the data in a haze of entropy. Since these results are extracted from emails that people save, we expact that very few of these emails are spam.

There is also no significant reaction to a change in email habits after the termination in October. It could be hypothesized that people did not change their email traffic patterns but merely changed the topics of conversation. We did not do content analysis so we cannot discount that hypothesis. However it is clear from the right middle plot in Figure 8 that there was a significant increase in staff visiting each other in their offices. It seems safe to assert that they may have been discussing topics that did not seem prudent to commit to email. This finding highlights an important power of social network tools based on physical behavior since people may use different modalities for different purposes.

6 Summary

We have looked inside an organization of several hundred people and watched it react to disruptions over the course of almost a year in a way that has never been done before. We attempted to formulate a quantitative measure of the level of social disruption in such a building-centered setting. It is virtually impossible to obtain the exact transcription of daily events for the data set collected on this scale. Instead we focused on a sort of a ranking system that allows us to identify time periods where the dynamics of the group behavior significantly deviated from the "usual".

These results demonstrate that it is possible to measure and model the dynamics of social structures within an organization without instrumenting individuals or installing invasive sensors in the environment, although much works is left to be done refining this method. This building-centered, non-invasive approach allows us to collect comprehensive datasets that seamlessly include everyone in an organization: including transient staff and visitors. By focusing on the building we have created a system that is very human-centered: inclusive while remaining economical, and sensitive to the social fabric without sacrificing the privacy of the individuals who weave that fabric.

References

[1] Gregory Abowd, Aaron Bobick, Irfan Essa, Elizabeth Mynatt, and Wendy Rogers. The aware home: Developing technologies for successful aging. In *Proceedings of AAAI Workshop on Automation as a Care Giver*, 2002.

[2] Ryan Aipperspach, Elliot Cohen, and John Canny. Modeling human behavior from simple sensors in the home. In *Proceedings Of The IEEE Conference On Pervasive Computing*, 2006.

[3] T.J. Allen. Architecture and communication among product development engineers. In *Proceedings of the Engineering Management Society*, pages 153–158. IEEE, 2000.

[4] M. Athans and C. B. Chang. Adaptive estimation and parameter identification using multiple model estimation algorithm. Technical Report 1976-28, Massachusetts Institute of Technology Lincoln Laboratory, Lexington, Massachusetts, USA, June 1976. Group 32.

[5] H. Russell Bernard and Peter D. Killworth. Informant accuracy in social network data ii. *Human Communications Research*, 4(1):3–18, 1977.

[6] T. Choudhury and A. Pentland. Characterizing social networks using the sociometer. In *Proceedings of the North American Association of Computational Social and Organizational Science (NAACSOS)*, 2004.

[7] Sunny Consolvo and Miriam Walker. Using the experience sampling method to evaluate ubicomp applications. In *Pervasive Computing*, pages 24–31. IEEE, 2003.

[8] A. Culotta, R. Bekkerman, and A. McCallum. Extracting social networks and contact information from email and the web. In *Conference on Email and Spam*, 2004.

[9] N. Eagle and A. Pentland. Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4):255–268, 2006.

[10] Yuri Ivanov, Alexander Sorokin, Christopher Wren, and Ishwinder Kaur. Tracking people in mixed modality systems. In *Visual Communications and Image Processing*, volume EI123. IS&T/SPIE, January 2007.

[11] Michael E. Kuhl and James R. Wilson. Modeling and simulating poisson processes having trends or nontrigonometric cyclic effects. *European Journal of Operational research*, 133:566–582, 2001.

[12] A. McCallum, A. Corrada-Emmanuel, and X. Wang. Topic and role discovery in social networks. In *19th Joint Conference on Artificial Intelligence*, 2005.

[13] Thomas B. Moeslund and Erik Granum. A survey of computer vision-based human motion capture. *Computer Vision and Image Understanding*, 81:231–268, 2001.

[14] E. Munguia Tapia, S. S. Intille, L. Lopez, and K. Larson. The design of a portable kit of wireless sensors for naturalistic data collection. In *Proceedings of PERVASIVE 2006*, Dublin, Ireland, 2006. Springer-Verlag.

[15] Ziad Rached, Fady Alajaji, and L. Lorne Campbell. The kullback-leibler divergence rate between markov sources. *IEEE Transactions on Information Theory*, 50(5), May 2004.

[16] C. Reynolds and R. Picard. Evaluation of affective computing systems from a dimensional metaethical position. In *First Augmented Cognition International Conference*, Las Vegas, NV, July 2005.

[17] Carson J. Reynolds and Christopher R. Wren. Worse is better for ambient sensing. In *Pervasive: Workshop on Privacy, Trust and Identity Issues for Ambient Intelligence*, May 2006.

[18] J. P. Scott. *Social Network Analysis: A Handbook*. SAGE Publications, 1991.

[19] Henry Stark and John W. Woods. *Probability, Random Processes, and Estimation Theory for Engineers*. Prentice Hall, 2 edition, 1994.

[20] N. M. Tichy, M. L. Tushman, and C. Fombrun. Social network analysis for organizations. *The Academy of Management Review*, page 27, 1979.

[21] J. Tyler, D. Wilkinson, and B. A. Huberman. *Communities and Technologies*, chapter Email as Spectroscopy: Automated Discovery of Community Structure within Organizations. Kluwer Academic, 2003.

[22] M. van Alstyne and J. Zhang. Emailnet: A system for automatically mining social networks from organizational email communication. In *Annual Conference of the North American Association for Computational Social and Organizational Sciences*, 2003.

[23] Daniel H. Wilson and Chris Atkeson. Simultaneous tracking & activity recognition (star) using many anonymous, binary sensors. In *The Third International Conference on Pervasive Computing*, pages 62–79, 2005.