Occupancy Detection with Implicit Sensing for Energy Savings and More

Bruce Nordman, Lawrence Berkeley National Laboratory

ABSTRACT

Implicit sensing is obtaining information useful for managing energy from hardware that already exists in buildings that does not have energy or sensing as a primary function. The principal source of this data is information technology (IT) network systems, which typically have fine-grained visibility into the existence, status, and activity of many devices in buildings. The immediate opportunity for implicit sensing is information about human occupancy, as revealed in data about or from IT devices and networks. This produces time-series data of indicators of building occupancy. This paper describes many types of implicit sensing, and their advantages, disadvantages, and complexities. The mechanisms vary in many factors including latency of data availability, time frequency, resolution, and privacy and security concerns. Since no new hardware is involved, an ideal deployment of implicit sensing would combine data from multiple sources, to get a better picture than any single source could provide. That said, a clear initial first choice is data from Wi-Fi access points on the number of devices (mostly phones) that are connected to them. The Wi-Fi data are available in most buildings (residential and commercial), cover a large number of devices, and have low latency of devices appearing or disappearing from the network to data being available. The Wi-Fi method is also easier to explain to building owners, occupants and others, than many other types of implicit sensing data.

Introduction

Improved management of building energy use increasingly relies on the ability to sense information about the physical world to use to make better decisions about how to operate systems and deliver services. Traditionally this has required installing dedicated sensors that need to be installed, maintained, powered, and provided with communications infrastructure. This is expensive and so often done only to a limited degree, or not at all.

Implicit sensing originated with the realization that information technology (IT) devices leave traces on computer devices and networks that provide data useful for estimating the energy use of the IT devices themselves (Norford et al., 1990). As many IT devices are highly interactive with human beings, and in recent years move with them all the time, it later became clear that device status on a network can indicate human activity, and therefore occupancy. Occupancy detection is a principle application of implicit sensing.

The information that implicit sensing could deliver to buildings already exists in hundreds of millions of buildings today. The challenge is not to install hardware in buildings to create new useful information; the challenge is to understand what data already exists, how to extract and process it, and how it might be used. That is, implicit sensing works in reverse from normal processes that identify a need, and then install hardware to fit that need. As information technology integrates into buildings and energy in more, new, and increasingly unexpected ways, implicit sensing is of interest not only for what it can do directly, but the lessons it holds for where new efficiency technology opportunities can arise, and how to create them.

Since most energy in buildings is used to provide services to occupants, such as lighting and space conditioning, understanding in detail the dynamics of building occupancy can help

drive the delivery of these services to be more aligned with the presence of people in space and time. Traditional occupancy sensors have been expensive to install and maintain, and so used only in niche applications, as in restrooms or other sparsely used locations. With this deployment pattern, overall building occupancy remains unknown.

Typical building controls operate with fixed daily schedules that may be differentiated by day of the week. The mechanics of changing building controls, and arranging for someone to constantly know what changes to make and to manually enter them, are daunting. The usual result is that building system operational patterns are static, even in cases of varying occupancy.

Improved understanding of occupancy requires solutions that are low-cost, widely applicable, and highly granular (Brambley et al, 2005). Only 10% of U.S. commercial buildings use an energy management and control system (EMCS) and these tend to be the larger. As a result, using approaches that do not rely on an EMCS existing can be more broadly applicable in the market (Katipamula et al, 2012). High installation and other costs as well as poor interoperability and proprietary systems create substantial market barriers to EMCS use.

A principal way to use implicit sensing data is to feed it directly into dynamic building operation. The largest opportunity here is climate control, so that buildings can be run on the basis of *actual* occupancy, rather than fixed schedules of expected occupancy. For example, a commercial building could only initiate workday temperature and ventilation levels only when a nominal fraction of the building occupants have arrived (perhaps 5%), or start normal operation early, but cease it if expected occupancy does not occur. This approach automatically detects holidays and daylight saving time changes, and also accounts for anomalous total occupancy, whether it is more than or less than usual. For buildings that can vary the amount of ventilation, that amount can be based on the fraction of normal occupancy that occurs at each moment rather than a constant value as is more typical. When the occupancy data is differentiated by HVAC zone, the variations can be done by zone, as implicit sensing data are often zonal.

The paper is organized as follows: the first section explains what implicit sensing is, in an organized and formal way; the next section presents the variety of types of implicit sensing we have explored; the third section dives into the method we found most promising (counts of devices connected to Wi-Fi systems); this is followed by a section on miscellaneous other issues, and finally, a summary.

Implicit Sensing Definition

While conventional sensors use devices dedicated to that purpose, “implicit sensing” leverages hardware and communications infrastructure already installed in buildings, which have some other primary purpose (Melfi et al., 2011; Nordman et al., 2014). The sensing action is *implicit* in the behavior of a device not intended to be a sensor. Implicit sensing is often low-to-no cost to begin using. To use implicit sensing to save energy or help understand energy use patterns, the sensing data must be collected, processed, communicated, and then interpreted for use.

Implicit sensing extends our traditional sense of occupancy in several dimensions, as shown in Figure 1. A traditional occupancy sensor provides a single yes/no result, for a single location, and single point in time. Implicit sensing can extend this for people to give a count, identify individuals, and specify their activity; can provide time-series data for analysis; and can provide visibility across rooms or zones or a whole building.

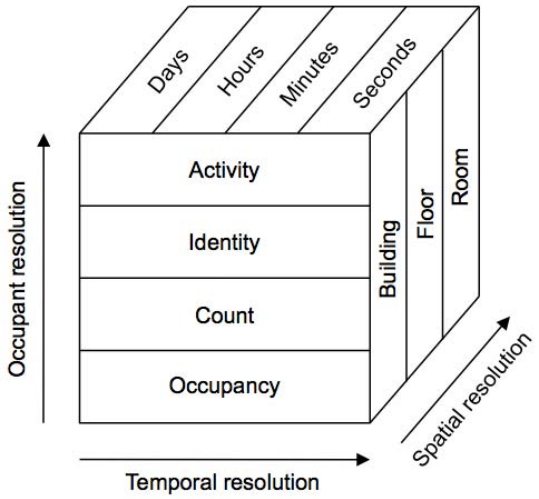

Figure 1. Implicit sensing characteristics (source: Melfi et al., 2011)

The base case of implicit sensing operation is systems that are used completely as-is; called Tier 1 and shown in Table 1. Tier 2 deviates from this by modifying software ib the device. An example could be software added to a wired telephone system to constantly monitor phone microphones (when not being used for calls) to assess where sounds indicate occupancy—by volume or content. Tier 3 extends this by also adding hardware to the device. An example is adding a USB temperature sensor to a PC. Most methods reviewed here are Tier 1.

Table 1. Tiers of Implicit Sensing (Nordman et al., 2014).

Tier 1	requires no modification to existing systems other than a data collection and processing point.
Tier 2	involves the addition of software to existing infrastructure to make existing occupancy related data available.
Tier 3	involves the addition of software and hardware to introduce new sources of occupancy data to existing systems.

Each implicit sensing method is based on some principle of operation and has limitations and variations. Methods may require additional data, require calibration, raise privacy concerns, or suggest research issues. Implicit sensing of occupancy generally involves the following steps:

- Some physical interaction occurs between a human and a device (active or passive on part of the person);
- The interaction is reflected in changes to data on the device or an IT system it is connected to;
- Those data are extracted and transferred to a central entity;
- The central entity may combine the data with data from other sources (static or dynamic);
- The processed data are delivered for use to final user—immediately for dynamic operation or in bulk for retrospective analysis.

Different implicit sensing methods work through these stages in different ways.

As an example, the first application of implicit sensing was to "ARP" data. ARP is the Address Resolution Protocol, one of the oldest and most fundamental elements of the Internet Protocol Suite (Plummer, 1982). The principle of operation of this method is that the power state of individual devices is highly correlated with the presence of the person that uses that device. The worker powers up the PC on arriving at work and it is powered down on leaving, either manually, or automatically based on a non-use timer. When the person arrives at work and powers up the PC, an early data packet sent by the PC onto the network will be one using the ARP protocol to announce its own network address or determine the network address of another device. The distribution of the ARP packet on the network will be noted by one or more local network routers. The list of devices seen recently can be obtained from the routers as frequently as desired. Typically the IP address of the device will indicate the building it is in, and sometimes even the particular floor. With data to associate individual device network addresses (IP or MAC) with people and rooms, a fine-grained view of occupancy can be provided.

ARP as a method of implicit sensing does have its problems though. While devices appear in ARP tables rapidly, to aid in efficient communication, they are dropped from them slowly, often taking hours to be removed after the device ceases to be present on the network. If the PC fails to power down when its user leaves, then no insight to occupancy can be gleaned. The data needed for localization within a building may be cumbersome to keep correct or be viewed as intruding on privacy.

Types of Implicit Sensing

We have identified more than a dozen potential data sources for implicit occupancy sensing in buildings, and collected sample data on eight of them from LBNL buildings. Specifically, we acquired data from LBNL's telephone system, its Wi-Fi infrastructure, and several sources from the IP network infrastructure. Each source has its own advantages, disadvantages, and peculiarities. A general feature of most sources is that data can be extracted as frequently as desired, and it is almost as easy to analyze results for many buildings as it is to do so for a single one. Since all hardware required is already present in buildings, the implementation cost is close to zero to add buildings. The technology appears to be highly replicable and scalable. In the primary study buildings, occupancy patterns are readily visible in the data, particularly arrival, departure, and lunch times, as are weekends and holidays.

One of the research goals was to understand the level of effort required to compile and process the data from each source to enable it to be useful. We also sought to understand the relative reliability, granularity, and latency of each source. Some sources are easier to obtain, and some are available in more buildings. Some data are retrieved only periodically, sometimes requiring manual effort, so suitable only for retrospective analysis, or future planning. The most useful data are available immediately and automatically.

The implicit sensing mechanisms we explored fall into the following categories; **bold text** indicates items for which we collected example data; *italicized text* identifies methods we have explored in some detail:

Internet Protocol Network Presence — **ARP**, **DHCP**, *Ping*, *Port Status*, **Wi-Fi networks**.

These methods rely on whether and how end-use devices are connected to IT networks, and retrieve data from network equipment (ARP and Port Status), network infrastructure servers (DHCP and Wi-Fi), or directly from end-use devices (Ping). The fact of each *device* being

connected to the network—or rather how that changes over time—can be an indicator of *human* occupancy. Depending on how these data are acquired, some methods may have delays on when presence is recognized; due to how protocols operate; other methods will keep devices on the list for times—often several hours—after the device has left the network. These limitations are a problem for dynamic building operation, but not for retrospective analysis (for which corrections can be done after the fact). These methods all, with additional information, provide granularity of occupancy data, often down to the individual device (and hence office). Some of them apply differently to wired and wireless network technologies.

IT Network Traffic — **DNS**, **Web browsing**, *direct traffic analysis*.

These methods use data from the network traffic itself to provide information indicating occupancy. This can be available from servers or from analysis of data traffic from individual devices. Direct traffic analysis operates in reverse of a network security firewall; a firewall analyses traffic to keeps bad data out of a local network—analysis for implicit sensing creates good data within the local network. Some data on the network only exists when a person is present and using a device or is found in much greater quantities when someone is present. Since these methods all involve active analysis of network traffic, the results can be available immediately. The beginning of occupancy is apparent immediately, but the end is only apparent when no activity is detected for an extended time. The spatial resolution is about the same as with network presence. These methods are all generic to any organization or building.

Enterprise Applications — Databases, email, *authentication*.

When someone uses an IT system, it creates traces of activity that indicate occupancy. These are usually specific to a particular company and while focused on the identity of the individual accessing the system, do provide a network address which can indicate a specific location. As with the network traffic methods, these involve active analysis so lack significant latencies.

Other IT Systems — *Access control*, **wired phones**, *cameras*.

These may or may not use the IP network, but do not fit into the previous three categories. Cameras could track both people entering and leaving a building; access systems may apply to both, or only to entry. Wired phones only provide data when they are used, so only useful when greatly aggregated as for whole-building estimates.

Building Infrastructure — **Electricity meters**, elevators, water, gas, chilled water, hot water.

These systems traditionally have not been connected to IT networks to make their data readily available, but that is changing. When people are present in a building, they use services which require resources that can be tracked[1]. These are commonly scoped to an entire building, but could be measured on a more fine-grained basis.

[1] In Price et al., 2015, water use data for a building showed that it was not occupied over night.

Wi-Fi Device Counts

Figure 2 shows an example series of implicit sensing data, from Wi-Fi access points (APs). It covers one typical work day and shows the number of "Wi-Fi associations" seen by the network management system on 10 minute intervals, for building 90 at LBNL. All devices that are connected to the Wi-Fi network are included in the total, primarily phones, with notebook PCs making up most of the rest. The base load of about 20 devices is likely PCs that are left on 24/7 (potentially including some desktop PCs that use Wi-Fi for convenience). The peak value of just under 300 corresponds to the number of people we expect to be in the building at any time. While the number of building occupants in principle is considerably higher—over 350—work travel, vacation, sick days, and telecommuting all reduce this, even when visitors are added. While some people have two Wi-Fi devices, a notebook PC may not be awake and connected to the network some of the time, and some people have no Wi-Fi device at all. The pattern over the day also matches our expectations, with significant variation in arrival and departure times, and a noticeable lunch dip. The graph also shows how the HVAC system in the building is operated, with it being either all-on or all-off on a fixed schedule, supplemented by an optimal-start period to ensure that the building reaches its intended thermal range by 8 a.m.

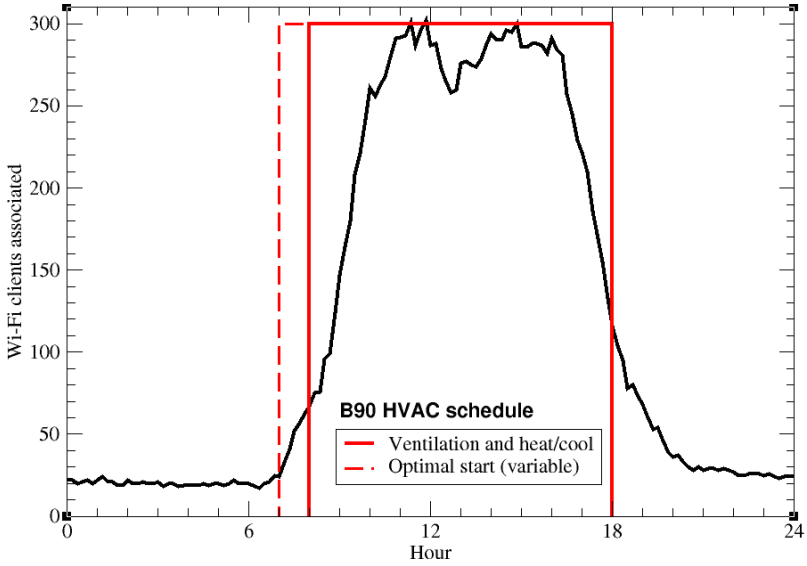

Figure 2. Wi-Fi Associations over an example day for building 90 (LBNL)

Why Wi-Fi

A key finding from our work is that device counts from Wi-Fi systems is the best single opportunity for implicit sensing at this time. Reasons for this conclusion include that it is:

- Easy to understand for people who don't understand network technology
- The most widespread method available—applicable to nearly any building type
- Simple to implement from an IT perspective
- Is covered by a modest number of key manufacturers (for commercial sector at least)
- Has a clear path to alleviating legitimate privacy and security concerns
- Has low latency of detecting arrival and departure of devices

- Offers clear zoning (albeit imperfect)

In commercial buildings (other than those so small that a single access point suffices), individual access points are scattered throughout each building to provide coverage. A central access point controller device coordinates their operation (including handoff of moving devices from one AP to another), provides a common point for authentication/security, and manages the entire collection, including archiving data about usage. The controller can provide implicit sensing data to building energy systems. The controller can be set to actively "push" the data out on a regular basis, or to respond to queries from the outside to "pull" the data out. Which methods are possible depends on the manufacturer.

Building 90 has 32 APs spread over its four main floors and basement. With about 350 occupants of the building, this is an average of about one AP per eleven people. The lab-wide AP controller can be queried any time for a list of devices connected to each AP. The LBNL network staff set up a system to do this every 10 minutes[2]. These data include many aspects of the device and its network connection, and notably the IP and MAC addresses. To not risk intruding on privacy, we only obtained a count of the number of devices per AP, not any data about each individual device.

There are at least two figures of merit potentially available. One is the number of devices that have actively connected to the Wi-Fi network or "authenticated". The other is just devices in the area that have made no attempt (or made a failed attempt) to connect, but are known to the AP to be in the area by virtue of how Wi-Fi works. A second source of data from UC Berkeley where both counts are available show they are not significantly different.

There is a quite consistent nighttime level of about 23 associated devices, when no one is present. This could include some desktop PCs using Wi-Fi because Ethernet was either unavailable or inconvenient. It could also include notebook PCs that are left fully on. Finally, there could be printers or other devices in the buildings that are always on the network. For purposes of occupancy detection, this base load can simply be subtracted from the current value.

Digging deeper

To understand broader patterns, it is necessary to apply some statistical analysis to the data. This is needed when longer time periods and more buildings are assessed. To simplify analysis we ignore the lunchtime dips, and so treat daytime occupancy as relatively flat. With this, the issues to determine are the peak occupancy level, the nature of transition from unoccupied to peak in the morning, and the nature of the reverse transition at the end of the day.

To explore occupancy variation in more detail, we focused on an example time surrounding the winter holiday break, which has more variable occupancy. Figure 3 shows two load duration curves of data from December and January that cover ordinary workdays, workdays impacted by the holiday break, holidays, and weekends. The black line shows all ten-minute time periods, sorted from highest to lowest. There are 144 periods per day; the red line takes only the peak period for each day. The lowest red days all are during the holiday break (a few PCs likely powered down making them lower than ordinary weekends). The next five days are ordinary weekend days. The highest two days in that section are Monday and Tuesday between Christmas and New Years; the lab is closed, but the data show a few people coming to the building. The days between the two roughly horizontal sections include the actual workdays

[2] For a short time we increased data collection to once every seven seconds.

during the two-week holiday period; many people take these as vacation days (one day is erroneous and reflects a part-day at the beginning of the analysis period). The ordinary day in this time period with the lowest peak value has 284 devices.

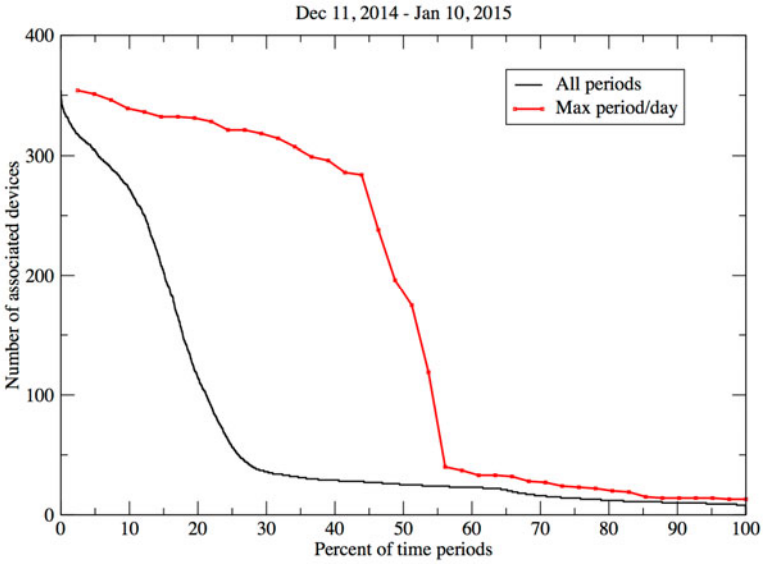

Figure 3: Load duration curve of Wi-Fi associations for LBNL Building 90

For the analysis below, of the same period covered above, 280 was taken as full occupancy, to identify statistical periods for each day, to assure that the fully-occupied state is reached during each regular workday. We calculated the values that are 10%, 25%, 50%, 75%, and 90% of the way between the night-time base and the value chosen for peak. We found the first time of the day that each of these percentile values was exceeded, as well as the last time during each day it was exceeded. A graph of these values is shown in Figure 4, with each horizontal point showing data for one day. It is notable that there is a significant group of people who arrive promptly at 8am each day, with never more than 10% arriving before that time and most of the time over 25% there before 8:10; this is shown by the red line being on top of the black line for arrival for most days in the graph. Leaving times are neither so regular nor so closely clustered. Figure 4 includes two holiday weeks with much lower occupancy. The dramatic dip in the second week of the graph is due to missing data.

The AP control system that covers this building reports the *type* of each device, if known. There are several ways that the system can determine this. For example, queries for web pages often include data about the hardware and operating system of the device; this can help the web server that creates the page to know the device's screen size and capabilities. Also, Media Access Control (MAC) addresses are allocated to companies in patterns that can be used to infer the manufacturer and device type. The AP can observe this and other data to identify or infer the type of many devices. A snapshot of device counts taken during one day found that over 50% were definitely phones, with many of the remainder possibly phones. Tablets and PCs made up most of the rest.

In more exploratory work, data on the time of associations to APs for one phone was obtained for many months. This showed it moving around building 90, as well as occasionally showing up at other lab buildings. It also showed up sometimes at the AP for a building on the vehicle route up to building 90, which also has a major shuttle bus stop. However, since the snapshots were only every 10 minutes, many short associations to that AP were not captured.

This ability to track individual behavior is intriguing but also raises significant privacy concerns.

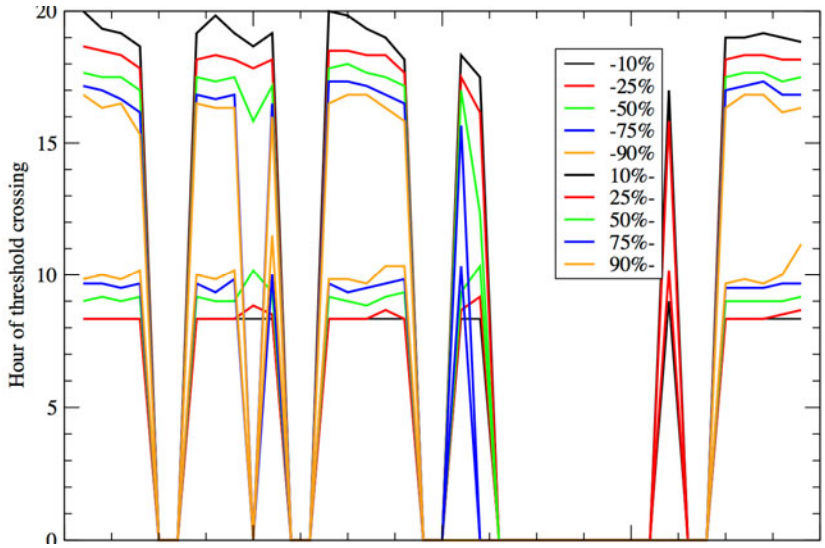

Figure 4. Percentile times for Wi-Fi associations for LBNL Building 90

One issue for this source of data is that devices stay in the list of associated devices for about 30 minutes after they have been last seen. To some degree this is needed as one doesn't want a device to be knocked off if it loses contact for a few seconds as a person moves around. That said, providing a count that well-reflects occupancy would likely benefit by using a shorter time period. AP systems could be set to only list devices seen within a shorter time period.

Other Issues

Ground Truth

Two key questions arise in interpreting such data. First, "How to assess how many people occupy a building in a general sense?". This has many potential answers. For building 90 for example, there are 357 wired phones; the lab's facilities department considers that there are 446 people assigned space; and there are about 530 offices and work stations. Issues include:

- Some spaces have no one occupying them;
- Many workspaces that can accommodate two people only have one;
- Some people have no wired phone;
- Some people, such as custodians, have no assigned workstation;
- Some people work part-time or mostly off-site, and may, or may not, share a space when they are present.
- Some are on extended vacations or work travels; and
- Some people come to meetings and so are in the building for a time without having any permanent association.

We are aware of no people who are regularly in the building in the middle of the night.

The second question is "How many people are *actually* in the building at a single given time?". This should be compared to the figure provided by the implicit sensing system. Such

"ground truth" data for a period of time (even a few hours during a morning or evening) could determine a conversion ratio from the value that implicit sensing provides to actual occupancy, and validates that such a ratio is stable over the course of a day. This needs to be done for a few buildings to validate implicit sensing in general, and to provide reference values for different building types. Developing inexpensive ways to collect episodic ground-truth data would be helpful. That said, for many purposes, it may not matter how accurate implicit sensing data are on an absolute basis, or how consistent over the course of the day. Even rough estimates may get most of the benefit that perfect data could provide, as the starting point is having no data.

Privacy and Security

An issue which quickly arises as one delves into implicit sensing is the potential conflict between obtaining data useful for energy purposes and perceived (and often very real) concerns with privacy and security. At its extreme, implicit sensing is about using our IT infrastructure to spy on people—to track and record their movements, activities, and electronic fingerprints. This could exceed bounds on what is allowable by policy or law, or be quite unsettling to people who work in places where it is used. That said, the goals of implicit sensing in general do not require the most sensitive data, so it is critical to delineate what is, and is not, being collected, be transparent about this, and seek to build systems wherein sensitive data are filtered out before being passed on. Research in this area does require some exploration of sensitive data, such as to confirm that masked or aggregated data reasonably reflect what the detailed data would indicate. Crossing the line to a limited and controlled degree can help establish where the line of limit should be for this technology.

For many purposes it isn't important to know *who* is in the building, but only *how many* people are present. In these cases, counts of devices are sufficient, and individually identifying information can be filtered out. In other cases, it may be necessary to be more specific. For example, the telephone of someone appearing on a campus Wi-Fi AP likely means that that person will soon appear in their office. This could be used to adjust HVAC, lighting, and electronics operation. However, it does require the implicit sensing system to have a persistent association between the device and the owner, which then allows the system to track the owner around the building and campus.

(mobile) Phones vs. (wired) Phone calls

One of the data sources that we obtained is a count of outgoing phone calls from wired office phones by hour. This was compared to the counts of Wi-Fi devices (mostly mobile phones) for a four week period. Two notable results became clear from this comparison.

- People these days make very few phone calls, at least on wired phones. Our source only counted outgoing calls to outside the laboratory, but even so, peak hours only had about 50 calls for the entire building of several hundred people. Since some people make several calls per day, many people must make no outside calls at all on these phones.
- In assessing the ratio of phone calls to Wi-Fi associations, we found that the ratio in the morning was about four times that in the afternoon (surprisingly different). Part of this is calls to the east coast all occurring in the morning.

Next Steps

Despite all these advantages, implicit sensing is not available on the market as there is no standard protocol to communicate such data between the sources of the information and the devices that could receive it. Creating such a standard is a near-term priority for future work. This could be used by any source of implicit sensing data.

DOE is currently supporting LBNL to help make implicit sensing data more readily available from systems from major manufacturer's, and to help demonstrate using implicit sensing data directly in dynamic building HVAC operation.

Several manufacturers of Wi-Fi access points sell hardware and software to obtain high-resolution tracking of individuals by monitoring the Wi-Fi footprint of phones, with retail shopping the primary target market. However, these systems generally require extra hardware and are relatively expensive.

Moving forward, implicit sensing data—and in fact all sensing data—should be forwarded to an "Occupancy Server" for each building. This entity would gather occupancy and related data from other devices, process and aggregate the data, and then provide it to devices that can use the information. This could be for dynamic building operation and retrospective analysis such as monitoring and verification analyses. Granularity needs over space and time vary so the occupancy server can provide data of the form needed by the requestor, combining the best insights of all sources, be resilient as possible to sources becoming unreliable or disappearing entirely, and easily add new sources. Ideally an occupancy server is not a stand-alone device, but rather a function of a device that already exists for some other purpose. This can make the cost (and energy use) of an occupancy server be low.

Summary

Implicit sensing is a promising low-cost way to provide information about occupancy. It need not and should not be the only sensing technology used, but can be a good complement and asset to other methods.

There is clear future work in this area. This includes to validate implicit sensing data with ground truth occupancy data, to compare the various methods with each other, and extract key metrics. Further work is also needed to use these and related occupancy data to better understand the linkages between occupancy and energy use for measurement and verification, fault diagnostics, energy forecasting, and improved building control.

Many systems already existing in buildings today have, or could have, data that could indicate occupancy. A challenge is to extract this data, process it, and combine it with other sources. Many of these methods are likely to be imperfect so that a combination of several will produce a better result. Concrete examples for each method will help better understand the operations necessary to obtain them, features or limitations, and ways that they could be made easier to obtain or more accurate.

Acknowledgments

The author would like to acknowledge the contributions and assistance of Marina Sofos and Joe Hagerman of the U.S. Department of Energy; Ken Christensen of the University of South Florida; as well as many LBNL colleagues including: Michael Smitasin, Mary Ann Piette, Rich Brown, Janie Page, Phil Price, Daniel Fuller, Christian Kohler and Stephen Czarnecki.

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

M. R. Brambley et al. *Advanced Sensors and Controls for Building Applications: Market Assessment and Potential R&D Pathways* (http://apps1.eere.energy.gov/buildings/ publications/pdfs/corporate/pnnl-15149_market_ assessment.pdf, 2005).

Li, N., G. Calis, and B. Becerik-Gerber. Measuring and Monitoring Occupancy with an RFID Based System for Demand-Driven HVAC Operations. *Automation in Construction* 24:89-99, 2012.

Katipamula, S., R. M. Underhill, J. K. Goddard, D. Taasevigen, M. A. Piette, J. Granderson et al. *Small- and medium-sized commercial building monitoring and controls needs: A scoping study*. Pacific Northwest National Lab., Richland, WA, USA, Tech. Rep. PNNL-22169. Oct. 2012.

Melfi R., B. Rosenblum, B. Nordman, and K. Christensen. Measuring Building Occupancy Using Existing Network Infrastructure. International Green Computing Conference and Workshops, July 2011.

Nordman, B., K. Christensen, R. Melfi, B. Rosenblum, and R. Viera. Using Existing Network Infrastructure to Estimate Building Occupancy and Control Plugged-in Devices in User Workspaces. *International Journal of Communication Networks and Distributed Systems*, Vol. 12, No. 1, pp. 4-29, January 2014.

Norford, L., A. Hatcher, J. Harris, J. Roturier, and O. Yu. 1990. "Electricity Use in Information Technologies." In *Annual Review of Energy 1990*. Edited by J. M. Hollander. Palo Alto, CA: Annual Reviews, Inc. pp. 423-53.

Price, P., M. A. Piette, J. Granderson, and J. Elliott. Automated Measurement and Verification of Transactive Energy Systems, Load Shape Analysis, and Consumer Engagement. In *GridWise® Architecture Council Transactive Energy Conference*. Portland, Oregon, 2015

Plummer, D. RFC 826, An Ethernet Address Resolution Protocol -- or -- Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware. Internet Engineering Task Force, Network Working Group. November 1982.