Sensor Networks Challenges for Intelligent Buildings

Alberto Eduardo Cerpa
University of California, Merced
acerpa@ucmerced.edu

1. Introduction

Sensor networks will play a fundamental role in future intelligent buildings. There are two basic areas that we envision this technology having a critical impact, these are energy conservation and security.

Most of the power consumption in most countries is done inside buildings. Furthermore, the largest consumers of energy are the richest most developed countries, and the proportion of energy used inside buildings is much larger in those nations than the rest. Regarding security, it is clear that having a large number of sensors in the building that can monitor human movement further increase the levels of security in the building. Having multiple sensing points can provide greater capability, as is leveraged commonly by sensor arrays such as in radar systems or binocular vision.

There are significant advantages for this technology to have an impact in intelligent buildings. First of all, sensor networks use small and non-intrusive devices. They have relatively low energy consumption at the system level, which implies long lifetime per deployment. Second, it is possible to retrofit old buildings to do sensing and monitoring with minimal changes, and even integrate them with HVAC systems for finer grained distributed control in all new and old buildings –if they have one–, as long as these systems provide an open API. Finally, a large number of parameters could be measured to capture spatial and temporal distributions at a much finer granularity than other technologies available. For example, a sensor network system consisting thousands of sensors in a building could capture temperature, humidity, light and human mobility at a spatial and temporal granularity not available today.

2. Challenges

The main challenges to move forward are:

Hardware Cost: the current cost of each individual sensor unit is still very high. Commercially available platforms cost in the order of $100 per unit with temperature, humidity and light sensors when bought in large quantities. Capable sensors able to track human mobility inside buildings are even more expensive with research prototypes of low power consumption cameras costing in the order of $300 per unit.

System Architecture: there is no unified system and networking architecture that is stable and mature enough to build different applications on top. Most of the applications and research prototypes are vertically integrated in order to maximize performance. This prevents further innovation at application layer, since currently each application has to reinvent the wheel at some portion of the lowest level of the stack.

Wireless Connectivity: having the capability of the sensor nodes to be able to communicate and to coordinate each other in a distributed manner completely un-tethered from the environment and the infrastructure is on of the main strengths of this technology. However, wireless communication in indoor environments is still quite unpredictable using low-power consumption RF transceivers, in particular in clutter environments common inside buildings, with many interfering electromagnetic fields, such as the one produced by elevators, machinery and computers, among others.

Programmability: manual re-programming of nodes is not viable when deploying systems in the order of thousands of nodes. While it is clear that some form of network reprogrammability is desirable, doing so in an energy and communication conservative form remains a challenge.

Security: there are security challenges at many levels. From the system point of view, it is critical that the information provided by the nodes be authenticated and the integrity verified, since this information provides the feedback loop to expensive equipment controlling power consumption in the building. From the users' point of view, it is also critical that this information cannot be easily spoofed and remains protected in the back end processor, since it may affect the privacy of users.

3. Research Needs

Each of the challenges above set the direction for many of the information technology research needs. We believe the most important research needs for the next 5 to 10 years are:

Design and production of new hardware: there are several research questions regarding how to build the next generation of sensor nodes. While some advances has been achieved with system integration and new designs, the most notable of them being the mote-in-a-chip 'spec' design, they remain at large research prototypes with limited programmability and still costly to obtain in large quantities in practice. Ironically, the most important issue facing the future may not be necessarily doing innovations in hardware architecture or the fabrication of components themselves, but perhaps integration issues that push the cost of each node to at least less than $10 per unit. Another important issue is the packaging of these units. Perhaps more critical for outdoors applications, packaging has been a no-man land in the last 5 years, with ingenious graduate students finding ad-hoc solutions that do not necessarily work for real deployments. While many of these issues are not pure research questions and they may be a matter of economies of scale, they remain a critical roadblock to make progress in other areas.

System Architecture: it is *unclear* what should be done at this point regarding the system and network architecture of the system; to the point that some researchers argue there should be none. While most of the applications and research prototypes are vertically integrated in order to maximize performance, we strongly believe there should be some significant thought spent into analyzing a wide range of applications and stacks, and understanding what are the fundamental components and services necessary to a large range of applications. This exercise would allow a system architecture design based on common needs of a group of applications, and would also allow the synthesis of common functionality at the lower levels of the stack. However, there are researchers that propose no basic architecture at all, for example in the context of declarative networks. We believe this is similar to the past history in computer architecture when at some point researchers were designing microprocessors being instruction set independent. We also believe that there should be more research on tiered architectures with heterogeneous type of nodes, since we believe that one size does not fit all.

Wireless Communication: there are no accepted and implemented wireless models in the community that are well understood so algorithm and application researchers could at least get simulation results with more confidence and accuracy. This process is similar to the one experienced by simulations tools used by the networking community in wired networks (e.g.: ns-2, optnet) and these tools still have significant problems when used to simulate wireless environments. Furthermore, as new more capable transceivers are introduced, there is a need to further improve upon previous work in other areas, like link quality estimation. For example, the effective nominal radio range of the new zigbee 802.15.4 radios significantly changes based on the bit rate used, so it is unclear how to perform link quality estimation for any parameter.

Security: while the community has produced several algorithms and schemes to provide cryptographic authentication and integrity at the packet transmission level under stringent energy constraints, the more general problem of system security has not been addressed extensively. These problems are critical if we are to move forward to real system deployments. For example, when the Soviet Union invaded Afghanistan in the '80s, they installed sensor motion detectors in the perimeters of all critical bases. The Afghan rebels attacked the system by releasing large number of rabbits several weeks before the attacks. The soviet security officers were so exhausted with the false alarm rates, that they simply disconnected the systems.

Programming: the more prevalent OS used in sensor networks for the smallest sensor platforms is TinyOS. However, as the number of nodes increases, other operating systems like SOS may have a more prevalent role since they are more efficient and use smaller quantity and size of packet updates. This could also mean that we may move towards a limited virtual machine paradigm, similar to Mate. However, at some scale, none of these schemes can get the job done with very large number of nodes. Research should be done to completely change the paradigm, from deterministic to may be more statistically oriented programming, perhaps inspired by statistical mechanics. We think there are other more pressing needs, so we left this issue to the last.

Short Bio

Alberto E. Cerpa received the PhD degree in computer science from the University of California, Los Angeles (UCLA, 2005), working under the supervision of professor Deborah Estrin, the MSc degree in computer science from the University of Southern California (USC, 2000), the MSc degree in electrical engineering (USC, 1998), and the engineer degree in electrical engineering from the Buenos Aires Institute of Technology, Argentina (ITBA, 1995). He is an assistant professor of computer science and engineering in the School of Engineering at the University of California, Merced. His interests lie broadly in the computer networking and distributed systems areas. His recent focus has been on systems research in wireless sensor networks, with emphasis on network self-configuration, radio channel measurement and characterization, programming models, and development of wireless testbeds. He is a member of the ACM and IEEE.