A CAUSE OF DATA PERTURBATIONS IN BUILDING SENSORS

Michael W. Bigrigg[1], and H. Scott Matthews[2]

ABSTRACT

The Critter temperature sensor device has been developed to understand the full nature of pervasive sensor networks. Off the shelf integrated sensor devices incorporate some amount of adaptation to make the devices more reliable. For instance, they use a weighted averaging scheme to account for large fluxuations in temperature, primarily considered to be sensor failures. A weighted average is used to limit the effects of data outliers, considered to be sensor failures. The Critter provides raw instantaneous readings including outlier data that may be considered anomalies or perturbations. We have deployed Critter data sensors pervasively through one academic building for almost 18 months. This paper explores the causes of temperature data perturbations, two temperature data readings taken within seconds of each other that differ by several degrees. Temperature sensor data perturbations are actually the effects of user activity within buildings. By capturing the raw data without automatic processing, we are able to show a correlation between time of the work day, and the frequency of data perturbation.

KEY WORDS

Networks, Building Management, Reliability, Temperature Sensor, Experimental Models.

INTRODUCTION

In order to pervasively deploy a large network of sensor nodes, we have developed the Critter Dataport Sensor. The Critter is an analog computer-attached sensor device developed at Carnegie Mellon University. It is an economical (and thereby pervasively deployable) alternative to all-in-one sensor devices, which include sensing, processing, memory, and networking. A sensor network is formed by the Critter-attached computers using the existing host networking. The Critter platform, while attached to a computer with significant resources from which it draws its power, is merely a guest on that host computer and cannot dominate its resources. By leveraging the computing, power, and communications resources of the host, we are able to leapfrog some of the issues that are a current focus of wireless sensor research and look towards future research issues that will be common of wired or wireless sensors pervasively deployed within a building environment.

The temperature-sensing version of the Critter uses a thermistor, a resistance-based temperature detector. A thermistor is made from a material for which its resistance decreases

[1] Project Scientist, Institute for Complex Engineered Systems, Carnegie Mellon University, Pittsburgh PA 15213, Phone +1 412/268-7668, bigrigg@cmu.edu

[2] Professor, Civil & Environmental Engineering and Engineering & Public Policy, Carnegie Mellon University, Pittsburgh PA 15213, hsm@cmu.edu

as its temperature increases. A current is passed through the thermistor and the resistance is measured. The Critters attach to the game port of a computer, which is then used to read the temperature sensor value. While game port sensors are not novel, the Critter uses them in conjunction with a CMU-developed auto-calibration mechanism that will automatically calibrate a resistance-based sensor without user interaction. This leads to extremely low installation and maintenance costs of the sensor system compared to existing technologies.

The Critter has been used to understand the full nature of pervasive sensor networks. Off the shelf integrated sensor devices incorporate some amount of adaptation to make the devices more reliable. For instance, they use a weighted averaging scheme to account for large fluxuations in temperature, primarily considered to be sensor failures. A weighted average is used to limit the effects of data outliers, considered to be sensor failures. The Critter provides raw instantaneous readings including outlier data that may be considered anomalies or perturbations.

We have deployed Critter data sensors pervasively through one academic building for almost 18 months. This paper explores the causes of temperature data perturbations, two temperature data readings taken within seconds of each other that differ by several degrees. Temperature sensor data perturbations are actually the effects of user activity within buildings. By capturing the raw data without automatic processing, we are able to show a correlation between time of the work day, and the frequency of data perturbation.

BACKGROUND

Building commissioning has emerged as the preferred method of ensuring that building systems are installed and operated to provide the performance envisioned by the designer. While most commissioning processes focus on bringing building operation to the original design intent, continuous commissioning focuses on optimizing the HVAC system and operation control for existing building condition. Continuous commissioning is an ongoing process to resolve operating problems, improve comfort, optimize energy use and identify retrofits for existing commercial and institutional buildings and central plant facilities (Continuous Commissioning Guidebook, 2002). The main focus of continuous commissioning, however, is to optimize the building's operations costs in terms of decreased energy costs.

Until recently, commissioning agents relied on data from building energy management systems (BEMS) to gather information about the indoor environment state. The BEMS data, at best, provided a crude approximation of the real building state due to lack of ubiquity of the sensors being used (Isakson et al., 2004).

The ability to continuously monitor the environment using sensors and to be able to store that information opens new opportunities in integrating the indoor environment state and thus increasing the benefits from optimizing occupant's comfort, in the continuous commissioning process. Current developments in sensing technologies have enhanced the ability to continuously collect information about facility performance and the indoor environment (Estrin et al., 1999, Fraser et al., 2003, Mehta, et al., 2002). There is a need for efficient storage, retrieval and processing mechanisms for the large number of distributed data streams and to build various real time monitoring applications and decision support systems (Gibbons, 2003).

CMU CRITTER SENSOR

In order to pervasively deploy a large network of sensor nodes, I have developed the Critter Dataport Sensor, shown in Figure 1. The Critter is an analog computer-attached sensor device developed at Carnegie Mellon University, and licensed to and manufactured by a Carnegie Mellon spin-off company, Pervasive Sensors. It is an economical (and thereby pervasively deployable) alternative to all-in-one sensor devices, which include sensing, processing, memory, and networking. A sensor network is formed by the Critter-attached computers using the existing host networking. The Critter platform, while attached to a computer with significant resources from which it draws its power, is merely a guest on that host computer and cannot dominate its resources. By leveraging the computing, power, and communications resources of the host, we are able to leapfrog some of the issues that are a current focus of wireless sensor research and look towards future research issues that will be common of wired or wireless sensors.

Figure 1. CMU Critter Sensor Attached to Desktop

The temperature-sensing version of the Critter uses a thermistor, a resistance-based temperature detector. A thermistor is made from a material for which its resistance decreases as its temperature increases. A current is passed through the thermistor and the resistance is measured. The Critters attach to the game port of a computer, which is then used to read the temperature sensor value. While game port sensors are not novel, the Critter uses them in conjunction with a CMU-developed auto-calibration mechanism that will automatically calibrate a resistance-based sensor without user interaction (patent pending). This leads to extremely low installation and maintenance costs of the sensor system compared to existing technologies. Similarly packaged Critter sensors have been built with other detectors including light, humidity, water, and door open/close. Assuming availability of ports, Critters can also be used on laptop or mobile computers.

The game port was selected in order to leverage the existing analog-to-digital electronics within a computer. This allows for a very economical production of these sensors - on the order of $5 compared to hundreds of dollars. These factors inspire cost-effective testbeds of thousands of nodes. Alternatively, a data acquisition device (PCI card, serial or parallel port adapter) would be used. There are also digital sensors, which are analog sensors with built-in electronics to do the analog-to-digital conversion and connect to either the parallel or serial port of the computer. As the Critter uses the game port to provide the basic analog-to-digital electronics, removing the need for additional electronics, they can leverage the existing substantial investment in information and communications technology. Economical, the

Critter has been able to be deployed en masse through the West Wing of Hamburg Hall occupied by the Institute for Complex Engineered Systems. This includes individual Critters in offices and laboratories, as well as multiple Critters installed to understand the value of sensor redundancy.

Our current system works such that once the Critter is attached to a computer; the installed host software will query the detector at periodic intervals. The host processor can send the data to a remote location on a network or store it locally for later retrieval. Additionally, the host processor can be used to process the data locally, depending on configuration.

The Critter has been used to understand the full nature of sensor failure. Off the shelf integrated sensor devices incorporate some amount of adaptation to make the devices more reliable. For instance, they use a weighted averaging scheme to account for large fluxuations in temperature due to user activity. The Critter provides raw instantaneous readings including anomalies and perturbations. This allows me to construct a more accurate fault model than can be obtained using only off the shelf sensor components.

TESTBED

We instrumented one wing of an academic building comprising laboratories, faculty and graduate student offices with fifteen Critter sensor devices. The floor layout is shown in Figure 2.

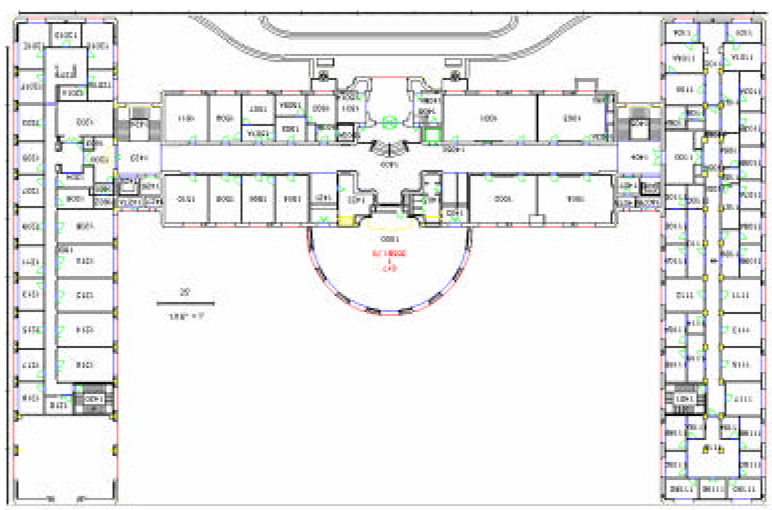

Figure 2. Floor Layout

The original goal was to understand the temperature variations based on the usage. These rooms are served by the same HVAC system. The Critter sensors were placed into the offices over a period of almost 18 months, from January 2004 to June 2005. The sensors ran continuously except for periods of downtime that included security patches, power failures, and network connection problems. The sensors were configured to collect a temperature reading every second. The total number of records by room is shown in Table 1.

Table 1. Total Number of Records By Room

Room Type and Number	Number of Records (in Millions)
Laboratory A1	76.6
Laboratory A2	87.5
Laboratory B1	83.8
Laboratory B2	61.3
Admin 1	58.6
Admin 2	72.5
Faculty 1	27.6
Faculty 2	58.9
Faculty 3	56.9
Faculty 4	32.9
Faculty 5	32.9
Faculty 6	6.7
Grad 1	85.3
Grad 2	80.7
Grad 3	72.5

While the devices collected a different number of records, the downtime of collection was uniformly distributed and not clustered at a consistent time of day. The Critter sensor device in Faculty Office #6 was removed after only a few months. One data stream from a laboratory is shown in Figure 3 as an example of the data collected. The temperature ranged from a low of 65.9 to a high of 84.8.

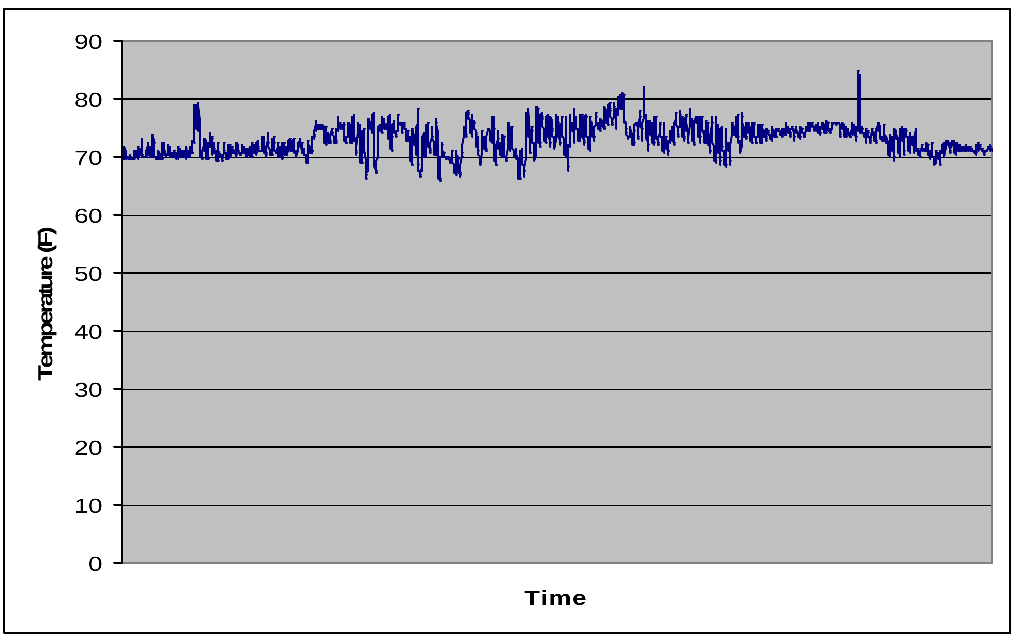

Figure 3. Example Temperature Data Stream

ANALYSIS

The original inspection of the data collected revealed what we first termed to be anomalies in the data. Two subsequent temperature readings taken a second apart could differ upwards of 4 degrees Fahrenheit. We considered these anomalies to be a reliability problem with our sensor device and considered pre-processing of the data to eliminate these anomalies. Unlike commercial sensors, the Critter does no processing of the data such as a weighted averaging.

Anomalies were extracted from the full data set of all 15 temperature sensor data streams for the entire 18 month period. This involved calculation of the difference between one temperature data point and the immediately following data point collected in the next second. An anomaly is determined as two subsequent readings being four or more degrees Fahrenheit apart. From this list, the hour was extracted regardless of the difference between readings. The hour was summed with the result is shown in Figure 4.

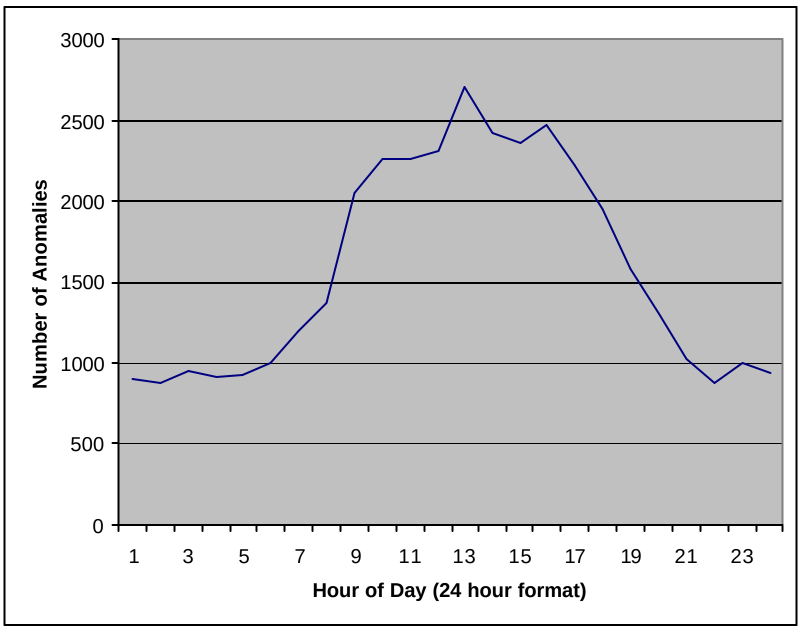

Figure 4. Anomalies Greater Than 4 Degrees

This graph shows a strong relationship between a work day and the number of anomalies. We believe that these differences are actually perturbations in the environment caused by human activity. There is a rise in the number of perturbations starting around 8AM and ending around 8PM. This corresponds with the human activity in a typical academic building that houses a mix of administrative personnel that work early in the day and research personnel that tend to work later in the day. The significant perturbations occur between 11AM and 3PM.

A second analysis was performed with the same method, only classifying anomalies as two subsequent readings of being nine or more degrees Fahrenheit apart. The result is shown in Figure 5.

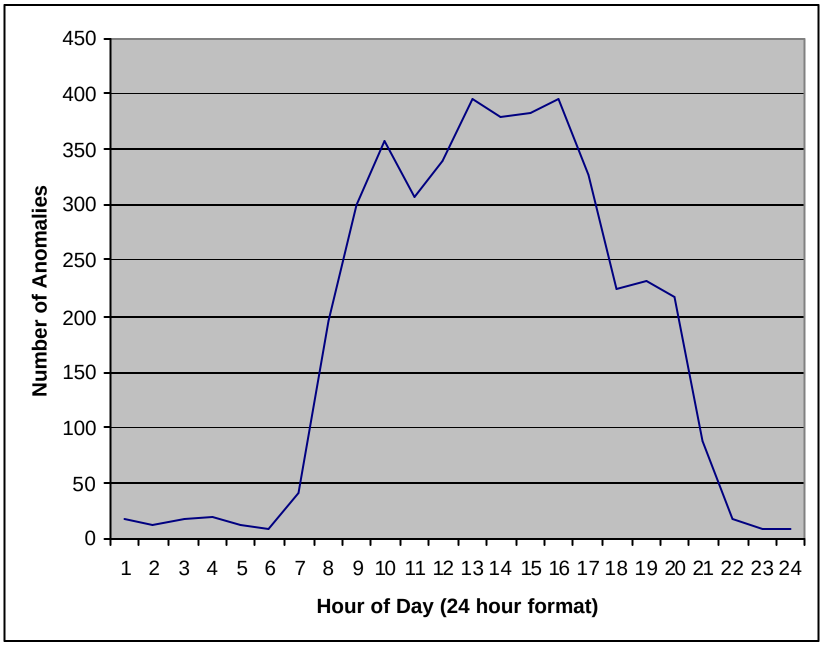

Figure 5. Anomalies Greater Than 9 Degrees

The same rise in the number of perturbations occurs at 8AM, ends around 8PM, with the significant number of perturbations centered in the middle of the day.

CONCLUSIONS

The Critter has been developed to understand the full nature of pervasive sensor networks. Off the shelf integrated sensor devices incorporate some amount of adaptation to make the devices more reliable such as a weighted averaging scheme to account for large fluxuations in temperature, primarily considered to be sensor failures. The Critter provides raw instantaneous readings including outlier data that may be considered anomalies or perturbations.

We have deployed Critter data sensors pervasively through one academic building for almost 18 months. This paper explored the causes of temperature data perturbations, two temperature data readings taken within seconds of each other that differ by several degrees. Temperature sensor data perturbations are actually the effects of user activity within buildings. By capturing the raw data without automatic processing, we are able to show a correlation between time of the work day, and the frequency of data perturbation.

ACKNOWLEDGMENTS

This work is supported in part by the Pennsylvania Infrastructure Technology Alliance, a partnership of Carnegie Mellon, Lehigh University, and the Commonwealth of

Pennsylvania's Department of Community and Economic Development (DCED). We want to thank Veritas Software for donating software that was used for data collection and storage, and Kustrim Kuqi who managed the data collection process.

REFERENCES

Continuous Commissioning Guidebook (2002), "http://www.eere.energy.gov/femp/operations_maintenance/commissioning_guidebook.cfm"

Estrin, D., Govindan, R., Heidemann, J., Kumar, S.,(1999) "Next Century Challenges: Scalable Coordination in Sensor Networks," Proceedings of the 5th Annual International Conference on Mobile Computing and Networks (MobiCOM '99), Seattle WA, August

Fraser, M., et al., (2003) "Elements of an Integrated Health Monitoring Framework," Proceedings of the 2003 SPIE Conference on Nondestructive Evaluation for Health Monitoring and Diagnostics, San Diego CA.

Gibbons, P., Karp, B., Ke, Y., Nath, S., Seshan, S., (2003), "IrisNet: An Architecture for a World Wide Sensor Web", IEEE Pervasive Computing.

Issakson, P., Wetterstrom, P., Carling, P., (2004), "BEMS Assisted Seasonal Functional Performance Testing in the Initial Commissioning of KIsta Entre and Katsan", International Conference for Enhanced Building Operations, Paris, France, 2004.

Mehta, V., and El Zarki, M., (2002), "Fixed Sensor Networks for Civil Infrastructure Monitoring – An Initial Study", First Annual Meditteranean Ad Hoc Networking Workshop, Med-hoc-Net 2002, Sardegna Italy.

.