Poster Abstract: Enabling Reliable and High-Fidelity Data Center Sensing

Chieh-Jan Mike Liang†, Jie Liu‡, Liqian Luo⋆, Andreas Terzis†

† Johns Hopkins University ‡ Microsoft Research ⋆ Google

ABSTRACT

RACNet is a sensor network that monitors a data center's environmental conditions at high temporal and spatial resolutions. The data RACNet collects can improve the energy efficiency of data centers, currently one of the fastest growing energy consumers in the U.S. RACNet overcomes the challenge of reliable and low-latency data gathering from dense networks deployed in harsh RF environments through rDCP, a novel network protocol that decouples data collection from topology control. Furthermore, rDCP automatically partitions the network to multiple routing trees, each operating at a different frequency channel. The combination of these features enables rDCP to scale up while maintaining high data yields. Preliminary results from a production deployment of 694 sensors (including 174 wireless nodes) show that rDCP achieves a data yield of 99%, while delivering 90% of the measurements in less than 30 seconds.

1. INTRODUCTION

Data center energy consumption has attracted global attention due to the fast growth of the IT industry and increasing concerns about carbon footprints and global climate change. Approximately 40% to 70% of the total energy that a typical data center uses is lost in the power distribution system or used by environmental control systems such as Computer Room Air Conditioning (CRAC) units, water chillers, and (de)humidifiers [1, 6]. A root cause for this low energy efficiency is the lack of visibility in the data center's operating conditions. For example, when servers issue thermal alarms, data center operators tend to decrease the CRAC's temperature settings, which might not alleviate the problem [4].

Wireless sensor network (WSN) technology offers many advantages for monitoring and control tasks. It is low-cost, non-intrusive, and can be easily re-purposed. Furthermore, WSNs require no additional network and facility infrastructure, simplifying their deployment in the already complex data center IT environment. At the same time, data center monitoring introduces new challenges for WSNs. Temperature can vary by as much as 5°C over a couple of meters inside a data center, which means that thousands of sensors might be necessary to cover a single colocation facility. Moreover, Liang et al. showed that 83% of the motes in a colocation are within one-hop distance from each other [3]. The combination of large scale and high density complicate the task of reliable and low-latency data collection. To make things worse, the RF environment inside a data center is harsh as it includes multiple metallic obstacles and interference from WiFi networks.

RACNet uses custom-made motes, called Genomotes, that form sensor chains using off-the-shelf USB cables. These cables deliver measurements, and they also distribute power from a server's USB port to all the sensors on the same server rack. Each chain consists of multiple sensing slaves and a master that is responsible for relaying the measurements over its radio. An added benefit of the chain configuration is that it reduces the number of nodes in the wireless network thereby reducing contention. Both mote types use an MSP430 microcontroller, while the master also has a TI CC2420 802.15.4 radio, 1MB of flash for data buffering, and a battery for continuous operation in case of server down time.

Figure 1 provides an architectural overview of RACNet's reliable Data Collection Protocol (rDCP). First, rDCP uses a distributed protocol for maintaining bi-directional routing trees (BiTrees) used for data collection. Multiple BiTrees can coexist, each rooted at a gateway that occupies a different frequency channel. Membership in individual trees dynamically adapts to the number of gateways in the network as well as link

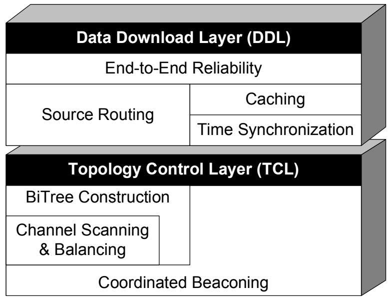

Figure 1: Reliable Data Collection Protocol (rDCP) architecture. The Topology Control Layer (TCL) maintains the tree topology while the Data Download Layer (DDL) reliably downloads data from individual master nodes.

qualities on each channel. Unlike centralized data collection protocols, such as Koala [5], the distributed topology control nature of rDCP allows it to promptly adapt to link quality changes and scale to large networks. Second, data downloads in rDCP are centrally controlled by the network's gateways. This is different from distributed data collection protocols, such as CTP [2], in that nodes wait for gateway requests before streaming measurements. This centralized coordination minimizes interference among network flows. At the same time, the pull-based data download approach allows the gateways to issue commands and negative acknowledgments used to implement reliability.

To minimize data latency, rDCP attempts to balance the load among all gateways, defined as the sum of hops necessary to reach all nodes in the tree. Once a gateway's load exceeds the network average by a predefined threshold, it calculates two probabilities according to its load relative to that of other gateways: A switch-out probability that determines the number of nodes in its own tree that should probe other channels, and a switch-in probability that determines the channel that those switched-out nodes should probe. The gateway then disseminates these probabilities to its tree's nodes. Upon receiving the command, each node probabilistically probes other channels and switches to the channel that provides the best parent link.

2. DEPLOYMENT RESULTS

We report results from one of the production deployments. This deployment consists of 694 Genomotes (including 174 wireless masters) in a 12,000 sq-ft Microsoft colocation facility. Slave Genomotes sample their temperature and humidity sensors every 30 seconds. The network uses up to four wireless channels. The system has been running for more than 7 months and collects more than 2.5 million measurement records per day.

Figure 2 presents the data yield over a 72-hour period. The percentage on the Y-axis is the ratio between the

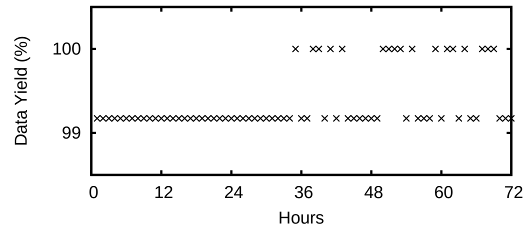

Figure 2: Average hourly data yield over a 72-hour period from the 694 sensors in the production deployment.

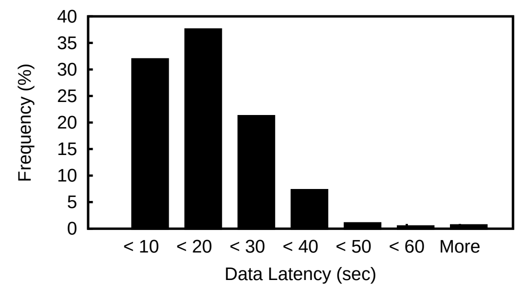

Figure 3: Data collection latency distribution of 10,000 data samples. The network had 3 gateways.

actual data received at the gateway during that hour and the expected amount of data, based on the sampling rate (i.e., 120 records per hour). One can see that the network has a data yield close to 100%. The minor data loss was the result of a misconfiguration where the base stations immaturely gave up on a record after only a few failed tries. Figure 3 shows that over 90% data have a latency of less than 30 seconds.

Last but not least, a considerable amount of effort was necessary to deploy thousands of sensors. One major pre-deployment overhead we have observed is assigning unique sensor addresses as this requires bookkeeping and multiple versions of the software that differ in only a single variable. To streamline the process, we include a 64-bit serial ID chip and a matching barcode on each Genomote. Since the collected measurements have spatial significance, the barcode minimizes human error at all stages of the deployment.

3. REFERENCES

[1] C. L. Belady. In the data center, power and cooling costs more than the it equipment it supports. *ElectronicsCooling magazine*, 3(1), February 2007.

[2] O. Gnawali, R. Fonseca, K. Jamieson, and P. Levis. Robust and efficient collection through control and data plane integration. Technical Report SING-08-02, 2008.

[3] C.-J. M. Liang, J. Liu, L. Luo, and A. Terzis. Racnet: Reliable acquisition network for high-fidelity data center sensing. Technical Report MSR-TR-2008-144, Microsoft Research, 2008.

[4] J. Liu, B. Priyantha, F. Zhao, C.-J. M. Liang, Q. Wang, and S. James. Towards fine-grained data center cooling monitoring using racnet. In *HotEmNets*, 2008.

[5] R. Musaloiu-E., C.-J. M. Liang, and A. Terzis. Koala: Ultra-Low Power Data Retrieval in Wireless Sensor Networks. In *IPSN*, 2008.

[6] The Green Grid. The green grid data center power efficiency metrics: PUE and DCiE, 2007.