Poster Abstract: Bandwidth Management in Wireless Sensor Networks

Bret Hull, Kyle Jamieson, and Hari Balakrishnan
MIT Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139
{bwhull, jamieson, hari}@lcs.mit.edu

Categories and Subject Descriptors

C.2.1 [**Computer Systems Organization**]: Computer-Communication Networks—*Network Architecture and Design*; C.2.4 [**Computer Systems Organization**]: Computer-Communication Networks—*Distributed Systems*

General Terms

Measurement, Performance, Design, Experimentation.

Keywords

Sensor networks, Flow Control.

1. INTRODUCTION

Wireless sensor networks are often used in remote monitoring and control applications. As sensor nodes get even smaller and sensor networks grow larger in size, we believe that bandwidth considerations will become increasingly imperative.

Our research addresses three important problems caused by constrained wireless link bandwidth in sensor networks:

1. How to allocate network bandwidth to sensor streams?
2. How to control network congestion?
3. How to improve forwarding capacity of the network?

The context for our work is a centrally administered, shared sensor network infrastructure, used concurrently by different applications. Due to the asynchronous nature of external events, sensor networks require a *bandwidth allocation* method by which nodes can decide how to allocate network bandwidth to sensor streams. The allocation method has to handle traffic that exhibits a high degree of spatial correlation and it must be able to change bandwidth allocations in response to observed phenomena.

Our solution to this problem employs *semantic priorities* that allow each node to map any sensor data stream to a desired traffic rate and class. This mapping may change over time depending on the sensor's values and other attributes (eg., the location of the sensed data). These semantic priorities work in conjunction with a static priority scheduling discipline that provides bandwidth separation between the traffic classes of the network, ensuring that high priority traffic will make it through the network.

Any network with constrained bandwidth requires *congestion control* to achieve good performance and avoid catastrophic congestion collapse when the offered load exceeds available network capacity. Since Internet-style, end-to-end congestion control protocols are not well-suited to our domain, we have developed a *hop-by-hop flow control* scheme that achieves good performance with low congestion-triggered packet loss rates.

Finally, sensor network applications traditionally forward the data they produce to a network bridge called a sensor access point (SAP). In order to increase the aggregate forwarding capacity of a sensor network it makes sense to deploy multiple SAPs. However, it is not clear how nodes should pick a SAP to avoid hot spots and realize this higher capacity. As a first step toward achieving higher utilization in networks with multiple SAPs, we have developed a new routing metric that uses hop-count, path load, and wireless path loss rate as the basis for selecting paths to SAPs.

2. DESIGN

Our architecture for bandwidth management in wireless sensor networks includes three main components: a semantic priority system to handle bandwidth allocation, hop-by-hop flow control to mitigate congestion, and a SAP selection metric for load balancing. The following sections briefly describe these components.

2.1 Semantic Priorities

Semantic priorities allow applications and network administrators to specify how data packets should be treated by nodes in the sensor network. Semantic priorities are expressed using rules, and describe an abstract function F that takes the sensor's data type, the value being reported, and other attributes (to the stream, etc.) as input and produces a rate and traffic class.

$$F(\text{type, value, attr}) \longrightarrow \{\text{rate, class}\}$$

Data is sent from a node to an SAP at a rate specified by the rule whose preconditions match the values and attributes of the sensor. The traffic class in a rule allows a network administrator as well as application writers to specify the relative importance of different data ranges for sensor streams. For example, one rule might specify that the temperature sensors provide a data rate of 4 packets/second,

SenSys'03, November 5–7, 2003, Los Angeles, California, USA.
ACM 1-58113-707-9/03/0011.

while a second more-specific rule might specify that temperature sensors that report values above 80 degrees from certain locations do so at 10 packets/second on a different traffic class that is treated at higher priority in the network.

Each node enforces semantic priorities in two ways: it implements a packet scheduling discipline to classify and forward packets in different traffic classes, and it implements a rate-control mechanism to throttle packets to the rates specified in their best-matching rule. The semantics we associate with a traffic class are *delay priority* rather than bandwidth sharing. That is, we choose a simple static priority forwarding scheme where packets are sent from the highest-priority queue that is not empty.

2.2 Hop-by-Hop Flow Control

Any network with competing and dynamically varying traffic requires some form of congestion control to avoid high packet loss rates, long queues, or congestion-triggered collapse. End-to-end flow control schemes like TCP are not well-suited to this domain. Our solution to this problem is to adapt the idea of *hop-by-hop flow control* to wireless sensor networks. The main thrust of hop-by-hop flow control is that a congested node provides immediate feedback to an upstream neighbor sending packets to it about the onset of congestion.

In our system, we send synchronous, link-layer NACKs (negative acknowledgments) when the instantaneous queue size at a node exceeds a high-water threshold. The packet causing this NACK will be enqueued, but the notification will cause the transmitting node to slow down. With this slow-down, transmitting nodes upstream from this part of the network will soon throttle back as queues fill up. By using the hop-by-hop scheme with synchronous NACKs, the expectation is that once a packet is admitted into the system, it is generally not dropped due to congestion.

An alternative design for rate control, previously explored by Woo and Culler [3], is to use implicit ACKs, whereby a packet is considered acknowledged if it is overheard being forwarded by a downstream neighbor. This scheme has several drawbacks, including an unknown delay between the time a packet is sent and the time implicit ACK is received (due to varying queue sizes) and an inability to handle application level suppression at transit nodes. Another congestion control design, CODA [2], uses backpressure messages either when nodes carrier-sense the channel to be busy more than a certain fraction of the time, or when packet queues start to grow.

2.3 SAP Selection

Every node in the sensor network needs to select a "good" path to send data toward an access point. We are interested in not the best path to a given destination node, but the best path to any one of several possible SAPs. Previous research has shown that a pure hop-count metric for path selection, standard in wired routing protocols, is not appropriate for wireless networks because picking a path with the smallest number of hops tends to pick marginal links that are maximally separated from each other, whose link quality is therefore bad [4].

In our network, nodes don't retransmit lost packets at the link-layer, as in 802.11. Furthermore, because our network uses hop-by-hop flow control to reduce congestion-related packet losses, proactively avoiding paths that have a higher traffic load is worthwhile. On any path, a measure of the *benefit* of using a path is proportional to the end-to-end packet success rate, and inversely proportional to the aggregate traffic load along the path. The *cost* of using the path, in the form of increased bandwidth used, is proportional to the number of hops it has. Together, the ratio of cost to benefit gives us the following path metric, M_P, for a path P:

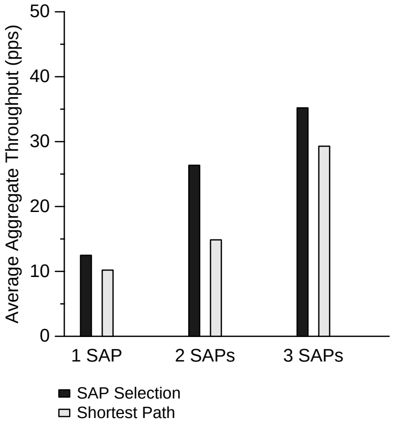

Figure 1: This figure shows the effects of the SAP selection metric on the aggregate throughput of a 40 node sensor network deployed throughout a floor of an office building.

$$M_P = \frac{\#\text{hops}(P) \cdot \max_P \text{Load}}{\text{path\_success\_prob}} \qquad (1)$$

Each node picks the path P with smallest value of M_P. Disseminating the information required to pick an appropriate path is easy to do using a distance-vector protocol, because the three components of the metric—"hops", "load", and "path success probability"—are all associative metrics. Each node propagates these three different quantities in distance-vector fashion corresponding to the path to the node's current estimate of the "best" SAP. Additionally, since nodes only need to keep track of the metric for their directly connected neighbors, per node state is kept quite low.

Figure 1 shows how our proposed metric compares to a simple shortest path metric on a 40 node sensor network deployed throughout one floor of an office building.

For a more detailed treatment of the work presented in this abstract, please see [1].

3. REFERENCES

[1] Hull, B., Jamieson, K., and Balakrishnan, H. Bandwidth management in wireless sensor networks. Tech. Rep. 909, MIT Laboratory for Computer Science, July 2003.

[2] Wan, C.-Y., Eisenmen, S., and Campbell, A. CODA: Congestion Control in Sensor Networks. To appear in *Proc. of ACM Sensys Conf.* (Los Angeles, CA, November 2003).

[3] Woo, A., and Culler, D. A Transmission Control Scheme for Media Access in Sensor Networks. In *Proc. of ACM MobiCom Conf.* (Rome, Italy, 2001).

[4] Yarvis, M., Conner, W., Krishnamurthy, L., Mainwaring, A., Chabra, J., and Elliott, B. Real-World Experiences with an Interactive Ad Hoc Sensor Network. In *Proceedings of the IEEE Intl. Conf. on Parallel Processing* (Vancouver, BC, August 2002), pp. 143–152.