Poster Abstract: A Case for All-Wireless, Dual-Radio Virtual Sinks

†Chieh-Yih Wan, †Andrew T. Campbell, ‡Jon Crowcroft
†Dept. of Electrical Engineering, Columbia University, New York, NY 10027
‡Computer Lab, Cambridge University Cambridge, CB3 0FD, UK
†{wan,campbell}@comet.columbia.edu, ‡Jon.Crowcroft@cl.cam.ac.uk

Categories and Subject Descriptors

C.2.1 [**Computer-Communications Networks**]: Network Protocols, Wireless Communications.

General Terms

Algorithms, Design, Performance.

Keywords

Wireless sensor networks, congestion avoidance.

1. INTRODUCTION

Wireless sensor networks are emerging technologies that offer low cost, distributed monitoring solutions for a wide variety of applications and systems. One application driving the development of sensor networks is the reporting of conditions within a region of interest where the environment can abruptly change due to an sudden event, such as enemy and target movements on the battlefield, or biochemical attacks, fires, etc. Our work focuses on sensor systems that need to efficiently deliver information during and immediately following an event that triggers such an abrupt change.

Congestion control and load balancing are critical issues in such sensor networks where the sensor field can move instantaneously from almost zero load to overload conditions. It is during these impulse or overload periods that the events in transit are of most importance and most likely to be lost due to congestion. Existing congestion control algorithms [1] [2] [3] are limited under these conditions because they rely on rate control or packet drop mechanisms at source or intermediate sensor nodes that can significantly impact the application's fidelity, as measured at one or more physical sinks.

We propose randomly distributing a small number of all-wireless dual radio *virtual sinks* throughout the sensor field that are capable of offering the existing low-power sensor network enhanced congestion avoidance support when persistent congestion is detected. In essence virtual sinks operate as safety valves in the sensor field to selectively siphon off high load traffic in order to maintain the fidelity of the application signal at the physical sink and to alleviate the funneling effect, as discussed in Section 2.1.

We call these specialized nodes virtual sinks to distinguish them from physical sinks, which typically have a wireline interface that provides a gateway to the Internet. Virtual sinks are equipped with a secondary long-range radio interface, such as IEEE 802.11, in addition to their primary low power mote radio. Virtual sinks are capable of dynamically forming a secondary ad hoc radio network that is rooted at a physical sink. Rather than rate controlling or dropping packets as in the case of the first generation congestion avoidance techniques [1], virtual sinks take the congested traffic off the low-powered sensor network (i.e., off the primary radio network) when persistent congestion is detected, and move it on to the secondary radio network, transiting it to a physical sink.

SenSys'04, November 3–5, 2004, Baltimore, Maryland, USA.
ACM 1-58113-879-2/04/0011.

2. DESIGN SPACE

A number of questions arise when studying the deployment of virtual sinks. What is the optimal number and distribution of virtual sinks to minimize congestion and energy consumption? Utilizing a longer-range radio is usually demanding in terms of energy consumption. Therefore, one should only activate the secondary long-range radio when it is needed. When does a virtual sink offer such hotspot services to local sensors? How do sensors discover local virtual sinks? When congestion or overload conditions occur which packets should be redirected onto the secondary long-range network? How can sensor networks automatically benefit from the existence of virtual sinks in their neighborhoods, but maintain uninterrupted services in their absences using the existing congestion avoidance mechanisms such as those discussed earlier [1] [2] [3]? What if the virtual sinks cannot form a fully connected network with the physical sink on their own? In what follows, we briefly explore these questions and discuss the technical considerations that underpin the design space.

2.1 Funneling Effect and Load Balancing

Conventional networks assume traffic flows in all directions. However, sensor networks exhibit a unique *funneling effect* where events are suddenly generated and then have to quickly move toward a relatively small number of physical sink points. Sensors within the range of an event epicenter generate impulse data that travels along a propagation funnel toward a sink when an event triggers. The funneling effect places heavier load on sensors that are closer to a sink point, causing them to use energy at a faster rate, significantly impacting the operational lifetime of the network. Meanwhile, traffic intensifies at the neck of the funnel toward the sink causing congestion, packet loss, and therefore, wasted energy and bandwidth. The aggregation of data

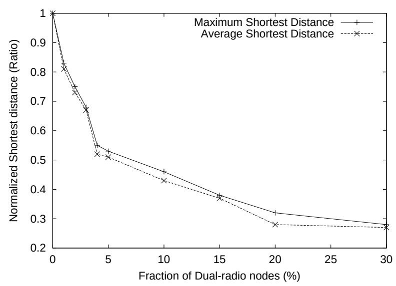

Figure 1: Reduction of average distance in a network with increasing percentage of dual-radio nodes.

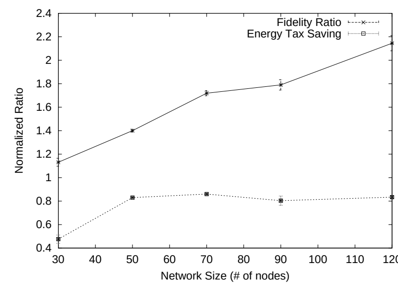

Figure 2: Fidelity and Energy Tax performance.

events can help offset congestion but it is unlikely that aggregation techniques alone can completely resolve the congestion problem and funneling effect. Because of the build up of traffic close to the sink, loss of aggregated data packets is more likely.

2.2 Small World Observations and Shortcuts

By using specialized dual-radio nodes as virtual sinks, the secondary long-range radio can serve the purpose of creating "shortcuts" in the sensor network among other virtual sinks and one or more physical sinks. While dual-radio sensor platforms (e.g., Stargate[1]) are feasible, the cost is still much higher than a single-radio platform (e.g., Berkeley motes series). Therefore, the cost of deploying a large number of dual-radio nodes in sensor networks is prohibitive. Recent "small worlds" studies conducted by Watts and Strogatz [4] has shown that a small fraction of shortcut nodes randomly distributed in a network is enough to effectively reduce the network diameter resulting in a fast distribution network. Following this, we conjecture that only a small fraction of shortcut nodes (i.e., virtual sinks) would be needed to create a fast secondary radio distribution network for overload traffic. Figure 1 examines this conjecture by simulating a sensor network of 100 nodes using the ns-2 simulator in an area of dimensions 350m x 350m with random sensor positions. The transmission radii of the sensors is $30m$ for the primary low power radio and $150m$ for the secondary long-range radio. Figure 1 shows that when only 5% of virtual sinks exist in the network, the average distance is halved.

2.3 Traffic Redirection

Virtual sinks resemble local sinks that "attract" part of the traffic from a congested neighborhood (hence the name), and sends it over a long-range radio link that provides a shortcut distribution path toward a physical sink. Utilizing long-range radio is usually expensive in terms of energy consumption, therefore, a virtual sink should only activate its secondary radio when it is needed and beneficial. Assuming nodes can measure and quantify their local congestion levels [1], we need to study at what point should sensors start (or stop) using local virtual sink services, assuming there is a virtual sink in the neighborhood.

An important aspect of traffic redirection is the algorithms for virtual sink discovery and visibility scope control. A virtual sink announce its existence to its neighbors to up to l hops away. With a larger value of l, more nodes are able to utilize a local virtual sink but these nodes may not necessary benefit from using the virtual sink because of possible longer paths that redirects traffic through other nodes to the virtual sink. On the other hand, with a smaller value of l, shorter redirecting paths can improve the delivery latency and energy consumption although fewer nodes would be able to utilize them. We need to study the tradeoffs and the optimum value for scope control.

[1] http://www.xbow.com/Products/Product-pdf-files

3. INITIAL RESULTS

We are currently developing algorithms for virtual sink discovery, selection, transit, and detection, and evaluating them using ns-2 simulation and through experimentation in a large-scale Stargate virtual sinks and Mica2 motes testbed. Figure 2 presents ns-2 simulation results for the fidelity ratio and energy tax savings [1] performance for our initial algorithms in a set of networks of different sizes. In each simulation, 5% of the nodes are randomly selected to be the virtual sinks. In this scenario, the virtual sinks are uniformly distributed across the sensor field and form a fully connected secondary network using long range radios. Always-on virtual sinks are utilized whenever they are able to deliver data events to the sink with lower delay. Figure 2 shows that our initial algorithms are able to obtain greater fidelity gain in a larger network, although the energy gain does not follow the same trend; it levels off for networks of 50 nodes in size providing up to 80% better performance in terms of energy tax savings. We are currently studying additional techniques that will push this performance benefit to much larger network configurations.

4. REFERENCES

[1] C-Y Wan, S. B. Eisenman, and A. T. Campbell. CODA: Congestion detection and avoidance in sensor networks. In *Proc. ACM SenSys 2003*, pages 266–279. Los Angeles, November 5-7 2003.

[2] Y. Sankarasubramaniam, O. Akan, and I. Akyildiz. Event-to-sink reliable transport in wireless sensor networks. In *Proc. of the 4th ACM Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc 2003)*, pages 177–188. Annapolis, Maryland, June 2003.

[3] B. Hull, K. Jamieson, and H. Balakrishnan. Bandwidth management in wireless sensor networks. In *Proc. ACM SenSys 2003*, pages 306–307. Los Angeles, November 5-7 2003.

[4] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393:440–442, 1998.