On Scalability in Sensor Networks

Upamanyu Madhow
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106, USA
Email: madhow@ece.ucsb.edu

***Abstract*— Three complementary approaches to scaling sensor networks up in the number of nodes and in energy efficiency are surveyed. Imaging Sensor Nets provide a massively scalable architecture for sensor networks with tens of thousands of nodes. Distributed Beamforming provides a gain in energy efficiency proportional to the number of nodes collaborating in the beamforming. Biased Resource Allocation provides a framework for scaling multihop sensor nets, combining in-network processing with joint routing/scheduling strategies that bias against connections that take up too many network resources.**

I. INTRODUCTION

A number of technical hurdles must be surmounted in order to attain the compelling vision of pervasive, networked sensor nodes that give us an ”up-close” view of our world. While there are a host of specific implementation challenges, the fundamental bottleneck to this vision is arguably one of scalability: (a) how to scale up the number of nodes so as to get adequate coverage of an area of interest, and (b) how to scale down the energy consumption per node so as to get by with energy scavenging, or to eke out an acceptable lifetime (months or years) with existing battery technology.

There is no one architecture that “solves” the problem of scalability, which manifests itself in different ways for different applications of sensor networks. In this paper, we describe three different approaches that solve specific subproblems related to scalability. While it is possible to combine the ideas in these different approaches, no attempt is made to do so here. We limit ourselves here to describing the “big picture” concepts. Details regarding the first two topics are available in some of our other publications, while we are in the processing of generating preliminary results on the third topic.

II. MASSIVELY SCALABLE IMAGING SENSOR NETS

Very large scale sensor nets with tens of thousands of randomly deployed nodes can have a fundamental impact on applications ranging from homeland security to interplanetary exploration. Conventional multihop wireless networking paradigms simply do not scale for data collection from such sensor nets. Furthermore, while the location of important events is a crucial element of the information we wish to gather, nodes in a large random deployment may not know their own locations, since localization techniques such as GPS may be unavailable either due to cost (e.g., of building in a GPS receiver) or environment (e.g. jamming or blockage of weak GPS signals). We use imaging as inspiration for jointly addressing the problems of scale and localization: the number of pixels in an image can scale to millions, being limited only by the size of the receiver’s “lens,” and the imaging device, if calibrated and location-aware, knows which location each pixel corresponds to. We combine the scalability of imaging with the generality of sensor networks by interpreting sensor nodes as pixels imaged by one or more sophisticated collector nodes. In one possible realization, a collector node scans the sensor field systematically using an RF beacon. Sensor nodes electronically reflect the beacon, possibly also modulating it with low-rate data. The collector processes the content and timing of the responses from the sensor nodes using radar, imaging and data demodulation techniques. By moving the complexity to the collector node, we eliminate the need for geolocation or networking in the sensor nodes. Indeed, the component sensor nodes need not even have a unique identity, since the association of their measurement with their location is performed by the collector. The elimination of network overhead implies that the data transmitted by the sensors can be drastically reduced, thus yielding link budgets that can sustain ranges of 10s or 100s of kilometers. This in turn enables an architecture in which one or more collectors directly communicate with sensors, bypassing the scaling problems associated with multihop transmission between sensors. Integration of the“views” obtained by multiple collectors (networked by conventional means) enhances coverage, spatial resolution, and estimation of sensor data.

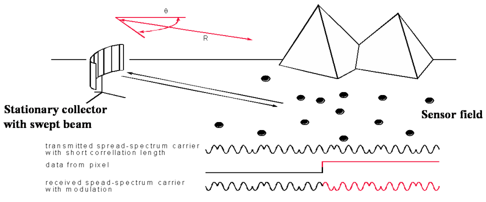

(waveforms designed for accurate localization and low-rate data collection)

Fig. 1. Imaging Sensor Net with Stationary Collector

Results on localization performance for an idealized model appear in [1], [2], [3], hence we do not repeat them here. Our current efforts focus on building a prototype imaging sensor net at millimeter (mm) wave frequencies, including integrated circuit realization of the sensor nodes and “brassboard” implementation of the collector node. The use of mm wave frequencies, which are an order of magnitude higher than those used in wireless communication systems today, increases imaging resolution, since we can realize highly directional transmission and reception using a moderate sized collector antenna, and reduces the sensor node form factor, since low-gain mm wave antennas can be realized as patterns of metal printed on the circuit board. The prototype architecture is depicted in Figure 1.

III. Energy Efficiency via Distributed Beamforming

Next, we present an approach to enhancing energy efficiency or communication range via collaboration between a cluster of sensor nodes. Extending the range of communication given power and bandwidth constraints is a fundamental problem for any wireless communication network. However, it is particularly critical for data collection from low-power sensor nodes deployed in an inaccessible area, if the link budget from a single sensor node to the collector node is not adequate. We consider distributed transmit and receive beamforming for “closing” the link in such situations, with a cluster of nodes agreeing upon a common message, and emulating an antenna array to produce a beam towards a distant receiver. Fixing the power emitted per node, collaboration between ten nodes, for example, has the potential of producing up to a ten-fold increase in range: the net transmitted power goes up by a factor of ten, so that, assuming a beamforming gain of ten, the received power goes up by a factor of hundred, resulting in a range extension of a factor of ten for free space propagation. However, the difficulty in achieving these gains lies in ensuring that the signals emitted from the different nodes add up in phase at the receiver because the delay between each transmitting node and the receiver is typically unknown. Even if such delay estimates were available, even small errors could result in large changes in the carrier phase. We have explored such difficulties, as well as potential solutions to them, in prior work based on an open-loop master-slave architecture in [4], [5], in which a master node controls the phase of slave nodes. However, in more recent work [6], we have realized that phase coherence is far easier to attain in a closed loop architecture in which there is feedback from the receiver.

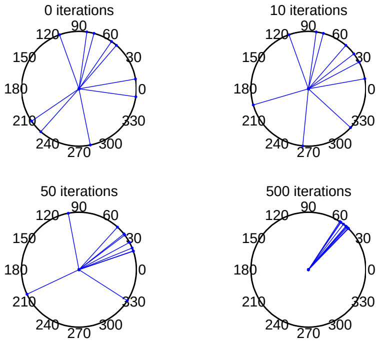

Fig. 2. Distributed beamforming with feedback control: alignment of phases using decentralized randomized probing (10 cooperating sensors)

The specific feedback strategy employed in [6] requires only one bit of feedback from the receiver to drive a completely decentralized adaptive algorithm, which ultimately results in the signals sent by the transmitters adding up coherently at the receiver. The transmitting nodes perturb their phases independently, and the receiver sends back a single bit indicating whether or not the net signal-to-noise ratio (SNR) has improved. If the SNR improves, then the transmitters keep the last perturbation; if not, they undo it. This is followed by another random perturbation of the phase. Our results show that a large fraction of the beamforming gain is attained quickly by this completely distributed algorithm. An analytical framework for optimizing the phase perturbations as a function of algorithm stage has been developed, and it has been shown that, when so optimized, the convergence time is linear in the number of nodes. A pictorial depiction of the evolution of the phases for the nodes is shown in Figure 2. See [6], [7] for details.

IV. Biasing Resource Allocation towards Scalability

For moderate scale sensor nets (e.g., with hundreds of nodes) with a conventional multihop architecture, the performance is limited by the negative results of Gupta and Kumar [8], which state that, if a source can choose its destination at random in a large network, then the throughput per node scales down with the network size. This result has been proved over and over again with various assumptions on the wireless environment in which the network operates, ranging from the classical narrowband collision channel to links operating at Shannon-theoretic limits. The reason why the result is so robust to the underlying assumptions has nothing to do with wireless, except that the broadcast nature of the wireless medium, and the rapid falloff in power with distance, forces us to use short hops. Once we decide to take short hops, Gupta-Kumar style results hold even if the links were replaced by wires. For example, for an N-node linear network, a randomly chosen source-destination pair is $O(N)$ hops apart. Elementary flow balance can then be used to show that the throughput attained under a maxmin fair allocation is $O(\frac{1}{N})$, and the sum throughput attained is $O(1)$. Given that the bottleneck is because typical connections in large networks take a large number of hops, an obvious approach to scalability is to bias against connections taking a large number of hops. Figure 3 shows a one-dimensional network, arranged in a ring to avoid edge effects. If we bias against long connections, using, for example, weighted round robin scheduling at the nodes with the weight of a connection inversely proportional to the number of hops it traverses, it can be shown that the throughput of an h-hop connection is approximately $\frac{1}{h}$, with a sum throughput of $\log N$. Comparing with a fair allocation, we realize that biasing against the long connections results in far better performance for the shorter connections, while reducing the throughput attained by the long connections only by roughly a factor of two.

In the context of sensor networks, it has been observed in prior work (e.g., [9]) that correlation in the sensor data can be

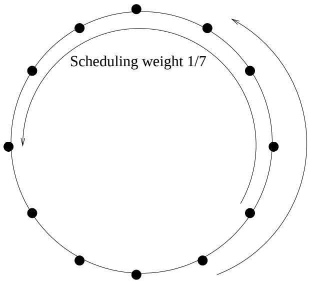

Fig. 3. Biased resource allocation for nodes arranged in a ring. Each link in the network uses weighted round robin, with the weight allocated to a connection inversely proportional to the number of hops it traverses.

exploited to potentially circumvent the Gupta-Kumar limits, since the latter assume independently generated traffic at each node. Our proposed framework would of course exploit such correlations using in-network processing involving relatively high-rate communication among clusters of nodes, followed by summaries being sent over a large number of hops to a data collection center. The specific biases implemented by joint routing and scheduling against long connections must therefore be adapted to the traffic generated by the specific sensing application, and and cross-layer optimization involving the application (sensing) layer, the network (routing) layer, and the link (scheduling) layer is required. This is very much work in progress, but the main point we wish to make is that it is possible to develop a systematic approach, abstracting away the vagaries of the wireless channel, for scaling multihop sensor networks, using the interplay of in-network processing and biased resource allocation.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation under grants EIA-0080134, ANI-0220118, CCF-0431205, and CNS-0520335, and by the Office of Naval Research under grants N00014-03-1-0090 and N00014-06-1-0066. The work surveyed here is in collaboration with a number of graduate students and colleagues at UCSB. The work on imaging sensor nets involves Bharath Ananthasubramaniam, Munkyo Seo, Prof. Mark Rodwell and Prof. B. S. Manjunath. The work on distributed beamforming involves Raghu Mudumbai, Dr. Gwen Barriac, and Prof. Joao Hespanha. The work on biased resource allocation involves Zhenwen Shao, Sumit Singh, and Prof. Elizabeth Belding.

REFERENCES

[1] B. Ananthasubramaniam, U. Madhow, “Virtual radar imaging for sensor networks,” *Proc. 3rd International Symposium on Information Processing in Sensor Networks (IPSN’04),* Berkeley, California, USA, April 2004.

[2] B. Ananthasubramaniam, U. Madhow, “Detection and Localization of Events in Imaging Sensor Nets,” *Proc. 2005 IEEE International Symposium on Information Theory (ISIT 2005),* Adelaide, Australia, September 2005.

[3] B. Ananthasubramaniam, U. Madhow, “On localization performance in imaging sensor nets,” submitted for publication (preprint available from author upon request).

[4] G. Barriac, R. Mudumbai, U. Madhow, “Distributed beamforming for information transfer in sensor networks,” *Proc. 3rd International Symposium on Information Processing in Sensor Networks (IPSN’04),* Berkeley, California, USA, April 2004.

[5] R. Mudumbai, G. Barriac, U. Madhow, “On the feasibility of distributed beamforming in wireless networks,” submitted for publication (preprint available from author upon request).

[6] R. Mudumbai, J. Hespanha, U. Madhow, G. Barriac, “Scalable Feedback Control for Distributed Beamforming in Sensor Networks,” *Proc. 2005 IEEE International Symposium on Information Theory (ISIT 2005),* Adelaide, Australia, September 2005.

[7] R. Mudumbai, J. Hespanha, U. Madhow, G. Barriac “Distributed transmit beamforming with feedback control,” submitted for publication (preprint available from author upon request).

[8] P. Gupta and P. R. Kumar, “The capacity of wireless networks,” *IEEE Trans. Information Theory,* vol. 46, no. 2., pp. 388-404, March 2000.

[9] A. Scaglione and S. Servetto, “On the interdependence of routing and data compression in multi-hop sensor networks,” *Proc. ACM Mobicom 2002.*