

Utility-Based Decision-Making in Wireless Sensor Networks

John Byers Gabriel Nasser

byers@cs.bu.edu comanche@cs.bu.edu

Computer Science Department
Boston University
Boston, MA 02215

Abstract—We consider challenges associated with application domains in which a large number of distributed, networked sensors must perform a sensing task repeatedly over time. We address issues such as resource constraints, utility associated with a sensing task, and achieving global objectives with only local information. We present a model for such applications, in which we define appropriate global objectives based on utility functions and specify a cost model for energy consumption. In the full version of this paper, we present algorithms and experimental results for this problem domain [2].

I. INTRODUCTION

In this paper, we argue that for many interesting applications of wireless sensor networks, a best-effort service model, in which nodes are expected to perform sensing operations and route data as best they can, may be too stringent. We adopt an economic decision model in which an activity is performed if its associated benefit outweighs its *opportunity cost*. A significant challenge here is the distributed nature of nodes in our networks, which implies that they do not have global information, making it unrealistic to expect nodes to accurately assess either the opportunity costs, or the relative benefits of a particular decision. Therefore, nodes in our model make heuristic assessments based on available local information in an attempt to optimize global objectives [1].

For the objectives we seek to address, computation in large-scale sensor networks will require scalable *coordination* amongst sensors to accomplish the desired tasks [3]. We consider global objective functions motivated by specific sensor network applications which are driven by *utility functions*, first studied in a networking context by Shenker [5]. Developing solutions which achieve these objectives are constrained in two primary ways: by the locality imposed by the distributed nature of the model, and by a resource constraint, namely the finite energy supply at sensor nodes. Our work develops a general model in which to study such problems and presents algorithmic results and experimental work in progress for a class of these problems.

While the objective functions and algorithms we propose are novel, they connect to a substantial body of work on ad-hoc routing protocols, fault tolerance, and energy conservation in sensor networks, which we survey in the full version of the paper. One work which considers several of the issues we consider here, including sensor fusion, or aggregating sensory information from multiple sources; load-balancing; and power conservation is the

Low-Energy Adaptive Clustering Hierarchy (*LEACH*) protocol [4]. The *LEACH* protocol [4] uses sensor fusion to compress datasets within the network, reducing the energy dissipated during the resulting transmission. One application-specific example they describe is *beamforming* algorithms, which combine a set of acoustic signals into a single signal without loss of relevant information. Our work applies the same general principle in advocating application-specific data aggregation as a technique for conserving energy.

II. MODEL

In this section, we present the network model we assume as well the objective functions we consider.

A. Network Model

We begin by assuming that every node i in our network, except for base stations, has a finite and non-replenishable reserve of energy p_i , a fixed transmission range R , a unique identifier, and that communication among nodes is *commutative*. The neighborhood of a node u is denoted by the set $N(u) = \{v \mid (u, v) \in E\}$. All transmissions from any node reach all the nodes in its neighborhood. For simplicity, we also assume that transmissions are *perfectly scheduled*, so that transmissions do not interfere with each other.

While we will generally assume that the nodes forming our ad hoc network are stationary for the algorithms we develop, this assumption is not an inherent limitation of our model.

We represent our network as an undirected graph $G = (V, E)$, where V is the set of all nodes, including the base stations. E is the set of edges in the network defined as follows, where $d(u, v)$ is the distance between nodes u and v .

$$E = \{(u, v) \mid u, v \in V \text{ and } d(u, v) \leq R\}$$

B. Sensing Model

We describe the sensing model we use in our network. We present the different costs associated with each operation, and explain our notion of *node specialization*, which is a network-adaptive role-based mechanism.

We propose to complement the role of adaptive algorithms with nodes that adapt their role, or *specialize*, as a response to changes not only in the virtual topology of the network but also to node power levels. In general, a node's typical role will be one of: idle, routing, sensing, or routing/sensing.

We assign the fixed costs c_s, c_t, c_r, c_a , to the four operations a node may perform: sensing, transmitting and receiving a fixed-size message, and aggregating sensory data, respectively. In our model, the ratio $\frac{c_a}{c_t}$, or the sense-to-transmit ratio, captures the relative importance of the two most important roles of nodes in the network in terms of energy consumption.

For many applications, it is not a requirement that the sensory data accumulated by the nodes of the network must be transmitted in full fidelity to the base station. Thus, we assume that data can be aggregated at each node with a fixed aggregation cost before it is transmitted upstream in the network.

C. Utility and Objective Functions

We associate each sensor domain with a monotonically non-decreasing utility function which maps the number of nodes participating in a sensory computation to a real value which measures the utility derived from output from a subset of sensors of that size. Shenker [5] motivated the use of utility functions in modeling a user's relative preference for a real-time stream encoded at varying levels of fidelity. Our motivation is similar – the user is the consumer of the output of the sensor network and the varying levels of fidelity correspond to increasingly detailed sensory output levels. Therefore we model the utility derived from a consumer of our sensor network resources by a monotone function $U : S \rightarrow [0, 1]$, which, for a network graph $G = (V, E)$, maps the *sensing subset* $S \subseteq V$, the set of all nodes in the graph that are sensing, to a real-valued interval. It is worth noting that for most applications, not all subsets of sensors of a given size are created equal – in many circumstances, having a geographically distributed set of reporting sensors is essential. Addressing such considerations is beyond the scope of this paper, but these considerations could be modeled by a more general (and non-monotone) utility function, $U : S^* \rightarrow [0, 1]$, mapping all possible subsets of nodes to values, where S^* denotes the power set of S .

First consider the utility curve to be represented by the step function on the left hand side of Figure 1. In this all-or-nothing case, useful data fusion is only possible when and only when the number of nodes participating in the sensing operation is at least as large as the threshold set by the function. In a more forgiving scenario, our utility function might resemble the inelastic curve on the right hand side of Figure 1, where we have some freedom in tuning the number of participating nodes to vary energy consumption in the network. This second curve has three regimes: when a very small number of nodes participate, the user derives little utility; at a certain threshold, the utility quickly increases dramatically; and then beyond a final inflection point, there are diminishing marginal returns and utility increases only very slowly. In both of these scenarios, ideally one would like to operate at the beginning of the third regime, just beyond the knee of the curve, to maximize utility relative to power consumption.

The objective function which we propose is one in which we maximize the sum, over the lifetime of the sensor network, of the utility of computations at intermediate timesteps. This objective reflects a natural goal – that of maximizing the total aggregated utility of the network over time. Let us define those vertices which elect to perform a sensing operation at any time




Fig. 1. Utility functions: Step function and inelastic utility function

t as the sensing subset $S_t \subseteq V$, and those transmitting data at any time t as the transmitting subset $R_t \subseteq V$. It follows that $S_t \subseteq R_t$, moreover, R_t must connect S_t and the base station.

Our objective function is the optimization problem:

$$\begin{aligned} & \text{maximize } \sum_t U(S_t) \text{ subject to:} \\ & \sum_t \sum_{i \in S_t} c_s + \sum_t \sum_{i \in R_t} c_t + (d_i - 1)c_r + c_a \leq p_i \\ & \forall t : R_t \subseteq V \end{aligned}$$

The first constraint of this formulation uses the cost model defined earlier to ensure that nodes cannot consume more power than they have available, where d_i denotes the degree of node i in the subgraph induced by R_t . The second constraint ensures that the data collected from all nodes who get credit for participating in the sensing subset at time t actually gets routed to the base station. This long-term strategy can be realized only through a combination of careful power management combined with distributed coordination on the part of the nodes in the sensor network in choosing their roles over time.

We provide algorithms and experimental results for this problem domain in the full version [2].

III. CONCLUSION

In a best-effort service model, nodes attempt to optimize the utilization of resources in the present without regard to future cost. In sensor networks, economic considerations, especially resource constraints, motivate objective functions which give substantial freedom in letting nodes choose their role over time. With the goal of optimizing the total utility derived over the lifetime of the network, the model we motivate enables nodes to discount current gains in lieu of future rewards; thereby optimizing their consumption of energy over time.

REFERENCES

- [1] Y. Bartal, J. Byers, and D. Raz. Achieving Global Objectives using Local Information with Applications to Flow Control. In *Proc. of the 38th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, Miami Beach, FL, October 1997.
- [2] J. Byers and G. Nasser. Utility-Based Decision-Making in Wireless Sensor Networks. Technical Report CS-TR-2000-014, Boston University, June 2000. URL = http://www.cs.bu.edu/fac/byers/pubs/sensor_utility_TR.ps.
- [3] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next Century Challenges: Scalable Coordination in Wireless Networks. In *Proc. of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, pages 263–270, Seattle, WA, 1999.
- [4] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-Efficient Communication Protocols for Wireless Microsensor Networks. In *Proc. of Hawaiian Int'l Conf. on Systems Science*, January 2000.
- [5] S. Shenker. Fundamental Design Issues for the Future Internet. *IEEE Journal on Selected Areas in Communications*, 13:1176–88, September 1995.