Title

Rethinking Data Fusion-Based Services in Tiered Sensor Networks

Authors

Dantu, Karthik
Sukhatme, Gaurav

Publication Date

2006-05-31

Rethinking Data Fusion-based Services in Tiered Sensor Networks

Karthik Dantu dantu@usc.edu
Gaurav Sukhatme gaurav@usc.edu

***Abstract*—** **Tiered sensor network architectures are gaining currency. In contrast with flat networks of impoverished nodes (the hitherto common assumption in sensor networking), such systems offer the promise of migrating computational load from sensing nodes to higher capability 'master' nodes. We argue that for certain data fusion-based services this means that compute intensive algorithms, often shunned as impractical for sensor networks, are in fact a viable possibility. Using localization as an example, we show how accurate results may be obtained by leveraging this capability without the use of specialized hardware or high configuration detail; both of which are standard approaches to the problem when computation is at a premium. Specifically, we propose a mathematical optimization-based framework for localization based on proximity constraints. Most variants of localization can be cast into this framework depending on the kinds of input available (e.g. ranging). We show accurate results, and exploit a technique from distributed optimization to divide the problem into pieces suitable for computation at the master-level nodes. We conclude with remarks on the general implications of this example for tiered systems, with pointers on how it is likely to be applicable to other problems such as power-aware routing.** [1]

I. Introduction

Tiered sensor network architectures are a natural platform on which to build compute-intensive systems where the computational load is resident on the higher capability 'master' nodes, rather than the simpler impoverished sensor nodes. Data fusion-based services such as localization, tracking, coverage, power-aware routing are all compute-intensive.

The Tenet architecture [1] provides a vision of the level of support future sensor network deployments will enjoy. The basic idea is to have a large number of 'mote-class' systems providing high density in sensing and a smaller number of 'master' nodes with more powerful radios and processors. This is clearly different from earlier approaches [2] which thought in terms of large numbers of mote-class systems deployed in a 'flat' manner. Tiered systems are finding their way into field deployments. As examples consider the Great Duck Island [3] and the James Reserve [4] deployments. In each case the number of master nodes is about 1-2 orders of magnitude smaller than the motes. Roughly, each master is 'responsible' for 10-100 motes.

Tiered networks have several advantages: they are easier to program and debug [1] because there is very little application logic in the mote-class system; they are easier to manage since there are few master-level nodes where the logical complexity resides; they constrain network diameter to minimize wireless link loss [5]; there is evidence that they lead to longer lifetimes when the master nodes are carefully placed [6]; Finally, and most importantly for us, tiered sensor networks are not limited by the processing available at each mote-class device, the master nodes do computation on behalf of the impoverished nodes.

Based on this background, we ask the following two questions.

- Can one demonstrate that intensive computation at the master nodes, coupled with extremely simple data collection at sensor nodes, produces data-fusion performance comparable to systems where the sensor nodes have been outfitted with specialized sensing hardware, or carefully configured, yet have access to relatively poor computational resources ?
- Are there convenient and efficient ways to distribute and manage computation across the master nodes ?

We answer both questions in the affirmative, in the special yet representative case of localization. We also sketch a preliminary example from power aware routing using the same formulation. In sensor networks the focus is on developing lightweight computationally tractable algorithms for data-fusion services suitable for a collection of computationally constrained nodes. In some cases, such as localization, several approaches rely on the use of specialized hardware either for ranging [7], beacons [8], or a 'super-node' which assists the nodes in the network for localization [9]. Other approaches carefully calibrate radio frequencies [10] along with accurate time synchronization to obtain accurate localization results.

Inspired by the computation available in tiered systems, we propose an approach to localization which requires no special hardware nor any configuration. The input to our system is strictly radio-based proximity. Based on these proximity constraints, we propose a mathematical optimization-based framework for localization. Our results

[1]This work is supported in part by NSF grants CNS-0540420, CNS-0520305, CNS-0325875 and CCR-0120778

are accurate; and techniques from distributed optimization are available to divide the problem into coarse pieces suitable for computation at the master-level nodes. Our system thus trades off computation to achieve localization accuracy with simple inputs.

Having established the technical details for a specific problem (localization), we argue that this result is promising for data fusion-based services in general because it provides evidence that compute intensive algorithms, often shunned as impractical for sensor networks, are in fact a viable possibility in tiered sensor networks. Specifically we give arguments to show how other forms of localization (e.g. using ranging) are easily dealt with in our framework. We also give a rationale for how other data fusion services such as power-aware routing could be implemented using our formalism.

In the next section we present a generic framework for low level services in tiered networks distributed based on mathematical optimization. In section III, we give the technical details for proximity-based localization as an example data fusion service which can be solved using our framework. Following that we discuss generalizations to other data fusion problems. We conclude with a summary and a sketch of ongoing and future work.

II. General Framework

This section describes the generic framework of our approach. We give specific examples of formulations of localization (Section III) and power-aware routing(Section IV) in later sections.

Consider the following optimization problem -

$$\begin{gathered}\min f(X)\\ \text{subject to } h(X)=0 \text{ and}\\ g(X)\le 0\end{gathered}$$

where h and g are matrices and $X \in R^n$.

This problem can easily be solved in a centralized manner by collecting all the relevant information at a single point.

We would like to split the computation among the master-class nodes in a tiered network while retaining a fixed number of constraints corresponding to each lower-tier node that spans multiple master nodes. This form of distribution has been previously studied in the context of parallel computing. One such paradigm called Hierarchical Overlapping Coordination [11] has been studied in the systems engineering and mathematical optimization literature. We adopt it here, with a minimal description of how it works reproducing a condensed version of the treatment in [11].

Represent the above equations using a Functional Dependence Table (FDT). In the FDT, the (i,j)-th entry of this table is one if the i-th constraint depends on the j-th variable, and zero otherwise. Efficient heuristic techniques exist [12] to partition the graph and thereby rearrange this FDT into a set of p_α decoupled blocks. Each block corresponds to a sub-problem.

The problem (denoted α) can be rewritten as

$$\begin{gathered}\min f_{\alpha_0}(X_{\alpha_0})+\textstyle\sum_{i=1}^{p_\alpha} f_{\alpha_i}(X_{\alpha_0},X_{\alpha_i})\\ \text{subject to } h(X_{\alpha_0},X_{\alpha_i})=0 \;\; i=1,\cdots,p_\alpha\\ \text{and } g(X_{\alpha_0},X_{\alpha_i})\le 0 \;\; i=1,\cdots,p_\alpha\end{gathered}$$

X_{α_0} is the vector of linking variables. If we fixed the linking variables $X_{\alpha_0}=d_\alpha$, $d_\alpha \in R^{n\alpha}$, then we have p_α problems of the form -

$$\begin{gathered}\text{For each } i=1,\cdots,p_\alpha,\\ \min_{x_{\alpha_i}} f_{\alpha_i}(d_\alpha,x_{\alpha_i}) \text{ subject to}\\ h_{\alpha_i}(d_\alpha,x_{\alpha_i})=0,\\ g_{\alpha_i}(d_\alpha,x_{\alpha_i})\le 0,\end{gathered}$$

Assume that we could perform another graph partitioning in a similar fashion and call that problem β.

The HOC algorithm [11] is as follows:

- Step 1: Fix linking variables X_{α_0}, and solve problem α by solving p_α independent sub-problems given above.
- Step 2: Fix linking variables X_{β_0} to their values as determined in step 1 and solve problem β by solving p_β independent problems.
- Goto Step 1 with the fixed values of α-linking variables determined in Step 2
- Repeat until convergence is achieved

In previous work, a hypergraph mapping has been proposed to partition the graph into numerous partitions [12]. However, one of the advantages of using this framework in the context of sensor networks is that we already have a graph based on connectivity. The partitions could be constructed by a simple flood initiated by the master nodes. Each mote picks as its master, the master node that is closest to it (in terms of number of hops). Ties could be resolved arbitrarily. Thus, a protocol could be designed to partition the graph suitably and easily. We can resolve ties in two possible ways to create the two different partitions (α and β) that are needed in the optimization.

III. Case Study: Proximity-based Localization

We now study the simple case of proximity-based localization. We will later relax some of the assumptions made and generalize to incorporate other forms of input like ranging.

A. Assumptions

We assume that there are N sensor nodes that are deployed in a square field in a uniform random fashion such a path exists in the network graph between any two nodes. We assume that a rough initial estimate of the locations of these nodes (this could be because the nodes were deployed by a human, robot or aerial vehicle which knows roughly where they were deployed). As we shall see, these initial estimates

(X_0,Y_0) need not be very accurate. Every node is assumed to have a unique ID. In the process of communicating with other nodes, each node assembles a list of its neighbors. We assume that the radio power levels are the same across all the nodes, and this determines the range of communication of each node. Let the radius of communication be R. Let us also assume that in the network graph the set of nodes is denoted V and the set of edges is denoted E.

B. Problem Formulation

Given nodes (i,j), $\forall(i,j) \in E$ we have a connectivity constraint of the form

$$(x_i - x_j)^2 + (y_i - y_j)^2 \leq R^2 \tag{1}$$

If the degree of i is d_i, we can have $(N - d_i)$ constraints for every node that is not a neighbor of i of the form

$$(x_i - x_j)^2 + (y_i - y_j)^2 \geq R^2 \tag{2}$$

We formulate the optimization problem as an unconstrained minimization. In such a formulation, each neighbor constraint of the form Eqn. [1] contributes cost $max(((x_i - x_j)^2 + (y_i - y_j)^2 - R^2),0)$ and each constraint of the form Eqn. [2] contributes cost $max((R^2 - (x_i - x_j)^2 - (y_i - y_j)^2),\ 0)$.

We minimize the following total cost.

$$\min \left\{ \begin{array}{c} \sum_{i=1}^{N} \max((x_i - x_j)^2 + (y_i - y_j)^2 - R^2),0) \\ \forall(i,j) \in E \\ +\sum_{i=1}^{N} max(R^2 - ((x_i - x_j)^2 + (y_i - y_j)^2),0)) \\ \forall(i,j) \notin E \end{array} \right\} \tag{3}$$

C. Scaling by Tiering

The second set of constraints (Eqn. 2) are over-constraining. An intelligent choice of a subset of these constraints could reduce the computational load significantly while preserving accuracy. We exploit this in the distributed version of the optimization. In the formulation (Eqn. 3), each new node added to the network adds N constraints. Hence, the number of constraints scales as N^2. Assume now, that there are k master-class nodes in the network. The network is divided into k partitions, each roughly containing N/k nodes. Now, each such node has $(N/k)^2$ constraints. This is a significant reduction in workload and also forms a trade off between the number of master-class nodes deployed and the computational load on each of them. We now apply the HOC algorithm to the problem and solve the distributed optimization problem across k partitions. For the simulations, we assumed $k = 4$.

Fig. 1 shows the localization error achieved with the distributed technique. The x-axis is the average degree of the network, and the y-axis shows the average localization error (distance of the localized position from the actual position). Although there is some improvement with increase in average network degree, it is interesting to note that good localization is achieved even in networks with low average degree of ≈6 (error ≈0.2*R). Empirical observation has led us to believe that much of this error is due to edge effects.

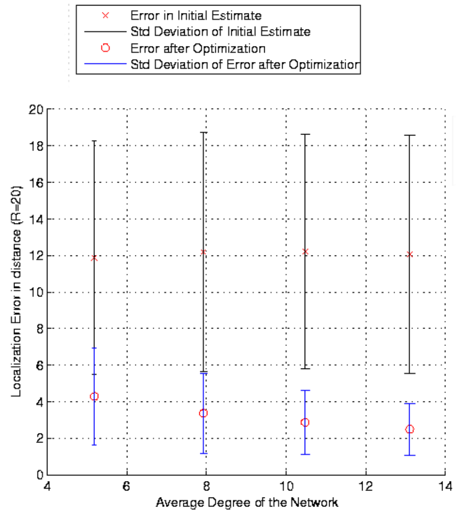

Fig. 1. Average error vs. average degree of connectivity. Localization error is measured as the magnitude of the distance between the estimated location of a node and its actual location. This is averaged over all nodes. The average degree of connectivity is the sum of the degree of each node in the network divided by the total number of nodes.

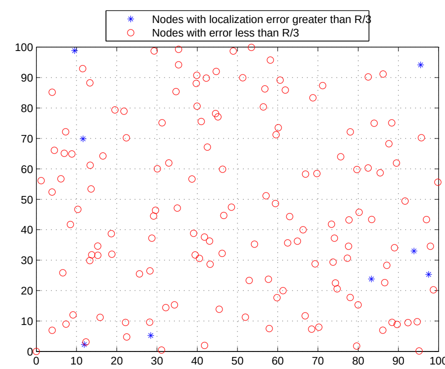

Fig. 2. Edge Effects: For a specifi c instance (with Num nodes = 150 and Radius of communication = 20) the empty circles denote nodes localized within (R/3) distance from actual location and the stars denote nodes localized with error greater than (R/3)

The second graph (Fig. 2) shows these edge effects. For a specific trial (with avg degree ≈8), it shows the nodes that were localized with error $\geq$ 0.33 * R as asterisks. As can be seen, they are at the periphery, and we intend to incorporate additional information (like signal strength) to mitigate this error in future work.

The third figure (Fig. 3) plots the decrease in localization error as a function of the optimization time (measured in

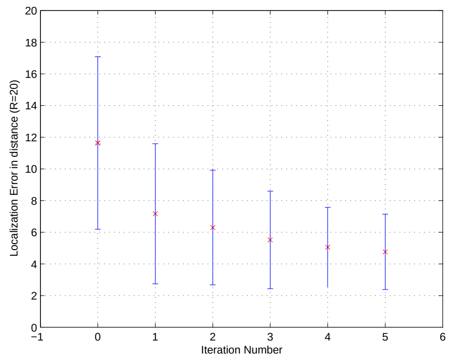

Fig. 3. Localization error vs. number of optimization iterations shown for the case R=20 and degree 8

iterations). The point corresponding to iteration zero is the error in the initial estimate and converges quickly. The error decreases with iterations. This shows that there is a trade off between computation and accuracy that can be exploited depending on the exact application at hand.

We have chosen a simple proximity-based localization example to illustrate two ideas. The first is that with no network configuration, and no special sensors, one can perform accurate localization if one is prepared to use compute-intensive methods such as optimization. Second, such methods are a natural fit to tiered systems since techniques exist to distribute them efficiently across multiple nodes. We now extend our model to incorporate other forms of input like ranging and radio interferometry for completeness.

D. Other Localization Techniques

If nodes i and j have ranging equipment using which they estimate their separation to be d, this includes a constraint in our optimization formulation

$$(x_i - x_j)^2 + (y_i - y_j)^2 = d^2 \tag{4}$$

This is easily added to our framework as a quadratic equality constraint and the optimization can proceed as before. Radio Interferometry is similar to ranging. It provides a linear combination of ranges as opposed to individual ones. They can be incorporated into our formulation similar to the above constraint. Finally, it is possible to tune the neighborhood (Eqn. 1) and non-neighborhood constraints (Eqn. 2) if signal strength measurements are present. This is likely to lead to better localization accuracy and may need fewer constraints to achieve the same accuracy as our present setup.

IV. Power-aware Routing

There is a wealth of literature on power-aware routing in wireless networks [13] examining various trade-offs. We show how this can be represented in our framework.

A. Assumptions

We assume that there are N sensor nodes (say motes) that are deployed in a square field in a uniform random fashion such that a path exists in the network graph between any two nodes. Let the set of edges on this network graph be denoted by A. The neighbor set of a node i is denoted by B_i. We assume that there are k master-class nodes also uniformly deployed in the network. Let us call the set of masters M and the set of motes V. Each mote produces sensor readings periodically (every t seconds) which it wants to relay to any of the master nodes (over multiple hops if need be). Each mote has a fixed amount of energy (battery capacity E_i for node i) which depletes over time. Let each mote consume e_r energy units to receive a packet and e_t energy units to transmit. The masters are assumed to have an infinite supply of energy (they have a renewable energy source like solar power, or are powered externally). From our assumptions, it is clear that shortest path routing will deplete the energies of nodes closest to the masters first (since they have to carry a bulk of the traffic), thus leading to disconnection of nodes farther away from the masters and the masters themselves. Ideally, we would want the network to gracefully deplete energies of all motes so that data collection lifetime is maximized.

B. Problem Formulation

We adapt the multicommodity flow formulation of the above problem [14]. Let us assume that $q_{(i,j)}$ is the flow from node i to node j where (i,j) are neighbors. Let T be the time interval for which the routes are calculated. Let each node generate Q packets of data during that time. By principle of conservation,

$$\sum_{j:i\in B_j} q_{(j,i)} + Q = \sum_{j\in S_i} q_{(i,j)} \forall (i,j) \in A \text{ and for every } i \in V$$

The lifetime of a node i is defined as

$$T_i = \frac{E_i}{\sum_{j\in B_i} e_t.q_{(i,j)} + \sum_{j\in B_i} e_r.q_{(j,i)}}$$

The system lifetime is

$$T_{system} = min(T_i) \;\; \forall i = 1, \ldots, N$$

Our objective is to maximize T_{system}. Hence the problem is of the form

$$\max T_{system}$$

such that

$$\sum_{j:i\in B_j} q_{(j,i)} + Q = \sum_{j\in S_i} q_{(i,j)} \forall (i,j) \in A \text{ and for every } i \in V$$

$$q_{(k,i)} = 0 \;\; \forall k \in M and \forall i \in B_k$$

$$q_{(j,i)} \leq 0$$

$$\sum_{j\in B_i} e_t.q_{(i,j)} + \sum_{j\in B_i} e_r.q_{(j,i)} \le E_i$$

This is very similar to our formulation of localization in eqn. 3. We can distribute the computation using the framework described in section II. Assuming we have k master nodes in the network, we can allocate the motes to "belong" to a master each. Let the set of masters be called M. Let S_l be the set of motes that belong to master l. Let L be the set of links that connect motes between sets. We could split the problem into k smaller problems, and rewrite the formulation as follows -

$$For\ \alpha = 1, \ldots, k$$

$$\max\ T^{\alpha}_{system}$$

such that

$$\sum_{j:i\in B_j} q^{\alpha}_{(j,i)} + Q = \sum_{j\in S_i} q^{\alpha}_{(i,j)} \forall (i,j) \in A \text{ and } \forall i \in S_\alpha$$

$$q^{\alpha}(k,i) = 0\ \forall k \in M$$

$$q^{\alpha}_{(j,i)} \le 0\ \forall (i,j) \in A \text{ and } \forall i \in S_\alpha$$

$$\sum_{j\in B_i} e_t.q^{\alpha}_{(i,j)} + \sum_{j\in B_i} e_r.q^{\alpha}_{(j,i)} \le E_i\ \forall (i,j) \in A \text{ and } \forall i \in (S_\alpha - M)$$

We can now use the algorithm discussed in section II to solve the above problem iteratively.

V. Conclusions

We have presented a framework for low-level services in sensor networks based on the idea that one can trade-off computation to achieve accuracy in tiered networks. We showed a simple proof-of-concept proximity-based localization study, and argued how we can extend it readily to other kinds of localization, as well as other services such as power-aware routing.

Our immediate goal is to implement our framework on a real network and ensure its feasibility. We also intend to investigate formally rigorous ways of partitioning the network to evenly distribute the load among master nodes. This is particularly relevant when the mote distributions are not uniform. One of the biggest problems in localization is that of graph isomorphism. Our initial formulation side-stepped this problem by overconstraining the system. Ultimately, it is important to make an educated choice of the constraints to avoid degrees of freedom which will cause error. We are embarking on a detailed study of the accuracy of our algorithm under various scenarios with real data from a network of motes/masters (thus violating the random placement assumption and the disc communication model). We also plan to compare the accuracy of this method with known range-free and range-aware localization techniques. In addition to localization, we intend to systematically study power-aware routing using our framework in simulation and on a mote/master testbed and compare it to existing power-aware routing techniques. Finally, we believe that our results may ultimately influence decisions about tiered network architecture design.

References

[1] Ramesh Govindan, Eddie Kohler, Deborah Estrin, Fang Bian, Krishnakant Chintalapudi, Om Gnawali, Ramakrishna Gummadi, Sumit Rangwala, and Thanos Stathopoulos, "Tenet: An architecture for tiered embedded networks," Tech. Rep. cens-tr-56, Center for Embedded Networked Sensing, UCLA, November 2005.

[2] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MOBICOM-00)*, N. Y., Aug. 6–11 2000, pp. 56–67, ACM Press.

[3] Robert Szewczyk, Alan Mainwaring, Joseph Polastre, John Anderson, and David Culler, "An analysis of a large scale habitat monitoring application," in *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, New York, NY, USA, 2004, pp. 214–226, ACM Press.

[4] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao, "Habitat monitoring: Application driver for wireless communications technology," Jan. 05 2001.

[5] Jerry Zhao and Ramesh Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *SenSys '03: Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, 2003, pp. 1–13.

[6] Mark Yarvis, Nandakishore Kushalnagar, Harkirat Singh, Anand Rangarajan, York Liu, and Suresh Singh, "Exploiting heterogeneity in sensor networks," in *Proceedings of IEEE INFOCOM 2005*, March 2005.

[7] David Moore, John Leonard, Daniela Rus, and Seth J. Teller, "Robust distributed network localization with noisy range measurements," In *SenSys '04: Proceedings of the 2nd International conference on Embedded networked sensor systems* [7], pp. 50–61.

[8] Nirupama Bulusu, John Heidemann, Deborah Estrin, and Tommy Tran, "Self-configuring localization systems: Design and experimental evaluation," New York, NY, USA, 2004, vol. 3, pp. 24–60, ACM Press.

[9] Radu Stoleru, Tian He, John A. Stankovic, and David Luebke, "A high-accuracy, low-cost localization system for wireless sensor networks," in *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, New York, NY, USA, 2005, pp. 13–26, ACM Press.

[10] Miklos Maroti, Peter Volgyesi, Sebestyen Dora, Kusye Branislav, Andras Nadas, Akos Ledeczi, Gyorgy Balogh, and Karoly Molnar, "Radio interferometric geolocation," in *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, New York, NY, USA, 2005, pp. 1–12, ACM Press.

[11] Nestor Michelena, Hyungju A. Park, Panos Papalambros, and Devadatta Kulkarni, "Hierarchical overlapping coordination for large-scale optimization by decomposition," June 29 1999.

[12] Nestor Michelena, Panos Papalambros, and G. G. Brown, "A hypergraph framework for optimal model-based decomposition of design problems," june 29 1997.

[13] Christine E. Jones, Krishna M. Sivalingam, Prathima Agrawal, and Jyh Cheng Chen, "A survey of energy efficient network protocols for wireless networks," *Wirel. Netw.*, vol. 7, no. 4, pp. 343–358, 2001.

[14] Jae-Hwan Chang and Leandros Tassiulas, "Maximum lifetime routing in wireless sensor networks," Piscataway, NJ, USA, 2004, vol. 12, pp. 609–619, IEEE Press.