Poster Abstract: Is Data-Centric Storage and Querying Scalable?

Joon Ahn and Bhaskar Krishnamachari *
Department of Electrical Engineering – Systems
University of Southern California, Los Angeles, CA 90089
joonahn@usc.edu, bkrishna@usc.edu

Abstract

The scalability of a wireless sensor network has been of interest and importance. We use a constrained optimization framework to derive fundamental scaling laws for both unstructured sensor networks (which use blind sequential search for querying) and structured sensor networks (which use efficient hash-based querying). We find that the scalability of a sensor network's performance depends upon whether or not the increase in energy and storage resources with more nodes is outweighed by the concomitant application-specific increase in event and query loads. We have figured out the theoretical scaling laws for the networks of 2 dimensional deployment in our previous work [2]. We report on our work-in-progress aimed at extending the scaling laws to networks of various dimensional deployment. As a recent achievement, we find that $m \cdot q^{1/2}$ must be $O(N^{\frac{d-1}{2d}})$ for unstructured networks, and $m \cdot q^{\frac{d}{d+1}}$ must be $O(N^{\frac{d-1}{d}})$ for structured networks, to ensure scalable network performance, where m is the number of events sensed by a network over a finite period of deployment, q is the number of queries for each event, d is the dimension of deployment, and N is the size of the network. These conditions determine (i) whether or not the energy requirement per node grows without bound with the network size for a fixed-duration deployment, (ii) whether or not there exists a maximum network size that can be operated for a specified duration on a fixed energy budget, and (iii) whether the network lifetime increases or decreases with the size of the network for a fixed energy budget. An interesting finding of this extension is that 3D uniform deployments are inherently more scalable than 2D uniform deployments, which in turn are more scalable than 1D uniform deployments.

Categories and Subject Descriptors: C.2.2 Computer Communication Networks: Network Protocols

General Terms: Design, Performance, Theory

Keywords: Modeling, Wireless Sensor Networks, Energy Efficiency, Theory, Performance Analysis, Querying, Scalability

*This work was supported in part by the NSF through the NETS-NOSS grant 0435505, CAREER grant 0347621, and ITR grant CNS-0325875.

SenSys'06, November 1–3, 2006, Boulder, Colorado, USA.

1 Introduction

Wireless sensor networks are envisioned to consist of large numbers of embedded devices that are each capable of sensing, communicating, and computing. While the network as a whole is required to provide fine resolution monitoring for an extended period of time, the individual embedded devices face some fundamental constraints. They are typically deployed with limited battery supplies and, because of their form factor and low cost, may also have limited data storage capability. The goal of our on-going research is to understand the conditions under which a query-based data-centric sensor network [1] deployed in various dimensions can be operated in a scalable manner despite these constraints on energy and storage; and to verify them with experiments using real wireless sensors.

We consider both unstructured and structured varieties of data-centric querying along with replicated storage in this research. In unstructured querying schemes, the node issuing the query does not know in advance where any copy of the requested event information can be found. The query dissemination is therefore a form of blind search (this can take the form of an expanding ring search or a sequential trajectory search). In structured querying schemes, a hash or index is used so that the querying node knows exactly where the nearest copy of the requested event information can be found. In such networks, there is a trade-off between the energy costs of replicated storage and querying that is determined by the number of replicas created for each event. A large number of replicas results in lowered query cost at the expense of greater storage cost, and vice versa. We can formulate an optimization problem whose aim is to select the optimum number of replicas that minimizes the total energy cost of querying and storage, subject to constraints on storage. We use this optimization problem as a tool to identify the conditions, in terms of the numbers of events and queries, under which query resolution can be performed in a scalable manner despite constraints on storage and energy.

We find that operating a network in a scalable fashion essentially requires that the traffic load due to additional events and queries be outweighed by the improvement in energy and storage resources obtained as the network size increases. Moreover, an interesting finding of this extension is that networks deployed in higher dimension has more scalability. Thus, 3D uniform deployments are inherently more scalable than 2D uniform deployments, which in turn are more scalable than 1D uniform deployments.

2 Extended Scaling Laws

The key assumptions in this work is same as those of our previous work [2] except that N nodes are deployed in a d-dimensional hypercube space. Notation is summarized as follows: N is the number of nodes in the network, R is the radio radius of a node, T is the fixed application-specific time duration for which the sensor network is deployed, m is the number of atomic events that are sensed in the environment, r_i is the number of copies of i-th event, and q_i is the number of queries for the i-th event.

The extended replication cost for d-dimensional deployment is as follows:

$$C_{replication} = c_1 \cdot (r-1) \cdot \sqrt[d]{N} \quad (1)$$

The extended search cost for structured and unstructured environment are as follows:

$$C_{search,structured} = c_2 \cdot \sqrt[d]{\frac{N}{r}} \quad (2)$$

$$C_{search,unstructured} = c_3 \cdot \frac{N}{r+1} \quad (3)$$

Note that the search cost for unstructured environment is independent of the dimension of deployment. We have derived these three expressions for the special case of two-dimensional deployment in [4, 3].

With these extended costs, we can extend the optimization problem in [2] and solve it using the method of Lagrange multiplier. The following theorems are of the resulting extended version corresponding to those in [2]:

THEOREM 3. *The total energy costs for unstructured networks grow with network size N as follows:*

$$C^*_{t,u} = \begin{cases} \Theta\left(m \cdot q^{1/2} \cdot N^{\frac{d+1}{2d}}\right), & \text{(inactive)} \\ \Theta\left(N^{\frac{d+1}{d}} + m^2 \cdot q\right), & \text{(active)} \end{cases} \quad (4)$$

THEOREM 4. *The total energy costs for structured networks grow with network size N as follows:*

$$C^*_{t,s} = \begin{cases} \Theta\left(m \cdot q^{\frac{d}{d+1}} \cdot N^{1/d}\right), & \text{(inactive)} \\ \Theta\left(N^{\frac{d+1}{d}} + m^{\frac{d+1}{d}} \cdot q\right), & \text{(active)} \end{cases} \quad (5)$$

THEOREM 5. *For unstructured networks, the energy requirement per node is bounded if and only if*

$$m \cdot q^{1/2} \text{ is } O\left(N^{\frac{d-1}{2d}}\right)$$

THEOREM 6. *For structured networks, the energy requirement per node is bounded if and only if*

$$m \cdot q^{\frac{d}{d+1}} \text{ is } O\left(N^{\frac{d-1}{d}}\right)$$

THEOREM 7. *For unstructured networks, given fixed average per-node energy e (i.e., the total energy allocated optimally among the nodes in the network grows linearly with the network size as $E = e \cdot N$), the following statements describe the conditions on the network size N, the number of events m and the number of queries per event q that ensure that the network can be operated successfully.*

1. *If $m \cdot q^{1/2}$ is $o(N^{\frac{d-1}{2d}})$, then there exists a minimum network size $N_{min}(e)$ beyond which it can always be operated successfully.*
2. *If $m \cdot q^{1/2}$ is $\Theta(N^{\frac{d-1}{2d}})$, then there exists an average per-node energy e^* such that for all $e < e^*$, it is not possible to operate a network of any size successfully, while for all $e \geq e^*$ it is possible to operate a network of any size successfully.*
3. *If $m \cdot q^{1/2}$ is $\omega(N^{\frac{d-1}{2d}})$, but $o(N)$, then there exists a maximum network size $N_{max}(e)$ beyond which the network cannot be operated successfully. Further N_{max} is a convex function of e.*
4. *If $m \cdot q^{1/2}$ is $\Theta(N)$, then N_{max} increases linearly with e.*
5. *If $m \cdot q^{1/2}$ is $\omega(N)$, then N_{max} increases as a concave function of e.*

THEOREM 8. *For structured networks, the same conditions as in the above Theorem 7 hold, with just the following substitutions:*

- $m \cdot q^{1/2} \rightarrow m \cdot q^{\frac{d}{d+1}}$
- $N^{\frac{d-1}{2d}} \rightarrow N^{\frac{d-1}{d}}$ *for condition 1, 2, and 3*
- $N \rightarrow N^{\frac{2d}{d+1}}$ *for condition 3, 4, and 5*

THEOREM 9. *For unstructured networks, with a fixed average per-node energy budget of e, so long as the number of events and queries scale temporally so that $m \cdot q^{1/2}$ is an increasing function of time, the lifetime of deployment T over which the network can operate successfully scales with the network size as per the following conditions:*

1. *if $m \cdot q^{1/2}$ is $o(N^{\frac{d-1}{2d}})$ then T increases with N.*
2. *if $m \cdot q^{1/2}$ is $\Theta(N^{\frac{d-1}{2d}})$ then T is constant with respect to N.*
3. *if $m \cdot q^{1/2}$ is $\omega(N^{\frac{d-1}{2d}})$ then T decreases with N.*

THEOREM 10. *For structured networks, the same conditions as in the above Theorem 9 hold, with just the following substitutions:*

- $m \cdot q^{1/2} \rightarrow m \cdot q^{\frac{d}{d+1}}$
- $N^{\frac{d-1}{2d}} \rightarrow N^{\frac{d-1}{d}}$

3 Conclusions and Future Work

We have investigated the fundamental scaling behavior of storage and querying in wireless sensor networks deployed d-dimensional space. The main take away from this study is that the event and query rates must scale sufficiently slowly with the network size if scalable performance is desired; and the network deployed in higher dimension has more scalability. In particular, an important scaling condition is ensuring that $m \cdot q^{1/2}$ be $O(N^{\frac{d-1}{2d}})$ for unstructured networks, and that $m \cdot q^{\frac{d}{d+1}}$ be $O(N^{\frac{d-1}{d}})$ for structured networks. Satisfying this condition ensures that adding nodes to the network is beneficial in that the energy and storage resources they bring outweigh the additional event and query activity they induce. This can be seen from many perspectives: satisfying this condition implies that (i) sensor networks require bounded energy and storage per node, (ii) arbitrarily large networks can be operated successfully with a limited energy budget, and (iii) that the network lifetime increases with network size for a given energy budget.

In our study we have not explicitly considered bandwidth capacity; we have implicitly assumed that the energy constraints will be more severe than bandwidth constraints in the system. However, if energy constraints are not significant (consider as an extreme case if all nodes could be wired for power), bandwidth issues could be the dominant consideration. This is a topic for future work.

As an on-going research, we would like to undertake realistic simulations and large-scale experiments to validate the analytical results presented in this proposal.

4 References

[1] R. Govindan, "Data-centric Routing and Storage in Sensor Networks," in *Wireless Sensor Networks*, Eds. Raghavendra, C.S., Sivalingam, K.M., and Znati, T., Kluwer Academic Publishers, 2004, pp. 185-206.

[2] J. Ahn and B. Krishnamachari, "Fundamental Scaling Laws for Energy-Efficient Storage and Querying in Wireless Sensor Networks," *ACM MobiHoc*, May 2006.

[3] B. Krishnamachari, J. Ahn, "Optimizing Data Replication for Expanding Ring-based Queries in Wireless Sensor Networks", *International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WIOPT '06)*, April 2006, Boston, USA.

[4] J. Ahn and B. Krishnamachari, "Derivations of the Expected Energy Costs of Search and Replication in Wireless Sensor Networks", *USC Computer Engineering Technical Report CENG-2006-3*, April 2006.

USC Autonomous Networks Research Group

Is Data-Centric Storage and Querying Scalable?

Joon Ahn and Bhaskar Krishnamachari
(joonahn@usc.edu, bkrishna@usc.edu)
Department of Electrical Engineering,
Viterbi School of Engineering, University of Southern California

Introduction:

Background

- Wireless sensor networks are envisioned to provide fine-grained resolution monitoring for large areas.
- There are fundamental constraints on energy and storage resources available at each node.
- An important paradigm is performing in-network data-centric storage, using queries to obtain event information on-demand.

Objective

- Is it possible to mathematically characterize the application conditions under which data-centric storage and querying in wireless sensor networks can be performed in a scalable fashion?
- The analysis considers structured and unstructured sensor networks with grid and random deployments in *d*-dimension.

Problem Formulation: The Minimum Expected Total Cost and Its Scalability

The Total Cost

$$C_t = \sum_{i=1}^{m} q_i C_s(r_i) + \sum_{i=1}^{m} C_r(r_i)$$

- ***m***: the total number of events
- q_i: the query rate for i-th event
- r_i: the number of copies of i-th event
- $C_s(r_i)$: the expected minimum search cost of i-th event
- $C_r(r_i)$: the expected replication cost of i-th event
- The total cost can be optimized globally using the method of Lagrange multipliers.

Optimization Formulation

$$\text{Minimize} \quad C_t = \sum_{i=1}^{m} q_i C_s(r_i) + \sum_{i=1}^{m} C_r(r_i)$$

$$\text{s.t} \quad \sum_{i=1}^{m} r_i \le S$$

- ***S***: the total amount of storage available in the network
- It is solvable using the method of Lagrange multiplier.

The Search Cost

Structured:

$$C_{s,st} = c_1 \sqrt[d]{\frac{N}{r}}$$

Unstructured:

$$C_{s,un} = c_2 \frac{N}{r+1}$$

Note that there is no variable *d* here.

- *N*: the number of nodes
- *d*: dimension
- *r*: the number of copies of an event

The Replication Cost

$$C_r = c_3 (r-1) \sqrt[d]{N}$$

- The replication cost is same for both structured and unstructured networks

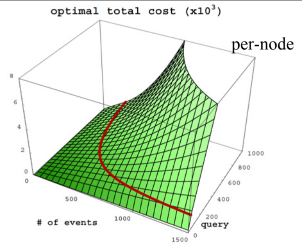

Unstructured network

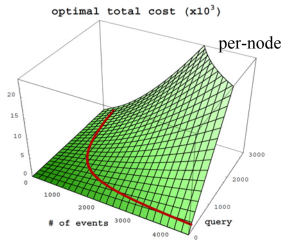

Structured network

Analysis and Results: Scaling Conditions and their Implications

Scaling Conditions
(for Structured Networks*)

Result 1. The optimal total energy costs grow as follows:

$$C^*_{t,s} = \Theta\left(m \cdot q^{\frac{d}{d+1}} \cdot \sqrt[d]{N}\right)$$

Result 2. The energy requirement per node is bounded if and only if:

$$m \cdot q^{\frac{d}{d+1}} = O\left(N^{\frac{d-1}{d}}\right)$$

- Energy constraints are stricter than storage constraints.

Result 3. Given fixed average per-node energy *e*,

1. If $m\, q^{d/(d+1)} = o(N^{(d-1)/d})$, then the network of size $N > N_{min}(e)$ can always be operated successfully.
2. If $m\, q^{d/(d+1)} = \Theta(N^{(d-1)/d})$, then there exists a threshold e^* such that for all $e < e^*$, the network cannot be operated successfully, otherwise the network can always be operated successfully.

If $m\, q^{d/(d+1)} = \omega(N^{(d-1)/d})$, then the network of size $N > N_{max}(e)$ cannot be operated successfully. Further,

3. If $m\, q^{d/(d+1)} = o(N^{2d/(d+1)})$, then $N_{max}(e)$ is increasing convex.
4. If $m\, q^{d/(d+1)} = \Theta(N^{2d/(d+1)})$, then $N_{max}(e)$ is increasing linear.
5. If $m\, q^{d/(d+1)} = \omega(N^{2d/(d+1)})$, then $N_{max}(e)$ is increasing concave.

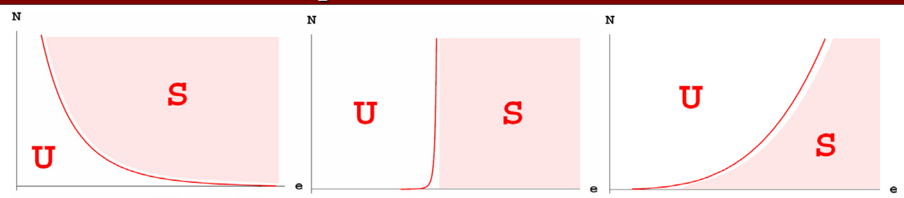

case 1: $m \cdot q^{2/3}$ is $o(N^{1/2})$ case 2: $\Theta(N^{1/2})$ case 3: $\omega(N^{1/2})$ and $o(N^{4/3})$

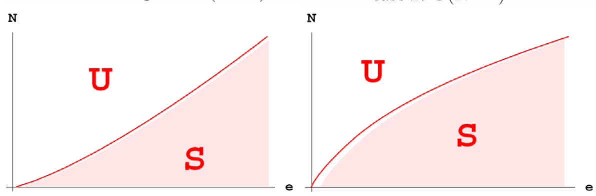

case 4: $\Theta(N^{4/3})$ case 5: $\omega(N^{4/3})$

Network size conditions for 2D deployment (Result 3)

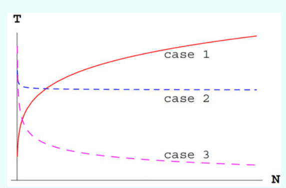

The network lifetime for 2D deployment (Result 4)

Result 4. With a fixed average per-node energy budget *e*, so long as $m\, q^{d/(d+1)}$ is an increasing function of time, the lifetime of deployment *T* with *N* nodes scales as follows:

1. If $m\, q^{d/(d+1)} = o(N^{(d-1)/d})$, then *T* increases with *N*.
2. If $m\, q^{d/(d+1)} = \Theta(N^{(d-1)/d})$, then *T* is constant with *N*.
3. If $m\, q^{d/(d+1)} = \omega(N^{(d-1)/d})$, then *T* decreases with *N*.

Fundamental Implication

Together, these results imply that

1. **only certain classes of applications can be sustained in arbitrarily large sensor networks.**
2. **The network of higher dimension has more scalability.**

*** Results apply to both grid and random topologies. The results of unstructured networks for 2D deployment can be found in the full paper: J. Ahn and B. Krishnamachari, "Fundamental Scaling Laws for Energy-Efficient Storage and Querying in Wireless Sensor Networks,"** ***ACM MobiHoc*** **'06**

This work has been supported in part by the National Science Foundation through the NETS-NOSS grant 0435505, CAREER grant 0347621, and ITR grant CNS-0325875