Brief Announcement: Locality-Based Aggregate Computation in Wireless Sensor Networks*

Jen-Yeu Chen
National Dong Haw University
Hualien, 974, Taiwan, R.O.C.
jenyeu@ieee.org

Gopal Pandurangan
Purdue University
W. Lafayette, IN 47907, USA
gopal@cs.purdue.edu

Jianghai Hu
Purdue University
W. Lafayette, IN 47907, USA
jianghai@purdue.edu

ABSTRACT

We present DRR-gossip, an energy-efficient and robust aggregate computation algorithm in wireless sensor networks. We prove that the DRR-gossip algorithm requires $O(n)$ messages and $O(\frac{n^{3/2}}{\log^{1/2} n})$ one-hop wireless transmissions to obtain aggregates on a random geometric graph. This reduces the energy consumption by at least a factor of $\frac{1}{\log n}$ over the standard uniform gossip algorithm. Experiments validate the theoretical results and show that DRR-gossip needs much less transmissions than other gossip-based schemes.

Categories and Subject Descriptors: C.2.4 [Computer Communication Networks]: Distributed Systems, F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Design, Performance, Theory

Keywords: wireless sensor networks; aggregate computation; gossip; distributed algorithm; randomized algorithm; random geometric graph

1. INTRODUCTION

A wireless sensor network is abstracted as a connected undirected graph $G(V, E)$ with all the sensor nodes as the set of vertices V, and all the bi-directional wireless communication links as the set of edges E. The total number of sensor nodes is n, i.e., $|V| = n$. This underlying graph can be arbitrary depending on the deployment of sensor nodes. Let each sensor node i be associated with an initial observation or measurement value denoted as $v_i \in \mathbb{R}$. The assigned values over all vertices form a vector $\mathbf{v}$. The goal is to compute aggregate functions such as average, sum, max, min etc. on the vector of values $\mathbf{v}$ [2].

In this paper, a novel approach named the DRR-gossip algorithm is presented. A Distributed Random Ranking (DRR) algorithm is used to first build a forest of (disjoint) localized trees over the sensor network. The height of a tree is small, thus the aggregates of its nodes can be quickly obtained at its root. All the roots then perform the uniform gossip algorithm on their respective aggregates to reach distributed consensus on the global aggregate. Compared with the standard uniform gossip, our DRR-gossip algorithm requires $O(n)$ messages and $O(\frac{n^{3/2}}{\log^{1/2} n})$ transmissions to reach consensus, reducing the energy consumption by at least a factor of $1/\log n$ on a Poisson random geometric graph. The DRR-gossip algorithm is inspired from the efficient gossip ([5]) idea of reducing the number of nodes participating in the gossip process in order to decrease the numbers of messages and transmissions for achieving consensus. Further, the DRR-gossip algorithm takes advantage of the locality of the trees to further decrease the number of messages to $O(n)$ while still maintains a running time of $O(\log n)$, whereas the efficient gossip algorithm needs $O(\log n \log \log n)$ time and $O(n \log \log n)$ messages. For further details on the analytical and experimental results of the DRR-gossip algorithms, we refer to the full version of this paper [3].

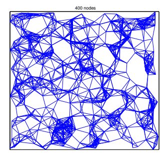

(a) network graph.

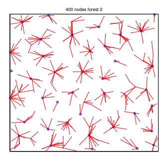

(b) the forest.

Figure 1: An instance of a forest.

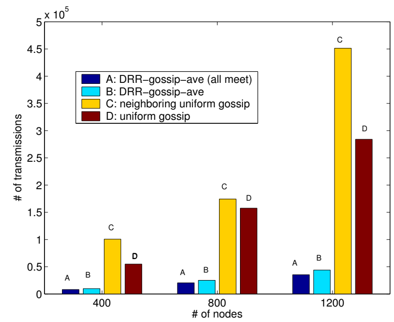

Figure 2: Simulation results: the numbers of transmissions for the computation of average aggregate (Ave) on Poisson random geometric graphs with various network sizes.

2. ALGORITHMS

The DRR-gossip algorithm proceeds in phases. In phase one, all the sensor nodes run the DRR algorithm to construct a forest of disjoint trees. In phase two, within each tree, the local aggregate (e.g. sum or maximum) of a tree is computed by a convergecast process and obtained at the root of the tree. In phase three, all the roots of the trees utilize a suitably modified version of the uniform gossip algorithm [6] to obtain the global aggregate. Finally, if necessary, a root can forward the global aggregate along the tree to all the other nodes of the tree.

Phase I: The DRR algorithm

The DRR algorithm is implemented in the following way. On an undirected connected graph $G = (V, E)$, every node $i \in V$ *uniformly at random* generates a random value, $rank(i) \in [0, 1]$ (its "rank") and then sends its rank to all its (one-hop) neighboring nodes. Each node compares the ranks it receives with its own and then connects (i.e., establishes a tree edge) to the node of the highest rank among all of its neighbors and itself. If a tie between two

*J.-Y. Chen was supported in part by NSC Award 97-2218-E259-003 and G. Pandurangan was supported in part by NSF Award CCF-0830476.

PODC'09, August 10–12, 2009, Calgary, Alberta, Canada.
ACM 978-1-60558-396-9/09/08.

Table 1: DRR-gossip vs. other gossip-based algorithms on Poisson random geometric graph.

Algorithm	time (steps)	messages	transmissions
neighboring gossip [1]	$O(n^2)$	$O(n^2)$	$O(n^2)$
geographic gossip [4]	$O(n \log^2 n \cdot \hat{R})$	$O(n \log n)$	$O(n \log^2 n \cdot \hat{R})$
efficient gossip [5]	$O(\log n \log\log n \cdot \hat{R})$	$O(n \log\log n)$	$O(n \log\log n \cdot \hat{R})$
uniform gossip [6]	$O(\log n \cdot \hat{R})$	$O(n \log n)$	$O(n \log n \cdot \hat{R})$
DRR-gossip [this paper]	$O(\log n \cdot \hat{R})$	$O(n)$	$O(n \cdot \hat{R})$

On a Poisson random geometric graph $G(n, r(n))$, $\hat{R} = 1/r(n) = O((n/\log n)^{1/2})$ by geographic routing.

ranks happens, nodes can break the tie by their identifiers. Through these connections, many *disjoint* trees are established on the graph. These trees together constitute a forest $\mathbb{F}$, which is a subgraph of G, i.e., $\mathbb{F}(V', E') \subseteq G(V, E)$, where $V' = V$ and $E' \subseteq E$. If a node is of the highest rank among its neighbors and itself, it becomes the root. Since every node, except the root nodes, connects to a node with a higher rank, there are no loops. (Fig. 1. is an example to illustrate the forest $\mathbb{F}$ constructed on a Poisson random geometric graph.) In this phase, $O(1)$ time steps, $O(n)$ messages and $O(n)$ one-hop wireless transmissions are required. The trees constructed by the DRR algorithm have the following advantages: (1) locality; (2) robustness to adversary; and (3) easy load balancing. We prove the following theorem.

THEOREM 1 (LOCALITY OF TREES). *On a Poisson random geometric graph $G(n, r)$, the height of the ranking trees constructed by the DRR algorithm is bounded by $O(\log n)$ with high probability (w.h.p).*

Because the height of a ranking tree is small, the local aggregate of a tree can be quickly obtained at the root of a tree in $O(\log n)$ time steps by the Converge-cast algorithm.

Phase II: Converge-cast

In the second phase of our algorithm, the local aggregate of each tree will be obtained at the root by the Convergecast algorithm — an aggregation process starting from the leaf nodes and proceeding upward along the tree to the root node. For example, to compute the local max/min, all leaf nodes simply send their values to their parent nodes. An intermediate node collects the values from its children, compares them with its own value and sends its parent node the max/min value among all received values and its own. A root node then obtains the local max/min value of its tree. By Theorem 1, the local aggregates within a tree can be computed in $O(\log n)$ time steps and by $O(n)$ messages and $O(n)$ one-hop wireless transmissions.

Phase III: Gossip

In the third phase, all roots of the trees compute the global aggregate by performing the uniform gossip algorithm on an abstract overlaying graph $\tilde{G} = clique(\tilde{V})$, where $\tilde{V} \subseteq V$ is the set of roots and $|\tilde{V}| = m$. To traverse through an edge of $\tilde{G}$, a message may need a multi-hop relay, i.e., several one-hop transmissions, on the original physical network graph G. We proved that both the global maximum and the global average can be obtained within $O(\log n)$ rounds of the uniform gossip algorithm. The following theorem is used to bound the running time of the DRR-gossip algorithm for computing the global average.

THEOREM 2. *W.h.p., there exists a time $T_{ave} = O(\log m + \alpha \log n) = O(\log n)$, $\alpha > 0$, such that for all time $t \geq T_{ave}$, the relative error of the estimate of average aggregate on the root of the largest ranking tree, z, is at most $\frac{2}{n^{\alpha-1}}$, where the relative error is $\frac{|\hat{x}_{ave,t,z} - x_{ave}|}{|x_{ave}|}$ and the global average is $x_{ave} = \sum_j x_j$ (assume that all x_j have the same sign).*

We also proved that the global maximum (or minimum) can be obtained after $O(\log n)$ rounds of the gossip-max algorithm, a version of uniform gossip subtly revised for computing global maximum, in phase 3.

THEOREM 3. *Running the gossip-max algorithm on the overlay complete graph $\tilde{G} = clique(\tilde{V})$ of $G(V, E)$, after $O(\log n)$ rounds of the sampling procedure of the gossip-max algorithm, w.h.p., all the roots will know the global maximum, Max.*

3. PERFORMANCE COMPARISONS

3.1 Analytical comparisons

To compare the DRR-gossip algorithms with other representative gossip-based algorithms, we provide a table (Table 1) for their performances on a Poisson random geometric graph. The neighboring gossip [1] is a variant of the standard uniform gossip: nodes only communicate with its one-hop direct neighbors. All the messages in neighboring gossip requires only local transmissions as the DRR algorithm. For both uniform gossip [6] and efficient gossip [5], every message requires $\hat{R}$ multi-hop transmissions. However, for DRR-gossip, only messages in phase three need multi-hop routing. In the other two phases, each message requires only one "local one-hop transmission." It is clear that, in all performance categories, the DRR-gossip-algorithm is at least equal to or better than all the other representative gossip-based algorithms.

3.2 Simulation results

Fig. 2 shows that the DRR-gossip-ave algorithm (the DRR-gossip algorithm for computing the global average) saves a tremendous number of transmissions, compared to the uniform gossip [6] and the neighboring gossip [1].

4. CONCLUSION

In this paper, we presented a set of DRR-gossip algorithms to compute aggregates in a wireless sensor network. Our analyses show that the proposed DRR-gossip algorithm, on a Poisson random geometric graph, requires the same time complexity as the standard uniform gossip algorithm but reduces the energy consumption by a factor of $\frac{1}{\log n}$ over the standard uniform gossip algorithm.

5. REFERENCES

[1] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE/ACM Trans. on Netw.*, vol. 14, no. SI, 2006.

[2] J.-Y. Chen, G. Pandurangan, and D. Xu. "Robust Computation of Aggregates in Wireless Sensor Networks: Distributed Randomized Algorithms and Analysis," *IEEE Trans. on Parallel and Distributed Systems*, 17(9), 2006, 987-1000.

[3] J.-Y. Chen, G. Pandurangan, and J. Hu, "http://www.cs.purdue.edu/homes/gopal/locality-sensor-agg.pdf."

[4] A. G. Dimakis, A. D. Sarwate, and M. J. Wainwright, "Geographic gossip: efficient aggregation for sensor networks," in *IPSN*, 2006.

[5] S. Kashyap, S. Deb, K. V. M. Naidu, R. Rastogi, and A. Srinivasan, "Efficient gossip-based aggregate computation," in *PODS*, 2006.

[6] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proc. FOCS*, 2003.