Brief Announcement: Convergence Analysis of Scalable Gossip Protocols

Stacy Patterson[1], Bassam Bamieh[2], and Amr El Abbadi[1]

[1] Department of Computer Science, University of California, Santa Barbara
{sep, amr}@cs.ucsb.edu
[2] Department of Mechanical Engineering, University of California, Santa Barbara
bamieh@engineering.ucsb.edu

Introduction. We present a simple, deterministic gossip protocol for solving the distributed averaging problem. Each node has an initial value and the objective is for all nodes to reach consensus on the average of these values using only communication between neighbors in the network. We first give an analysis of the protocol in structured networks, namely d-dimensional discrete tori and lattices, and show that in an n node network, the number of rounds required for the protocol to converge to within ϵ of the average is $O(|\log(\epsilon)|\ n^{2/d})$. We then extend our results to derive upper and lower bounds on convergence for arbitrary graphs based on the dimensions of spanning supergraphs and subgraphs.

Analyzing Convergence. The network is represented by an undirected graph $G = (V, E)$ where V is the set of nodes in the network, with $|V| = n$, and E is the set of communication channels between them. The neighbor set of a node i, denoted N_i, is the set of nodes $j \in V$ such that $(i, j) \in E$. Every node i has an initial value $x_i(0)$, and the average of all values in the system is $x_{ave} = \frac{1}{n}\sum_{i=1}^{n} x_i(0)$. In each round k, every node sends an equal fraction β of its current value to each of its neighbors and sends the remaining fraction, $\alpha = 1 - |N_i|\beta$, to itself. Each node updates its current value $x_i(k+1)$ to be the sum of all values received in round k. The desired goal of the protocol is for the system to converge to an equilibrium where $x_i(k) = x_{ave}$ for all $i \in V$.

We measure how far the current state of the system is from the average state using the "deviation from average" vector, with each component defined by

$$\tilde{x}_i(k) \ := \ x_i(k) \ - \ \frac{1}{n}\left(x_1(k) \ + \ \ldots \ + \ x_n(k)\right).$$

When the vector $\tilde{x}$ equals $\mathbf{0}$, the vector of all 0's, the vector x equals $x_{ave}\mathbf{1}$, where $\mathbf{1}$ is the vector of all 1's. The rate at which $||\tilde{x}||$ approaches 0 determines the rate at which the nodes reach consensus at x_{ave}. We define the *ϵ-consensus time* to be the number of rounds required for $\frac{||\tilde{x}(k)||}{||\tilde{x}(0)||} \leq \epsilon$ for a given ϵ.

If $x(k)$ denotes the vector of current values in round k, the distributed averaging protocol can be represented as an $n \times n$ matrix $A = [a_{ij}]$ that transforms $x(k)$ to $x(k+1)$, where $x(k+1)$ is the vector of values in round $k+1$. a_{ij} is the fraction of j's current value that j sends to node i in each round. The evolution of the vector x is given by following recursion equation.

S. Dolev (Ed.): DISC 2006, LNCS 4167, pp. 540–542, 2006.
© Springer-Verlag Berlin Heidelberg 2006

$$x(k+1) = A\, x(k)$$

It can be shown that the evolution of the vector x depends on the second largest eigenvalue of A, λ_2. More specifically, an upper bound on the number of rounds required for ϵ-consensus is given by

$$k \geq \frac{\log(\epsilon)}{\log(|\lambda_2|)}. \tag{1}$$

Therefore, the ϵ-consensus time depends upon the inverse of $|\lambda_2|$. In general, it is not possible to determine λ_2 analytically. However, for certain graph structures, we can derive λ_2 and thus derive an asymptotic bound for the protocol.

In particular, we consider the protocol matrix for a d-dimensional discrete torus $\mathbb{Z}_s^d$. In each round, every node sends an equal fraction $\beta < \frac{1}{2d}$ of its current value to each of its $2d$ neighbors. In the 1-dimensional case, the protocol matrix A is a circulant matrix. In the general d-dimensional case, A is known as a circulant operator. The eigenvalues of A can be explicitly obtained using the multi-dimensional Discrete Fourier Transform [1]. Using this and Equation (1), we can determine the ϵ-consensus time of the protocol in any d-dimensional discrete torus. For large n, the asymptotic convergence of the protocol in a d-lattice is the same as for a d-dimensional torus, so the same analysis applies.

Theorem 1. *The ϵ-consensus time of the distributed averaging protocol in a discrete d-dimensional torus or d-lattice with n nodes is $O(|\log(\epsilon)|\ n^{2/d})$.*

This result shows that the dimensionality of the torus determines the convergence rate of the averaging protocol. In tori, the dimension is closely related to the connectivity of the graph; a higher dimensional torus has greater connectivity and a faster convergence rate than a lower dimensional torus. If β is chosen carefully, a similar relationship between connectivity and convergence can be derived for arbitrary graphs.

Theorem 2. *Let $G_1 = (V_1, E_1)$ be an undirected, connected graph. Let $G_2 = (V_2, E_2)$ be a spanning subgraph of G_1. For $\beta \leq \frac{1}{2\Delta(G_1)}$, where $\Delta(G_1)$ is the maximum degree of G_1, the convergence rate of the protocol on G_1 is greater than or equal to that on G_2.*

According to the Theorem 2, if we start with a d-dimensional torus and add edges, the convergence rate of the protocol on the new graph will be at least as fast as that on the original graph. Similarly, if we take an arbitrary graph and add edges to form a d-dimensional torus, the convergence rate of the protocol on the original graph will be less than or equal to that of the protocol on the torus. This result is stated precisely in the following corollary.

Corollary 1. *Let $G = (V, E)$ be an arbitrary graph with $|V| = n$.*

1. *If D is the dimension of the largest dimensional torus that is a spanning subgraph of G, the consensus protocol reaches ϵ-consensus on G in $O\left(|\log(\epsilon)|\ n^{2/D}\right)$ rounds.*

2. *If d is the dimension of the smallest dimensional torus for which G is a spanning subgraph, the consensus protocol reaches ϵ-consensus on G in $\Omega\left(|\log(\epsilon)|\ n^{2/d}\right)$ rounds.*

This corollary is a direct result of Theorems 1 and 2.

Details on the above results can be found in [2].

Related Work. Stochastic gossip protocols for solving the distributed averaging problem have been proposed [3,4]. In these protocols, each node selects a neighbor at random in each round and averages its own value with the neighbor's value. Bounds derived for convergence rates of these protocols are probabilistic. A deterministic protocol in which a node can send a different fraction of its current value to each of its neighbors in each round have also been studied [5]. The authors use offline analysis based on global information to determine the optimal fraction to send along each edge and give localized online heuristics. However, the selection of fractions does not affect the asymptotic convergence rate of the protocol. Distributed averaging has also been studied in the context of anonymous networks [6]. In this work, each node must know the size and topology of the network. The role of dimensionality in optimal error bounds in sensor networks was studied in [7]. Our work characterizes the relationship between network dimensionality and the convergence rate of a deterministic gossip protocol for distributed averaging that requires only local information, making it well-suited for large scale P2P systems, sensor networks, and mobile ad-hoc networks.

References

1. Salapaka, S., Peirce, A., Dahleh, M.: Analysis of a circulant based preconditioner for a class of lower rank extracted systems. In: Numerical Linear Algebra with Applications. Volume 12. (2005) 9–32
2. Patterson, S., Bamieh, B., El Abbadi, A.: Convergence analysis of scalable gossip protocols. Technical Report 2006-09, Department of Computer Science, University of California, Santa Barbara (2006)
3. Kempe, D., Dobra, A., Gehrke, J.: Gossip-based computation of aggregate information. In: IEEE Symposium on Foundations of Computer Science. (2003) 482
4. Boyd, S., Ghosh, A., Prabhakar, B., Shah, D.: Gossip algorithms: Design, analysis and applications. In: INFOCOM. Volume 3. (2005) 1653–1664
5. Xiao, L., Boyd, S.: Fast linear iterations for distributed averaging. In: Systems and Control Letters. Volume 53. (2005) 65–78
6. Kranakis, E., Krizanc, D., van den Berg, J.: Computing boolean functions on anonymous networks. Information and Computation **114** (1994) 214–236
7. Barooah, P., Hespanha, J.: Graph effective resistances and distributed control: Part ii – electrical analogy and scalability. In: 45th IEEE Conference on Decision and Control. Feb. 2006 (submitted). (2006)