

Distributed Data Aggregation

MAP-I – Module for the Thematic Seminar on Global Computing

Carlos Baquero and Paulo Sérgio Almeida – CCTC-DI – U. Minho

Summary

This document describes a proposal for a 8 hours tutorial module to be included on the MAP-I Thematic Seminar on Global Computing, for 2008-09.

1 Context

This module provides an introduction to the state-of-the-art in data aggregation algorithms for distributed networks. Data aggregation is a core technique in the design of efficient sensor networks and scalable systems in general. An introduction to the subject area can be found on Robbert Van Renesse paper on the *The Importance of Aggregation*

... We define aggregation as the ability to summarize information. In the area of sensor networks it is also referred to as data fusion. It is the basis for scalability for many, if not all, large networking services. For example, address aggregation allows Internet routing to scale. Without it, routing tables would need a separate entry for each Internet address. Besides a problem of memory size, populating the tables would be all but impossible. DNS also makes extensive use of aggregation, allowing domain name to attribute mappings to be resolved in a small number of steps. Many basic distributed paradigms and consistency mechanisms are based on aggregation. For example, synchronization based on voting requires votes to be counted.

One application of aggregation is decentralized network size estimation algorithms. Systems, such as Skype, estimate the number of active users by centralized algorithms, enacting a high demand on central resources and creating single points of failure. Aggregation introduces scalable decentralized algorithms that allow counting (or summing in the general case) over a distributed network of nodes. The results can be used in a broad class of distributed algorithms, including Distributed Hash Tables, Multicast Algorithms, Quorum Building and Load Balancing.

A particular application is found on sensor networks, where data collection is often only practical if aggregation is performed. Aggregation-based techniques can be used to detect how many sensors are deployed in a given area or what is the average temperature in the covered area. In these systems data aggregation must be performed along the communication paths, as it would be impractical, energy-wise, to collect all measured values.

2 Objectives

The goal of this module is to provide a solid background on the emerging subject of data aggregation in distributed networks. The course will present the abstractions that model data aggregation and basic network topologies (concentrating on mesh, random and small-world networks); it will introduce the state-of-the-art in data aggregation algorithms; and will discuss new research directions in the area.

The course will show what can be currently achieved with the most recent results and hopefully motivate graduate students towards new directions of research in this recent field.

3 Learning Outcomes

Upon successful completion of this course, students should be able to:

- know the main aggregation techniques and classes of aggregation algorithms;
- discuss the properties and complexity tradeoffs of aggregation algorithms;
- analyze a given aggregation task and decide how to achieve it with existing tools;
- extract information from scientific papers in the area;
- address new problems beyond the state-of-the-art.

4 Pre-Requisites

The course assumes familiarity with distributed systems and networking topics at undergraduate level.

5 Format

Tutorial module. Dedicated lecture notes will be provided.

6 Grading

The assessment will be made by a written exam about one or more research papers covered in the module.

7 Content

1. Aggregation: Basic Functions, Function Decomposition;
2. Network Models: Topologies, Metrics;
3. Network Size Estimation;
4. Aggregation of Sums: FM Sketches, Separable Functions, Extrema Propagation;
5. Averaging Techniques: Push-Pull, Push-Sum, Distributed Random Grouping, Delta Adjusting.

8 References

- [1] R. Van Renesse. The Importance of Aggregation. *Future Directions in Distributed Computing*. Springer LNCS, volume 2584, pages 87-92. 2003.
- [2] D. Kempe, A. Dobra and J. Gehrke. Gossip-based Computation of Aggregate Information. *44th Annual IEEE Symposium on Foundations of Computer Science*. 2003.
- [3] M. Jelasity, A. Montresor and O. Babaoglu. Gossip-based Aggregation in Large Dynamic Networks. *ACM Transactions on Computer Systems*. 23(3):219-252. 2005.
- [4] J. Chen, G. Pandurangan and D. Xu. Robust Computation of Aggregates in Wireless Sensor Networks: Distributed Randomized Algorithms and Analysis. *IEEE Transactions on Parallel and Distributed Systems*. 17(9):987-1000. 2006.
- [5] M. Durand and P. Flajolet. Loglog Counting for Large Cardinalities. *ESA*. Springer LNCS, volume 2832, pages 605-617. 2003
- [6] M. Bawa, H. Garcia-Molina, A. Gionis and R. Motwani. Approximate Aggregation Techniques for Sensor Databases. *ICDE, IEEE*, pages 449-460. 2004.
- [7] K. Horowitz and D. Malkhi. Estimating Network Size from Local Information. *Information Processing Letters*. 88(5):237-243. 2003.
- [8] P. S. Almeida, C. Baquero, N. Preguiça, D. Hutchison. Scalable Bloom Filters. *Information Processing Letters*. 101(6):255-261. 2007.

A Instructors

The teaching team consists of two members of the Distributed Systems Group (GSD) of the Informatics Department of Minho University. Both lecturers have more than 10 years of experience of teaching and research in distributed systems.

The teams research on Data Aggregation is still recent: In 2005 there was a focus on improving Bloom Filters; this lead to a Journal publication in Information Processing Letters (Elsevier) [9]; in 2006 a new probabilistic technique was developed for aggregation of sums; the results are under submission and available as Technical Report since May 2006. Also in 2006, a new research line was started on providing message loss tolerance to aggregation algorithms.

Paulo Sérgio Almeida is a lecturer at the Department of Informatics of Minho University, and a researcher member of CCTC. His scientific research activities are centered in distributed systems. The two main topics of research have been time/version stamping mechanisms and distributed data aggregation algorithms. The mains results of late have been Dynamic Version Stamps, Bounded Version Vectors and Scalable Bloom Filters.

Carlos Baquero is a lecturer at the Department of Informatics of Minho University, and a researcher member of CCTC. His research interests are focused on distributed systems, in particular in causality tracking, peer-to-peer systems and mobile computing. Recent research is focused on highly dynamic distributed systems, both in internet P2P settings and in mobile and sensor networks.

Selected Publications

- [1] P. S. Almeida, C. Baquero, N. Preguiça, D. Hutchison, Scalable Bloom Filters, Information Processing Letters 101 (2007) 255-261, Elsevier.
- [2] J. B. Almeida, P. S. Almeida, C. Baquero. Bounded version vectors. In Rachid Guerraoui, editor, Proceedings of DISC 2004: 18th international symposium on distributed computing, number 3274 in LNCS, pages 102–116. 2004. Springer Verlag.
- [3] P. S. Almeida, C. Baquero, V. Fonte. Version stamps – decentralized version vectors. Proceedings of the 22nd international conference on distributed computing systems (ICDCS), pages 544–551. 2002. IEEE Computer Society.