

Topic For Directed Study:

Fundamental Research Issues in Performance Analysis of Runtime Status Collection and Aggregation

Stefan Witt, WSU ID 10539995
stefan.witt@wsu.edu

EE 595, Summer 2002
Instructor: David Bakken

“Mr. Fusion” is a data fusion middleware framework [AMB02]. It is used to collect data from multiple client replicas, and categorize these data as correct or incorrect according to a set of well-defined rules. The incorrect data are then stored as erroneous values in a multidimensional hierarchical database, a data warehouse. Data fusion is performed by aggregating the collected data. Queries on the data can be answered by the data warehouse.

Mr. Fusion consists of two basic components, the Fusion Core, and the Fusion Status Service [AMB02]. While the former component is responsible for collecting data from the replicas and categorizing them, the latter is the data warehouse that aggregates the collected values and answers the queries.

In this project, Mr. Fusion will be used to detect slowdown failures of subscribers to an atomic multicast. In [BHO+99] it has been shown that the multicast throughput dramatically decreases if even one client suffers from a slowdown failure. This is due to the fact that ACKs are sent late by the slow receiver, and the multicast system as a whole slows down because of that.

In order to be able to detect slowdown failures, the Fusion Core needs information about the response times of the clients. This can be inferred from ACKs that are sent by the clients or by ballots of one data fusion session. Then the Fusion Status Service can be used to query for slow clients by some appropriate query.

The project consists of three basic steps:

1. Becoming familiar with a Java multicast system that is total and atomic. This multicast will then be used to send data from the clients to Mr. Fusion.
2. Examining the throughput performance of the multicast with and without one slow client. Furthermore, use Mr. Fusion to exclude the slow client automatically.
3. Writing a project report that explains the background, describes the approach for the performance analysis, and presents the results in text and diagrams.

References

- [BHO+99] K. P. Birham, M. Hayden, O. Ozkasap, Z. Xiao, M. Buidu: *Bimodal Multicast*, ACM Transactions on Computer Systems, Vol. 17, No. 2, May 1999, pp. 41–88
- [AMB02] A. Franz, R. Mista, D. Bakken, C. Dyreson, M. Medidi: *Mr. Fusion: A Programmable Data Fusion Middleware Subsystem with a Tunable Statistical Profiling Service*, to appear in Proceedings of the International Conference on Dependable Systems and Networks (DSN-2002), IEEE/IFIP, June 23–26 2002, Washington, DC