M.I.T. Laboratory for Computer Science June 18, 1976
Computer Systems Research Division Request for Comments No. 116

Trip report: The Berkeley Workshop on Distributed Data Management and Computer Networks. May 25 - 26, 1976

by Harry Forsdick

I recently attended the Berkeley Workshop on Distributed Data Management and Computer Networks. The following list outlines the topics discussed at the conference:

1. An Overview of Existing Distributed Networks
2. Network Graphics Protocol
3. Mini-host Front Ends to the ARPA Net
4. Validating Software on Distributed Systems
5. High Level Protocols
6. Distributed Data Bases
7. Teleconferencing and Network Mail
8. Data Translation and Exchange Standards
9. Gateways and Local Networks
10. Network Requirements for Data Base Support
11. Data Access and Manipulation Languages
12. Data Management Machines

While the people at the conference were from many different backgrounds, there was a noticeable attendance by members of the ERDA community. ERDA has apparently begun to put emphasis on computer networks for interconnecting computers at the various ERDA laboratories. One of the purposes of this conference was to explore the issues in distributed data management and computer networks that people entering the field will encounter.

I attended the sessions on topics 1, 2, 4, 6, 8, 10, and 12.

There were several talks that stand out in my mind and I will give a brief description of them.

In the session on network graphics, Robert Sproull of Xerox-PARC presented a talk entitled "Network Graphics Isn't Networking." He presented a brief perspective on graphics with respect to the Network Graphics Protocol. Before the NGP, graphics had reached the point where very exotic devices were being used that required high bandwidth paths and quick response from a processor. However, the exotic devices were expensive and graphics was not reaching a wide enough market. What was required was a lowering of sights in graphics. At about the same time, storage scopes were being marketed at a price suitable for wide spread use; they required a similar lowering of sights. Out of the desire to use both existing applications and cheaper devices, the notion of separating the device independent portion of graphics from the device dependent portion arose. The use of networks increased the desire for standardization on the device independent aspect of graphics. The device dependent portion could run locally while the device independent portion could run on a processor foreign to the graphics device. In this way, the limited bandwidths of network paths would only have to be used to transmit information across the interface of the two parts. This separation is one of the basics of the NGP. The unfortunate aspect of the NGP is that it is oriented to exotic, high performance devices. It is hard to see the simple options available for modest applications. Sproull's main point was that

the NGP was not so much an advance in using computer networks as it was a realization that graphics needed to be standardized.

Peter Alsberg of the University of Illinois, Center for Advanced Computation, gave two talks in the sessions on Distributed Data Bases and Network Requirements for Data Base Support. In the first, entitled "Data Distribution Strategies," he outlined four different schemes for storing data in a distributed system:

1. Single copy, no backup.
2. Single copy, backup by keeping off-line copy and journal of changes; reconstruction on failure.
3. Primary copy, one or more on-line slave backups that can serve as alternative sources of data if the primary is unavailable.
4. Multiple copy with dynamic updating at any copy.

It was Alsberg's opinion that scheme 3 was superior to the rest when managing distributed data bases that were frequently modified and where high availability and minimal cost were primary requirements. Scheme 4, he argued, results in disastrous communication costs if general types of operations (e.g., test and set) are allowed. Alsberg also suggested that scheme 2 can result in lower availability than scheme 1 because of the delay introduced byonstruction. In his second talk, "Synchronization and Resiliency in Network Data Access," Alsberg described a discipline, referred to as "N-Host Resiliency," for insuring that the failure of up to N hosts still allows the rest of a distributed system to proceed as normal. The basic idea is that in an N-host resilient system, an agreement by N hosts that an

operation will be completed is needed before the operation is considered successfully completed. Alsberg presented an argument based on probability that even a 2-host resilient system has a very low probability of failure.

The last talk of the conference was given by Stuart Schuster of the University of Toronto on "The Case for a Parallel-Associative Approach to Data Base Machine Architectures." Schuster argued that many operations on data bases are really associative in nature and a data base machine ought to reflect this. He described a system under development that employed multiple processors which were connected to circulating memories (e.g. bubble). The instructions of the processors used associative addresses to select records in the data base. There was a fair amount of good discussion on the merits and feasibility of this approach.

Several other people made some observations in their talks that should be noted. Arie Shoshani of the Systems Development Corporation in a talk entitled "Common Standards for Distributed Data Bases" emphasized that the particular standard chosen is not important, but rather what you choose to standardize; in addition it is important to standardize the minimum amount possible. James White of the Stanford Research Institute spoke about common attributes of third level protocols (function oriented protocols) which are typically implemented by user and server processes. Every time such a service is programmed, many of the same functions are reprogrammed. White argued that a network virtual

programming environment should support data types common to third level protocols as well as a non-blocking procedure call that looks a lot like network wide interprocess communication. John McCarthy of Stanford University suggested, among other things, that it is useless to try to get people to agree on standard organizations of data bases and as a result, the real emphasis ought to be on methods for describing other people's data bases to your own programs. In a talk on the Datacomputer, Jerry Farrell of the Computer Corporation of America pointed out that the semantics of the language used to access the Datacomputer is influenced by the fact that there are communication delays in the network links to the Datacomputer. It is inefficient to go to the Datacomputer every time you want to retrieve a record and thus Datalanguage attempts to optimize around common high level operations. Finally, several people including Jon Postel of SRI and Harvard Holmes and Michael Stonebreaker of Lawrence Berkeley Laboratory talked about the need to divide systems up into parts that execute locally and parts that execute remotely. The term "Frontend-Backend" was used to describe this division. Descriptions of how several existing systems (the data base management system INGRES and the Berkeley Data Base Management System) would be split up to operate in a distributed environment were presented. Typically parsing, formatting and even some query operations based on small amounts of data in a cache memory are delegated to the local frontend processor (intelligent terminal) while data retrieval, long term storage and large

operations on the entire data base are performed in the backend processor.

A copy of the conference proceedings should arrive shortly.