The PLUS Accelerator

Roberto Bisiani, Principal Investigator

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Our plan is to demonstrate a workstation accelerator (called PLUS) that not only executes the current speech recognition systems in real time, but that will also support the increasing computational demands of speech recognition systems for the next few years. The accelerator will also be useful for other compute bound applications. The accelerator will be replicated and made available to other sites.

Our approach is to connect single-processor machines of roughly the performance of high-end workstations by means of a very fast and powerful communication mechanism. This makes the memory on each node visible to every other node as if the system were a single shared-memory multiprocessor. The prototype, for example, uses 25MHz Motorola 88000 processors with up to 32 Mbytes of DRAM memory and 512 Kbytes of SRAM in each node; the interconnection is based on a MOSIS-built mesh router with 40Mbytes/sec bandwidth.

PLUS's research contribution comprises new mechanisms to support shared-memory and synchronization between processors. The emphasis is on low latency and low overhead synchronization.

The PLUS prototype is interfaced to host computer systems by means of a SCSI interface, this interface can also perform speech input/output and limited signal processing. Interfacing PLUS to a host does not require any operating system change on the host.

As of December 1990, 40 PLUS nodes have been fabricated and are awaiting testing. A three node PLUS system has been assembled and debugged and is currently running the Sphinx system in batch mode. The software environment now includes MACH C-threads which provides C language support for multiprocessor programming. This is the same multiprocessor programming environment that is used in the standard MACH operating system. Therefore developers can create applications on their MACH workstations that are easily portable to a PLUS system.

Our future plans include the construction of 8 and 16 node PLUS systems to support the development of very large vocabulary speech recognition. We are also developing a parallel implementation of Viterbi beam search to take full advantage of the PLUS architecture.