



Technical Progress Report 11/1/93 - 2/1/94
Construction of a Connectionist Network Supercomputer
University of California, Berkeley
ONR URI Grant No. N00014-92-J-1617

DTIC
ELECTR
FEB 14 1994
S C D

1 Abstract

Work has progressed on many fronts this quarter:

- Significant efforts continue in the VLSI implementations of the first Torrent processor, T0, and CNS-1 network interface chip, Hydrant.
- Several novel board level technologies for the CNS-1 have been verified.
- We have made good progress in the development of support software for the Torrent processor.
- Work continued in CNS-1 performance evaluation and architecture refinement.
- We have worked on adapting speech algorithms for the CNS-1.
- We have reached a significant milestone in the use of analog preprocessors for speech recognition.

The CNS-1 project continues to have a significant effect on the education of graduate and undergraduate students at our institution. There are currently 16 Ph.D., 1 M.S., and 2 B.S. students associated with the project (some are paid through supporting agencies other than the ONR). Also, many of the design principles, VLSI building blocks, and CAD tools developed as part of the implementation of the T0 processor are now used in CS250, Graduate VLSI Systems Design, here at Berkeley.

2 Technical Status

2.1 Software

The software effort has made considerable progress in several areas. The main emphasis for the last quarter has been on developing a stable software environment for development of code to run on the Torrent processors. Extensive comparisons have been made between the instruction set and register-level T0 simulators and both now produce identical results for all test programs. The assembler, C compiler and binary utility programs are now reasonably stable, and several large C programs have been run on the simulators. Finally, a single-tasking kernel has been written and debugged using the simulators. This is initially for

94-04866

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

the SPERT single board system, but much of the code, including floating point instruction emulation and host system I/O support, will be used in the CNS-1 operating system.

At the higher levels, the Sather language system also showed good progress. The 0.5 version of Sather was released with our Australian partners and provides an important stepping stone to the 1.0 version. A full 1.0 system will be completed this quarter and released later this year. The parallel Sather project was marked by the Ph.D. completion of C.Lim and the production of a complete language definition. The current effort focuses on a portable pSather that will be the basis for the CNS-1 version.

A new version of the "Boxes of Boxes" (BoB) simulation package was released internally and is in use. Parallelism in BoB is achieved either by subdividing vectors among processors, or by assigning a subgraph of one more interconnected BOX objects to a subset of the processors. The BoB software environment has been further developed and a number of users are now providing feedback for its initial implementation on the RAP.

The ICSIM simulator was recoded in Sather 1.0 leading to a considerable simplification. We are currently specifying a parallel version of ICSIM in pSather and this will be a major milestone towards mapping the system to CNS-1.

2.2 Performance Evaluation and Applications

The analytical studies of CNS-1 performance on various key computations continues to yield fruitful results. We have also analyzed more complex, less regular tasks and extended the analysis to cover different memory strategies, including the case of static memory. Two additional technical reports will be released soon. We have also begun mapping the IUE environment of the ARPA image understanding program to the CNS-1 architecture.

We have been conducting experiments in parallelization of network training for speech. The particular approach we have focused on is to train separate networks for each individual speaker, and then merge the resulting networks for speaker independent recognition by computing a weighted average of the phonetic probabilities from each net. Our initial experiments use a uniform weight across each gender, as we have found that cross-gender prediction is extremely poor. This work is in a preliminary phase, but its success would mean that we could drastically reduce the communication requirements for training our recognizers on large speech corpora.

2.3 Hardware Development

Testing of the interface test chip (fabricated by MOSIS last quarter) proceeded. This chip mimics many of the features of the Rambus interface to transfer data over short distances at a 250 MHz rate. This work will influence the CNS-1 network hardware interface.

A circuit board containing two of the interface chips (transmitter and receiver) and aux-

iliary circuits was designed and fabricated. This board incorporates two features expected to be used on the SPERT board design:

1. Chip-on-board. Also known as MCM-L (Multi-chip Module - Laminate), this is the most cost-effective way for us to obtain high performance for limited production runs. The die is attached directly to the circuit board and the wire-bonds are made between the chip and gold plated pads on the board.
2. Elastomeric test connector. To avoid adding conventional connectors to each circuit board, an array of test points is contacted using a flexible Z-axis connector material. (Similar material is used in calculators and watches to attach the display to the circuit board.) This method mimics the expensive "bed-of-nails" fixtures used for production testing of circuit boards.

Both of these features have proven successful for the interface test board, and operation of the 1.2 micron chips has been verified. The maximum frequency of operation is lower than expected, and additional testing is underway.

The SPERT board design has stabilized, and will be laid out and fabricated after the silicon design of T0 is finished. Two new features were added to the board design during this reporting period, 1) a variable speed clock, and 2) a temperature limit sensor.

2.4 Analog VLSI pre-processors

This quarter has been devoted to the initial evaluation of analog VLSI auditory pre-processors in speech recognition systems. As outlined in previous reports, a silicon auditory model of spectral shape, with on-chip support for efficient communications and parameter storage, has been designed, fabricated, and tested. Also outlined in previous reports has been the design and coding of a software environment for the evaluation of this pre-processor for pattern recognition processing.

Building on these efforts, we have finished a prototype system that connects the analog pre-processor with a commercial speech-recognition software library. Using this system, we are conducting preliminary experiments using an isolated-word, telephone-quality, speaker-independent digit database. These early experiments do not attempt to exploit the unique characteristics of auditory models; instead, we use the auditory chip as a filterbank, and use conventional techniques for converting filterbank outputs into features suitable for speech recognition. Our intention is to use the recognition results of these experiments as a baseline, with which to evaluate later experiments that fully exploit the structure of auditory models.

Also in this quarter, we have written an article on this analog VLSI processor, and submitted it to the technology magazine *IEEE Micro*; the article is now undergoing peer review. We have also publicized this work at a tutorial given at the Neural Information Processing Systems conference, and at talks at Stanford University and the Xerox Palo Alto Research Center.

3 Presentations

Ben Gomes, "The Connectionist Network Supercomputer (CNS-1)," NIPS*93 Workshop, Connectionist Modelling and Parallel Architectures, Vail, Colorado, Dec. 4, 1993.

J. Lazzaro, "A VLSI Implementations Tutorial," Neural Information Processing Systems Conference, Denver, CO, Nov. 29, 1993.

J. Lazzaro, "Silicon Auditory Processors as Computer Peripherals," CRRMA Auditory Colloquium, Stanford University, Palo Alto CA, Dec. 16, 1993.

J. Lazzaro, "Modeling Hearing in Analog VLSI," PARC Forum (a public seminar series), Xerox Palo Alto Research Center, Palo Alto CA, Jan 6, 1993.

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>Rec A273327</i>	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
<i>A-1</i>	