

19970717 091

1

Technical Progress Report 8/1/95 – 10/31/95
Construction of a Connectionist Network Supercomputer
University of California, Berkeley
ONR URI Grant No. N00014-92-J-1617

1 Abstract

This report presents a summary of the technical status for the period 8/1/95–10/31/95. A major milestone was passed during this period with the running of a complete neural network training on our SPERT neurocomputing board. With this advancement, our multi-year effort to apply digital VLSI chip design to speech recognition algorithm development has come to fruition. In addition, our effort in the area of analog auditory preprocessors for speech recognition continues to show good progress.

2 Technical Status

2.1 The T0 Processor and SPERT board

The prototype SPERT boards described in our previous status report are in daily use, primarily for speech recognition research. The next run of 30 boards is nearing completion, and will greatly enhance our computational capacity. Nearly half of these boards are committed to collaborators at other sites. Before taking delivery of their board, each site will send a researcher to Berkeley for training on the Torrent architecture and SPERT tools. By the end of this quarter, we expect to have an active international community of Torrent users.

With more than six months of experience since the first T0 die was tested, we are pleased to report that no functional errors have been found in the processor design.

2.2 Software

At the beginning of the quarter, our new MLP library, "Quicknet", was made available to speech researchers at ICSI. Quicknet provides a high level interface for training and running neural nets on either SPERT boards or workstations. Several programs based upon the library are now being used for training MLPs for speech recognition tasks.

Additional experiments have been performed to determine the effects of using fixed point for speech recognition training. Our latest results show no significant differences in the cross validation error rate when comparing training runs using floating point computation on workstations with fixed point SPERT boards. There is a significant difference in training

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

DTIC QUALITY INSPECTED B



DEPARTMENT OF THE NAVY
OFFICE OF NAVAL RESEARCH
SEATTLE REGIONAL OFFICE
1107 NE 45TH STREET, SUITE 350
SEATTLE WA 98105-4631

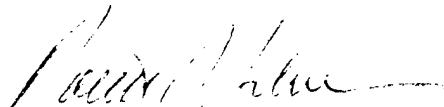
IN REPLY REFER TO:

4330
ONR 247
11 Jul 97

From: Director, Office of Naval Research, Seattle Regional Office, 1107 NE 45th St., Suite 350,
Seattle, WA 98105
To: Defense Technical Center, Attn: P. Mawby, 8725 John J. Kingman Rd., Suite 0944,
Ft. Belvoir, VA 22060-6218

Subj: RETURNED GRANTEE/CONTRACTOR TECHNICAL REPORTS

1. This confirms our conversations of 27 Feb 97 and 11 Jul 97. Enclosed are a number of technical reports which were returned to our agency for lack of clear distribution availability statement. This confirms that all reports are unclassified and are "APPROVED FOR PUBLIC RELEASE" with no restrictions.
2. Please contact me if you require additional information. My e-mail is silverr@onr.navy.mil and my phone is (206) 625-3196.


ROBERT J. SILVERMAN

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY • DAVIS • IRVINE • LOS ANGELES • RIVERSIDE • SAN DIEGO • SAN FRANCISCO



SANTA BARBARA • SANTA CRUZ

COLLEGE OF ENGINEERING
DEPARTMENT OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCES
COMPUTER SCIENCE DIVISION

BERKELEY, CALIFORNIA 94720

November 6, 1995

Dr. Clifford Lau
Electronics Division, ONR 312
Office of Naval Research
Scientific Office Code 1114SE
800 North Quincy Street
Arlington, VA 22217-5660

Reference: ONR URI Grant No.: 144-92-J-1617

Dear Dr. Lau:

Enclosed you will find the quarterly progress report for the research project for Professor John Wawrzynek: "Construction of a Connectionist Network Supercomputer," ONR URI Grant No.: N00014-92-J-1617. This report covers the research done during the performance period 8/1/95-10/31/95. Please direct any mail or questions regarding this contract to me.

Thank you for your continued support of this project.

Sincerely,

A handwritten signature in black ink, appearing to read "Theresa Lessard-Smith".

Theresa Lessard-Smith
Project Coordinator

cc: Administrative Grants Officer, Seattle
Director, Naval Research Laboratory
Defense Technical Information Center
Sponsored Projects Office - UC Berkeley
file

time, however, with the SPERT board between 6X and 20X faster than a SparcStation 20 on complete training runs, including the overhead of initialization, disk I/O and data format conversions.

The Torrent fixed point, floating point and integer libraries have been improved. Considerable time was spent working on the system tools and enhancing code portability.

There was also significant progress in the area of high level programming for parallelism. Sather release 1.0.8 is the first to include an extensive Fortran interface which significantly improves the usability of the system for scientific computation. An article on this appeared in the September 1995 issue of *Computers in Physics*. Our group was also awarded a grant from Lawrence Livermore National Lab supporting use of their Meiko CS2 super computer.

The Myrinet system for parallel Sather has been installed and is fully functional. This is being used as a platform for language and algorithm development and also as a testbed for possible deployment in larger CNS systems. A first parallel version of the ICSIM connectionist simulator is nearing completion.

2.3 Analog VLSI pre-processors.

As described in the previous progress report, we have been applying our three-chip sound pre-processing system to a speaker independent speech recognition problem. We are conducting our experiments using a 200 speaker, telephone quality, isolated word database (1-9, "oh", "zero", "yes", "no"), combining our pre-processing system with a Hidden Markov Model based speech recognizer using Multilayer Perceptrons for phoneme classification.

In the last quarter, we have continued to improve the recognition performance on the system, achieving a 4.1% error rate. While this score still falls short of the best software-based front-ends (1.8% error, using the same database and testing methods), the recognizer is comparable to systems current in use: commercial isolated-word telephone-quality digit recognizers achieve 2-5% error performance in the field, handling a vocabulary that is slightly smaller vocabulary (11 words) than our 13-word task.

Our research plans for future work on the recognition system require major changes to our experimental framework, including switching to a more challenging database, and the application of new algorithms for speech recognition. Before this work begins, we believe it is appropriate to disseminate information on our current recognition system, including representation choices, feature extraction methods, experimental results, and formulation of an agenda for continued research into auditory models. Much of the work this quarter has been in preparing experimental results suitable for publication, and on the preliminary drafting of a journal paper describing the work.

Also in the past quarter, we received word of the acceptance of a paper we submitted to the Neural Information Processing Systems conference about the chip design work for our three-chip sound pre-processing system. We are now preparing our spotlight presentation and poster for this early December meeting, as well as an 8 page paper describing the chip

design for the conference proceedings.

3 Recent Publications

Krste Asanović, James Beck, Bertrand Irissou, Brian E. D. Kingsbury, Nelson Morgan and John Wawrzynek, "The T0 Vector Microprocessor," Proceedings of Hot Chips VII, August 1995.

Bourlard, H., Konig, Y., and Morgan, N., "REMAP: Recursive Estimation and Maximization of a Posteriori Probabilities in Connectionist Speech Recognition," Proceedings of Eurospeech 1995, pp. 1663-1666, Madrid, Spain.

Lazzaro, J. and Wawrzynek, J., "Silicon Models for Auditory Scene Analysis," accepted by Neural Information Processing Systems (NIPS 95), December 1995.

David Stoutamire and Matthew Kennel, "Sather Revisited: A High-Performance Free Alternative to C++," Computers in Physics, Vol.9, No. 5, Sep/Oct 1995.

Wawrzynek, J., Asanović, K., Kingsbury, B., Beck, J., Johnson, D., Morgan, N., "SPERT-II: A Vector Microprocessor System and its Application to Large Problems in Backpropagation Training," accepted by Neural Information Processing Systems (NIPS 95), December 1995.