

REPORT DOCUMENT

AD-A265 004

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to be 1 hour per response, including the time for reviewing the instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters and the Office of Management and Budget, Paperwork Reduction Project (0704-0188).


actions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information, including suggestions for reducing this burden, to Washington Headquarters and the Office of Management and Budget, Paperwork Reduction Project (0704-0188).

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE March 1993		3. REPORT TYPE AND DATES COVERED Professional Paper	
4. TITLE AND SUBTITLE HETEROGENEOUS SUPERCONCURRENT PARALLELISM (HeSCP)				5. FUNDING NUMBERS PR: ZW65 PE: 0601152N WU: DN302054	
6. AUTHOR(S) R. F. Freund				8. PERFORMING ORGANIZATION REPORT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Command, Control and Ocean Surveillance Center (NCCOSC) RDT&E Division San Diego, CA 92152-5001				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Chief of Naval Research Independent Exploratory Development Program (IED) OCNR-20T Arlington, VA 22217				11. SUPPLEMENTARY NOTES	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) HeSCP (Heterogeneous SuperConcurrent Parallelism, pronounced "hee-skip") is the <i>tuned</i> use of diverse processors to solve distinct computational needs, in which codes or code portions are directed to the processor(s) best suited for their execution. Implicit in this definition is the idea that different tasks or different portions of these tasks indicate different architectural requirements or different balances of requirements. The processors themselves are generally parallel or other HPC machines. Typically the processors are used concurrently and it is often the case that tools and methodologies used for parallel and vector processing, e.g., code profiling, flow analysis, data dependencies, etc., can be extended to the larger, meta-parallel processor. The <i>super</i> in SuperConcurrent refers to the traditional supercomputing aims of HeSCP, as well as the heterogeneous concurrency superimposed on the various parallel processors, and finally the aim of superlinear performance from the <i>tuned</i> orchestration of the heterogeneous components.					
93-11960 93 5 26 090					
Published in Newsletter of the Parallel Processing Technical Committee, Vol. 1, No. 1, Oct 1992, pp 4-7.					
14. SUBJECT TERMS high performance computing heterogeneous processing superconcurrency				15. NUMBER OF PAGES	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED				18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	
19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED				20. LIMITATION OF ABSTRACT SAME AS REPORT	

UNCLASSIFIED

21a. NAME OF RESPONSIBLE INDIVIDUAL R. F. Freund	21b. TELEPHONE (include Area Code) (619) 553-4071	21c. OFFICE SYMBOL Code 423

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.

Call for Newsletter Contributions

TCPP *Newsletter* will serve as a forum for scientists and engineers in the area of parallel processing to express their opinions. The emphasis of the newsletter will be to expedite exchange of information among its readers. Short articles in all areas related to parallel processing will be considered for publication. Items of interest include (but not limited to):

- technical summaries of research results,
- announcements of new parallel processing systems, technologies and standards,
- reports on projects under development in academia, industry and national laboratories,
- opinions on current and future trends in parallel processing,
- reviews of new books, and
- information on new and existing courses in parallel processing.

Opinions expressed in the contributions will be those of the individual authors rather than the official position of TCPP, the IEEE Computer Society, or organizations with which the authors may be affiliated. Besides the articles, the newsletter will cover announcements of general interest: calendar of events about forthcoming conferences, workshops and symposia, call for papers (including special issues of journals).

In its first year, TCPP is planning to bring out three issues of the newsletter. Short articles (in camera-ready form or $\text{\LaTeX}$ format) can be sent to either of the co-editors listed on the back cover of this newsletter.

Heterogeneous SuperConcurrent Parallelism (HeSCP)

Richard F. Freund

Naval Research and Development Center, San Diego, CA 92152-5000

WHAT IS HeSCP?

HeSCP (Heterogeneous SuperConcurrent Parallelism, pronounced "hee-skip") is the *tuned* use of diverse processors to solve distinct computational needs, in which codes or code portions are directed to the processor(s) best suited for their execution. Implicit in this definition is the idea that different tasks or different portions of these tasks indicate different architectural requirements or different balances of requirements. The processors themselves are generally parallel or other HPC machines. Typically the processors are used concurrently and it is often the case that tools and methodologies used for parallel and vector processing, e.g., code profiling, flow analysis, data dependencies, etc., can be extended to the larger, meta-parallel processor. The *super* in SuperConcurrent refers to the traditional supercomputing aims of HeSCP, as well as the heterogeneous concurrency superimposed on the various

Distribution /	
Availability Code	
Dist	Avail and/or Special
A-1	

parallel processors, and finally the aim of superlinear performance from the *tuned* orchestration of the heterogeneous components.

RELATIONSHIP TO OTHER HETEROGENEOUS EFFORTS

A number of other projects and products have described themselves as heterogeneous processing. Many of these employ clusters or workstations using the principle of opportunistic load-balancing (OLB) for assigning processes to processors. In these cases the term heterogeneity is partially justified because the workstations may be different and/or the connectivity between them may vary in either topology or bandwidths. However this approach, while often useful for cycle-soaking, consists of processing components essentially homogeneous with respect to computational characteristics (the different workstations). Thus this approach lacks the full computational spectrum of HeSCP. Another common type of heterogeneous endeavor is represented by the Cray supercomputer. The heart of computational units is the vector processor, but several processors can be arrayed in parallel, achieving a high-level MIMD parallelism. Thus the Cray is really a mixed-mode, partly heterogeneous processor. However this heterogeneity is static and allows for little, if any, "tuning", and therefore also does not display the full-range of heterogeneity potential.

FUNDAMENTAL HeSCP HEURISTIC

The cycle soaking philosophy of using OLB among clusters of workstations is effective only because the work stations are essentially homogeneous, i.e., it really doesn't make any performance difference which workstation is selected for any particular process. However the OLB approach is clearly inadequate on its own in a very heterogeneous environment, e.g., a CM-2 and a Convex. Just by chance, for example, one might assign a process to the CM-2, because it was the first available machine, that would cause it to limp along on only one (of its 64K) processors! OLB is simple and leads to the machines being busy, but it does not necessarily lead to faster computation of any program or set of programs. OLB should be secondary to (near-)optimal matching first. This leads us to the Fundamental HeSCP Heuristic: First evaluate suitability of tasks to processor types, then load-balance among selected machines for final assignment.

GOALS OF HeSCP

One way to look at some of the goals of HeSCP is from a mathematical programming point of view. In a very abstract sense, we are aiming to minimize the objective function of compute time, subject to a fixed constraint, e.g., cost, C . If C is set at some medium level, e.g., $C = \$1M$, we could say we are tuning our HeSCP environment for cost-effective computing. Typically this would stem from using a few mini-supers and small parallel machines, at one site, in place of one mid-sized supercomputer. On the other hand, if C is very large, say $C = \$10M$, then we are aiming to tune our environment for the most effective supercomputing. Typically this would take the form of several, different HPC machines, at different sites, linked by gigabit networks. Another dimension to look at is whether we are tuning our HeSCP system to get maximal throughput for a set of processes, or maximum performance on a single job.

CONNECTIVITY, BANDWIDTH, AND GRANULARITY

HeSCP is an abstract model with potentially many different goals and instantiations. Within this model, no specific connectivity, bandwidth, or granularity is either required or ex-

cluded. Clearly higher bandwidth or less data to transmit imply less transfer latency penalty and therefore greater ease to share subprocesses amongst different processors and thus potentially a finer granularity of effort. For example, a heterogeneous environment on a global WAN with only medium bandwidth may have to settle for coarse granularity of effort and therefore often only throughput effectiveness rather than superlinear performance on most individual tasks. Increasing the bandwidth dramatically could dramatically change the optimal strategy, and, in particular, imply far greater likelihood of optimal performance on individual tasks, as is also typically the case for Heterogeneous LANs. Since data is generally far larger than instructions, another level of heuristic is to first determine (when possible and by whatever means) where to place data initially so that it is least often moved.

HIGH LEVEL ORCHESTRATION TOOLS

Heterogeneity relies heavily on the development of network orchestration tools, e.g., PVM. However, by the Fundamental Heuristic, these tools in themselves are insufficient to lead to effective heterogeneous solutions. A missing ingredient is the rationale for assigning different subtasks to different processing components, i.e., the theoretical and empirical understanding of what kinds of computation fit what kind of architectures. The network orchestration tools and the rationale together, aka Distributed Intelligent Network System or DINS, satisfy the two components of the Fundamental Heuristic. DINS-like tools are necessary in any case to satisfy the anticipated future explosion in HPC requirements. Most HPC projects today rely on one or more computer scientists in addition to applications programmers. DINS would allow applications experts to develop HPC programs. This step is clearly a necessity if we are to process ten to a hundred times more computationally complex HPC projects in the 90's than those of the 80's.

PROGRAMMING PARADIGMS FOR HeSCP

Heterogeneous environments require heterogeneous programming paradigms. One level of this is the tuning of existing languages, e.g., HP FORTRAN or SISAL for heterogeneous environments. Other paradigms, expressly designed with heterogeneity in mind, e.g., HAsP and Cluster-M (see reference), have also been proposed. full-fledged heterogeneous environments will soon require a menu of potential paradigms to be available. A more advanced level of such heterogeneity in programming is the CHASM concept, in which different language instructions are intermingled in the broadcast meta- instruction stream. The combination of DINS above, and new programming paradigms represent great technical challenges. However I would argue they are both necessary for future HPC problems and will substantially ease the burden of applications development in future highly networked and heterogeneous software/hardware environments.

NHeSCP-C

The NHeSCP-C (National HeSCP Consortium) is a proposed group of academic and government scientists (at Emory, Kent State, NJIT, NRaD, Purdue, UCSD, and USC) aiming to study in full generality the varieties of tools needed in heterogeneous environments, develop a prototype DINS, and demonstrate the potential of this approach on several Grand Challenge problems, using existing HPC consortia.

REFERENCE

Currently the best single reference on HeSCP is the Proceedings of the First Workshop on Heterogeneous Processing held in March 1992 (IEEE order number 3007, ISBN 0-8186-3007-8). In addition, HAsP - Heterogeneous Associative Processing by Freund and Potter will appear in Computing Systems in Engineering.

UPCOMING EVENTS

The second Heterogeneous Processing Workshop (HPW2) of the International Parallel Processing Symposium (IPPS7) will be held in Newport Beach, California in April 93 (please contact R. F. Freund, freund@superc.nosc.mil for HPW2 and Viktor Prasanna, prasanna@halcyon.usc.edu for IPPS7). In addition there will be a Special Issue of IEEE Computer on Heterogeneous Processing to be published in June 1993 as well as a Special Issue of Journal of Parallel and Distributed Computing (JPDC) on Heterogeneous Processing to be published in January 1994. The deadline for submission to the Computer Special Issue has passed. Oct 19, 1992 is the deadline for the JPDC issue (contact R. F. Freund, freund@superc.nosc.mil or Vaidy Sunderam, vss@mathcs.emory.edu).

Executive Summary of the Report of the NSF-Sponsored Purdue Workshop on Grand Challenges in Computer Architecture for the Support of High Performance Computing*

*H. J. Siegel and S. Abraham
Purdue University, Workshop Co-Chairs*

B. Bain, Intel Corp. SSD	K. E. Batcher, Kent State Univ.
T. L. Casavant, Univ. of Iowa	D. DeGroot, Texas Instruments
J. B. Dennis, M.I.T.	D. C. Douglas, Thinking Machines Corp.
T-Y. Feng, Penn. State Univ.	J. R. Goodman, Univ. of Wisconsin
A. Huang, AT & T Bell Labs.	H. F. Jordan, Univ. of Colorado
J. R. Jump, Rice Univ.	Y. N. Patt, Univ. of Michigan
A. J. Smith, UC-Berkeley	J. E. Smith, Cray Research, Inc.
L. Snyder, Univ. of Washington	H. S. Stone, IBM T.J. Watson Research Ctr
R. Tuck, MasPar Computer Corp.	B. W. Wah, Univ. of Illinois
Z. Zalcstein, NSF Liaison	

The Purdue Workshop on Grand Challenges in Computer Architecture for the Support of High Performance Computing was held at Purdue University on December 12 and 13, 1991. The workshop was sponsored by the National Science Foundation to identify critical research

*This work was supported by the National Science Foundation Division of Computer and Computation Research Systems Program under grant number CCR-9200735.

topics in computer architecture as they relate to high performance computing. This is an executive summary of the workshop report.

After a wide-ranging discussion of the computational characteristics and requirements of the grand challenge *applications*, four major architectural challenges were identified as crucial to advancing the state of the art of high performance computation in the coming decade. These *computer architecture* grand challenges are summarized concisely below.

Challenge 1: Idealized Parallel Computer Models

A parallel computer model provides the interface between parallel hardware and parallel software. It is the idealization of computation that computer architects strive to support with the greatest possible performance. Although a single model may not fulfill the requirements of all effective architectures and application domains, the multitude of alternatives must be reduced to a small number to support portability of programs and reusability of program parts.

Challenge 2: Usable Peta-Ops Performance

This challenge addresses the need for usable computer performance orders of magnitude greater than both the giga-ops performance available today and the tera-ops performance that may be achieved soon. This computer performance cannot be obtained by simply interconnecting massive quantities of existing processor, memory, and I/O resources. Such a system would be unmanageable to program and would ineffectively utilize its processors. The challenge is to (1) dramatically improve and (2) effectively harness the base technologies impacting processors, memory, and I/O into a computer system such that the grand challenge applications programmer has available peta-ops (10^{15} ops) of usable processing performance.

Challenge 3: Computers in an Era of HDTV, Gigabyte Networks, and Visualization

Emerging technologies are providing an opportunity to support startling new communication intensive applications, such as digital video workstations that treat images as easily as characters are treated today. How can computer architecture and new communications technology evolve to enable such applications?

Challenge 4: Infrastructure for Prototyping Architectures

Testing a new idea in computer architecture has been a difficult process requiring large investments in building design tools and providing a suitable software environment for an experimental machine. Prototype development involves not only hardware, but also software in the form of compilers and operating systems. An infrastructure is needed to facilitate the study of the effects of new hardware technologies and machine organizations against different application requirements.

These grand challenges in computer architecture are inherently multidisciplinary and will require team efforts crossing boundaries from software to hardware to applications. While it is crucial that the above challenges be addressed, it is important to stress that the viability and usability of parallel computers is also a function of the supporting software systems. Thus, a substantial effort must be devoted to advancing the software aspects of high performance computing. This involves improving the software interface supported by parallel computers and developing languages, compilers, and software tools that simplify the task of parallelizing, mapping, and optimizing algorithms for efficient execution on parallel computers. In the arena

of high performance parallel computers, it is more important than ever for computer architects to consider the issues of system software, application needs, and usability when designing and implementing machines. To address the interaction of this full range of subjects is a challenge in itself.

The workshop report will be available in three printed forms.

A. Conference Paper: A summary of the report will appear in the proceedings of *Frontiers '92: The Fourth Symposium on the Frontiers of Massively Parallel Computation*, sponsored by the IEEE Computer Society and NASA, October 1992 (proceedings published by the IEEE Computer Society Press, Los Alamitos, CA).

B. Journal Article: This entire report is scheduled to appear in the November 1992 issue of the *Journal of Parallel and Distributed Computing* (published by Academic Press, Orlando, FL).

C. Technical Report: The entire report is also available as "Purdue University, School of Electrical Engineering, Technical Report Number 92-26" by sending a request to:

Technical Reports
School of Electrical Engineering
1285 Electrical Engineering Building
Purdue University
West Lafayette, IN 47907-1285
Fax: 317-494-6440

A Report on ICPP '92

Robert H. B. Netzer

Dept. of Computer Science, Brown University, Box 1910, Providence, RI 02912-1910

The 21st International Conference on Parallel Processing was held during August 17-21 this year in St. Charles, IL. In addition to the technical program, high points of the conference included the keynote address by Henry Burkhardt of Kendall Square Research, a lively panel that addressed the question "Are networks of workstations tomorrow's supercomputers," and six tutorials offered on a wide range of topics.

As usual, the conference contained three parallel tracks (architecture, software, and applications) with papers across the spectrum of parallel processing. This year, Trevor Mudge (of the University of Michigan) acted as conference chair; Quentin Stout and Kang Shin (also of the University of Michigan) were Co-chairs. About 500 papers were submitted, 30 Unlike previous years, the submissions were spread fairly evenly across the three tracks (with the architecture track receiving the most submissions and the software track the fewest). In addition to the technical program (and several well-stocked beer, wine, and cocktail parties), six tutorials were offered. Tom Casavant (of the University of Iowa) discussed visualization

tools for parallel programs. Yale Patt (of the University of Michigan) and Wen-mei Hwu (of the University of Illinois) covered hardware and software approaches to exploiting concurrency. Chuan-lin Wu (of UT-Austin) overviewed the current state of parallel computing. Jesse Fang (of HP Labs) and Lionel Ni (of Michigan State University) presented compiler techniques for automatic parallelization. Peter Kogge (of IBM) and Hanad Ghose (of SUNY at Binghamton) focused on pipelining and parallelism. Finally, H.J. Siegel (of Purdue University) spoke about using tightly coupled parallel machines.

In the opening session, several awards were given for outstanding and original papers. Daniel Slotnick received the most-original paper award for his paper "Adaptive, deadlock-free packet routing in torus networks with minimal storage," co-authored with Robert Cypher and Luis Gravano. Outstanding paper awards were presented to Arun Nanda and Lionel Ni for their paper "MAD kernels: An experimental testbed to study multiprocessor memory system behavior," and to Shridhar Shukla, Brian Little and Amr Zaky for their paper "A compile-time technique for controlling real-time execution of task-level data-flow graphs."

The opening session also included the keynote address by Henry Burkhardt III, founder of Kendall Square Research. He discussed the motivation behind the KSR1 (their controversial shared-memory multiprocessor) and its hardware and software architecture. KSR's goal was to create a scalable, general-purpose, easy-to-use supercomputer. The ALLCACHE memory architecture is novel: it organizes the entire memory as a collection of local caches (no distinction is made between local and shared memory). Processors are 64-bit custom CMOS, interconnected via a two-level hierarchy of rings. The machine's software is intended to be familiar and support standard scientific, business, and personal computing. The KSR1 runs traditional UNIX tools: OSF/1 UNIX, C, C++, FORTRAN, COBOL, plus a parallel version of the ORACLE relational database.

During the first evening of the conference a panel discussed the question "Are networks of workstations tomorrow's supercomputers?" The participants were Peter Kogge (of IBM), Tilak Agerwala (also of IBM), Lionel Ni (of Michigan State Univ.), Yale Patt (of the Univ. of Michigan), H.J. Siegel (of Purdue Univ.), and Salim Hariri (of Syracuse Univ.). The panel began with a lively talk by H.J. Siegel, with the analogy that supercomputing is like renting a fast car, but using networks of workstations is like hitchhiking. He contrasted the two approaches with respect to cost, security, availability, response time, variance of response time, reliability, and applicability. One example is cost: obtaining performance from rental cars costs lots of money, but hitchhiking requires only a pair of shoes. Similarly, obtaining high performance with supercomputers requires spending millions of dollars, but using networks of workstations requires less investment (especially since everyone has workstations, just like shoes). H.J.'s closing remark was that networks of workstations will probably not replace supercomputers in the future (and that if all people were hitchhikers, no one would pick them up!). The viewpoints of the other panelists were mixed: some basically agreed with H.J.'s point, but others argued that networks of workstations are more cost-effective for some problems. In the end, better arguments were made that future networks of workstations will not *replace* supercomputers, but will certainly be an *cheap, important source of computation* (appropriate for many important problems).

Let's Think About Parallel Processing Education

Russ Miller

Dept. of Computer Science, State University of New York, Buffalo, NY 14260

While the agenda for The IEEE Technical Committee on Parallel Processing is still evolving, it is clear that one of our missions should be to disseminate information pertaining to the education of future generations of scientists in terms of parallel processing. Therefore, we solicit information from you regarding current and proposed mechanisms for education. This includes formal graduate and undergraduate courses, whether or not experimental, regularly offered seminars and independent studies, interdisciplinary programs that involve parallel processing, as well as other programs and offerings that are aimed at parallel processing education. Initially, we plan to collate the information and provide a summary in some subsequent issue of the newsletter.

Secondly, we would like to solicit confidential comments, criticisms, and suggestions on books and monographs involving various aspects of parallel computing. Again, we plan to provide a summary.

Please send information via either E-mail or snailmail to:

Prof. Russ Miller, Department of Computer Science, 226 Bell Hall, State University of New York, Buffalo, NY 14260, USA, Ph: (716) 645-3295, E-mail: miller@cs.buffalo.edu.

Additional comments or suggestions involving educational issues should also be addressed to Russ Miller.

Announcement of a New Journal

Sajal K. Das

Dept. of Computer Science, University of North Texas, Denton, TX 76203-3886

A new journal **Parallel Algorithms and Applications (PAA)** is being planned to be published by Gordon and Breach Science Publishers, Reading, UK, with its first volume scheduled in September of 1992. The journal will be publishing papers relating to Parallel and Multiprocessor computer systems covering the following areas:

Parallel Algorithms: Design, Analysis and Usage in Numerical Analysis, Discrete Mathematics, Non-numerical, Geometric, Graphics, Genetic, Optimization, Pattern Recognition, Simulation, Signal/Image Processing, Systolic Algorithm.

Parallel Applications: Usage in the areas of Artificial Intelligence, Systems Software and Compilers, CAD/CAM, Databases, Expert Systems, Information retrieval, Neural Networks, Industrial, Scientific and Commercial Applications for Pipelined, Vector, Array, Parallel and Distributed Computers.

Prof. David J. Evans, Director, Parallel Algorithms Research Centre, Loughborough University of technology, Loughborough, Leicestershire, LE11 3TU, UK, telephone: (0509) 222670, FAX: (0509) 211586, is the **Editor-in-Chief** for the journal.

The **Editorial Review Board** consists of the following members: S. G. Akl (Canada), Y. P. Boglaev (Russia), M. Clint (UK), S. K. Das (USA), L. Dehne (Canada), C. C. Douglas (USA), C. S. Jeong (Korea), R. C. T. Lee (China), G. Loizou (UK), K. G. Margaritis (Greece), G. H. Megson (UK), S. Olariu (USA), N. Petkov (The Netherlands), M. J. Quinn (USA), V. Raymond-Smith (UK), I. Stojmenovic (Canada), M. Thuné (Sweden), T. Tollenare (Belgium), M. Vajteric (Czechoslovakia), S. A. Zenios (USA).

Contributors should submit their papers to the Editor in Chief at the above address. Further information about the submission process can be obtained from Prof. Sajal K. Das, Center for Research in Parallel and Distributed Computing (CRPDC), Department of Computer Science, P.O. Box 13886, University of North Texas, Denton, TX 76203-3886, USA. Tel: (817) 565-4256, E-mail: das@ponder.csci.unt.edu, FAX: (817) 565-2799.

Calls for Papers/Participation

Symposium on Integrated Systems, Previously the Conference on Advanced Research in VLSI, March 17-19, 1993, Seattle, Washington, USA. Submit 5 copies of the paper, not exceeding 15 pages, with a 500-word abstract postmarked by **Tuesday, October 6, 1992** to University of Washington Conference, 114 Sieg Hall, FR-35, Seattle, WA 98195, USA. Contact Kay Beck at the above address or at kbeck@cs.washington.edu.

The 13th Int'l Conference on Distributed Computing Systems, May 25-28, 1993, Pittsburgh, Pennsylvania, USA. Submit 6 copies of double-spaced manuscript (maximum of 20 pages) with an abstract and keywords by **Thursday, October 15, 1992** to Prof. Larry Wittie, Z4400 Computer Science, SUNY at Stony Brook, Stony Brook, NY 11794-4400, USA. Tel: (516)-632-8456, Fax: (516)-632-8334, E-mail: lw@sbcs.sunysb.edu. Contact Prof. Benjamin W. (Ben) Wah, Coordinated Science Laboratory, University of Illinois, MC228, 1101 W. Springfield Avenue, Urbana, IL 61801-3082, Tel: (217)-333-3516, Fax: (217)-244-7175, E-mail: b-wah@uiuc.edu.

Second Workshop on Heterogeneous Processing, April 13, 1993, Newport Beach, California, USA. Send 5 copies of complete paper (not to exceed 15 single-spaced, single-sided pages) to: HP '93 - EEB 200, Dept. of EE-Systems, University of Southern California, Los Angeles, CA 90089-2562. E-mail: hp93@halcyon.usc.edu. Manuscripts must be received by **October 15, 1992**.

Special Issue of the Journal of Parallel and Distributed Computing (JPDC) on Heterogeneous Processing to be published on January 1994. Five copies of complete double-spaced manuscripts should be received by either of the following Guest Editors by **October 19, 1992**: Dr. Richard F. Freund, Code 423, Naval Research and Development Center, San Diego, CA 92152-5000, USA, Tel: (619) 553-4071, Fax: (619) 553-5136, E-mail: freund@superc.nosc.mil; Prof. Vaidy Sunderam, Department of Mathematics and Computer Science, Emory University, 1300 Clifton Rd, Atlanta, GA 30322, USA, Tel: (404) 727-5926, Fax: (404) 727-5611, E-mail: vss@mathcs.emory.edu.

The 20th Annual Int'l Symposium on Computer Architecture, May 16-19, 1993, San Diego, California, USA. Send 8 copies (double-spaced and not exceeding 6000 words in length) of the paper to Prof. John Hennessy, Computer Systems Laboratory, Stanford University,

Stanford, CA 94305, USA, E-mail: hennessy@sierra.stanford.edu. Papers will be accepted for considerations until **November 1, 1992**.

Workshop on Parallel and Distributed Real-Time Systems, April 13, 1993, Newport Beach, California, USA. Submit five copies not exceeding 2500 words (about 10 pages) of the paper to either of the following program Co-chairs. Submissions should be received by **November 30, 1992**. Prof. Alexander D. Stoyenko or Prof. Lonnie R. Welch, The Real-Time Computing Laboratory, Dept. of Computer and Information Science, New Jersey Institute of Technology, Newark, NJ 07102, USA; Ph: (201) 596-5765, alex@vienna.njit.edu; Ph: (201) 596-5683, welch@vienna.njit.edu; Fax: (201) 596-5777.

Performance '93, Int'l Symposium on Computer Performance Modeling, Measurement, and Evaluation, September 29 - October 1, 1993, Roma, Italy. Contact Prof. Bruno Ciciani, Universita' di Roma "Tor Vergata", Dip. di Ingegneria Elettronica, via Della Ricerca Scientifica, I-00173, Roma, Italy, Tel + 39-6-7259.4478, Fax +39-6-2020519, E-mail: PERF93@irmias.bitnet or TUCCI@utovrm.it. Submit 6 copies of papers (not exceeding 20 doubled-spaced pages) by **November 30, 1992**.

North-America authors submit papers to: Dr. Stephen S. Lavenberg, IBM T. J. Watson Research Center, P. O. Box, 704 Yorktown Heights, N. Y. 10598, USA, Tel + 1-914-784.7573. E-mail: SSLAVEN@ watson.ibm.com.

Europe and others submit papers to: Prof. Giuseppe Iazeolla, Universita' di Roma "Tor Vergata", Dip. di Ingegneria Elettronica, via Della Ricerca Scientifica, I-00173, Roma, Italy. Tel + 39-6-7259.4486, E-mail: IAZEOLLA@irmias.bitnet.

Joint Symposium on Parallel Processing (JSPP '93), May 17-19, 1993, Waseda University, Tokyo, Japan. Send 8 copies of extended abstract (not exceeding 2000 words, 4 single-spaced pages), no later than **December 4, 1992**, to Prof. Hiroshi Nakashima, Dept. of Information Science, Faculty of Engineering, Kyoto University, Yoshida Hon-machi, Sakyo-Ku, Kyoto, 606 Japan, Tel: +81-75-753-5383, Fax: +81-75-753-5379, E-mail: jssp93@kuis.kyoto-u.ac.jp.

Second Int'l Symposium on High Performance Distributed Computing, (HPDC-2), July 21-23, 1993, Spokane, Washington, USA. Submit five copies of the manuscript (not exceeding 20 double-spaced pages) by **January 15, 1993** to either: Prof. C. S. Raghavendra, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164-2752, USA, Ph: (509)335-8246, E-mail: raghu@eecs.wsu.edu; or Prof. Salim Hariri, Electrical and Computer Engineering Department, Syracuse University, 111 Link Hall, Syracuse, NY 13244, USA, Ph: (315)443-4282, E-mail: hariri@cat.syr.edu.

Calendar of Events

October 12-15, 1992, Fifth Int'l Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-V), Boston, Massachusetts, USA. Contact Darlene Forsyth at (508) 436-5450 or asplos@apollo.hp.com.

October 19-21, 1992, The Fourth Symposium on the Frontiers of Massively Parallel Computation, Mclean, Virginia, USA. Contact Pearl Wang, Dept. of Computer Science,

George Mason Univ, Fairfax, Va 22030, Ph: (703)-993-1527, E-mail: f92info@cs.gma.edu.

November 16-18, 1992, Supercomputing '92, Minneapolis, Minnesota, USA. Contact Supercomputing '92, Conference Office, SCD/NCAR, PO Box 3000, Boulder, CO 80307, USA. E-mail: sc92info@ncar.ucar.edu.

December 2-4, 1992, Workshop on Cluster Computing, Supercomputer Computations Research Institute, Florida State University, Tallahassee, Florida 32306, USA. Contact the organizers at (904)-644-1010 or send E-mail to cluster-workshop@scri.fsu.edu.

December 7-9, 1992, Fifth Australian Supercomputing Conference, Melbourne, Australia. Contact Fifth Australian Supercomputing Conference, Royal Melbourne Institute of Technology, Computer Centres - Building 5, 124 Latrobe Street, Melbourne Vic 3000, AUSTRALIA, E-mail: 5asc@rmit.edu.au, Fax +61 3 663 5652.

December 7-9, 1992, Fourth IEEE Symposium on Parallel and Distributed Processing, Dallas, Texas, USA. Contact Prof. Kai Hwang, Dept. of EE-Systems, University of Southern California, Los Angeles, CA 90089-0781, USA, E-mail: kaihwang@panda.usc.edu.

January 5-8, 1993: Twenty-sixth Hawaii Int'l Conference on System Sciences (HICSS-26), Koloa, Hawaii. Contact Pamela S. Harrington, Center for Executive Development, University of Hawaii at Manoa, 2404 Maile Way, B-101, Honolulu, HI 96822, Tel: (808) 956-7396, Fax (808) 956-3766, E-mail: hicss@uhunix.uhcc.hawaii.edu.

January 20-22, 1993: Working Conference on Architectures and Compilation Techniques for Fine- and Medium-Grain Parallelism, Orlando, Florida, USA. Contact Prof. Jean-Luc Gaudiot, Univ. of Southern California, Dept. of EE-Systems, Los Angeles, CA 90089-0781, USA, Tel: (213)-740-4484, Fax: (213)-740-4449, E-mail: gaudiot@usc.edu.

March 17-19, 1993 (*): Symposium on Integrated Systems, previously known as the Conference on Advanced Research in VLSI.

April 13-16, 1993: International Parallel Processing Symposium (IPPS '93), Newport Beach, California, USA. Contact Prof. Viktor K. Prasanna, Univ. of Southern California, Dept. of EE-Systems, Los Angeles, CA 90089-0781, USA, Fax: (213)-740-4449, E-mail: ipps93@halcyon.usc.edu.

April 13, 1993 (*): Second Workshop on Heterogeneous Processing.

April 13, 1993 (*): Workshop on Parallel and Distributed Real-Time Systems.

May 16-19, 1993 (*): The 20th Annual Int'l Symposium on Computer Architecture.

May 17-19, 1993 (*): Joint Symposium on Parallel Processing, (JSPP '93).

May 25-28, 1993 (*): The 13th Int'l Conference on Distributed Computing Systems.

July 21-23, 1993 (*): Second Int'l Symposium on High Performance Distributed Computing, HPDC-2.

Note: A (*) indicates the detailed information is included in Calls for Papers/ Participation section.

Lets hear from you

Contributions should be sent to either of the newsletter Co-editors in camera-ready form or $\text{\LaTeX}$ format. For details, please refer to the *Call for Newsletter Contributions* section in this issue. If you wish to be on the TCPP mailing list to receive announcements and newsletters, please send a note to tcpp@halcyon.usc.edu.

Newsletter Co-editors

Ms. Sally Jelinek
Electronic Design Associates
3184-H, Airway Ave.
Costa Mesa, CA 92626
Ph: (714) 751-5992
Fax: (714) 979-7488

Prof. D. K. Panda
Dept. of Computer and Info. Science
The Ohio State University
Columbus, OH 43210-1277
Ph: (614) 292-5199
FAX: (614) 292-2911
E-mail: panda@cis.ohio-state.edu
