Hindawi Publishing Corporation
Mathematical Problems in Engineering
Volume 2016, Article ID 2632306, 3 pages
http://dx.doi.org/10.1155/2016/2632306

Editorial

Advances in High Performance Computing and Related Issues

Veljko Milutinović,[1] Borko Furht,[2] Zoran Obradović,[3,4,5] and Nenad Korolija[1]

[1] *Department of Computer Engineering and Information Theory, School of Electrical Engineering, University of Belgrade, Bulevar Kralja Aleksandra 73, 11120 Belgrade, Serbia*
[2] *NSF I/UCRC CAKE, FAU Site, Department of Computer & Electrical Engineering and Computer Science, College of Engineering and Computer Science, Florida Atlantic University, 777 Glades Road, Boca Raton, FL 33431, USA*
[3] *Center for Data Analytics and Biomedical Informatics, Temple University, 1925 N. 12th Street (SERC: 035-02), Philadelphia, PA 19122, USA*
[4] *Computer and Information Sciences Department, Temple University, 1925 N. 12th Street (SERC: 035-02), Philadelphia, PA 19122, USA*
[5] *Statistics Department, Fox School of Business (Secondary Appointment), Temple University, 1925 N. 12th Street (SERC: 035-02), Philadelphia, PA 19122, USA*

Correspondence should be addressed to Veljko Milutinović; vm@etf.rs

Received 28 September 2016; Accepted 28 September 2016

1. Introduction

With the current trend of increasing the number of cores in modern computers, society faces a problem of high energy dissipation, while increased performance is limited by communication delays. Further major advances in computing are possible only if some important notions are taken seriously into consideration. In our opinion, the major four notions of interest for the time to come are coming from the research of four Nobel Laureates who worked in science or economy. These are Richard Feynman, Ilya Prigogine, Daniel Kahneman, and Andre Geim. The work of Feynman et al. [1] suggests that those solutions are better than minimizing data movements and using shorter communications distances. The work of Nicolis and Prigogine [2] suggests that energy injections could lower the entropy of a computing system, which makes it easier for compiler optimizations. The work of Kahneman and Tversky [3] suggests that some applications do not require optimal computing, meaning that the saved resources could be reinvested elsewhere, for much higher benefits in another domain. The work of Geim [4] suggests that solutions that tolerate latencies could bring benefits in other more important domains. The quality of the research selected for presentation in this special issue could be judged also by how much the above-mentioned notions were taken care of.

Hardware dataflow computing solves problems of control-flow computing by executing more instructions in parallel on the same chip die. This approach makes communication channels shorter and the execution faster, even with much lower frequencies, while reducing energy consumption. However, the full benefits of dataflow computing are obtained only if combined with control-flow machines and if appropriate advances are done in both the architectures and algorithms for control-flow machines.

Historically, the number of instructions executed per second on processors approximately has doubled every two years [5]. Technology limits are nearly reached and it is not possible to increase the speed of the CPU any more even if the growing cost of energy is not an issue. Multicore and many core paradigms address this problem by introducing many cores instead of one, resulting in a paradigm shift from designing fastest single-core applications to parallelizing program execution. However, communication delays and energy consumption limit these paradigms as well. Hardware dataflow paradigm [6–9] naturally solves these problems by treating computer processing as factory production lines instead of one or many specialized workers. The approach is based on hardware being configured for executing certain sets of instructions, where all instructions could be executed in parallel. Since high performance computing algorithms normally also include instructions that are not repeatedly executed, dataflow computers are often combined with a control-flow processor. For CPU-demanding and data-demanding applications, one can design data flow and program a dataflow computer accordingly. Until recently, dataflow computer programmers would have to be hardware specialists. By noticing the advantage of pipelining execution of loops, industries have developed various tools that facilitate the process of transforming Von Neumann architecture applications into dataflow applications.

2. Problem Statement

However, programming dataflow computer architectures comes with a cost. When programming control-flow computer architectures, programmers only consider executing instructions sequentially. In order to cope with data dependencies in the case of dataflow programming, redesigning algorithms is needed and this requires new mathematical models. New paradigms are usually hard to accept, and therefore redesigning algorithms is needed before programmers change their programming preferences. As the percentage of dataflow computers rises, it is expected that additional algorithms will become available in open literature.

Towards this objective, authors of this special issue presented their original research articles that seek to combine existing and new paradigms in order to achieve better execution performances, lower power consumption, and therefore dissipation, as well as lower hardware costs.

3. Reducing Computation and Power Consumption

In the article entitled "Elimination of the Redundancy Related to Combining Algorithms to Improve the PDP Evaluation Performance" by F. Deng et al., a novel method is proposed for eliminating redundant policies loaded on the policy decision point (PDP) in the authorization access control model and therefore improves both storage usage and evaluation performance of the PDP. Experimental results show that the evaluation performance of the PDP can be prominently improved by eliminating the redundancy related to combining algorithms.

An article entitled "The Role of High Performance Computing and Communication for Real-Time Biofeedback in Sport" by A. Umek and A. Kos regards the main technological challenges of real-time biofeedback in sports. A multiuser signal processing in a football match is recognized as a high performance application that needs high-speed communication and high performance remote computing. Dataflow computing is found to be a good choice [10] for real-time biofeedback systems with large data streams.

In "A Protocol for Provably Secure Authentication of a Tiny Entity to a High Performance Computing One" by S. Tomović et al., the problem of developing authentication protocols dedicated to a specific scenario where an entity with limited computational capabilities should prove the identity to a computationally powerful verifier is addressed. It is shown that the proposed protocol is secure against active attacking scenarios and so-called GRS man-in-the-middle (MIM) attacking scenarios.

A manuscript by B. Zhang et al. entitled "Probabilistic Analysis of Steady-State Temperature and Maximum Frequency of Multicore Processors considering Workload Variation" presents a probabilistic method to analyze the temperature and maximum frequency for multicore processors based on workload variations. Experimental results provide evidence that hotspot temperatures of multicore processors are not deterministic and have significant variations, and the number of active cores and running frequency simultaneously determine the probabilistic distribution of hotspot temperatures.

In a paper by L. Verdoscia and R. Giorgi entitled "A Data-Flow Soft-Core Processor for Accelerating Scientific Calculation on FPGAs," a new type of soft-core processor called the "Data-Flow Soft-Core" is introduced that can be implemented through FPGA technology with adequate interconnected resources, eliminating partial data and instructions as traffic for load and store activities. The proposed design aims at combining the performance of a fine-grained dataflow architecture with the flexibility of reconfiguration, without requiring a partial reconfiguration or a new bitstream for reprogramming it.

An article entitled "An Encryption Technique for Provably Secure Transmission from a High Performance Computing Entity to a Tiny One" by M. J. Mihaljević et al. proposes an encryption/decryption approach dedicated to a one-way communication between a transmitter which is a computationally powerful party and a receiver with limited computational capabilities based on stream ciphering and simulation of a binary channel which degrades channel inputs by inserting random bits. It is shown that deliberate and secret key controlled insertion of random bits into the basic ciphertext provides security enhancement of the resulting encryption scheme.

Potentially enormous computing resources available in a distributed system are exploited effective in J. Yang et al. paper entitled "A Hierarchical Load Balancing Strategy Considering Communication Delay Overhead for Large Distributed Computing Systems." In this study a hierarchical load balancing strategy is proposed based on a generalized neural network (HLBSGNN) for hiding communication delays and therefore achieving scalability.

A paper entitled "A Novel CSR-Based Sparse Matrix-Vector Multiplication on GPUs" by G. He and J. Gao addresses the problem of irregular memory access patterns that occur when multiplying sparse matrices based on the compressed sparse row (CSR) format. In this study a CSR-based SpMV on the GPU is proposed which involves two kernels and a middleware in order to allow CSR arrays accesses in a fully coalesced manner.

In "Buckling Instability Behavior of Steel Bridge under Fire Hazard" by Y. Wang and M. Liu critical buckling stress of a bridge under fire hazard and a thermal analysis model of a steel bridge is characterized by a Fire Dynamics Simulator. Thermal parameters of the steel are determined by a polynomial fitting method and finite element software ANSYS.

Finally, in "Mining the IPTV Channel Change Event Stream to Discover Insight and Detect Ads" by M. Kren et al., it is analyzed how the data stream of the user-generated channel change events received from the entire IPTV network can be mined to obtain insights about the content. This study also predicts the occurrence of TV ads with high probability and shows that the approach could be extended to model the user behavior and classify viewership in multiple dimensions.

4. Conclusion

The control-flow computing paradigm assumes processing units capable of executing all instructions defined by computer architecture. However, it could execute only a few instructions simultaneously. Hardware dataflow computers solve this problem by designing hardware in such a manner that each instruction that it should execute is made in hardware and this part of hardware is connected to those parts of hardware that are responsible for dependent instructions. Therefore, executing instructions in a row is done by a flow of electrical signals through the hardware. As soon as one instruction is executed, input for a new instruction could be given. In this way, thousands of instructions can be executed in parallel.

In this special issue, authors present new trends in high performance computing, with the goal of reducing computer execution time, even when using dataflow hardware with an order of magnitude lower frequency than modern control-flow processors. A major part of price in high performance computing is the cost of electrical power. Apart from performance improvement, authors reduced power dissipation and therefore the overall cost of computing.

It is expected that future computers would include both hardware dataflow and control-flow processors. In this way, one could utilize high frequencies of modern control-flow processors, but also the parallel execution capabilities with low power consumption available in dataflow computing for instructions that should be executed over and over again. As of 2016, Maxeler dataflow is available through AWS, which is a cornerstone for the development of new horizons in modern computing.

Veljko Milutinović
Borko Furht
Zoran Obradović
Nenad Korolija

References

[1] R. Feynman, A. Hey, and R. Allen, *Feynman Lectures on Computation*, Addison-Wesley Longman, 1998.

[2] G. Nicolis and I. Prigogine, *Self-Organization in Nonequilibrium Systems*, Wiley-Interscience, New York, NY, USA, 1977.

[3] D. Kahneman and A. Tversky, "Prospect theory: an analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–292, 1979.

[4] K. S. Novoselov, A. K. Geim, S. V. Morozov et al., "Electric field in atomically thin carbon films," *Science*, vol. 306, no. 5696, pp. 666–669, 2004.

[5] G. E. Moore, "Lithography and the future of Moore's law," in *Proceedings of the Advances in Resist Technology and Processing XII*, vol. 2438 of *Proceedings of SPIE*, Santa Clara, Calif, USA, February 1995.

[6] J. Sharp, *Data Flow Computing: Theory and Practice*, Intellect Books, 1992.

[7] P. C. Treleaven, R. P. Hopkins, and P. W. Rautenbach, "Combining data flow and control flow computing," *The Computer Journal*, vol. 25, no. 2, pp. 207–217, 1982.

[8] R. Nikhil, "Can dataflow subsume von Neumann computing?" *ACM SIGARCH Computer Architecture News*, vol. 17, no. 3, 1989.

[9] R. Iannucci, "Toward a dataflow/von Neumann hybrid architecture," *ACM SIGARCH Computer Architecture News*, vol. 16, no. 2, pp. 131–140, 1988.

[10] A. Kos, V. Ranković, and S. Tomažič, "Sorting networks on Maxeler dataflow supercomputing systems," in *Dataflow Processing*, vol. 96 of *Advances in Computers*, pp. 139–186, Elsevier, Academic Press, Amsterdam, The Netherlands, 2015.

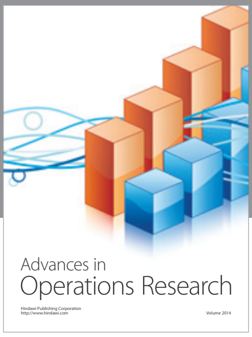
Advances in
Operations Research

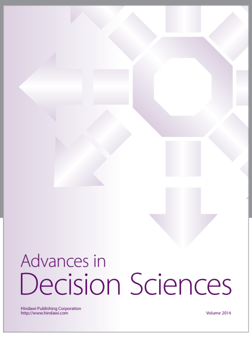
Advances in
Decision Sciences

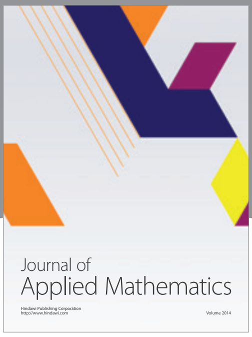
Journal of
Applied Mathematics

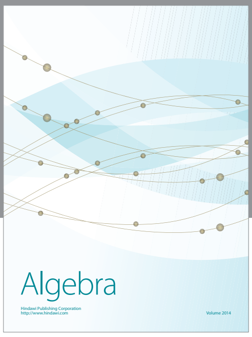
Algebra

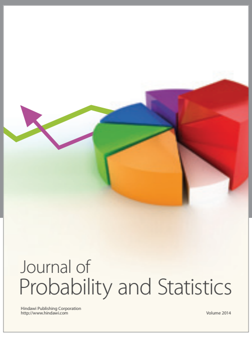
Journal of
Probability and Statistics

The Scientific
World Journal

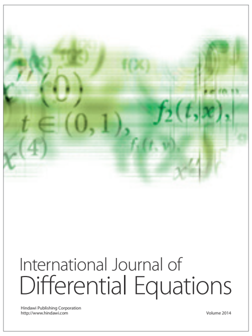
International Journal of
Differential Equations

International Journal of
Combinatorics

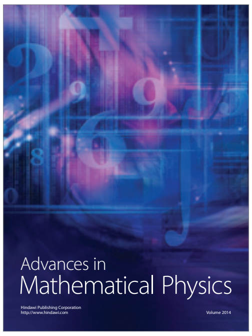
Advances in
Mathematical Physics

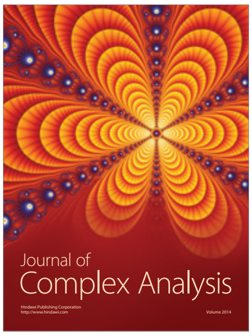
Journal of
Complex Analysis

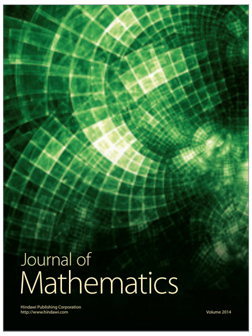
Journal of
Mathematics

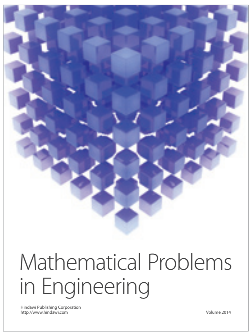
Mathematical Problems
in Engineering

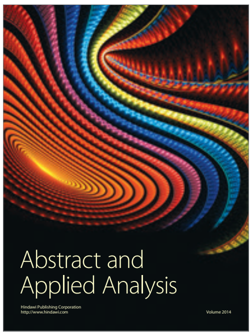
Abstract and
Applied Analysis

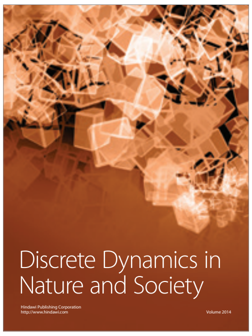
Discrete Dynamics in
Nature and Society

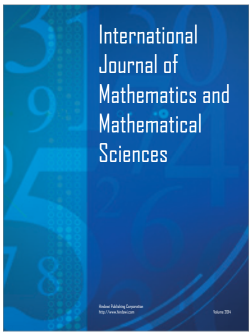
International
Journal of
Mathematics and
Mathematical
Sciences

Journal of
Discrete Mathematics

Journal of
Function Spaces

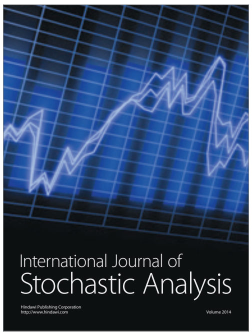
International Journal of
Stochastic Analysis

Journal of
Optimization