sudheer@anl.gov

Sudheer Chunduri

Research Interests

High Performance Computing, Performance Modeling and Analysis

Work Experience

Post-doctoral Appointee, Performance Engineering Group, Leadership Computing Facility, Argonne National Laboratory, June 2016 - Present
Researcher, High Performance Computing, IBM India Research Lab, India, July 2013 - May 2016
Assistant Professor, Sri Sathya Sai Institute of Higher Learning, India, July 2011 - May 2013

Educational background

Ph.D., Computer Science, Sri Sathya Sai Institute of Higher Learning, India, April 2013
- Advisor: *Prof. Ashok Srinivasan, Florida State University*
- Thesis: *Topology and Routing Aware Mapping on Parallel Processors*

M.Tech., Computer Science, Sri Sathya Sai Institute of Higher Learning, India, March 2006, GPA: 4.85/5
- Thesis Advisor: *Shakti Kapoor, STSM, STG, IBM Austin*

B.Tech., Information Technology, RVR&JC College of Engineering, India, April 2004

Peer Reviewed Publications

Journal Publications

Dynamic Load Balancing for Petascale Quantum Monte Carlo Applications: The Alias Method, **C.D. Sudheer**, S. Krishnan, A. Srinivasan, and P. R. C. Kent. Computer Physics Communications, Volume 184, Issue 2, February 2013, Pages 284292, Impact Factor : 3.268, (5-year Impact Factor: 2.812).

Static and Dynamic Frequency Scaling on Multicore CPUs, Wenlei Bao, Changwan Hong, Sriram Krishnamoorthy, **C. D. Sudheer**, Louis-Noel Pouchet, P. Sadayappan, Accepted for publication in ACM Transactions on Architecture and Code Optimization (TACO) 2016.

A59-Self-Propelled Pedestrian Dynamics Model for Studying Infectious Disease Propagation during Air-Travel., Sirish Namilae, Ashok Srinivasan, **C.D. Sudheer**, Anuj Mubayi, Robert Pahle, and Mathew Scotch., Journal of Transport & Health 3, no. 2 (2016): S40.

Reducing the Disk IO Bandwidth Bottleneck through Fast Floating Point Compression using Accelerators, Ajith Padyana, **Devi Sudheer**, Pallav Kumar Baruah, Ashok Srinivasan, International Journal of Advanced Computer Research (ISSN (print): 2249-7277 ISSN (online): 2277-7970), Volume-4 Number-1 Issue-14 March-2014.

Conference Publications

Optimizing Massively Parallel Simulations of Infection Spread Through Air-Travel for Policy Analysis, Ashok Srinivasan, **C. D. Sudheer**, Sirish Namilae, accepted for publication at CCGrid 2016.

Efficient Barrier Implementation on the POWER8 Processor, **C. D. Sudheer**, Ashok Srinivasan, pp. 165-173, 2015 Proceedings of the IEEE 22nd International Conference on High Performance Computing, December 2015, Bangalore, India.

A Semi-Discrete Matrix Free Spectral Element Adjoint Model of 3D Elastic Wave Equation, Steven Moore, Lior Horesh, Sergiy Zhuk, **Devi Sudheer Chunduri**, Tigran Tchrakian, Albert Akhriev, Alberto Costa Nogueira Junior and Andrew Rawlinson, 2015 SIAM Conference on Mathematical & Computational Issues in the Geosciences, June 29 to July 2, 2015, Stanford University, Stanford, California USA.

Semi-Discrete Matrix-Free Formulation of 3D Elastic Full Waveform Inversion Modeling, Stephen Moore, **Devi Sudheer Chunduri**, Sergiy Zhuk, Tigran Tchrakian, Ewout van den Berg, Albert Akhriev, Alberto Costa Nogueira Junior, Andrew Rawlinson and Lior Horesh, Euro-Par 2015, LNCS 9233, pp. 507-518,2015.

Semi-discrete Matrix Free Formulation of 3D Full Waveform Elastic Modeling and Inversion, Stephen Moore, Sergiy Zhuk, **Devi Sudheer Chunduri**, Tigran Tchrakian, Ewout van den Berg, Albert Akhriev, Alberto Costa Nogueira Junior, Andrew Rawlinson and Lior Horesh, WS10 - Full Waveform Inversion for Near-surface Characterization, 77th EAGE Conference 2015, Madrid, Spain.

Scaling up the training of Deep CNNs for Human Action Recognition, M. Sai Rajeswar, A. Ravi Sankar, Vineeth N. Balasubramanian, **C.D. Sudheer**, PARLEARNING workshop 2015, IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW), pp. 1172 - 1177, 25-29 May 2015.

Optimizing MPI Collectives on Intel MIC Through Effective Use of Cache, Pinak Panigrahi, Sriram Kanchiraju, Ashok Srinivasan, **C.D. Sudheer**, 2014 International Conference on Parallel, Distributed and Grid Computing (PDGC), pp. 88 - 93, 11-13 Dec. 2014, (Best Paper Award).

Parallel Learning of Deep Convolutional Neural networks and its Application to Action Recognition, M. Sai Rajeswar, A. Ravi Sankar, Vineeth N. Balasubramanian, **C.D. Sudheer**, Proceedings of the IEEE International Conference on High Performance Computing - Student Research Symposium, Goa, India, 2014, (NVIDIA SRS Best GPU paper award).

Optimization of the Hop-Byte Metric for Effective Topology Aware Mapping, **C.D. Sudheer**, Ashok Srinivasan, Proceedings of the 19th IEEE International Conference on High Performance Computing (HiPC), 2012, (Acceptance rate: 25%).

Optimizing Assignment of Threads to SPEs of the Cell BE Processor, **C.D. Sudheer**, T. Nagaraju, P.K. Baruah, Ashok Srinivasan, 10th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC), Proceedings of the 23rd International Parallel and Distributed Processing Symposium, IEEE, 2009, (Citations: SG - 5).

High Throughput Compression of Floating point numbers in GPUs, Ajith Padyana, **C.D Sudheer**, P.K. Baruah, Ashok Srinvasan, Proceedings of the 2nd IEEE International Conference on Parallel, Distributed and Grid Computing - 2012 Himachal Pradesh, December 2012.

Investigating Algorithmic Techniques for Enhancing Application Performance on Multicore Processors, **C.D. Sudheer** (Advisor: Ashok Srinivasan), PhD Forum at IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2009.

A Communication Model for Determining Optimal Affinity on the Cell BE processor, **C.D. Sudheer**, Sriram, S.: Proceedings of the 16th IEEE International Conference on High Performance Computing (HiPC), Student Research Symposium, Dec 2009.

Research Projects at IBM

Scaling DL_MESO on to POWER8+GPUs

- The initial goal of this project is to port a meso scale particle simulation package called DL_MESO on to the GPUs. Later, the code need to be tuned such that both the CPU (POWER8) and GPU are utilized optimally.

Full Seismic Waveform Modeling and Inversion

- Ported the 2D FDFD Acoustic Forward modeling code written in MATLAB to C, and also implemented a salable parallel version of the code using PETSc library.
- Worked on performance optimization of the 3D SEM Elastic Forward modeling code using a hybrid MPI + OpenMP based approach. Implemented core level optimizations using QPX intrinsics on BG/Q and AVX intrinsics on Intel SandyBridge based nodes. Also ported the code to make use the Intel Xeon Phi accelerators using MIC intrinsics.

Large Scale Parallel Community Detection

- The algorithm is based on the weighted label propagation, which works solely on local information, thus giving it the scalability advantage over conventional approaches. Our parallel algorithm employs load balancing mechanism across the compute nodes, along with communication-computation overlap via asynchronous communication to achieve the scalable performance.
- Experimental results on well-known real life massive scale graphs such as uk-2002 (261M edges) and also on RMAT graphs with 10M - 100M edges, demonstrate the superior performance and scalability of our algorithm compared to the well-known approaches for community detection.
- The code is also ported to Intel MIC systems with a hybrid approach to parallelism at the node level using both the CPU and the MIC accelerator for label processing.

Research Projects

Topology and routing aware mapping tool for massively parallel processors

- Developed general mapping techniques by posing the hop-byte metric as a quadratic assignment problem (QAP). Rather than using the metric just for evaluation of the mapping quality, the idea is intuitive in optimizing the metric itself.
- A metric based on the idea of minimizing the number of bottleneck links, called the maximum contention metric, requires the routing information along with the topology details. We showed that our heuristics for optimizing this metric are more effective in reduction communication costs.
- Tools/programming methods used: C, C++, Cray MPI, Cray Scheduler, Python, OSU MPI benchmarks and large scale code profiler - HPC toolkit.

Topology aware implementation of Global Arrays data management for QMCPACK

- An automated data management approach is generally used to enable existing QMCPACK applications using the GA library to significantly enhance the range of problem sizes that can be handled.
- In this work, the topology and routing aware grouping of GA groups was implemented such that the contention created by the communication in the GA calls is minimized.
- Tools/programming methods used: C, C++, Global Arrays and MPI.

Optimal dynamic load balancing algorithm for large scale codes involving near identical computational tasks

- A new load balancing algorithm for Quantum Monte Carlo and showed that it scales well, especially when compared with existing implementations.
- We also theoretically analyze its performance characteristics under a variety of metrics, in order to provide more insight into its strengths and limitations.
- An important feature of the new algorithm is that the load can be perfectly balanced with each process receiving at most one message. It is also optimal in the maximum size of messages received by any process.
- Tools/programming methods used: C, MPI, OSU MPI benchmarks and profilers such TAU, IPM.

Optimizing assignment of threads to SPEs on the Cell BE Processor

- The actual bandwidth obtained for inter-SPE communication is strongly influenced by the assignment of threads to SPEs (thread-SPE affinity) in many realistic communication patterns. We identify the bottlenecks to optimal performance and use this information to determine good

affinities for common communication patterns.
- Our solutions improve performance by up to a factor of two over the default assignment. We also discussed the optimization of affinity on a Cell blade consisting of two Cell processors, and provided a software tool to help with this.
- Tools/programming methods used: Cell BE programming involving signals, DMAs and mailboxes, C and MPI.

Reducing the disk IO bandwidth bottleneck through fast floating point compression using accelerators
- We proposed a compression technique based on time-series analysis, and investigate its effectiveness on floating point data from a variety of applications.
- We show that significant reduction in IO time can be achieved, even accounting for the compression overhead, on a Cell BE processor. In our experiments, the typical improvement was around 30%, varying from a slight loss in performance to a factor of seven improvement, depending on the type of application.
- The contribution of this work lies in demonstrating the potential of floating point compression in reducing the IO bandwidth bottleneck for important classes of scientific applications.
- Tools/programming methods used: Cell BE programming, GPU CUDA programming and C.

Design and implementation of an optimized Proc file system for a message passing operating system
- The aim is to design a process(proc) file system in the Minix OS by first understanding the implementation of it in Linux and then implement it for Minix OS running on PowerPC 405GP board.
- Implementation involved the debugging of Minix on PowerPC using IBM RISCWatch debugger.
- Tools/programming methods used: MINIX OS, PowerPC ISA and C.

Mini software projects
- Implementation of TFTP (Trivial File Transfer Protocol) in Java. It was also provided with a GUI.
- Implementation of a primitive web server / proxy / browser in UNIX.
- Tools/programming methods used: Java, C and UNIX programming internals.

Teaching

High Performance Computing with Accelerators, 2012. http://dmacssite.github.io

- Students of this course learn skills to produce scalable parallel programs targeting the unique requirements for obtaining high performance on GPUs and multi-cores. We will compare and contrast parallel programming style for GPUs and conventional multi-cores.
- *This course was highly successful resulting in 9 student papers, more than one third of the total papers, accepted for Student Research Symposium, at HiPC 2012. And also, 3 out of the 4 awards constituted for Best Presentation and Best Poster were secured by the students of this course.*

Programming for Performance, 2011, 2012 and 2013. http://progforperf.github.io

- To obtain the high level of performance needed in scientific computing, it is necessary for programs to exploit many of the features of modern computer architectures. In this course, we study the performance-critical features of state-of-the-art architectures such as Intel MIC, Intel SandyBridge and AMD Opetron, and discuss how applications can take advantage of them to obtain high performance.
- This is not a course on software tricks; rather, the emphasis is on abstractions of computer architecture, understanding performance, and obtaining performance when you need it.

Parallel Computing, 2013. http://parallelcomp.github.io

- This course is a comprehensive exploration of parallel programming paradigms, examining core concepts, focusing on a subset of widely used contemporary parallel programming models, and

providing context with a small set of parallel algorithms.

Computer Organization and Design, 2010, 2011 and 2012.

- This course provides a programmer's view of how computer systems execute programs, store information, and communicate. It enables students to become more effective programmers, especially in dealing with issues of performance, portability and robustness.
- It also serves as a foundation for courses on compilers, networks, operating systems, and computer architecture, where a deeper understanding of systems-level issues is required.
- Topics covered include: machine-level code and its generation by optimizing compilers, performance evaluation and optimization, computer arithmetic, memory organization and management and IO.

Processor Architecture and its Applications, 2008.

- The objective of this course is to teach intricate details of processor system Architecture. The idea is to design and develop an instruction set simulator for a simple in-order single issue RISC processor.
- It involves designing an orthogonal ISA for a hypothetical architecture, simulation routines for pipeline functional units, assembler, physically indexed cache, and then translation.

Operating Systems Design and Implementation, 2007.
Systems Programming using MINIX Operating System, 2006.

- The aim is to explain the interaction between the Operating System and the Architecture. MINIX OS code for the PowerPC architecture is used in the course.
- Various topics covered in the course: context switching (system call interrupt), OS message passing mechanism, keyboard interrupt, Programmable Interrupt Timer interrupt, reset interrupt, Address Translation (TLB interrupt) etc.

Course management tools

- Maintained course websites locally over e-learning portal and also in public domain over Git.
- Incorporated the use of Piazza, a web based tool for an efficient way to manage class Q&A.

Invited Tutorials

Tutorial on Parallel Programming and Performance Optimization on GPUs jointly with Prof. P Sadayappan, Ohio State University, at International Symposium on Computational Sciences (ISCS-2015), Sri Sathya Sai Institute Of Higher Learning, Prasanthi Nilayam, India, December 12-15, 2015; http://iscs-sssihl.github.io/2015/gpututorial.html

Invited Presentations

An Overview of the Global Arrays Toolkit, Five-days Technology Workshop on Heterogeneous Computing - Many Core/ Multi GPU - Performance of Algorithms, Application Kernels (HeMPa), 2011, at CMSD, UoHYD by C-DAC Pune & CMSD.
Programming for Performance on Cell BE processor, Performance Enhancement on Emerging Parallel Processing Platforms Workshop (PEEP), 2008, jointly organized by C-DAC and IUCAA.

Competitive Honors/Awards

Certificate of Appreciation from the Director of IBM Research India for contribution towards a O&G customer project, 2015.
Best Paper Award, PDGC conference 2014.
NVIDIA SRS Best GPU paper award, HiPC conference 2014.

Professional Service

Reviewer, Journal of Parallel and Distributed Computing.
Reviewer, IEEE Transactions on Cloud Computing.

Sub-reviewer, ESPM2 2016: Second International Workshop on Extreme Scale Programming Models and Middleware, SC 2016.
Sub-reviewer, The 7th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS16), SC 2016.
Reviewer (Computational Readiness aspect), DOE Leadership Computing Allocation (INCITE) proposals, 2016.
Technical Program Committee member, International Conference on High Performance Computing and Communications (HPCC), 2011, 2012.
TPC member, The 11^{th} IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA), 2013.
TPC member, The 12^{th} International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP), 2012.
TPC member, Student Research Symposium, IEEE International Conference on High Performance Computing (HiPC), 2012, 2014, 2015, 2016.
TPC member, International Symposium on Computational Sciences (ISCS-2015), 2015.
Reviewer, Computing, Springer Journal.

Computer Access Time Grants

PRAC Allocation for project VIPRA on NCSA Blue Waters, 650,000 SUs, PI: Prof. Ashok Srinivasan, FSU.
XSEDE Research Allocation: Scaling Communication Performance for Massively Parallel Applications, 800,000 SUs, PI: Prof. Ashok Srinivasan, Florida State University.
Teragrid Startup Allocation, PI: Prof. P. Sadayappan, Ohio State University.
XSEDE Education Allocation:Programming for Performance on multicore and many-core processor, PI: Prof. Ravi Mukkamala, Old Dominion University. (Have used this allocation effectively for teaching graduate level courses. Had access to the following supercomputers: TACC systems Ranger, Lonestar and Longhorn, SDSC systems Gordon and Trestles, PSC Blacklight.)