

Yogish Sabharwal

email: ysabharwal@in.ibm.com,

web: <http://www.cse.iitd.ernet.in/~yogish>

Phone: +91-9811301726

Date of Birth: February 21st, 1974



Executive Summary:

Ph. D. Computer Science and Engineering, IIT Delhi, 2006-07
“Approximation algorithms for proximity & clustering problems”

Master of Computer Applications Delhi University, 1997

Bachelors(Honours) in Mathematics Delhi University, 1994

Research Interests: High Performance Computing, Parallel Algorithms, Algorithms (Theory and applications), Computational Geometry

Teaching: Combinatorial Optimization, Delhi University (2010, '11, '12, '13, '14, '15)
Special Topics in Parallel Computing, IIT Delhi (2014, '15, '17, '18, '21)
Data Structures and Algorithms, IIT Delhi (co-taught 2020)
Introduction to Parallel Programming with OpenMP, NPTEL (2017, '18, '19, '21)
Introduction to Parallel Programming with OpenMP and MPI, NPTEL (2022)
Students (thesis) advised: 1 Ph.D., 10+ Masters

Publications/Patents: 85+ publications/15+ Patents

Technical Participation: TPC of IPDPS 2011, 2012, 2013, 2014, 2015
TPC of Supercomputing 2022
TPC of HiPC 2016, 2017, 2019, 2020, 2021
HiPC Program Vice Chair (Applications Track) 2018 and 2022
TPC of FSTTCS 2013
TPC of ICDCN 2012
TPC of WALCOM 2011
Organizing committee of ScaDL 2019, 2020, 2021

Experience:

24 years experience in High Performance Computing and Networking & Communication technologies.

High Performance Systems: Blue Gene Supercomputer, Cell multicore processor, Power processors, PowerPC, Intel

Languages: C, C++, Java, Python, Pytorch

Networking Technologies: IP, PPP, MATIP, ATM, IPOA, LANE, Frame Relay, X.25, 3G

Operating Systems: Solaris, Linux, VxWorks, Windows, GAME

IBM India Research Lab 2003-Present (STSM & Manager - High Performance Computing)

I currently manage the HPC Team at IBM Research India, with a primary focus on performance optimization and tuning of a wide variety of benchmarks and scientific applications such as the HPC Challenge benchmarks, numerical linear algebra methods, environmental science applications, cognitive computing frameworks and financial applications on massively parallel multicore architectures such as Blue Gene Supercomputers, Cell processor and

Power processor systems. The optimizations include designing new algorithms, data structures and other intricate techniques for distributed memory and multicore architectures. Projects include

- *Optimization of Cognitive Computing Platforms*

Training in deep learning networks is extremely compute intensive. GPUs and other accelerators are commonly used to speed up the training process. I lead the IBM India Research Lab team working on the optimization of deep learning frameworks such as Torch, Tensorflow and Theano on distributed memory systems comprising of GPU enabled nodes.

- *HPC in Environmental Science Applications*

Modeling and Simulation of environmental science applications such as weather modelling, flood modelling and climate modelling on high performance computing platforms (Blue Gene and Power7 systems). Technical co-lead of the UBD|IBM collaboration with University of Brunei. Leading the team enabling the modeling and simulation for weather, flood forecasting and climate research; this involves development and optimization of parallel numerical kernels underlying these physical models. Successfully enabled an operational weather model for Brunei on the Blue Gene system. Led the design of a parallel hydrological numerical model on the Blue Gene system. Also co-lead the relationship with UBD, helping in defining the research agenda for the collaboration and capacity building at the center.

- *Optimization of the Graph 500 benchmarks on the Blue Gene Supercomputer*

Designed parallel scaleable algorithms for the Graph 500 Breadth First Search and Single Source Shortest Path benchmarks. The BFS algorithm **won the Graph500 Competition in 2012 and 2013**. The SSSP algorithm won the **best paper award at IPDPS 2014**.

- *Optimization of the HPC Challenge benchmarks on the Blue Gene Supercomputer*

Our algorithm for the RandomAccess benchmark is the **winner of the HPCC Class I award at Supercomputing for 6 years in a row**. Associated publications detailing the algorithm include a **best paper award nomination at Supercomputing 2006**, a **best paper award at IPDPS 2009** and a **best paper award at HiPC 2010**. Also optimized Fast Fourier Transform benchmark and Transpose benchmark on Blue Gene.

HPC Challenge is a suite of benchmarks that measure the performance of various subsystems of supercomputers. The RandomAccess benchmark measures the performance of the memory subsystem and the Transpose benchmark measures the system interconnection bandwidth.

- *Optimization of Application Performance on Power Systems*

Developing algorithms for mapping and allocation of processors for computational tasks on Power systems with dragonfly type interconnection network in order to minimize network congestion.

- *Optimization of ADVC (a finite element analysis code) on the Blue Gene/L Supercomputer*

The optimizations allowed for scaling the code to larger number of nodes resulting in unparalleled performance. This work was the **finalist for the Gordon Bell award** at Supercomputing 2006.

- *Parallelization of Financial Engineering methods*

Parallelized the Least-Squares Monte Carlo (LSM) algorithm, proposed by Longstaff and Schwartz for pricing of American options – an important problem in quantitative finance.

- *Optimization of BLAS (Basic Linear Algebra Subroutines) on the Cell processor: Architected and designed the strategy for parallelization of the linear algebra sub-routines on the multi-core processor.*

Hughes Software Systems, India 2001 – 2003 (Sr Technical Leader- 3G Technologies)

Lead a team of 10 engineers responsible for requirements study, architecture, system design, integration, implementation and testing of the 3G-SGSN (Serving GPRS Support Node) – a network entity responsible for managing the mobility of a GPRS UE in a 3G mobile network.

Nortel Networks, India 1998 – 2001 (S/W Designer- Routers & Remote Access Servers)

Lead multiple project teams of 2-3 engineers involved in preparing the functional specifications, designing, implementation, testing and/or maintenance of various protocols, drivers and features for the Nortel Router and Remote Access Server products. The projects included (*not exhaustive*)

- Inverse Multiplexing for ATM (IMA) driver (splitting a cell stream over multiple physical links in a cyclical fashion and retrieving the original stream at the far end) on the Nortel routers
- Protocol Priority Queuing (PPQ) services (which allow for prioritization/re-ordering of packets) for ATM and X.25 protocols on the Nortel routers
- Enhancement of Differentiated Services Queuing Management & Scheduling for ATM services
- Support for Inter-working of VRRP (Virtual Router Redundancy Protocol) and Multi-protocol Over ATM for ATM LAN Emulation services

- RADIUS - a remote authentication and accounting service for dial in users on Nortel Remote Access servers.

Hughes Software Systems, India 1997 (Software Engineer- Internet Technologies)

Part of a 3-member team involved in preparing the functional specifications, designing, implementation and testing of a tool that monitors web-sites and keeps the administrator informed of the status and the load on the web-site.

Awards and Recognition:

- **IBM Outstanding Technical Achievement Award** for contributions to Model Compression (2020)
- **IBM Research Division Award** for contributions to entity analytics (2018)
- **IBM Outstanding Technical Achievement Award** for contributions to randomized algorithms in data mining (2017)
- **IBM Research Division Award** for contributions to approximation algorithms for facility location (2017)
- **IBM India Research Lab Award** for outstanding publication for the work on scalable subgraph matching (2015).
- **IBM India Research Lab Award** for outstanding publication for the work on scalable single source shortest path computation (2014).
- **Best Paper Award at IPDPS (2014).**
- **IBM Research Division Award** for contributions to Graph500 and HPCC benchmarks (2013)
- **IBM Outstanding Technical Achievement Award** for UBD-IBM Project (2012).
- **IBM Client OTA Award** for UBD-IBM Project (2011).
- **IBM Award** in recognition of contributions to Graph 500 benchmark (2011).
- **IBM Research Accomplishment** for contributions to UBD-IBM Collaboration (2011).
- **IBM Research Accomplishment** for contributions to KAUST Collaboration (2011).
- Our Randomaccess algorithm is the **winner of the HPCC award at Supercomputing 6 years in a row (2005-10).**
- **IBM Research Accomplishment** for contributions to Blue Gene/P (2010).
- **IBM Best People Manager Award** (2010).
- **Best Paper award at HiPC (2010).**
- **Best Paper award at IPDPS (2009).**
- **IBM Award** in recognition of receiving best paper at IPDPS (2009).
- Member of the team that received the **US National Medal of Technology and Innovation** (2009).
- **IBM First Patent Plateau Achievement award** (2008).
- **IBM Bravo Award** for contributions to Blue Gene/P Supercomputer (2008).
- **IBM Research Accomplishment** for contributions to Basic Linear Algebra Library on the Cell (2007).
- **Best Presentation Award (for Ph.D. work) at IRISS 2007 (6th Annual Indian Inter-Research Institute Student Seminar in Computer Science)**
- **IBM Bravo Award** for contributions to the HPC Challenge benchmarks (2007)
- **IBM Research Division Award** for contributions in Blue Gene/L Supercomputer (2006)
- **IBM Research Accomplishment** for contributions to Blue Gene/L (2005).
- **IBM Award for Impact on the Business and Significant Achievement** for contributions to HPC Challenge Benchmark on the Blue Gene/L Supercomputer (2005)
- **Hughes Team Award** for technical excellence in design and development of 3G SGSN at Hughes (2001)
- **Nortel Appreciation Award** for consistent performance in sustaining of router product of Nortel (2000)
- **Nortel Excellence Award** for leadership in design and development of MATIP at Nortel (1999)
- **Nortel Appreciation Award** for excellence in design and implementation of X.25 Priority Queuing at Nortel (1998)
- **Hughes Snap Award** for excellence in development of BazaarWatch at Hughes Software Systems (1997)

Selected Publications:

1. *Efficient Scaling of Dynamic Graph Neural Networks, with Venkatesan T. Chakaravarthy, Shivmaran S. Pandian, Saurabh Raje, Toyotaro Suzumura and Shashanka Ubaru, ACM/IEEE SC2021 Conference on High Performance Computing, Networking, Storage and Analysis (SC).*
2. *PoWER-BERT: Accelerating BERT Inference via Progressive Word-vector Elimination, with Saurabh Goyal, Anamitra Roy Choudhury, Saurabh Raje, Venkatesan T. Chakaravarthy and Ashish Verma, International Conference on Machine Learning (ICML), pp.3690-3699, 2020.*
3. *On optimizing distributed non-negative Tucker decomposition, with Venkatesan T. Chakaravarthy, Shivmaran S. Pandian and Saurabh Raje, International Conference on Supercomputing (ICS) 2019.*

4. *On Optimizing Distributed Tucker Decomposition for Sparse Tensors*, with Venkatesan T. Chakaravarthy, Jee W. Choi, Douglas J. Joseph, Prakash Murali, Shivmaran S. Pandian and Dheeraj Sreedhar, *International Conference on Supercomputing (ICS) 2018*.
5. *On Optimizing Distributed Tucker Decomposition for Dense Tensors*, with Venkatesan T. Chakaravarthy, Jee W. Choi, Douglas J. Joseph, Xing Liu, Prakash Murali and Dheeraj Sreedhar, *31st International Symposium on Parallel & Distributed Processing (IPDPS)*, 2017.
6. *Scaleable Single Source Shortest Path Algorithms for Massively Parallel Systems*, with Venkatesan Chakaravarthy, Fabio Checconi, Prakash Murali & Fabrizio Petrini, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2017.
7. *Subgraph Counting: Color Coding Beyond Trees*, with Venkatesan Chakaravarthy, Michael Kapralov, Prakash Murali, Fabrizio Petrini, Xinyue Que and Baruch Scheiber, *30th International Symposium on Parallel & Distributed Processing (IPDPS)*, 2016.
8. *Reusable Resource Scheduling via Colored Interval Covering*, with Venkatesan Chakaravarthy, Sreyash Kenkre, Sakib A. Mondal and Vinayaka Pandit, *30th International Symposium on Parallel & Distributed Processing (IPDPS)*, 2016.
9. *Facility Location with Knapsack or Matroid Constraints*, with Ravishankar Krishnawamy, Amit Kumar, Vishwanath Nagarajan and Barna Saha, *Mathematics of Operations Research*, Volume 40, 2015.
10. *Scaleable Single Source Shortest Path Algorithms for Massively Parallel Systems*, with Venkatesan Chakaravarthy, Fabio Checconi & Fabrizio Petrini, *28th International Symposium on Parallel & Distributed Processing (IPDPS)*, 2014.

Best Paper Award.

11. *Distributed Algorithms for Scheduling on Line and Tree Networks with Non-uniform Bandwidths*, with Venkatesan Chakaravarthy, Anamitra R. Choudhury and Sambuddha Roy, *27th International Symposium on Parallel and Distributed Processing (IPDPS)*, pp. 973-984, 2013
12. *Looking under the hood of the IBM Blue Gene/Q network*, with Dong Chen, Noel Easley, Philip Heidelberger, Sameer Kumar, Amith Mamidala, Fabrizio Petrini, Robert Senger, Yutaka Sugawara, Robert Walkup, Burkhard Steinmacher-Burow, Anamitra R. Choudhury, Swati Singhal and Jeffrey J. Parker, *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, pp. 69, 2012
13. *A divide and conquer strategy for scaling weather simulations with multiple regions of interest*, with Preeti Malakar, Thomas George, Sameer Kumar, Rashmi Mittal, Vijay Natarajan, Vaibhav Saxena and Sathish Vadhiyar, *ACM/IEEE SC2012 Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, pp. 37

Best Student Paper Finalist

14. *Breaking the speed and scalability barriers for graph exploration on distributed-memory machines*, with Fabio Checconi, Fabrizio Petrini, Jeremiah Willcock, Andrew Lumsdaine and Anamitra R. Choudhury, *ACM/IEEE SC2012 Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, pp. 13
15. *Distributed algorithms for scheduling on line and tree networks*, with Venkatesan Chakaravarthy and Sambuddha Roy, *31st ACM Symposium on Principles of Distributed Computing (PODC)*, 2012
16. *Collective Algorithms for sub-communicators*, with Anshul Mittal, Nikhil Jain, Thomas George and Sameer Kumar, *International Conference on Supercomputing (ICS) 2012*.
17. *The Matroid Median Problem*, with Ravishankar Krishnaswamy, Amit Kumar, V. Nagarajan and Barna Saha, *Symposium on Discrete Algorithms (SODA) 2011*.
18. *Minimum Cost Resource Allocation for meeting job requirements*, with Venkatesan T. Chakaravarthy, Amit Kumar, Gyana Parija and Sambuddha Roy, *International Parallel and Distributed Processing Symposium (IPDPS) 2011*.
19. *Improved Algorithms for the Distributed Trigger Counting Problem*, with Venkatesan T. Chakaravarthy and Anamitra R. Choudhury, *International Parallel and Distributed Processing Symposium (IPDPS) 2011*.
20. *Efficient Algorithms for Global Snapshots in Large Distributed Systems*, with Rahul Garg and Vijay K. Garg, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2010.
21. *Linear time approximate clustering in any dimension*, with Amit Kumar & Sandeep Sen, *Journal of the ACM (JACM) 2010*.
22. *Optimal Bucket Algorithms for large MPI Collectives on Torus Interconnects*, with Nikhil Jain, *International Conference on Supercomputing (ICS) 2010*.
23. *Varying Bandwidth Resource Allocation Problem with Bag Constraints*, with Venkatesan T. Chakaravarthy, Vinayaka Pandit and Deva P. Seetharam, *International Parallel and Distributed Processing Symposium (IPDPS) 2010*.
24. *Approximating Decision Trees with Multiway Branches*, with Venkatesan T. Chakaravarthy, Vinayaka Pandit and Sambuddha Roy, *International Colloquium on Automata, Languages and Programming (ICALP) 2009*.
25. *Analysis of sampling techniques for association rule mining*, with Vinayaka Pandit and Venkatesan T. Chakaravarthy, *International Conference on Database Theory (ICDT) 2009*.

26. *HPCC RandomAccess Benchmark for Next Generation Supercomputers*, with Vikas Aggarwal, Rahul Garg and Philip Heidelberger, *International Parallel and Distributed Processing Symposium (IPDPS) 2009*.
Best Paper Award
27. *Performance Analysis and Optimization of All-to-all communication on the Blue Gene/L Supercomputer*, with Sameer Kumar, Rahul Garg and Philip Heidelberger, *International Conference on Parallel Processing (ICPP) 2008*.
28. *Optimizations in Financial Engineering : The Least-Squares Monte Carlo method of Longstaff and Schwartz*, with Anamitra Choudhury, Alan King and Sunil Kumar, *Intn. Parallel and Distributed Computing Symposium (IPDPS) 2008*.
29. *Software Routing and Aggregation of Messages to Optimize the Performance of the HPCC Randomaccess Benchmark*, with Rahul Garg, *ACM/IEEE SC2006 Conference on High Performance Computing, Networking, Storage and Analysis (SC) 2006*.
Best Paper Award Finalist.
30. *Large Scale Drop Impact Analysis of Mobile Phone Using ADVC on Blue Gene/L*, with H. Akiba, T. Ohyama, Y. Shibata, K. Yuyama, Y. Katai, R. Takeuchi, T. Hoshino, S. Yoshimura, H. Noguchi, M. Gupta, J. Gunnels, V. Austel, R. Garg, S. Kato, T. Kawakami, S. Todokoro and J. Ikeda, *ACM/IEEE SC2006 Conference on High Performance Computing, Networking, Storage and Analysis (SC) 2006*.
Gordon Bell Finalist.
31. *Scalable Algorithms for global snapshots in distributed systems*, with Rahul Garg and Vijay K. Garg, *International Conference on Supercomputing (ICS) 2006*.
32. *Linear Time Algorithms for Clustering Problems in any Dimensions*, with Amit Kumar and Sandeep Sen, *International Colloquium on Automata, Languages and Programming (ICALP) 2005*.
33. *A Simple linear time $(1+ \epsilon)$ -approximation algorithm for k-means clustering in any dimensions*, with Amit Kumar and Sandeep Sen, *Foundations of Computer Science (FOCS) 2004*.