Bibliography

[1] Gcc. The gnu compiler collection. `http://gcc.gnu.org`.

[2] Robert Hundt, Sandya Mannarswamy, and Dhruva Chakrabarti. Practical structure layout optimization and advice. In *Proceedings of the International Symposium on Code Generation and Optimization*, CGO '06, page 233–244, USA, 2006. IEEE Computer Society.

[3] Shruti Jadon, Pradyuman Kavedi Kannan, Urmil Kalaria, K. R. Varsha, Karthik Gupta, and Prasad B. Honnavalli. A comprehensive study of load balancing approaches in real-time multi-core systems for mixed real-time tasks. *IEEE Access*, 12:53373–53395, 2024.

[4] Matheus da Silva Serpa Philippe Olivier Alexandre Navaux, Arthur Francisco Lorenzon. Challenges in high-performance computing. *Journal of the Brazilian Computer Society*, 29(1):51–62, 2023.

[5] Tiago Augusto Engel, Andrea Schwetner Charao, Manuele kirsch Pinbeiro, and Luiz-Angelo Steffenel. Performance improvement of data mining in weka through multicore GPU acceleration. *Ambient Intell Human comput*, 6:377–390, 2015.

[6] Wei Jiang and Gagan Agrawal. MATE-CG: A map reduce-like framework for accelerating data-intensive computations on heterogeneous clusters. *IEEE 26th. International Parallel Distributed Processing Symposium (IPDPS)*, pages 644–655, 2012.

[7] Peter Tröger. The multi-core era - trends and challenges. *ArXiv*, abs/0810.5439, 2008.

[8] Xiaohong Qiu, Geoffrey fox, Huapeng Yuan, George Chrysanthakopoulos, and

Henrik Nielsen. Parallel data mining on multicore clusters. *Seventh International Conference on Grid and Cooperative Computing*, pages 41–49, 2008.

[9] Jack Dongarra Ed. The promise and perils of the coming multicore revolution and its impact. *CTWatch Quarterly*, 3(1):1–35, 2007.

[10] EE Schadt, MD Linderman, J Sorenson, L Lee, and GP Nolan. Computational solutions to large-scale data management and analysis. *Nature Rev Genetics*, 11(9):647–657, 2010.

[11] Jaewon Kwon, Yongju Lee, Hongju Kal, Minjae Kim, Youngsok Kim, and Won Woo Ro. Mccore: A holistic management of high-performance heterogeneous multicores. In *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO '23, page 1044–1058. Association for Computing Machinery, 2023.

[12] Khaleghzadeh Hamidreza, Manumachu Ravindranath Reddy, and Lastovetsky Alexey. A novel data-partitioning algorithm for performance optimization of data-parallel applications on heterogeneous HPC platforms. *IEEE Transactions on Parallel and Distributed Systems*, 29(10):2176–2190, 2018.

[13] Chen Xiaomeng, Zhang Hui Bai, Hanli Yang, Chunming, Zhao Xujian, and Li Bo. Runtime prediction of high-performance computing jobs based on ensemble learning. *Association for Computing Machinery*, 29(7):56–62, 2020.

[14] F. Yao, Y. Yao, L. Xing, and H. Chen. An intelligent scheduling algorithm for complex manufacturing system simulation with frequent synchronizations. *Clouds[J] Memtic Computing*, 11:357–370, 2019.

[15] K. Gurleen, B. Anju, and C. Inderveer. An intelligent regressive ensemble approach for predicting resource usage. *Journal of Parallel and Distributed Computing*, 123:1–12, 2019.

[16] Xinnian Zheng, Haris Vikalo, Shuang Song, Lizy K. John, and Andreas Gerstlauer. Sampling-based binary-level cross-platform performance estimation. *In Proceedings of the Conference on Design, Automation and Test in Europe. European Design and Automation Association*, pages 1713–1718, 2017.

[17] Junichi Sakamoto, Daisuke Fujimoto, Riku Anzai, Naoki Yoshida, and Tsutomu Matsumoto. High-throughput bilinear pairing processor for server-side FPGA applications. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, pages 1–14, 2024.

[18] Sathwika Bavikadi, Purab Ranjan Sutradhar, Amlan Ganguly, and Sai Manoj Pudukotai Dinakarrao. Reconfigurable processing-in-memory architecture for data intensive applications. In *2024 37th International Conference on VLSI Design and 2024 23rd International Conference on Embedded Systems (VLSID)*, pages 222–227, 2024.

[19] Shashikiran Venkatesha and Ranjani Parthasarathi. Survey on redundancy based-fault tolerance methods for processors and hardware accelerators - trends in quantum computing, heterogeneous systems and reliability. *ACM Comput. Surv.*, 2024.

[20] Aashish Phansalkar, Ajay Joshi, , and Lizy K. John. Analysis of redundancy and application balance in the SPEC CPU2006 benchmark suite. *In Proceedings of the 34th Annual International Symposium on Computer Architecture. ACM*, pages 412–423, 2007.

[21] Bao Huynh and Bay Vo. An efficient method for mining erasable itemsets using multicore processor platform. *Complexity*, page 9, 2018.

[22] Nathan Beckmann and Daniel Sanchez. Modeling cache performance beyond LRU. *22nd International Symposium on High Performance Computer Architecture (HPCA),*, pages 1–12, February 2016.

[23] Xiaodong Wang and Jose F. Mart ́ ́ınez. Xchange: A market-based approach to scalable dynamic multiresource allocation in multicore architectures. *In Proceedings of the International Symposium on High Performance Computer Architecture (HPCA'15)*, PP:113–125, 2015.

[24] Ching-Hwa Cheng. Design example of useful memory latency for developing a hazard preventive pipeline high-performance embedded-microprocessor. *Advanced VLSI Design Methodologies for Emerging Industrial Multimedia and Communication Applications, Hindawi*, pages 1–10, July 2013.

[25] J. Wu and B. Hong. Collocating CPU-only jobs with GPU-assisted jobs on gpu-assisted HPC. *13th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing*, pages 19–24, 2013.

[26] J. Zhang, J. Zhai, W. Chen, and W. Zheng. Process mapping for MPI collective communications. *European Conference on Parallel Processing. Springer*, pages 81–92, 2009.

[27] Dimitrios Ziakas, Allen Baum, Robert A. Maddox, and Robert J. Safranek. Intel® quickpath interconnect architectural features supporting scalable system architectures. *18th IEEE Symposium on High Performance Interconnects*, pages 1–6, 2010.

[28] Mulya Agung, Muhammad Alfian Amrizal, Kazuhiko Komatsu, Ryusuke Egawa, and Hiroyuki Takizawa. A memory congestion-aware MPI process placement for modern NUMA systems. *IEEE 24th International Conference on High Performance Computing (HiPC)*, pages 152–161, 2017.

[29] Mulya Agung, Muhammad Alfian Amrizal, Ryusuke Egawa, and Hiroyuki Takizawa. An automatic MPI process mapping method considering locality and memory congestion on NUMA systems. *IEEE 13th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoC)*, pages 17–24, 2019.

[30] Franck Cappello and Daniel Etiemble. Mpi versus MPI+OpenMP on the IBM SP for the NAS benchmarks. *Proceedings of the IEEE/ACM SC2000 Conference (SC'00)*, pages 1–12, 2000.

[31] D.H. Bailey, E. Barszcz, J.T. Barton, D.S. Browning, R.L. Carter, L. Dagum, R.A. Fatoohi, P.O. Frederickson, T.A. Lasinski, R.S. Schreiber, H.D. Simon, V. Venkatakrishnan, and S.K. Weeratunga. The NAS parallel benchmarks. *International Journal of High Performance Computing Applications*, pages 63–74, September 1991.

[32] Fabien Gaud, Baptiste Lepers, Justin Funston, Mohammad Dashti, Alexandra Fedorova, Renaud Lachaize Vivien Quema, and Mark Roth. Optimizing NUMA systems applications with carrefour. *Communications of the ACM*, 58(12):59–66, December 2015.

[33] Tan Li a, Yufei Ren, Dantong Yu, and Shudong Jin. Analysis of NUMA effects in modern multicore systems for the design of high-performance data transfer applications. *Future Generation Computer Systems*, 74(5):41–50, 2017.

[34] H.S. Stone, J. Turek, and J.L. Wolf. Optimal partitioning of cache memory. *IEEE Transactions on Computers*, 41(9):1054–1068, 1992.

[35] Chenjie Yu and Peter Petrov. Off-chip memory bandwidth minimization through cache partitioning for multi-core platforms. In *Proceedings of the 47th Design Automation Conference*, DAC '10, page 132–137, New York, NY, USA, 2010. Association for Computing Machinery.

[36] Mahmut Kandemir, Taylan Yemliha, and Emre Kultursay. A helper thread based dynamic cache partitioning scheme for multithreaded applications. In *2011 48th ACM/EDAC/IEEE Design Automation Conference (DAC)*, pages 954–959, 2011.

[37] Ying Ye, Richard West, Zhuoqun Cheng, and Ye Li. Coloris: A dynamic cache partitioning system using page coloring. In *2014 23rd International Conference on Parallel Architecture and Compilation Techniques (PACT)*, pages 381–392, 2014.

[38] Henry Cook, Miquel Moreto, Sarah Bird, Khanh Dao, David A. Patterson, and Krste Asanovic. A hardware evaluation of cache partitioning to improve utilization and energy-efficiency while preserving responsiveness. *SIGARCH Comput. Archit. News*, 41(3):308–319, 2013.

[39] Jiang Lin, Qingda Lu, Xiaoning Ding, Zhao Zhang, Xiaodong Zhang, and P. Sadayappan. Gaining insights into multicore cache partitioning: Bridging the gap between simulation and real systems. In *2008 IEEE 14th International Symposium on High Performance Computer Architecture*, pages 367–378, 2008.

[40] Dimitris Kaseridis, Muhammad Faisal Iqbal, and Lizy Kurian John. Cache friendliness-aware managementof shared last-level caches for high performance multi-core systems. *IEEE Transactions on Computers*, 63(4):874–887, 2014.

[41] Yuejian Xie and Gabriel H. Loh. Pipp: promotion/insertion pseudo-partitioning of multi-core shared caches. *SIGARCH Comput. Archit. News*, 37(3):174–183, 2009.

[42] Nauman Rafique, Won-Taek Lim, and Mithuna Thottethodi. Architectural support for operating system-driven cmp cache management. In *2006 International*

Conference on Parallel Architectures and Compilation Techniques (PACT), pages 2–12, 2006.

[43] Miao Zhou, Yu Du, Bruce Childers, Daniel Mosse, and Rami Melhem. Symmetry-agnostic coordinated management of the memory hierarchy in multicore systems. *ACM Trans. Archit. Code Optim.*, 12(4), jan 2016.

[44] Heechul Yun and Prathap Kumar Valsan. Evaluating the isolation effect of cache partitioning on cots multicore platforms. In *Proceedings of the 11th Annual Workshop on Operating Systems Platforms for Embedded Real-Time Applications*, (OSPERT'15), New York, NY, USA, 2015. Association for Computing Machinery.

[45] Harshad Kasture and Daniel Sanchez. Ubik: efficient cache sharing with strict qos for latency-critical workloads. *SIGPLAN Not.*, 49(4):729–742, feb 2014.

[46] Daniel Sanchez and Christos Kozyrakis. Vantage: Scalable and efficient fine-grain cache partitioning. In *2011 38th Annual International Symposium on Computer Architecture (ISCA)*, pages 57–68, 2011.

[47] Hong Jiang Dongyuan Zhan and Sharad C. Seth. Clu: Co-optimizing locality and utility in threadaware capacity management for shared last level caches. *IEEE Trans. Comput.*, 63(7):1656–1667, 2014.

[48] Sai Prashanth Muralidhara, Mahmut Kandemir, and Padma Raghavan. Intra-application cache partitioning. In *2010 IEEE International Symposium on Parallel Distributed Processing (IPDPS)*, pages 1–12, 2010.

[49] Lei Liu, Yong Li, Zehan Cui, Yungang Bao, Mingyu Chen, and Chengyong Wu. Going vertical in memory management: handling multiplicity by multi-policy. *SIGARCH Comput. Archit. News*, 42(3):169–180, jun 2014.

[50] Ramazan Bitirgen, Engin Ipek, and Jose F. Martinez. Coordinated management of multiple interacting resources in chip multiprocessors: A machine learning approach. In *2008 41st IEEE/ACM International Symposium on Microarchitecture*, pages 318–329, 2008.

[51] Taha Abdelazziz Rahmani, Ghalem Belalem, Sidi Ahmed Mahmoudi, and Omar Rafik Merad-Boudia. Machine learning-driven energy-efficient load balancing for real-time heterogeneous systems. *Cluster Computing*, 2024.

[52] J. Li, G. Michelogiannakis, B. Cook, D. Cooray, and Y. Chen. Analyzing resource utilization in an HPC system: A case study of NERSC's perlmutter. In *High Performance Computing. ISC High Performance*, volume 13948 of *LNCS*, 2023.

[53] Z. Majo and T. R. Gross. Memory management in NUMA multicore systems: trapped between cache contention and interconnect overhead. *International Symposium on Memory Management, ACM*, pages 11–20, 2011.

[54] F. Pellegrini and J. Roman. Scotch: A software package for static mapping by dual recursive bipartitioning of process and architecture graphs. *International Conference on High-Performance Computing and Networking (HPCN 1996). Springer*, pages 493–498, 1996.

[55] David Lo, Liqun Cheng, Rama Govindaraju, Parthasarathy Ranganathan, and Christos Kozyrakis. Heracles: Improving resource efficiency at scale. In *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, pages 450–462, 2015.

[56] William Hasenplaugh, Pritpal S. Ahuja, Aamer Jaleel, Simon Steely Jr., and Joel Emer. The gradient-based cache partitioning algorithm. *ACM Trans. Archit. Code Optim.*, 8(4), jan 2012.

[57] Vineeth Mekkat, Anup Holey, Pen-Chung Yew, and Antonia Zhai. Managing shared last-level cache in a heterogeneous multicore processor. In *Proceedings of the 22nd International Conference on Parallel Architectures and Compilation Techniques*, pages 225–234, 2013.

[58] Emmanuel Jeannot, Guillaume Mercier, and Francois Tessier. Process placement in multicore clusters: Algorithmic issues and practical techniques. *IEEE Transactions on Parallel and distributed Systems*, 25(4):993–1002, April 2014.

[59] Stijn Eyerman, Kenneth Hoste, and Lieven Eeckhout. Mechanistic-empirical processor performance modeling for constructing CPI stacks on real hardware. *In Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'11), IEEE*, pages 216–226, 2011.

[60] Nguyen X, Hoang T, Nguyen H, Inoue K, and Pham C. An FPGA-based hardware accelerator for energy-efficient bitmap index creation. *IEEE Access*, 6:16046–16059, 2018.

[61] W Fang, KK Lau, M Lu, X Xiao, CK lam, PY Yang, B he, Q Luo, and PVSK Yang. Parallel datamining on graphics processors: Technical report. *Department of Computer Science and Engineering, Hong Kong University of Science and Technology*, pages 1–10, 2008.

[62] MS Perez, A. sanchez, P. Herrero, V. Robles, and JM Pena. Adapting the weka data mining toolkit to a grid based environment. *Lecture Note in Computer Science, Springer, Lodz, Polonia*, 3528:492–497, 2005.

[63] Benjamin C. Accurate and efficient regression modeling for microarchitectural performance and power prediction. *In Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems. ACM*, pages 185–194, 2006.

[64] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.

[65] Karan Singh, Engin İpek, Sally A. McKee, Bronis R. de Supinski, Martin Schulz, and Rich Caruana. Predicting parallel application performance via machine learning approaches. *Concurrency and Computation: Practice and Experience*, 19(17):2219–2235, 2007.

[66] Ang Li, Xuanran Zong, Srikanth Kandula, Xiaowei Yang, and Ming Zhang. CloudProphet: Towards application performance prediction in cloud. *In Proceedings of the ACM SIGCOMM 2011 conference (SIGCOMM'11)*, pages 426–427, 2011.

[67] Ren Wu, Bin Zhang, and Meichun Hsu. GPU-accelerated large scale analytics. *Technical Report HPL, HP Labs*, 38:2–11, 2009.

[68] Sameh Sharkawi, Don Desota, Raj Panda, Rajeev Indukuru, Stephen Stevens, Valerie Taylor, and Xingfu Wu. Performance projection of HPC applications using SPEC CFP2006 benchmarks. *In Proceedings of the IEEE International Symposium on Parallel and Distributed Processing, IEEE*, pages 1–12, 2009.

[69] Allan Hartstein and Thomas R. Puzak. The optimum pipeline depth for a microprocessor. *In Proceedings of the 29th Annual International Symposium on Computer Architecture (ISCA'02). IEEE Computer Society*, pages 7–13, 2002.

[70] Newsha Ardalani, Clint Lestourgeon, Karthikeyan Sankaralingam, and Xiaojin Zhu. Cross-architecture performance prediction (XAPP) using CPU code to predict GPU performance. *In Proceedings of the 48th International Symposium on Microarchitecture. ACM*, pages 725–737, 2015.

[71] Xinnian Zheng, Lizy K. John, and Andreas Gerstlauer. Accurate phase-level cross-platform power and performance estimation. *In Proceedings of the 2016 53rd ACM/EDAC/IEEE Design Automation Conference. IEEE*, pages 1–6, 2016.

[72] Kaiyang Liu, Jingrong Wang, Zhiming Huang, and Jianping Pan. Sampling-based multi-job placement for heterogeneous deep learning clusters. *IEEE Transactions on Parallel and Distributed Systems*, 35(6):1029–1043, 2024.

[73] Aashish Phansalkar, Ajay Joshi, Lieven Eeckhout, and Lizy Kurian John. Measuring program similarity: Experiments with SPEC CPU benchmark suites. *In IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'05), IEEE*, pages 10–20, 2005.

[74] Daniel Shelepov, Juan Carlos Saez Alcaide, Stacey Jeffery, Alexandra Fedorova, Nestor Perez, Zhi Feng Huang, Sergey Blagodurov, and Viren Kumar. HASS: A scheduler for heterogeneous multicore systems. *ACM SIGOPS Operating Systems Review*, 43:66–75, 2009.

[75] D. Shepelow and A. Fedorova. Scheduling on heterogeneous multicore processors using architectural signatures. *In Proceedings of the WIOSCA Workshop of the 35th Annual International Symposium on Computer Architecture*, pages 1–9, 2008.

[76] Christina Delimitrou and Christos Kozyrakis. Quasar: Resource-efficient and qos-aware cluster management. *In Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems. ACM*, pages 127–144, 2014.

[77] Jiawei Wang, Yutao Liu, Ming Fu, Hermann Härtig, and Haibo Chen. Brief announcement: Work stealing through partial asynchronous delegation. *36th ACM Symposium on Parallelism in Algorithms*, page 281–283, 2024.

[78] Shruti Jadon, Pradyuman Kavedi Kannan, Urmil Kalaria, K. R. Varsha, Karthik Gupta, and Prasad B. Honnavalli. A comprehensive study of load balancing

approaches in real-time multi-core systems for mixed real-time tasks. *IEEE Access*, 12:53373–53395, 2024.

[79] W. Li, H. Cheng, Z. Lu, Y. Lu, and W. Liu. HASpMV: Heterogeneity-aware sparse matrix-vector multiplication on modern asymmetric multicore processors. In *2023 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 209–220, 2023.

[80] Ravi Reddy Manumachu and Alexey Lastovetsky. Parallel data partitioning algorithms for optimization of data-parallel applications on modern extreme-scale multicore platforms for performance and energy. *IEEE Access*, 6:69075–69106, 2018.

[81] Hamza Omar, Halit Dogan, Brian Kahne, and Omer Khan. Multicore resource isolation for deterministic, resilient and secure concurrent execution of safety-critical applications. *IEEE Computer Architecture Letters*, 17(2):230–234, 2018.

[82] R. Shaw, E. Howley, and E. Barrett. An energy efficient anti-correlated virtual machine placement algorithm using resource usage predictions. *Simulation Modelling Practice and Theory*, 93:322–342, 2019.

[83] Neill Richard, Drebes Andi, and Pop Antoniu. Automated analysis of task-parallel execution behavior via artificial neural networks. *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 647–656, 2018.

[84] Dhouha Mejri, Mohamed Limam, and Claus Weihs. A new dynamic weighted majority control chart for data streams. *Soft Computing*, 22, 01 2018.

[85] Carlos Villavieja, Isaac Gelado, Alex Ramirez, and Nacho Navarro. Memory management on chip-multiprocessors with on-chip memories. In *Proc. Workshop on the Interaction between Operating Systems and Computer Architecture*, 2008.

[86] Sam Van Den Steen, Stijn Eyerman, Moncef Mechri Sander De Pestel, Trevor E. Carlson, David Black-Schaffer, Erik Hagersten, and Lieven Eeckhout. Analytical processor performance and power modeling using micro-architecture independent characteristics. *In IEEE Transaction on Computers*, 65(12):3537–3551, December 2016.

[87] J. Zhu, J. Wang, Zhang, and Y. et al. Virtual machine migration method based on load cognition. *Soft Computing*, 23:9439–9448, 2019.

[88] N.H. Shahapure and P. Jayarekha. Virtual machine migration based load balancing for resource management and scalability in cloud environment. *Int. j. inf. tecnol.*, 12:1331–1342, 2020.

[89] Ivanoe De Falco, Eryk Laskowski, Richard Olejnik, Umberto Scafuri, Ernesto Tarantino, and Marek Tudruj. Extremal optimization applied to load balancing in execution of distributed programs. *Applied Soft Computing*, pages 501–513, 2015.

[90] N. Hu, C. Yan, Z. Chen, and Y. Lu. FinD: Fine-grained dynamic task scheduling with lightweight threads on many-core processors. In *2023 IEEE Intl Conf on Parallel and amp; Distributed Processing with Applications, Big Data and amp; Cloud Computing, Sustainable Computing and amp; Communications, Social Computing and amp; Networking (ISPA/BDCloud/SocialCom/SustainCom)*, pages 270–278, 2023.

[91] Ahmed Eleliemy and Florina M. Ciorba. A distributed chunk calculation approach for self-scheduling of parallel applications on distributed-memory systems. *Journal of Computational Science*, 51:101–284, 2021.

[92] F. Yao, Z. Zhang, Zeyu. Ji, Bin Liu., and Haoyuan Gao. LBB: load-balanced batching for efficient distributed learning on heterogeneous GPU cluster. *The Journal of Supercomputing*, 80:12247–12272, 2024.

[93] Kaige Yan, Yanshuang Song, Tao Liu, Jingweijia Tan, Xiaohui Wei, and Xin Fu. Hsas: Efficient task scheduling for large scale heterogeneous systolic array accelerator cluster. *Future Generation Computer Systems*, 154:440–450, 2024.

[94] Yu Wang, Victor Lee, Gu-Yeon Wei, and David Brooks. Predicting new workload or CPU performance by analyzing public datasets. *ACM Trans. Archit. Code Optim.*, 15(4), 2019.

[95] G. Bosilca, C. Foyer, E. Jeannot, G. Mercier, and G. Papaure. Online dynamic monitoring of mpi communications. *European Conference on Parallel Processing, Springer*, pages 49–62, 2017.

[96] Hongyan Liu, Fei Lei, Chen Tong, Chunji Cui, and Li Wu. Visual smoke detection based on ensemble deep CNNs. *Displays*, 69(3):102020, 2021.

[97] N. Mirghafori, M. Jacoby, and D. Patterson. Truth in spec benchmarks. *Computer Architecture News*, 23(5):34–42, 2020.

[98] Nelson M, Sorenson Z, Myre JM, Sawin J, and Chiu D. GPU acceleration of range queries over large data sets. *Proc. of the 6th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies.Association for Computing Machinery, New York.*, pages 11–20, 2019.

[99] J.P.Marcelo and A.D.Carreara. New model-based methods and algorithms for performance and energy optimization of data parallel applications on homogeneous multicores clusters. *IEEE Transactions on Parallel and distributed Systems*, 28(4):1119–1133, 2017.

[100] Ru Huang, Lei Ma, Jianhua He, and Xiaoli Chu. T-GAN: A deep learning framework for prediction of temporal complex networks with adaptive graph convolution and attention mechanism. *Displays*, 68:102023, 2021.

[101] Georg Hager and Gerhard Wellein. Introduction to high performance computing for scientists and engineers. *CRC Press*, pages 1–356, July 2010.

[102] A. Degomme, A. Legrand, G. Markomanolis, M. Quinson, M. Stillwell, and F. Suter. Simulating MPI applications: The SMPI approach. *IEEE Transactions on Parallel and Distributed Systems*, 28(8):2387–2400, 2017.

[103] Xuan-Thuan Nguyen, Trong-Thuc Hoang, Hong-Thu Nguyen, Katsumi Inoue, and Cong-Kha Pham. An FPGA-based hardware accelerator for energy-efficient bitmap index creation. *IEEE Access*, 6:16046–16059, 2018.

[104] Xi E. Chen and Tor M. Aamodt. Hybrid analytical modeling of pending cache hits, data prefetching, and MSHRs. *ACM Transactions on Architecture and Code Optimization (TACO)*, 8(3):1–10, 2011.

[105] Danny Segev. An approximate dynamic-programming approach to the joint replenishment problem. *Mathematics of Operations Research*, 39:432–444, 2014.

[106] G.M. Weber and N. N. MATOG: Array layout auto-tuning for CUDA. *ACM Trans Archit Code Optim (TACO)*, 14(3):28, 2017.

[107] M.S. Engin Ipek, Bronis R De Supinski, and S.A. McKee. An approach to performance prediction for parallel applications. pages 196–205, 2005.

[108] Benjamin C. Lee and David M. Brooks. Accurate and efficient regression modeling for micro-architectural performance and power prediction. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 185–194. ACM, 2006.

Appendix A

List of Publications

Journals

1. Navin Mani Upadhyay, Ravi Shankar Singh, Shri Prakash Dwivedi. "Prediction of multicore CPU performance through parallel data mining on public datasets", **Displays, Elsevier**, **[SCI]**, **[Impact Factor: 4.3]**, DOI: 10.1016/j.displa.2021.102112, **(Published.)**

2. Navin Mani Upadhyay, Ravi Shankar Singh. "An effective scheme for memory congestion reduction in multi-core environment", **Journal of King Saud University - Computer and Information Sciences, Elsevier**, **[SCIE]**, **[Impact Factor: 6.9]**, DOI: 10.1016/j.jksuci.2020.05.011, **(Published.)**

3. Navin Mani Upadhyay, Ravi Shankar Singh, Shri Prakash Dwivedi. "Dependency Prediction of Long-Time Resource Uses in HPC Environment", **IEEE Access**, **[SCI]**, **[Impact Factor: 3.9]**, DOI: 10.1109/ACCESS.2023.3341046, **(Published.)**

Conference Proceedings

1. Navin Mani Upadhyay, Ravi Shankar Singh. "Performance Evaluation of Classification Algorithm in Weka using Parallel Performance Profiling and Computing Technique", **Fifth International Conference on Parallel, Distributed**

and Grid Computing (PDGC),**[2018 Fifth International Conference]** DOI: 10.1109/PDGC.2018.8745940, **(Published.)**