

JAROSLAW ZOLA

University at Buffalo, SUNY
Department of Computer Science and Engineering
Department of Biomedical Informatics
335 Davis Hall, Buffalo, NY, 14260-2500

Phone: (716) 645-3187
jzola@buffalo.edu
<http://www.jzola.org/>
<http://www.score-group.org/>

Research Interests

Addressing challenging problems, and accelerating scientific discovery through:

- Development of scalable algorithms and software for big data in life sciences
- Novel bioinformatics methods in the context of next-gen sequencing technologies
- High performance, parallel and distributed computing on emerging architectures
- Data driven large-scale scientific computing

Education

- 2002 – 2005:** Ph.D. in Computer Science (these en cotutelle)
Grenoble Institute of Technology, Grenoble, France
- 2001 – 2005:** Ph.D. in Computer Science
Czestochowa University of Technology, Czestochowa, Poland
- 2000:** Study abroad (Computer Science, Erasmus Program)
Grenoble Institute of Technology, Grenoble, France
- 1996 – 2001:** M.Sc. in Computer Science
Czestochowa University of Technology, Czestochowa, Poland
Specialization: Software Engineering and Information Systems

Professional Experience

- Aug 2014 – present:** Assistant Professor
Department of Computer Science and Engineering
Department of Biomedical Informatics
University at Buffalo, Buffalo, NY

- Oct 2012 – Aug 2014:** Associate Research Professor
Department of Electrical and Computer Engineering
Rutgers Discovery Informatics Institute
Rutgers University, Piscataway, NJ
- Aug 2010 – Oct 2012:** Research Assistant Professor
Department of Electrical and Computer Engineering
Iowa State University, Ames, IA
- Jun 2006 – Aug 2010:** Postdoctoral Research Associate
Department of Electrical and Computer Engineering
Iowa State University, Ames, IA
- Nov 2001 – Jun 2006:** Research Assistant
Institute of Computer and Information Sciences
Czestochowa University of Technology, Czestochowa, Poland

Honors and Awards

- 2012:** IEEE Senior Member
- 2007:** Best paper award nomination, ACM/IEEE Supercomputing Conference
- 2002:** Ph.D. scholarship award of the French government

Research Proposals and Grants

Current

Grants

1. PI, UB IMPACT, “OneBioStore: Distributed Smart Storage and Scalable Algorithms for Collaborative Biomedical Discovery.” Duration: 12 mos. Starting date: 09/01/2016. Amount: \$18,560.
2. PI, NSF/CCF, “Student Travel Support: ACM International Workshop on Big Data in Life Sciences, Seattle, WA, October 2, 2016.” Duration: 12mos. Starting date: 07/01/2016. Amount: \$10,000.
3. Co-I, NIH, “Buffalo Clinical and Translational Research Center.” Duration 43 mos. Starting date: 08/12/2015. Amount: \$13,151,902.
4. PI, NSF/CBET, “Collaborative Research: CDS&E: Sculpting fluid flow using a programmed sequence of micro-pillars.” Duration: 36 mos. Starting date: 09/01/2013. Amount: \$59,842 plus \$8,000 REU Supplement.

Resource Allocations

1. PI, AWS Cloud Credits for Research, “Long DNA reads analysis in Amazon cloud.” Resource allocation on the Amazon cloud, Duration: 12 mos. Starting date: 04/01/2016. Amount: \$10,000.

Completed

1. Co-PI, NSF/IIS, “BIGDATA: Mid-Scale: DA: Collaborative Research: Genomes galore – Core techniques, libraries, and domain specific languages for high-throughput DNA sequencing.” Duration: 36 mos. Starting date: 01/01/2013. Amount: \$1,300,000.
2. PI, Argonne Leadership Computing Facility, “ELaSTIC – Similarity graphs from large-scale biological sequence collections.” Resource allocation on the Mira supercomputer, Duration: 12 mos. Starting date: 01/01/2015. Amount: 2M core-hours.
3. PI, NSF/CCF, “Collaborative Research: Student travel support: International Workshop on Big Data in Life Sciences, Newport Beach, CA, September 20, 2014.” Duration: 12 mos. Starting date: 09/01/2014. Amount: \$10,000.
4. PI, NSF/IIP (sub-award from Optimal Solutions, Inc.), “STTR Phase I: Using big data to support supply chain analytics and optimization.” Duration: 12 mos. Starting date: 01/01/2014. Amount: \$112,500.
5. Researcher, Polonium Programme, “Grid computing for solving large problems in bioinformatics.” Joint project: Grenoble Institute of Technology (France), Poznan University of Technology (Poland) and Czestochowa University of Technology (Poland), 2005-2006.
6. Researcher, Laboratoire Franco-Mexicain d’Informatique, “Phylogenetic tree reconstruction on large scale distributed systems and grids.” Joint project: Grenoble Institute of Technology (France) and Centro de Investigacion Cientifica y de Education Superior de Ensenada (Mexico), 2005-2006.

Publications and Talks

† – joint first authors.

Journal Publications

1. O. Wodo, **J. Zola**, B.S.S. Pokuri, P. Du, B. Ganapathysubramanian, “Automated, High Throughput Exploration of Process-structure-property Relationships Using the MapReduce Paradigm,” *Materials Discovery*, 1, pp. 21-28, 2015. DOI: 10.1016/j.md.2015.12.001
2. N. Clauvelin, P. Lo, O. I. Kulaeva, E.V. Nizovtseva, J. Diaz-Montes, **J. Zola**, M. Parashar, V.M. Studitsky, W.K. Olson, “Nucleosome Positioning and Composition Modulate In Silico Chromatin Flexibility,” *Journal of Physics: Condensed Matter*, 27(6):064112, 2015. DOI: 10.1088/0953-8984/27/6/064112 PMID: 25564155

3. S.K. Samudrala, **J. Zola**, S. Aluru, B. Ganapathysubramanian, "Parallel Framework for Dimensionality Reduction of Large-Scale Datasets," *Scientific Programming*, 2015.
DOI: 10.1155/2015/180214
4. J. Diaz-Montes, B. Ganapathysubramanian, M. Parashar, I. Rodero, Y. Xie, **J. Zola**: "Federated Computing for the Masses – Aggregating Resources to Tackle Large-scale Engineering Problems," *IEEE Computing in Science & Engineering*, 16(4), pp. 62–72, 2014.
DOI: 10.1109/MCSE.2013.134
5. O. Nikolova, **J. Zola**, S. Aluru: "Parallel Globally Optimal Structure Learning of Bayesian Networks," *Journal of Parallel and Distributed Computing*, 73(8), pp. 1039–1048, 2013.
DOI: 10.1016/j.jpdc.2013.04.001
6. X. Yang[†], **J. Zola**[†], S. Aluru: "Large-Scale Metagenomic Clustering on Map-Reduce Clusters," *Journal of Bioinformatics and Computational Biology*, 11(1):1340001, 2013.
DOI: 10.1142/S0219720013400015 PMID: 23427983
7. M. Aluru[†], **J. Zola**[†], D. Nettleton, S. Aluru: "Reverse Engineering and Analysis of Large Genome-Scale Gene Networks," *Nucleic Acids Research*, 41(1):e24, 2013.
DOI: 10.1093/nar/gks904 PMID: 23042249
8. A. Sarje[†], **J. Zola**[†], S. Aluru: "Accelerating Pairwise Computations on Cell Processors," *IEEE Transactions on Parallel and Distributed Systems*, 22(1), pp. 69–77, 2011.
DOI: 10.1109/TPDS.2010.65
9. X. Yu, L. Li, **J. Zola**, M. Aluru, H. Ye, A. Foudree, H. Guo, S. Anderson, S. Aluru, P. Liu, S. Rodermel, Y. Yin: "A Brassinosteroid Transcriptional Network Revealed by Genome-Wide Identification of BES1 Target Genes in *Arabidopsis thaliana*," *Plant Journal*, 65(4), pp. 634–646, 2011. DOI: 10.1111/j.1365-313X.2010.04449.x PMID: 21214652
10. **J. Zola**, M. Aluru, A. Sarje, S. Aluru: "Parallel Information-Theory-Based Construction of Genome-Wide Gene Regulatory Networks," *IEEE Transactions on Parallel and Distributed Systems*, 21(12), pp. 1721–1733, 2010. DOI: 10.1109/TPDS.2010.59
11. M. Aluru, **J. Zola**, A. Foudree, S.R. Rodermel: "Chloroplast Photooxidation Induced Transcriptome Reprogramming in *Arabidopsis* *immutans* White Leaf Sectors," *Plant Physiology*, 150(2), pp. 904–923, 2009. DOI: 10.1104/pp.109.135780 PMID: 19386811
12. M. Ott, **J. Zola**, S. Aluru, A.D. Johnson, D. Janies, A. Stamatakis: "Large Scale Phylogenetic Analysis on Current HPC Architectures," *Scientific Programming*, 16(2-3), pp. 255–270, 2008. DOI: <http://dx.doi.org/10.3233/SPR-2008-0247>
13. D. Trystram, G. Parmentier, **J. Zola**: "Large Scale Multiple Sequence Alignment with Simultaneous Phylogeny Inference," *Journal of Parallel and Distributed Computing*, 66(12), pp. 1534–1545, 2006. DOI: 10.1016/j.jpdc.2006.03.003

Peer-reviewed Conference Publications

14. F. Schoeneman, S. Mahapatra, V. Chandola, N. Napp, **J. Zola**, "Error Metrics for Learning Reliable Manifolds from Streaming Data," Accepted to SIAM International Conference on Data Mining, 2017 (preprint). arXiv: 1611.04067 [stat.ML]

15. S. Karan, **J. Zola**: “Exact Structure Learning of Bayesian Networks by Optimal Path Extension,” In Proc. of IEEE Int. Conference on Big Data (BigData), pp. 48–55, 2016.
DOI: 10.1109/BigData.2016.7840588
16. R. Nihalani, **J. Zola**, S. Aluru: “Resolving Read Assignment Ambiguities in Metagenomic Clustering,” In Proc. of ISCA Int. Conference on Bioinformatics and Computational Biology (BI-CoB 2013), 2013.
17. X. Yang, **J. Zola**, S. Aluru: “Parallel Metagenomic Sequence Clustering via Sketching and Maximal Quasi-Clique Enumeration on Map-Reduce Clouds,” In IEEE Proc. of Int. Parallel and Distributed Processing Symposium (IPDPS 2011), pp. 1210–1220, 2011.
DOI: 10.1109/IPDPS.2011.116
18. **J. Zola**[†], A. Sarje[†], S. Aluru: “Constructing Gene Regulatory Networks on Clusters of Cell Processors,” In IEEE Proc. of Int. Conference on Parallel Processing (ICPP 2009), pp. 108–115, 2009. DOI: 10.1109/ICPP.2009.35
19. O. Nikolova, **J. Zola**, S. Aluru: “Parallel Algorithm for Exact Bayesian Network Inference,” In IEEE Proc. of Int. Conference on High Performance Computing (HiPC 2009), pp. 342–349, 2009. DOI: 10.1109/HIPC.2009.5433194
20. **J. Zola**, M. Aluru, S. Aluru: “Parallel Information-Theory-Based Construction of Gene Regulatory Networks,” In Proc. of Int. Conference on High Performance Computing (HiPC 2008), LNCS vol. 5374, pp. 336–349, 2008. DOI: 10.1007/978-3-540-89894-8_31
21. M. Ott, **J. Zola**, S. Aluru, A. Stamatakis: “Large Scale Maximum Likelihood-Based Phylogenetic Analysis on the IBM BlueGene/L,” In Proc. of ACM/IEEE Supercomputing Conference (SC 2007), 2007 (**best paper finalist**). DOI: 10.1145/1362622.1362628
22. **J. Zola**, X. Yang, S. Rospondek, S. Aluru: “Parallel T-Coffee: A Parallel Multiple Sequence Aligner,” In Proc. of ISCA Int. Conference on Parallel and Distributed Computing and Systems (PDCS 2007), pp. 248–253, 2007.
23. D. Trystram, **J. Zola**: “Parallel Multiple Sequence Alignment with Decentralized Cache Support,” In Proc. of Int. European Conference on Parallel and Distributed Computing (Euro-Par 2005), LNCS vol. 3648, pp. 1217–1226, 2005. DOI: 10.1007/11549468_133
24. G. Parmentier, D. Trystram, **J. Zola**: “Cache-Based Parallelization of Multiple Sequence Alignment Problem,” In Proc. of Int. European Conference on Parallel and Distributed Computing (Euro-Par 2004), LNCS vol. 3149, pp. 1005–1012, 2004. DOI: 10.1007/978-3-540-27866-5_135
25. **J. Zola**, L. Lacinski, R. Wyrzykowski: “An Object-Oriented Library for Evolution Programs with Applications for Partitioning of Finite Element Meshes,” IUTAM Symp. on Evolutionary Methods in Mechanics, Solid Mechanics and its Applications, vol. 117, pp. 351–360, 2004.
26. R. Wyrzykowski, A. Tomas, K. Karczewski, **J. Zola**, T. Kuczynski: “Development, Administration and Use of Eltoro Cluster,” In Proc. of Pionier 2002, pp. 273–280, 2002 (in Polish).
27. **J. Zola**, R. Wyrzykowski: “EPL-Julia, the High Performance Library for Evolutionary Computations,” In Proc. of Int. Conference on Parallel Processing and Applied Mathematics (PPAM 2001), LNCS vol. 2328, pp. 633–639, 2002. DOI: 10.1007/3-540-48086-2_72

Peer-reviewed Workshop Publications

28. **J. Zola**: “Constructing Similarity Graphs from Large-scale Biological Sequence Collections,” In Proc. of IEEE Int. Workshop on High Performance Computational Biology (HiCOMB 2014), pp. 500–507, 2014. DOI: 10.1109/IPDPSW.2014.63
29. J. Diaz-Montes, Y. Xie, I. Rodero, **J. Zola**, B. Ganapathysubramanian, M. Parashar: “Exploring the Use of Elastic Resource Federations for Enabling Large-Scale Scientific Workflows,” In Proc. of Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS 2013), 2013 (**best paper finalist**).
30. **J. Zola**, D. Trystram, A. Tchernykh, C. Brizuela: “Parallel Multiple Sequence Alignment with Local Phylogeny Search by Simulated Annealing,” In Proc. of IEEE Int. Workshop on High Performance Computational Biology (HiCOMB 2006), 2006. DOI: 10.1109/IPDPS.2006.1639536

Book Chapters

31. **J. Zola**: “Parallel Computing for Gene Networks Reverse Engineering,” Chapter in “Approaches in Integrative Bioinformatics – Towards Virtual Cell,” M. Chen, R. Hofstaedt (eds.), Springer, 2014. DOI: 10.1007/978-3-642-41281-3_12
32. **J. Zola**, S. Aluru: “Network Inference in Systems Biology,” Chapter in “Encyclopedia of Parallel Computing,” D. Padua (ed.), Springer, 2011. DOI: 10.1007/978-0-387-09766-4_466
33. A. Sarje, **J. Zola**, S. Aluru: “Pairwise Computations on the Cell Processor with Applications in Computational Biology,” Chapter in “Scientific Computing with Multicores and Accelerators,” J. Kurzak, D.A. Bader, J. Dongarra (eds.), Chapman & Hall/CRC, 2010.
34. D. Trystram, **J. Zola**: “Multiple Sequence Alignment and Phylogenetic Inference,” Chapter in “Grids for Bioinformatics and Computational Biology,” A. Zomaya, E. Talbi (eds.), John Wiley & Sons, 2007. DOI: 10.1002/9780470191637.ch6

Edited Proceedings

35. A. Kalyanaraman, **J. Zola**, D.A. Bader, S. Aluru (eds.): “Proceedings of the Fifteenth IEEE International Workshop on High Performance Computational Biology (HiCOMB 2016),” IEEE Computer Society, 2016.
36. **J. Zola**, D.A. Bader, S. Aluru (eds.): “Proceedings of the Twelfth IEEE International Workshop on High Performance Computational Biology (HiCOMB 2013),” IEEE Computer Society, 2013.
37. D.A. Bader, D. Trystram, **J. Zola** (eds.): “Proceedings of the Parallel Bio-Computing Workshop (PBC 2005),” Springer Verlag, 2005.

Technical Reports

38. J. Diaz-Montes, B. Ganapathysubramanian, M. Parashar, I. Rodero, Y. Xie, **J. Zola**: “Federated Computing for the Masses – Aggregating Resources to Tackle Large-scale Engineering

Problems,” Technical Report, Rutgers Discovery Informatics Institute, TR-RDI2-20130515-0, 2013.

Invited Talks and Lectures

1. Invited talk: “Sketching Biological Sequences for Storage and Computation.” Presented at Workshop on Parallel Software Libraries for Sequence Analysis (pSALSA) in Seattle, October 2016.
2. Invited talk: “Parallel Computing and Bioinformatics: Practitioner’s Point of View.” Presented at Biomedical Informatics Grand Rounds at University at Buffalo, March 2016.
3. Invited talk: “Parallel Computing for Large-scale Computational Biology.” Presented at Pharmacology and Toxicology Seminar Series at University at Buffalo, October 2015.
4. Invited talk: “Similarity Detection in Large Sequence Collections.” Presented at Workshop on Parallel Software Libraries for Sequence Analysis (pSALSA) in Atlanta, September 2015.
5. Three-hours tutorial: “How to Find Who Lives in Your Belly Button: Not Too Short Introduction to Metagenomics.” Presented at DIMACS Center for Discrete Mathematics and Theoretical Computer Science, June 2014.
6. Invited talk: “Constructing Similarity Graphs from Large-scale Biological Sequence Collections.” Presented at Drexel University, March 2013.
7. Panelist: “Big Data Standards and the Potential Long Term Benefits for Research and Clinical Development.” Webinar organized by NGSLeaders, December 2012.
8. Invited talk: “Large Scale Metagenomic Clustering.” Presented at International Conference on the Bioinformatics of Genome Regulation and Structure/Systems Biology in Novosibirsk, Russia, June 2012.
9. Three-hours tutorial: “Not Too Short Introduction to Metagenomics.” Presented at Czestochowa University of Technology, Poland, December 2011.
10. Invited talk: “Computational Biology: Can We Do Without Parallel Computing?” Presented at Corporacion Universitaria para el Desarrollo de Internet meeting in Puebla, Mexico, April 2009.
11. Invited lectures: “Introduction to Bioinformatics.” and “High Performance Computational Biology.” Presented at Pannon University, Hungary, May 2006.
12. Two-hours talk with Denis Trystram: “Bio-Computing Today: What to Expect from Parallel and Distributed Systems?” Presented at 4th Franco-Mexican School on Distributed Systems in Grenoble, France, July 2005.
13. Seminar: “Parallel Server for Large Scale Multiple Sequence Alignment.” Presented at Poznan University of Technology, Institute of Computing Science, Poland, May 2005.
14. Invited talk with Roman Wyrzykowski and Tomasz Olas: “Parallel FEM Computations on PC-Based Cluster Using Myrinet.” Presented at Myrinet User’s Group Conference in Vienna, Austria, May 2002.

Teaching

Curriculum Development

1. CSE 470/570 "Introduction to Parallel and Distributed Processing," University at Buffalo.
2. BMI 503 "Biomedical Informatics Systems, Databases and Software Methods," University at Buffalo.

Courses Taught

1. BMI 503 "Biomedical Informatics Systems, Databases and Software Methods," University at Buffalo, Fall 2015, Fall 2016.
2. CSE 250 "Data Structures in C+," University at Buffalo, Spring 2015, Fall 2016.
3. CSE 603 "Parallel and Distributed Processing," University at Buffalo, Fall 2015.
4. CprE 426/526 "Introduction to Parallel Algorithms and Programming" (co-instructor), Iowa State University, Spring 2011 and 2012.
5. CS Graduate Level "Introduction to High Performance Computing," Czestochowa University of Technology, Spring 2006.
6. CS Undergraduate Level "Computer Languages and Programming Methods," Czestochowa University of Technology, Spring 2001-2005.
7. CS Undergraduate Level "Selected Topics in Combinatorial Optimization," Czestochowa University of Technology, Fall 2001.

Mentoring

Postdoctoral Supervisor

1. Vamsi Krishna Potluru (co-supervised with Manish Parashar), Rutgers University, 2014.

Ph.D. Advisor

2. Subhadeep Karan, University at Buffalo, 2015-present.
3. Frank Schoeneman, University at Buffalo, 2015-present.
4. Vicky Zheng, University at Buffalo, 2016-present.
5. Li Zhang (co-advised with Tevfik Kosar), University at Buffalo, 2016-present.

M.Sc. Advisor

6. Vinay Ashokkumar, independent study, University at Buffalo, 2016.
7. Kushal Bhandari, independent study, University at Buffalo, 2016.

8. Jeban Ephrim Gnanaraj Kanagarajan, independent study, University at Buffalo, 2016.
9. Ajay Sudhakar Deshpande, independent study, University at Buffalo, 2015-2016.
10. Dhanasekar Karuppasamy, independent study, University at Buffalo, 2015-2016.
11. Piotr Dziubecki, Thesis: "A User-Oriented Grid Portal For Large Scale Multiple Sequence Alignment," Czestochowa University of Technology, 2005-2006.
12. Adrian Rospondek, Thesis: "Parallel*-Coffee: Parallel Implementation of 3D/T-Coffee Method," Czestochowa University of Technology, 2005-2006.
13. Malgorzata Sikorska, Thesis: "A Practical Comparison of 2PC and 3PC Mechanisms for Distributed Databases," Czestochowa University of Technology, 2005-2006.
14. Monika Zagala, Thesis: "Rational Arithmetic of High Precision for C++," Czestochowa University of Technology, 2005-2006.

Thesis Committee

1. Luigi Di Tacchio, University at Buffalo.
2. Suchismit Mahapatra, University at Buffalo.
3. Pavan Kumar Behara, University at Buffalo.
4. Sai Kiranmayee Samudrala, Iowa State University, 2012.

Outreach Activities

Undergraduate Mentoring

1. Jon Goodrum, Research Internship, University at Buffalo, 2017-present.
2. Matthew Eichhorn, Research Internship, University at Buffalo, 2016-present.
3. Blake Hurlburt, Research Internship, University at Buffalo, 2016-present.
4. Grant Iraci, Research Internship, University at Buffalo, 2016-present.
5. Jacob Ekstrum, Research Internship, University at Buffalo, 2016-present.
6. Adithya Narayanan, Research Internship, University at Buffalo, 2016.
7. Ashish Tyagi, Research Internship, University at Buffalo, 2015-2016.
8. Paul Kowalski, NSF REU, University at Buffalo, 2015-2016.
(CSE Undergraduate Research Award, 2016)
9. Vicky Zheng, Research Internship, University at Buffalo, 2015-2016.
(CSE Undergraduate Research Award, 2016)

10. William Spoth, Research Internship, University at Buffalo, 2015.
11. David Bryant, Research Internship, University at Buffalo, 2015.
12. Hannah Graesser, Summer Research Internship, Canisius College, 2015.
13. Manikandan Sundararajan, Research Internship, University at Buffalo, 2015.
14. Vidita Gawade, Aresty Summer Research Program, Rutgers University, 2014.
15. Kush Oza, NSF REU, Rutgers University, 2014.
16. Dylan Quinta, NSF REU, Carnegie Mellon University, 2014.
17. Mariam Tsilosani, NSF REU, Rutgers University, 2014.
18. Walter Hummel, Aresty Research Assistant Program, Rutgers University, 2013-2014.
19. Alexio Mota, Aresty Research Assistant Program, Rutgers University, 2013-2014.

Early Education and K12 Outreach

1. Presenter, The Eric Pitman Annual Summer Workshop in Computational Science, Center for Computational Research, 2015 and 2016.
2. Instructor, Liberty Partnerships Program, University at Buffalo, 2015 and 2016.
3. Group Leader, "Science is Elementary" Outreach Program Westminster Community Charter School, Buffalo, 2014-2015.
4. Advisor, Chaitanya Aluru, "Research Internship in Bioinformatics and Computational Biology," Ames High School, IA (Extended Learning Program), 2011-2012.

Professional Service and Activities

Advisory Functions

1. Panelist, Gryphon Scientific/AAAS/FBI Roundtable on Enhancing the Research and Educational Environment at U.S. Universities, 2015.
2. Member, AAAS/FBI Working Group on Big Data in Life Sciences and National Security, 2014.

Organizing Committees

1. Founding co-chair, ACM International Workshop on Big Data in Life Sciences (BigLS).
2. Founding co-chair, Parallel Bio-Computing Workshop (PBC).
3. Advancing Algorithms and Methods track co-chair, ACM Conference on Bioinformatics, Computational Biology and Biomedical Informatics (ACM BCB), 2017.

4. Student travel award co-chair, ACM Conference on Bioinformatics, Computational Biology and Biomedical Informatics (ACM BCB), 2016.
5. Program chair, IEEE International Workshop on High Performance Computational Biology (HiCOMB), 2013 and 2016.
6. Co-organizer, ACM/IEEE Supercomputing workshop: Computational and Data Challenges in Genomic Sequencing, 2015.
7. Algorithms track co-chair, International Conference on Contemporary Computing (IC3), 2015.
8. Publicity chair, IEEE International Conference on Computational Advances in Bio and Medical Sciences (ICCABS), 2013.
9. Publicity co-chair, ISCA International Conference on Bioinformatics and Computational Biology (BICoB), 2013.
10. Co-organizer, IEEE Cluster workshop: Parallel Programming and Applications on Accelerator Clusters (PPAAC), 2010.

Program Committees

1. IEEE International Conference on High Performance Computing (HiPC), 2015 and 2017.
2. International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), 2013, 2016 and 2017.
3. International Conference on Parallel Processing and Applied Mathematics (PPAM), 2013, 2015 and 2017.
4. ACM Conference on Bioinformatics, Computational Biology and Biomedical Informatics (ACM BCB), 2013, 2015 and 2016.
5. IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2015 and 2016.
6. Workshop on Accelerator-Enabled Algorithms and Applications in Bioinformatics (WACEBI), 2016.
7. International Workshop on Data Intensive Distributed Computing (DIDC), 2016.
8. ACM/IEEE Supercomputing, Posters Committee, 2015.
9. IEEE Cluster 2015.
10. IEEE International Workshop on Foundations of Big Data Computing, 2015.
11. International Conference on Parallel Processing (ICPP), 2012 and 2014.
12. ASE/IEEE International Conference on Big Data, 2013 and 2014.
13. International Conference on Contemporary Computing (IC3), 2011 and 2014.
14. IEEE International Conference on Parallel and Distributed Systems (ICPADS), 2012 and 2013.
15. ISCA International Conference on Bioinformatics and Computational Biology (BICoB), 2012.
16. Workshop on Using Emerging Parallel Architectures (WEPA) 2009, 2010 and 2012.
17. ACM/IEEE Supercomputing, 2008.

Referee

Proposals

1. National Science Foundation, 2016, 2017.
2. Great Lakes Consortium for Petascale Computation, 2015, 2016, 2017.
3. AAAS Peer Review Committee for KACST, 2014.
4. Austrian Science Fund, 2013.
5. Polish National Science Centre, 2013.

Journals

1. IEEE Transactions on Parallel and Distributed Systems
2. Journal of Parallel and Distributed Computing
3. ACM Transactions on Parallel Computing
4. Parallel Computing
5. Future Generation Computer Systems
6. International Journal of Parallel Programming
7. Concurrency and Computation: Practice and Experience
8. International Journal of High Performance Computing
9. IEEE/ACM Transactions on Computational Biology and Bioinformatics
10. Bioinformatics
11. PLOS ONE
12. IET Systems Biology
13. International Journal of Data Mining and Bioinformatics
14. BMC Bioinformatics
15. BMC Research Notes
16. Software: Practice and Experience

Professional Societies Membership

1. Institute of Electrical and Electronics Engineers (IEEE), 2006-present (senior member since 2012).
2. IEEE Computer Society, 2006-present.
3. Association for Computing Machinery (ACM), 2007-present.
4. ACM Special Interest Group on High Performance Computing (SIGHPC), 2014-present.
5. ACM Special Interest Group on Bioinformatics, Computational Biology, and Biomedical Informatics (SIGBio), 2014-present.

6. American Association for the Advancement of Science (AAAS), 2013-present.

Professional Societies Service

1. ACM Special Interest Group on Bioinformatics, Computational Biology, and Biomedical Informatics, Director of Web Communications, 2017-present.
2. IEEE Computer Society Technical Committee on Parallel Processing, Student Travel Awards Chair, 2016-present.

University Service – University at Buffalo

University Level

1. Faculty Advisory Committee, Center for Computational Research, 2014-present.

Department of Computer Science and Engineering

1. Member, Teaching Effectiveness and TA Training Committee, 2016-present.
2. Member, Undergraduate Affairs Committee (chair of Data-Driven Decisions sub-committee), 2015-present.
3. Member, Student Awards Committee, 2015-present.
4. Member, Colloquium Committee, 2014-2016.
5. Member, Graduate Admission Committee, 2014.

Department of Biomedical Informatics

1. Member, Digital Outreach Committee, 2016-present.

Media Coverage

1. “Kudos: Zola member of ‘big data’ working group” – “UB Reporter,” A short story highlighting my participation in the AAAS/FBI Working Group on Big Data in Life Sciences and National Security.
2. “Big Biological Impacts From Big Data” – “Science Magazine,” Story about how big data impacts life sciences, featuring my comments on the subject.
3. “Our Bodies, Our Data” – “Quanta Magazine” and “Wired.com,” Story about the current status of big data in life sciences, featuring my comments on the topic.
4. “Sculpting Flow” – “NSF News From the Field” and “Science 2.0,” Brief report on the computational aspect of sculpting fluid flows in microfluidic channels. The report covers the HPC experiment I led at Rutgers. It has been featured on the main NSF News web page.

5. “Understanding Fluid Flow in Microchannels” – “Digital Manufacturing Report” and “HPC in the Cloud,” Article about using a cloud-like federation of HPC resources to solve pressing engineering problems. The story is based on the experiment I led at Rutgers to analyze a parameter space in microfluidic flow.

Software Development

All projects are open source.

1. **SABNA**: High performance, exact structure learning of Bayesian networks under different scoring criteria. The framework implements several novel algorithms and data structures to enable BN learning on large data. Developed together with Subhadeep Karan (Ph.D. student) – <https://gitlab.com/SCoRe-Group/SABNA-Release>.
2. **ELaSTIC**: Similarity graphs construction from massive biological sequence collections. The method is design to detect pairwise similarities in data sets with millions of DNA/RNA or protein sequences without aligning all sequences against each other – <http://www.jzola.org/elastic>.
3. **CLOSET**: Cloud-enabled framework based on map-reduce for large scale metagenomic data clustering. The software handles data sets consisting of millions of reads and can be used to classify data at different taxonomic levels – <http://www.jzola.org/closet>.
4. **TINGe**: Systems biology parallel framework for gene networks inference. The software can handle whole genome microarray expression data, and it has been successfully used to reconstruct the largest to date gene network of *Arabidopsis thaliana* – <http://www.jzola.org/tinge>.
5. **GeNA**: Cytoscape plugin and standalone application for analysis of gene regulatory networks using approach akin to the PageRank algorithm – <http://www.jzola.org/tinge>.
6. **Parallel T-Coffee**: Parallel multiple sequence aligner able to align thousands of sequences using the T-Coffee scoring method (more than 400 downloads, web portal built by researchers from Spain and England) – <http://www.jzola.org/ptc>.
7. **Libpnorm**: A C/C++ high performance library for scheduling pairwise computations on Cell processors taking into account local memory limitations – <http://www.jzola.org/libpnorm>.