

Topics in Parallel and Distributed Computing: Introducing Algorithms, Programming, and Performance within Undergraduate Curricula^{*†‡}

Chapter 1 – Editors Introduction and Road Map

Sushil K. Prasad, Anshul Gupta, Arnold L. Rosenberg,

Alan Sussman, and Charles Weems

*How to cite this book: Prasad, Gupta, Rosenberg, Sussman, and Weems. Topics in Parallel and Distributed Computing: Enhancing the Undergraduate Curriculum: Performance, Concurrency, and Programming on Modern Platforms, Springer International Publishing, 2018, ISBN : 978-3-319-93108-1, Pages: 337.

†Free preprint version of this book: https://grid.cs.gsu.edu/~tcpp/curriculum/?q=cder_book_2

‡Free preprint version of volume one: https://grid.cs.gsu.edu/~tcpp/curriculum/?q=cedr_book

Our premise is that every computer science (CS) and computer engineering (CE) undergraduate student should achieve a specified skill level in parallel and distributed computing (PDC) because of required coursework. Recognizing that both instructors and students need suitable textual material, a book project has been initiated to address the lack of suitable textbooks to integrate PDC topics into the lower level core courses (CS1, CS2, Systems, Data Structures and Algorithms, Logic Design, etc.). This book is a companion to our first book of the CDER Book Project* and has evolved organically based on contributions received in response to calls for book chapters in 2016. We invited proposals for chapters on teaching parallel and distributed computing topics, suitable for either instructors or students, specifically on topics from the current TCPP/CDER curriculum guidelines for introductory courses[†] that have not been addressed by the chapters in the earlier volume. Subsequently, we saw good community interest in writing chapters for higher level elective courses as well. To address this, we extended the scope of this book to both lower-level core courses and more advanced, specialized topics in parallel and distributed computing that are targeted at students in upper level classes. All contributions have been rigorously reviewed internally by the editors and externally by experts.

*Prasad, Gupta, Rosenberg, Sussman, and Weems. 2015. Topics in Parallel and Distributed Computing: Introducing Concurrency in Undergraduate Courses, 1st Edition, Morgan Kaufmann, ISBN : 9780128038994, Pages: 360. http://grid.cs.gsu.edu/~tcpp/curriculum/?q=cedr_book

[†]Prasad, S. K., Chtchelkanova, A., Dehne, F., Gouda, M., Gupta, A., Jaja, J., Kant, K., La Salle, A., LeBlanc, R., Lumsdaine, A., Padua, D., Parashar, M., Prasanna, V., Robert, Y., Rosenberg, A., Sahni, S., Shirazi, B., Sussman, A., Weems, C., and Wu, J. 2012. NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing - Core Topics for Undergraduates, Version I, Online: <http://www.cs.gsu.edu/~tcpp/curriculum/>, 55 pages.

This material is based upon work partially supported by the National Science Foundation under Grants IIS 1143533, CCF 1135124, CCF 1048711 and CNS 0950432. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

1.1 Why These Books?

Parallel and Distributed Computing (PDC) now permeates most computing activities. It is no longer sufficient for even novice programmers to acquire only traditional sequential programming skills. In recognition of this technological evolution, a working group was organized under the aegis of the US National Science Foundation (NSF) and the IEEE Computer Society Technical Committee on Parallel Processing (TCPP) to draft guidelines for the development of curriculum in PDC for low-level (core) undergraduates courses in computational subjects such as Computer Science (CS), Computer Engineering (CE), and Computational Science and Engineering. The NSF/TCPP working group proposed a set of core PDC topics, with recommendations for the level of covering each in undergraduate curricula. The resulting curricular guidelines have been organized into four tables, one for each of computer architecture, computer programming, algorithms, and cross-cutting topics[†].

The initial enthusiastic reception of these guidelines led to a commitment within the working group to continue to develop the guidelines and to foster their adoption at a broad range of academic institutions. Toward these ends, the Center for Curriculum Development and Educational Resources (CDER) was founded, with the five editors of this volume comprising the initial Board of Directors. A new working group has taken up revision of the 2012 TCPP/CDER curriculum during 2016-18 – specially looking from the aspects of Big Data, Energy, Distributed Computing, and Exemplars – to create version-II of the TCPP/CDER curriculum. CDER has initiated several activities toward the end of fostering PDC education.

1. A *courseware repository* has been established for pedagogical materials – sample lectures, recommended problem sets, experiential anecdotes, evaluations, etc. This is a living repository: CDER invites the community to contribute existing and new material to it[‡].
2. An *Early Adopter Program* has been established to foster the adoption and evaluation of the guidelines. This activity has fostered educational work on PDC at more than 100 educational

[‡]CDER Courseware Repository: [http://www.cs.gsu.edu/~tcpp/curriculum/?q=courseware_](http://www.cs.gsu.edu/~tcpp/curriculum/?q=courseware_management)
management

institutions in North and South America, Europe, and Asia. The Program has thereby played a major role in establishing a worldwide community of people interested in developing and implementing PDC curricula.

3. The *EduPar workshop series* has been established. The original instantiation of EduPar was as a satellite of the International Parallel and Distributed Processing Symposium (IPDPS). EduPar was – and continues to be – the first education-oriented workshop at a major research conference. The success of EduPar led to the development of a sibling workshop, EduHPC, at the Supercomputing Conference (SC). In August, 2015, EduPar and EduHPC will be joined by a third sibling workshop, Euro-EduPar, which will be a satellite of the International Conference on Parallel Computing (EuroPar). CDER has also sponsored panels and BOF sessions at the ACM Conference on Computer Science Education (SIGCSE).

The preceding activities have succeeded not only as individual activities, but even more exciting, by spawning a vibrant international community of educators who are committed to PDC education. It is from this community that the desirability of the *CDER Book Project*, a series of volumes to support both instructors and students of PDC, became evident. This is the first volume in the series.

Motivation and Goals of the CDER Book Project: Curricular guidelines such as those promulgated by both the CDER Center and the CS2013 ACM/IEEE Computer Science Curriculum Joint Task Force[‡] are an essential first step in propelling the teaching of PDC-related material into the 21st century. But such guidelines are only a first step: both instructors and students will benefit from suitable textual material to effectively translate guidelines into curriculum. Moreover, experience to this point has made it clear that the members of the PDC community have much to share with each other and with aspiring new members, in terms of creativity in forging new directions and experience in evaluating existing ones. The Book Project's goal is to engage the community to address the need for suitable textbooks and related textual material to integrate PDC topics into the lower level core courses (which we affectionately, and hopefully transparently, refer to as CS1, CS2, Systems, Data Structures and Algorithms, Logic Design, etc.). The current edited book series intends, over time, to cover all of these proposed topics. This first volume in the projected series

has two parts.

Part I - For instructors: These chapters are aimed at instructors to provide background, scholarly materials, insights into pedagogical strategies, and descriptions of experience with both strategies and materials. The emphasis is on the basic concepts and references on what and how to teach PDC topics in the context of the existing topics in various core courses.

Part 2 - For students: These chapters aim to provide supplemental textual material for core courses which students can rely on for learning and exercises.

Print and Free Web Publication: While a print version through a renowned commercial publisher will foster our dissemination efforts in a professional format, the preprint versions of all the chapters will be freely available on the CDER Book Project website[§].

Organization: This introductory chapter is organized as follows. Section 1.2 gives brief outlines of each of the ten subsequent chapters. Section 1.3 provides a roadmap for the readers to find suitable chapters and sections within these which are relevant for specific courses or PDC topics. Section 1.3 contains biographical sketches of the editors and authors.

[‡]The ACM/IEEE Computer Science Curricula 2013 document (<http://www.acm.org/education/CS2013-final-report.pdf>) explicitly refers to our proposed curriculum on page 146 as follows:

As multi-processor computing continues to grow in the coming years, so too will the role of parallel and distributed computing in undergraduate computing curricula. In addition to the guidelines presented here, we also direct the interested reader to the document entitled "NSF/TCPP Curriculum Initiative on Parallel and Distributed Computing - Core Topics for Undergraduates", available from the website: <http://www.cs.gsu.edu/~tcpp/curriculum>.

[§]CDER Book Project - Free Preprint Version: http://cs.gsu.edu/~tcpp/curriculum/?q=CDER_Book_Project

1.2 Chapter Introductions

1.2.1 Part I - For instructors

In Chapter 2, titled *What do we need to know about parallel algorithms and their efficient implementation?*, Voevodin et al. discuss

In Chapter 3, titled *Modules for Teaching Parallel Performance Concepts*, Apan Qasem discusses three teaching modules targeting parallel performance concepts. The first module discusses fundamental concepts in parallel computing performance and mainly targets a CS1 course, highlighting parallel programming tools and performance metrics, and provides several sample exercises. The second module targets an upper-level operating systems class and focuses on communication and synchronization and how they affect performance for parallel applications, introducing the concepts of data dependences, synchronization, race conditions, and load balancing, again providing several sample exercises. The third module focuses on performance measurement and estimation of parallel systems, targeting compiler and computer architecture classes. This module reviews basic performance concepts and discusses advanced concepts such as strong and weak, scaling, linear and super-linear speedup, and latency vs. bandwidth measurements, in the context of OpenMP, and provides two sample exercises.

Chapter 4, titled *Scalability in Parallel Processing*, by Yanik Ngoko and Denis Trystram

The objective of this chapter is to discuss the notion of *scalability*. We start by explaining the notion with an emphasis on modern (and future) large scale parallel platforms. We also review the classical metrics used for estimating the scalability of a parallel platform, namely, speed-up, efficiency and asymptotic analysis. We continue with the presentation of two fundamental laws of scalability: Amdahl's and Gustafson's laws. Our presentation considers the original arguments of the authors and reexamines their applicability in today's machines and computational problems. Then, the chapter discusses more advanced topics that cover the evolution of computing fields (in term of problems), modern resource sharing techniques and the more specific issue of reducing energy consumption. The chapter ends with a presentation of a statistical approach to the design of scalable algorithms. The approach describes how scalable algorithms can be designed by using

a “cooperation” of several parallel algorithms solving the same problem. The construction of such cooperations is particularly interesting while solving hard combinatorial problems. We provide an illustration of this last point on the classical satisfiability problem SAT.

This chapter is intended to be used in intermediate-advanced courses on the design and analysis of parallel algorithms. The material covers data parallelism, performance metrics, performance modeling, speedup, efficiency, Amdahl’s law, Gustafson’s law, and isoefficiency. It also presents an analysis of Amdahl’s and Gustafson’s laws when considering resource sharing techniques, energy-efficiency and problem types. The analysis could be too advanced for a CS2 student because it requires a background in modern parallel systems and computer architectures.

In Chapter 5, titled *Energy Efficiency Issues in Computing Systems*, Krishna Kant

In Chapter 6, titled *Scheduling for fault-tolerance*, Guillaume Aupy and Yves Robert

1.2.2 Part 2 - For students

In Chapter 7, titled *MapReduce The Scalable Distributed Data Processing Solution*, Bushra Anjum provides students with an overview of how to process large datasets using the MapReduce programming model. Along with multiple examples of MapReduce applications, the chapter provides an outline of the basic functions that must be written to build a MapReduce application, and also discusses how the map and reduce steps in a distributed MapReduce system (i.e. Hadoop) execute a MapReduce application with scalable performance. The chapter also discusses the strengths and limitations of the MapReduce model, addressing scalability, flexibility, and fault tolerance. Finally, the chapter discusses higher level services built on top of the basic Hadoop MapReduce system.

Chapter 8, titled *The Realm of Graphical Processing Unit (GPU) Computing*, by Vivek Palipuram and Jinzhu Gao

In this chapter, the authors provide an introduction to general-purpose graphical processing unit (GPGPU) computing using the Compute Unified Device Architecture (CUDA) programming model. The chapter extensively covers the GPGPU architecture as viewed by a CUDA programmer and CUDA concepts including CUDA thread management, memory management, and per-

formance optimization strategies. The chapter pedagogically reinforces the CUDA concepts using parallel patterns such as matrix-matrix multiplication and convolution. The chapter includes several active-learning exercises that engage students with the text. Throughout this chapter, students will develop an ability to write effective CUDA codes for maximum application performance. The chapter is intended for an upper-level undergraduate course with object-oriented programming and data structures using C++ as prerequisites. The chapter can also be used in a sophomore- or junior-level software engineering course, or in an undergraduate elective course dedicated to high-performance computing using a specialized architecture. Because the chapter covers the GPGPU architecture and programming in detail, a prior exposure to CS1/CS2 level programming with basic computer organization is desirable.

In Chapter 9, titled *Managing Concurrency in Mobile User Interfaces with Examples in Android*, Konstantin Läufer and George K. Thiruvathukal discuss parallel and distributed computing from a mobile application development perspective, specifically addressing concurrency in interactive, GUI-based applications on the Android platform. The chapter gives an overview of GUI-based applications and frameworks, then looks at implementing simple interactive application behavior in the Android mobile application development framework using a running example. More complex use cases are introduced that enable discussing event handling and timers, to further show how GUI applications display all the benefits and costs of concurrent execution. Finally, the chapter closes with a deeper exploration of long-running compute-bound applications, where the problem is to maintain responsiveness to user requests.

Chapter 10, titled *Parallel Programming for Interactive GUI Applications*, by Nasser Giacomman and Oliver Sinnen

In this chapter, the authors show students how to use Java threads to implement a graphical user interface (GUI) that is responsive even while computation is being done. Because this example of concurrency is concrete and visual, it can be introduced fairly early in the curriculum. If the first course in programming makes active use of the Java GUI framework, then this will be a modest extension of that coverage. By at least the second course in programming (again if GUI programming is already included), and certainly in a sophomore software engineering class, this

material can be used as a means to introduce many ideas that are basic to PDC, and get students thinking in terms of using explicit concurrency to take advantage of the capabilities of modern systems. The foundation that is laid by this material could easily be extended, for example, in a programming with data structures course, to introduce thread-safe processing of larger structures, including algorithms such as parallel merge sort.

Chapter 11, titled *Scheduling in Parallel and Distributed Computing Systems*, by Srishti Srivastava and Ioana Banicescu

1.3 How to Find a Topic or Material for a Course?

The following table lists the remaining chapters in the book, core/elective undergraduate courses they can be used for (see list below), and their prerequisites, if any. More detailed tables in the Appendix list the topics covered in each chapter.

CORE COURSES:

CS0: Computer Literacy for Non-majors

CS1: Introduction to Computer Programming (First Course)

CS2: Second Programming Course in the Introductory Sequence

Systems: Introductory Systems/Architecture Course

DS/A: Data Structures and Algorithms

CE1: Digital Logic (First Course)

ADVANCED/ELECTIVE COURSES:

Arch2: Advanced Elective Course on Architecture

Algo2: Elective/Advanced Algorithm Design and Analysis (CS7)

Lang: Programming Language/Principles (after introductory sequence)

SwEngg: Software Engineering

ParAlgo: Parallel Algorithms

ParProg: Parallel Programming

Compilers: Compiler Design

Networking: Communication Networks

DistSystems: Distributed Systems

OS: Operating Systems

Chap #	Short title	Primary core course	Other courses	Prerequisites
Part I				
2	Parallel algorithms and implementation	CS0	CS1, DS/A, ParAlgo	-
3	Parallel Performance Concepts	CS1 Systems	OS, DS/A	DS/A for advanced modules
4	Scalability	Systems	CS2, ParAlgo	Math maturity
5	Energy Efficiency	CE1	CS2, DS/A	Math maturity
6	Scheduling	Systems	CS2, DS/A	CS1, Probabilities
Part II				
7	MapReduce	DS/A	CS2, ParAlgo, DistSystems	CS0, CS1
8	GPU Computing	Systems	Arch 2, ParProg	CS1, CS2
9	Mobile User Interfaces	DS/A	OOP, ParProg	CS1, CS2
10	Interactive GUI Applications	CS2 DS/A		CS1, CS2
11	Scheduling	DS/A	ParAlgo, DistSystems	basic PDC terms and concepts

Editor and Author Biographical Sketches

1.3.1 Editors

Anshul Gupta is a Principal Research Staff Member in Mathematical Sciences department at IBM T.J. Watson Research Center. His research interests include sparse matrix computations and their applications in optimization and computational sciences, parallel algorithms, and graph/combinatorial algorithms for scientific computing. He has coauthored several journal articles and conference papers on these topics and a textbook titled "Introduction to Parallel Computing." He is the primary author of Watson Sparse Matrix Package (WSMP), one of the most robust and scalable parallel direct solvers for large sparse systems of linear equations.

Sushil K. Prasad (BTech85 IIT Kharagpur, MS86 Washington State, Pullman; PhD90 Central Florida, Orlando - all in Computer Science/Engineering) is a Professor of Computer Science at Georgia State University and Director of Distributed and Mobile Systems (DiMoS) Lab. He has carried out theoretical as well as experimental research in parallel and distributed computing, resulting in 140+ refereed publications, several patent applications, and about \$3M in external research funds as principal investigator and over \$6M overall (NSF/NIH/-GRA/Industry).

Sushil has been honored as an ACM Distinguished Scientist in Fall 2013 for his research on parallel data structures and applications. He was the elected chair of IEEE Technical Committee on Parallel Processing for two terms (2007-11), and received its highest honors in 2012 - IEEE TCPP Outstanding Service Award. Currently, he is leading the NSF-supported IEEE-TCPP curriculum initiative on parallel and distributed computing with a vision to ensure that all computer science and engineering graduates are well-prepared in parallelism through their core courses in this era of multi- and many-cores desktops and handhelds. His current research interests are in Parallel Data Structures and Algorithms, and Computation over Geo-Spatiotemporal Datasets over Cloud, GPU and Multicore Platforms. His home-page is www.cs.gsu.edu/prasad.

Arnold L. Rosenberg is a Research Professor in the Computer Science Department at Northeastern University; he also holds the rank of Distinguished University Professor Emeritus in the Computer Science Department at the University of Massachusetts Amherst. Prior to joining UMass, Rosenberg was a Professor of Computer Science at Duke University from 1981 to 1986, and a Research Staff Member at the IBM Watson Research Center from 1965 to 1981. He has held visiting positions at Yale University and the University of Toronto. He was a Lady Davis Visiting Professor at the Technion (Israel Institute of Technology) in 1994, and a Fulbright Senior Research Scholar at the University of Paris-South in 2000. Rosenberg's research focuses on developing algorithmic models and techniques to exploit the new modalities of "collaborative computing" (wherein multiple computers cooperate to solve a computational problem) that result from emerging computing technologies. Rosenberg is the author or coauthor of more than 170 technical papers on these and other topics in theoretical computer science and discrete mathematics. He is the coauthor of the research book "Graph Separators, with Applications" and the author of the textbook "The Pillars of Computation Theory: State, Encoding, Nondeterminism"; additionally, he has served as coeditor of several books. Dr. Rosenberg is a Fellow of the ACM, a Fellow of the IEEE, and a Golden Core member of the IEEE Computer Society. Rosenberg received an A.B. in mathematics at Harvard College and an A.M. and Ph.D. in applied mathematics at Harvard University. More details are available at <http://www.cs.umass.edu/~rsnrg/>.

Alan Sussman is a Professor in the Department of Computer Science and Institute for Advanced Computer Studies at the University of Maryland. Working with students and other researchers at Maryland and other institutions he has published numerous conference and journal papers and received several best paper awards in various topics related to software tools for high performance parallel and distributed computing, and has contributed chapters to 6 books. His research interests include peer-to-peer distributed systems, software engineering for high performance computing, and large scale data intensive computing. Software tools he has built with his graduate students have been widely distributed and used in many computational science applications, in areas such as earth science, space science, and medical

informatics. He is a subject area editor for the *Parallel Computing* journal and an associate editor for *IEEE Transactions on Services Computing*, and edited a previous book on teaching parallel and distributed computing. He is a founding member of the Center for Parallel and Distributed Computing Curriculum Development and Educational Resources (CDER). He received his Ph.D. in computer science from Carnegie Mellon University.

Charles Weems is co-director of the Architecture and Language Implementation lab at the University of Massachusetts. His current research interests include architectures for media and embedded applications, GPU computing, and high precision arithmetic, and he has over 100 conference and journal publications. Previously he led development of two generations of a heterogeneous parallel processor for machine vision, called the Image Understanding Architecture, and co-directed initial work on the Scale compiler that was eventually used for the TRIPS architecture. He is the author of numerous articles, has served on many program committees, chaired the 1997 IEEE CAMP Workshop, the 1999 IEEE Frontiers Symposium, co-chaired IEEE IPDPS in 1999, 2000, and 2013, was general vice-chair for IPDPS from 2001 through 2005, is on the steering committees of EduPar and EduHPC. He has co-authored twenty-eight introductory CS texts. He is a member of ACM, Senior Member of IEEE, a member of the Executive Committee of the IEEE TC on Parallel Processing, has been an editor for IEEE TPDS, Elsevier JPDC, and is an editor with *Parallel Computing*.

1.3.2 Authors

Bushra Anjum has a PhD in Computer Science from North Carolina State University and is currently serving as a Tech Lead for the Amazon Prime Program at Amazon, Inc. Alongside, she is also a visiting professor at the Computer Science Department of the California Polytechnic Institute, San Luis Obispo. Anjum has been extensively using Elastic MapReduce platform provided by Amazon Web Services for a few years now for job related tasks. Before joining industry, she served in academia full time both in the USA and in Pakistan for 5+ years. With unconventional career choices and international exposure, she brings the expertise of being an academician, a researcher and a practitioner at the same time.

Alexander Antonov is a leading researcher in Research Computing Center of Lomonosov Moscow State University (RCC MSU). His main research interests are related to research in such fields as parallel and distributed computing, performance and efficiency of computers, parallel programming, informational structure of algorithms and programs, application optimization and fine tuning, architecture of computers, benchmarks, etc. In 1999 Alexander Antonov received PhD degree on the subject of interprocedural analysis of programs. Alexander took part in a number of projects supported by the Ministry of Education and Sciences of the Russian Federation, Russian Foundation for Basic Research and Russian Science Foundation. Alexander is editor of Parallel.Ru Information analytical center for parallel computing. Alexander Antonov is one of the main developers of the AlgoWiki Open encyclopedia of parallel algorithmic features. At the present time Alexander Antonov takes part in different researches being conducted in RCC MSU that are devoted to efficiency analysis of parallel applications and supercomputer systems in general. He has published over 50 scientific papers with 4 books among them.

Guillaume Aupy received his PhD from ENS Lyon. He is currently a tenured researcher at Inria Bordeaux Sud-Ouest. His research interests include data-aware scheduling, reliability, energy efficiency in high-performance computing. He is the author of numerous articles, has served on many program committees. He was the technical program vice-chair of SC17.

Ioana Banicescu is a professor in the Department of Computer Science and Engineering at Mississippi State University (MSU). Between 2009 and 2017, she was also a Director of the NSF Center for Cloud and Autonomic Computing at MSU. Professor Banicescu received the Diploma in Electronics and Telecommunications from Polytechnic University - Bucharest, and the M.S. and the Ph.D. degrees in Computer Science from New York University Polytechnic Institute. Her research interests include parallel algorithms, scientific computing, scheduling and load balancing algorithms, performance modeling, analysis and prediction, and autonomic computing. Between 2004-2017, she was an Associate Editor of the Cluster Computing journal and the International Journal on Computational Science and Engineer-

ing. Professor Banicescu, served and continues to serve on numerous research review panels for advanced research grants in the US and Europe, on steering and program committees of a number of international conferences, symposia and workshops, on the Executive Board and Advisory Board of the IEEE Technical Committee on Parallel Processing (TCPP). She has given many invited talks at universities, government laboratories, and at various national and international forums in the United States and overseas.

Jinzhu Gao received the Ph.D. degree in Computer Science from The Ohio State University in 2004. From June 2004 to August 2008, she worked at the Oak Ridge National Laboratory as a research associate and then the University of Minnesota, Morris, as an Assistant Professor of Computer Science. She joined the University of the Pacific (Pacific) in 2008 and is currently a Professor of Computer Science at Pacific. Her main research focus is on intelligent data visual analytics. Over the past fifteen years, Dr. Gao has been working closely with application scientists and Silicon Valley technology companies to develop online data visual analytics and deep learning platforms to support collaborative science, mobile health, IoT data analytics, business operational visibility, and visual predicative analysis for industries. Her work has been published in top journals such as IEEE Transactions on Visualization and Computer Graphics, IEEE Transactions on Computers, and IEEE Computer Graphics and Applications.

Nasser Giacaman is a Senior Lecturer in the Department of Electrical and Computer Engineering at the University of Auckland, New Zealand. His disciplinary research interest includes parallel programming, particularly focusing on high-level languages in the context of desktop and mobile applications running on multi-core systems. He also researches Software Engineering Education by driving the development of tools and apps to help students learn difficult programming concepts.

Krishna Kant is a full professor in the department of computer and information science at Temple University. He has 37 years of combined experience in academia, industry, and government and has published in a wide variety of areas in computer science, telecommunications,

and logistics systems. His current research interests span energy management, data centers, wireless networks, resilience in high performance computing, critical infrastructure security, storage systems, database systems, configuration management, and logistics networks. He is a fellow of IEEE.

Konstantin Läufer is a full professor of computer science at Loyola University Chicago. He received a PhD in computer science from the Courant Institute at New York University in 1992. His research interests include programming languages, software architecture, and distributed and pervasive computing systems. His recent focus in research and teaching has been on the impact of programming languages, methodologies, frameworks, and tools on software quality. Konstantin has repeatedly served as a consultant in academia and industry and is a co-inventor on two patents owned by Lucent Technologies.

Yanik Ngoko received his B.Sc. in Computer Science from University of Yaoundé I (UYI), Cameroon, his M.Sc. in parallel and numerical computing also from UYI, and his doctorate in Computer Science from the Institut National Polytechnique de Grenoble, France (2010). From 2011 to 2014, he was a postdoctoral researcher, first at the university of São Paulo and then at the university of Paris 13. Since October 2014, he is a research scientist at Qarnot computing and an associate member of the Laboratoire d'Informatique de Paris Nord (University of Paris 13). His research interests include parallel and distributed computing, web services, cloud computing, applications of edge computing to IoT.

Vivek Pallipuram (B.Tech.2008 NIT Trichy, MS2010 Clemson University, Ph.D.2013 Clemson University) is an Assistant Professor of Computer Engineering at University of the Pacific, Stockton, California. His research interests include high-performance computing (HPC), heterogeneous architectures such as general-purpose graphical processing units (GPGPUs) and Xeon Phi co-processors, Cloud computing, image processing, and random signal processing. His interests also include promoting HPC education and scientific computing in primarily-undergraduate universities. His work has been published in journals such as the Journal of Supercomputing, and Concurrency and Computation: Practice and Experience;

and in top conferences such as IEEE Cluster and eScience. He is also a peer-reviewer for a number of international journals and conference proceedings. In the classroom, he strives to be a facilitator by engaging students using active-learning techniques. In addition to receiving information from the instructor, students interact with their peers via in-class group activities and gain valuable perspective. This process increases the influx of knowledge per student, promoting well-rounded and comprehensive learning. He enjoys teaching high-performance computing, computer systems and networks, random signals, and image processing.

Apan Qasem is an Associate Professor in the Computer Science Department at Texas State University. He received his PhD in 2008 from Rice University. Qasem directs the Compilers Research Group at Texas State where he and his students are working on a number of projects in the area of high-performance computing including developing intelligent software for improving programmer productivity and using GPUs for general-purpose computation. Qasem's research has received funding from the National Science Foundation, Department of Energy, Semiconductor Research Consortium (SRC), IBM, Nvidia and the Research Enhancement Program at Texas State. In 2012, he received an NSF CAREER award to pursue research in autotuning of exascale systems. Qasem has co-authored over 50 peer-reviewed publications including two that won best paper awards. He regularly teaches the undergraduate and graduate Compilers and Computer Architecture courses.

Yves Robert received the PhD degree from Institut National Polytechnique de Grenoble. He is currently a full professor in the Computer Science Laboratory LIP at ENS Lyon. He is the author of 7 books, 150 papers published in international journals, and 240 papers published in international conferences. He is the editor of 11 book proceedings and 13 journal special issues. He is the advisor of 30 PhD theses. His main research interests are scheduling techniques and resilient algorithms for large-scale platforms. He is a Fellow of the IEEE. He has been elected a Senior Member of Institut Universitaire de France in 2007 and renewed in 2012. He has been awarded the 2014 IEEE TCSC Award for Excellence in Scalable

Computing, and the 2016 IEEE TCPP Outstanding Service Award. He holds a Visiting Scientist position at the University of Tennessee Knoxville since 2011.

Oliver Sinnen graduated in Electrical and Computer Engineering at RWTH Aachen University, Germany. Subsequently, he moved to Portugal, where he received his PhD from Instituto Superior Tecnico (IST), University of Lisbon, Portugal in 2003. Since 2004 he is a (Senior) Lecturer in the Department of Electrical and Computer Engineering at the University of Auckland, New Zealand, where he leads the Parallel and Reconfigurable Computing Lab. His research interests include parallel computing and programming, scheduling and reconfigurable computing. Oliver authored the book “Task Scheduling for Parallel Systems”, published by Wiley.

Srishti Srivastava is an Assistant Professor of Computer Science at the University of Southern Indiana. She received her Ph.D. in Computer Science at Mississippi State University in May 2015. Her research interests include dynamic load balancing in parallel and distributed computing, performance modeling, optimization, and prediction, robustness analysis of resource allocations, and autonomic computing. Srishti has authored and co-authored a number of articles published in renowned IEEE and ACM conferences, journals, and book chapters. Srishti has served on the program committees of international conference workshops such as, EduHPC, and EduPar. She has also been a peer reviewer for a number of international journals, and conference proceedings. She is a professional member of the IEEE computer society, ACM, Society for Industrial and Applied Mathematics (SIAM), Computing Research Association (CRA, CRA-W), Anita Borg Institute Grace Hopper Celebration (ABI-GHC), and an honor society of Upsilon Pi Epsilon (UPE). She is also a 2014 young researcher alumna of the Heidelberg Laureate Forum, Germany.

George K. Thiruvathukal received his PhD from the Illinois Institute of Technology in 1995. He is a full professor of computer science at Loyola University Chicago and visiting faculty at Argonne National Laboratory in the Mathematics and Computer Science Division, where he collaborates in high-performance distributed systems and data science. He is the author

of three books, co-editor of a peer-reviewed collection, and author of various peer-reviewed journal and conference papers. His early research involved object-oriented approaches to parallel programming and the development of object models, languages, libraries, and tools (messaging middleware) for parallel programming, mostly based on C/C++ on Unix platforms. His subsequent work in Java resulted in the book *High-Performance Java Platform Computing*, Prentice Hall and Sun Microsystems Press, 2000. He also co-authored the book *Codename Revolution: The Nintendo Wii Platform* in the MIT Press Platform Studies Series, 2012. Recently, he co-edited *Software Engineering for Science*, Taylor and Francis/CRC Press, October 2016.

Denis Trystram is a Professor in Computer Science at Grenoble Institute of technology since 1991 and is now distinguished professor there. He was a senior member of Institut Universitaire de France from 2010 to 2014. He obtained in 2011 a Google research award in Optimization for his contributions in the field of multi-objective Optimisation. Denis is leading a research group on optimization of resource management for parallel and distributed computing platforms in a joint team with Inria. Since 2010, he is director of the international Master program in Computer Science at university Grenoble-Alpes. He has been elected recently as the director of the research pole in Maths and Computer Science in this university.

Vadim Voevodin is a senior research fellow in Research computing center of Lomonosov Moscow state university (RCC MSU). His main research interests are related to different aspects of high-performance computing: analysis of parallel program efficiency, development of system software, parallel programming, etc. Vadim Voevodin got his PhD in memory locality analysis in parallel computing. Also he was a main developer in a research devoted to the study of memory hierarchy usage. At the present time Vadim Voevodin is actively involved in different researches being conducted in RCC MSU that are devoted to efficiency analysis of parallel applications and supercomputer systems in general. One research is dedicated to detecting abnormal inefficient job behavior based on constant monitoring of supercomputer job flow. The other newly started research is aimed to develop a universal software tool suite

that will help common users to conduct both large-scale efficiency analysis of the entire set of applications and a professional in-depth analysis of individual parallel applications, based on many researches previously done in RCC MSU. Another major research area concerns the analysis of supercomputer resource utilization and efficiency of using application packages installed on a supercomputer.

Vladimir Voevodin is Deputy Director of the Research Computing Center at Lomonosov Moscow State University. He is Head of the Department Supercomputers and Quantum Informatics at the Computational Mathematics and Cybernetics Faculty of MSU, professor, corresponding member of Russian academy of sciences. Vl. Voevodin specializes in parallel computing, supercomputing, extreme computing, program tuning and optimization, fine structure of algorithms and programs, parallel programming technologies, scalability and efficiency of supercomputers and applications, supercomputing co-design technologies, software tools for parallel computers, and supercomputing education. His research, experience and knowledge became a basis for the supercomputing center of Moscow State University, which was founded in 1999 and is currently the largest supercomputing center in Russia. He has contributed to the design and implementation of the following tools, software packages, systems and online resources: V-Ray, X-Com, AGORA, Parallel.ru, hpc-education.ru, hpc-russia.ru, LINEAL, Sigma, Top50, OctoShell, Octotron, AlgoWiki. He has published over 100 scientific papers with 4 books among them. Voevodin is one of the founders of Supercomputing Consortium of Russian Universities established in 2008, which currently comprises more than 60 members. He is a leader of the major national activities on Supercomputing Education in Russia and General Chair of the two largest Russian supercomputing conferences.

Appendix: Chapters and Topics

The following tables list the topics covered in each chapter. The depth of coverage of each topic is indicated by the intended outcome of teaching that topic, expressed using Bloom’s taxonomy of educational objectives:

K = Know the term

C = Comprehend so as to paraphrase/illustrate

A = Apply it in some way

Chapter 2: What do we need to know about parallel algorithms and their efficient implementation?

PDC Concept	Chapter Section		
	2.3.1	2.3.2	2.3.3
Performance issues	C	C	C
Information structure	C	C	C
Data locality	C	C	
Computational intensity	K		
Resource of parallelism	C	C	C
Computational kernel		K	
Serial complexity		K	C
Parallel complexity		K	C
Load balancing		C	A
Determinacy of an algorithm		C	
Scalability		C	A
Efficiency		C	C
Race conditions			A

Chapter 3: Modules for Teaching Parallel Performance Concepts

PDC Concept	Chapter Section		
	3.2	3.3	3.4
Speedup	K		C
Efficiency	K		C
Linear and Super linear speedup	K		C
Strong and weak scaling	K		C
Amdahl's Law	C		A
Power vs. time trade-offs	K		A
Task granularity		A	
Load balancing		A	
Communication and synchronization		C	
Scheduling and thread mapping		A	
SMP and NUMA			C
Data locality			C

Chapter 4: Scalability in Parallel Processing

PDC Concept	Chapter Section				
	4.1	4.2	4.3	4.4	4.5
scalability	K	C			
speedup		C	C	C	A
efficiency		C	C	C	A
data parallelism			K		
isoefficiency		C			
Amdahl' law		K	C	C	A
Gustafson' law		K	C	C	A
strong scaling		C			C
weak scaling		C			C
resource sharing				C	A
energy efficiency		K		K	
P-completeness			K		
algorithm portfolio					A

Chapter 5: Energy Efficiency Issues in Computing Systems

PDC Concept	Chapter Section						
	5.1	5.2	5.3	5.4	5.5	5.6	5.7
Performance issues	C	C	C	C	C	C	C
Concurrency, sequential/parallel						A	
Scheduling				C	C		

Chapter 6: Scheduling for fault-tolerance

PDC Concept	Chapter Section					
	6.1	6.2	6.3	6.4	6.5	6.6
Why/What is Par/Dist Computing	K	K	K	K	K	K
Performance Issues, Computation	K	A	C	A	A	K
Cluster	K	K	K	K	K	K
Performance measures		A	A	A	A	
Basic Probabilities		C	A	C	C	
Programming SPMD		C				
Load balancing		K		K	A	
Scheduling		C		C	C	
Dynamic Programming					A	

Chapter 7: MapReduce The Scalable Distributed Data Processing Solution

PDC Concept	Chapter Section				
	7.1	7.2	7.3	7.4	7.5
Why/What is Par/Dist Computing	A	A	C	K	A
Concurrency	K	K	C		A
Cluster Computing	A	C	A	K	K
Scalability	A	C	A	K	A
Speedup	K	C	A		
Divide & Conquer (parallel aspects)	C	A	K		A
Recursion (parallel aspects)			K		A
Scan (parallel-prefix)	K	A			C
Reduction (map-reduce)	K	A		K	A
Time	C	A			A
Sorting	K	A			A

Chapter 8: The Realm of Graphical Processing Unit (GPU) Computing

PDC Concept	Chapter Section								
	8.1	8.2	8.3	8.4	8.5	8.6	8.7	8.8	8.9
Data parallelism	C								
GPGPU devices		C	A	A	A	A	A	A	
nvcc compiler			A						
Thread management				A	A				
Parallel patterns							A	A	
Performance evaluation							A		
Performance optimization							A	A	
CUDA		A		A	A	A	A	A	
Advancements in GPU computing									K

Chapter 9: Managing Concurrency in Mobile User Interfaces

PDC Concept	Chapter Section						
	9.1	9.2	9.3	9.4	9.5	9.6	9.7
Why and what is PDC	K						
Tasks and threads	K		C			A	A
Thread safety		K	C	C		A	
Race conditions		K	C	C		A	
Thread/task spawning			K			A	A
Synchronization			C	C		A	
Nondeterminism			C			A	
Deadlocks				K			

with Examples in Android

Chapter 10: Parallel Programming for Interactive GUI Applications

PDC Concept	Chapter Section					
	10.1	10.2	10.3	10.4	10.5	10.6
Concurrency	C	A	A	C	A	A
Race conditions		C	A			
Thread safety			C	C		
GUI concurrency				C	A	A

Chapter 11: Scheduling in Parallel and Distributed Computing Systems

PDC Concept	Chapter Section			
	11.1	11.2	11.3	11.4
MIMD architecture	K			
Multicore	C			
SMP	N			C
Topologies		N	N	
Latency		K		
Heterogeneous		K		K
Data Parallel		C		
Computation	C	C	C	C
Load balancing		C	C	C
Distributed memory			C	C
Client server			C	
Static		C	C	
Dynamic		C	C	C
Asymptotic	C			
Communication		C	C	
Synchronization		C	C	
Speedup			A	
Efficiency		A	A	
Makespan		C		C
Concurrency		C		
Performance modeling				K
Fault tolerance	K			