High Performance Computing

Modern Systems and Practices

Thomas Sterling
Matthew Anderson
Maciej Brodowicz
School of Informatics, Computing, and Engineering
Indiana University, Bloomington

Foreword by C. Gordon Bell

MK
MORGAN KAUFMANN PUBLISHERS
AN IMPRINT OF ELSEVIER

Morgan Kaufmann is an imprint of Elsevier
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, United States

Notices
Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-0-12-420158-3

Publisher: Katey Birtcher
Acquisition Editor: Steve Merken
Developmental Editor: Nate McFadden
Production Project Manager: Punithavathy Govindaradjane
Designer: Mark Rogers

Typeset by TNQ Books and Journals

CHAPTER

INTRODUCTION 1

CHAPTER OUTLINE

High Performance Computing. https://doi.org/10.1016/B978-0-12-420158-3.00001-0

Supercomputing, which means supercomputers and their application, is among the most important developments of the modern age, with unequaled impact across a vast diversity of fields of inquiry and practical effect. From the extremes of arcane sciences to the most immediate practical concerns, supercomputers play an essential role in the progress and advancement of human capabilities, environments, and understanding. No other single technology in the history of humanity has experienced a similar rate of growth, even in its relatively short existence. Within the span of a single human lifetime, supercomputers have expanded their ability to perform calculations by a factor of 10 trillion or 13 orders of magnitude, and this is a conservative estimate. From less than a 1000 basic operations per second in the late 1940s to today's performance in excess of a 100 quadrillion floating-point operations per second (over 100 petaflops), supercomputer speed has steadily improved by about a factor of 200 times every decade through a series of advances in technology, architecture, programming methods, algorithms, and system software (Fig. 1.1). High performance computing (HPC), synonymous with supercomputing, is a principal means of exploration complementing empirical methods used for more than 2 millennia and theory practiced in the age of enlightenment of the last 4 centuries. As the "third pillar" of investigation, supercomputing enables new paths of inquiry, new techniques of design, and new methods of operating process. Even discoveries correctly credited to other classes of tools and instrumentation, such as giant telescopes or particle accelerators, require the use of supercomputers as well to produce their final results through data analysis (sometimes referred to as "big data"). It can be asserted that supercomputing allows us to understand the past, to control the present, and in limited cases to predict the future.

The skills required to employ HPC are multiple and complex, while the means of acquiring such skills to a sufficient degree require potentially years of study and experience at least in normal practice.

FIGURE 1.1

The Titan petaflops machine fully deployed at Oak Ridge National Laboratory in 2013. It takes up more than 4000 sq ft and consumes approximately 8 MW of electrical power. It has a theoretical peak performance of over 27 petaflops and delivers 17.6 petaflops R_{max} sustained performance for the highly parallel Linpack (HPL) benchmark. This architecture includes Nvidia graphics processing unit accelerators.

Photo courtesy of Oak Ridge National Laboratory, US Dept. of Energy

This often means lengthy apprenticeships in research facilities in academia, industry, or national laboratories. There are many books written to teach particular programming languages; others describe in detail the structures and instruction sets of computer architectures; and still others discuss system software such as operating systems. But missing has been a single textbook that serves as an entry-level presentation of all these elements and their interrelationships in one place, combined with guided hands-on experience. This work, *High Performance Computing*, is developed as a carefully crafted synthesis of relevant elements of related disciplines, all of which contribute in critical ways to supercomputing and its use. This book presents the foundation concepts, in-depth relevant knowledge, and detailed skills that together will give you a meaningful understanding of HPC and an initial set of techniques to make you an effective, albeit incipient, practitioner in its use. Throughout this text the best practices employed by the community are presented with training, so you learn to do, the *how*, even as you are gaining understanding of the *what* and the *why*.

This textbook provides a comprehensive introduction to the field of HPC. It is presented in a form that will be both intellectually rewarding and practical in teaching useful basic skills. It combines perspectives about supercomputing concepts, knowledge about supercomputers, and techniques for using and programming supercomputers. But teaching a complex subject like HPC is challenging, in that just about everything is defined in terms of and relates to everything else. Yet by the nature of pedagogy, material must be presented in some sequential order. This first chapter is a brief introductory presentation of the essential elements of HPC to provide an overview of everything; a first pass that will allow successive in-depth chapters to be related to this broad context.

The chapter looks at the many facets which comprise HPC. The importance of the material is that it provides a complete, albeit simplified, perspective of HPC so that more detailed discussion of specifics can be understood within the full context. Because no piece makes sense without the others, almost all areas are briefly introduced in this chapter. To reinforce the interrelated broad-brush presentation of issues, this chapter concludes with a history of the field and its rapid evolution.

1.1 HIGH PERFORMANCE COMPUTING DISCIPLINES

As previously noted, HPC is really a collection of multiple interrelated disciplines, each providing an important aspect of the total field. To master HPC as a useful tool is to develop an understanding and associated skills in each of these corresponding areas. These broad areas are described here, including a formal definition of "high performance computing" that applies throughout the treatment of the field, end-user application problems that are the intended purpose of HPC across a wide range of science, engineering, societal, and security domains, the core concept of performance which is the distinguishing characteristic of HPC compared to other forms of computing, the hardware and software components that make up an HPC system, environments, tools, application programming, and the interfaces used. Each of these is presented in some detail in the following sections and together form a major portion of the concepts, knowledge content, and skills comprising this textbook.

1.1.1 DEFINITION

HPC is a field of endeavor that relates to all facets of technology, methodology, and application associated with achieving the greatest computing capability possible at any point in time and technology. It engages a class of electronic digital machines referred to as "supercomputers" to perform a wide array

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.