
A Load Balancing Framework For Clustered Storage Systems

Heterogeneous Metacomputing, Load Balancing and Programming in WebCom
4th IFIP TC 12 International Conference, ICCIDS 2021, Chennai, India, March 18-20, 2021, Revised Selected Papers
Engineering Societies in the Agents World X
ICWiCOM 2019
Ride the Performance Tiger
11th International Conference, Eugene, OR, USA, June 30 -- July 3, 2014, Revised Selected Papers
Practical Load Balancing
Machine Learning and Wireless Communications
Proceedings of International Conference on Wireless Communication
Parallel and Distributed Computing for Symbolic and Irregular Applications
High Performance Computing - HiPC 2008
International Conference, San Francisco, CA, USA, May 28-30, 2001. Proceedings, Part II
Proceedings of International Conference on Computational Intelligence and Data Engineering
Middleware 2006
Proceedings
The Architecture of Scientific Software
11th International Conference, VECoS 2017, Montreal, QC, Canada, August 24-25, 2017, Proceedings
Parallel Science and Engineering Applications
A Framework for Supporting Application Level Load Balancing in Distributing Systems
ICCIDE 2021
ACM/IFIP/USENIX 7th International Middleware Conference, Melbourne, Australia, November 27 - December 1, 2006, Proceedings
15th International Conference, Bangalore, India, December 17-20, 2008, Proceedings
A Framework for Dynamic Load Balancing and Physical Reorganization in Parallel Database Systems
Advances in Web-based GIS, Mapping Services and Applications
High Performance Computing
Application Delivery and Load Balancing in Microsoft Azure
Parallel and Distributed Processing
A Framework for
Scheduling in Distributed Computing Environment Using Dynamic Load Balancing
High-Performance Scientific Computing
A Framework for Solving Load Balancing Problems in Unique-path Circuit Switched Mesh-oriented Networks
Dynamic Load Balancing in Distributed Content-based Publish/subscribe
Load Balancing Servers, Firewalls, and Caches
Parallel Computing
Verification and Evaluation of Computer and Communication Systems
First JARA-HPC Symposium, JHPCS 2016, Aachen, Germany, October 4-5, 2016, Revised Selected Papers
Computational Science - ICCS 2001
IFIP TC2/WG2.5 Working Conference on the Architecture of Scientific Software October 2-4, 2000, Ottawa, Canada

BENJAMIN REEVES

Heterogeneous Metacomputing, Load Balancing and Programming in WebCom Springer

Scientific applications involve very large computations that strain the resources of whatever computers are available. Such computations implement sophisticated mathematics, require deep scientific knowledge, depend on subtle interplay of different approximations, and may be subject to instabilities and sensitivity to external input. Software able to succeed in this domain invariably embeds significant domain knowledge that should be tapped for future use. Unfortunately, most existing scientific software is designed in an ad hoc way, resulting in monolithic codes understood by only a few developers. Software architecture refers to the way software is structured to promote objectives such as reusability, maintainability, extensibility, and feasibility of independent implementation. Such issues have become increasingly important in the scientific domain, as software gets larger and more complex, constructed by teams of people, and evolved over decades. In the context of scientific computation, the challenge facing mathematical software practitioners is to design, develop, and supply computational components which deliver these objectives when embedded in end-user application codes. The Architecture of Scientific Software addresses emerging methodologies and tools for the rational design of scientific software, including component integration frameworks, network-based computing, formal methods of abstraction, application programmer interface design, and the role of object-oriented languages. This book comprises the proceedings of the International Federation for Information Processing (IFIP) Conference on the Architecture of Scientific Software, which was held in Ottawa, Canada, in October 2000. It will prove invaluable reading for developers of scientific software, as well as for researchers in computational sciences and engineering.
4th IFIP TC 12 International Conference, ICCIDS 2021, Chennai, India, March 18-20, 2021, Revised Selected Papers Anchor Academic Publishing

This book constitutes the refereed proceedings of the 10th International Workshop on Engineering Societies in the Agents World, ESAW 2009, held in Utrecht, The Netherlands, in November 2009. The 13 revised full papers and 5 short contributions presented together with two invited talks were carefully selected from 31 submissions. The papers are organized in topical sections on self-organization, software-engineering and architectures, social aspects of agent societies, organization and autonomy. This proceedings concludes with the extended abstracts of 6 contributions to a demonstration session on agent-based technologies and works.

Engineering Societies in the Agents World X IOS Press

As a concept, Concurrent Engineering (CE) initiates processes with the goal of improving product quality, production efficiency and overall customer satisfaction. Services are becoming increasingly important to the economy, with more than 60% of the GDP in Japan, the USA, Germany and Russia deriving from service-based activities. The definition of a product has evolved from the manufacturing and supplying of goods only, to providing goods with added value, to eventually promoting a complete service business solution, with support from introduction into service and from operations to decommissioning. This book presents the proceedings of the 20th ISPE International Conference on Concurrent Engineering, held in Melbourne, Australia, in September 2013. The conference had as its theme Product and Service Engineering in a Dynamic World, and the papers explore research results, new concepts and insights covering a number of topics, including service engineering, cloud computing and digital manufacturing, knowledge-based engineering and sustainability in concurrent engineering.

ICWiCOM 2019 Springer Science & Business Media

Developed in the context of science and engineering applications, with each abstraction motivated by and further honed by specific application needs, Charm++ is a production-quality system that runs on almost all parallel computers available. Parallel Science and Engineering Applications: The Charm++ Approach surveys a diverse and scalable collecti

Ride the Performance Tiger CRC Press

This book constitutes the proceedings of the 11th International

Conference International Conference on Verification and Evaluation of Computer and Communication Systems (VECoS 2017), held at Concordia University, Montreal, Canada, in August 2017. The 13 full papers, together with 3 abstracts in this volume were carefully reviewed and selected from 35 submissions. The aim of the VECoS conference is to bring together researchers and practitioners in the areas of verification, control, performance and dependability evaluation in order to discuss state-of-the-art and challenges in modern computer and communication systems in which functional and extra-functional properties are strongly interrelated. Thus, the main motivation for VECoS is to encourage the cross-fertilization between various formal verification and evaluation approaches, methods and techniques, and especially those developed for concurrent and distributed hardware/software systems.

11th International Conference, Eugene, OR, USA, June 30 - July 3, 2014, Revised Selected Papers IOS Press

Research on artificial life is critical to solving various dynamic obstacles individuals face on a daily basis. From electric wheelchairs to navigation, artificial life can play a role in improving both the simple and complex aspects of civilian life. The Handbook of Research on Investigations in Artificial Life Research and Development is a vital scholarly reference source that examines emergent research in handling real-world problems through the application of various computation technologies and techniques. Examining topics such as computational intelligence, multi-agent systems, and fuzzy logic, this publication is a valuable resource for academicians, scientists, researchers, and individuals interested in artificial intelligence developments.

Practical Load Balancing Springer Nature

This volume contains the proceedings from the workshops held in conjunction with the IEEE International Parallel and Distributed Processing Symposium, IPDPS 2000, on 1-5 May 2000 in Cancun, Mexico. The workshops provide a forum for bringing together researchers, practitioners, and designers from various backgrounds to discuss the state of the art in parallelism. They focus on different aspects of parallelism, from runtime systems to formal methods, from optics to irregular problems, from biology to networks of personal computers, from embedded systems to

programming environments; the following workshops are represented in this volume: { Workshop on Personal Computer Based Networks of Workstations { Workshop on Advances in Parallel and Distributed Computational Models { Workshop on Par. and Dist. Comp. in Image, Video, and Multimedia { Workshop on High-Level Parallel Prog. Models and Supportive Env. { Workshop on High Performance Data Mining { Workshop on Solving Irregularly Structured Problems in Parallel { Workshop on Java for Parallel and Distributed Computing { Workshop on Biologically Inspired Solutions to Parallel Processing Problems { Workshop on Parallel and Distributed Real-Time Systems { Workshop on Embedded HPC Systems and Applications { Reconfigurable Architectures Workshop { Workshop on Formal Methods for Parallel Programming { Workshop on Optics and Computer Science { Workshop on Run-Time Systems for Parallel Programming { Workshop on Fault-Tolerant Parallel and Distributed Systems All papers published in the workshops proceedings were selected by the program committee on the basis of referee reports. Each paper was reviewed by independent referees who judged the papers for originality, quality, and consistency with the themes of the workshops.

Machine Learning and Wireless Communications Springer
This book constitutes the refereed post-conference proceedings of the Fourth IFIP TC 12 International Conference on Computational Intelligence in Data Science, ICCIDS 2021, held in Chennai, India, in March 2021. The 20 revised full papers presented were carefully reviewed and selected from 75 submissions. The papers cover topics such as computational intelligence for text analysis; computational intelligence for image and video analysis; blockchain and data science.

Proceedings of International Conference on Wireless Communication IOS Press

Heterogeneous clusters with different kinds of compute elements are more and more a standard in the High Performance Computing Community. Efficiently utilizing all the available compute elements poses several challenges for a programmer. One of the major issues is load balancing which can for example result from the nature of the algorithm that needs to be executed or from the varying performance of the different compute elements. In this work a tasking framework is proposed that copes

with the architectural limits of a heterogeneous cluster. The initial implementation provides a simple interface, shows almost linear scaling and perfect load balancing for different case studies as for example a SpGEMM implementation. The final goal of this work was to design a framework that gives a simple task definition for tasks that can be efficiently executed on a GPU or any other compute elements available. In such heterogeneous systems the framework should allow to run tasks on all the CPU cores and the GPUs simultaneously and also can be used to decide if utilizing the GPU is beneficial or not. We further explore the possibilities and limitations for task execution and optimization for the GPU and what meta data and benchmarking is needed to decide if a GPU execution of a task turns out beneficial or if an execution on the CPU should be advised. Also the possibility of overlapping CPU task with the execution of a GPU task on the core that also drives the GPU was investigated. Furthermore the framework should still be able to efficiently load balance across the cluster as well as within the node or if the application allows for, inside the GPUs. Besides having a mostly positive impact on the overall performance, the framework also simplifies access to all compute elements and makes the implementation independent from variations in the topology of the cluster.

Parallel and Distributed Computing for Symbolic and Irregular Applications Springer

There are many scientific applications for which the computational load varies throughout the execution and causes uneven distribution of workload during run-time. One such class of applications is Adaptive Mesh Refinement (AMR) applications. AMR is a type of multiscale algorithm that achieves high resolution in localized regions of dynamic, multidimensional numerical simulations. A typical AMR application may require enormous computing resources, which usually cannot be satisfied by a single-processor machine, thereby requiring parallel and distributed systems. One of the key issues related to AMR is dynamic load balancing (DLB), which allows large-scale adaptive applications to run efficiently on parallel and distributed systems. In investigating DLB schemes, we first complete a detailed analysis of structured AMR (SAMR) applications, identifying the unique characteristics that impose severe challenges on DLB schemes. The results indicate that most of the available DLB schemes are not appropriate for SAMR applications due to their

unique adaptive characteristics. Thus, we propose a novel dynamic load balancing scheme for SAMR applications on parallel systems (denoted as parallel DLB). It integrates a grid-splitting technique with direct grid movements, for which the objective is to reduce the parallel execution time. Further, our experiment shows that simply moving a DLB scheme designed for parallel systems to distributed systems will introduce significant overhead. Therefore, we propose a framework for dynamic load balancing on distributed systems (denoted as distributed DLB). It takes into consideration: (1) heterogeneity of processors, (2) heterogeneity of networks, (3) shared nature of networks, and (4) adaptive characteristics of the applications. For SAMR applications, the distributed DLB incorporates the proposed parallel DLB during the load balancing process. Both parallel DLB and distributed DLB were implemented in the ENZO code, a parallel implementation of SAMR in astrophysics and cosmology. Experiments show that the proposed DLB schemes can significantly improve the performance of SAMR applications on both parallel and distributed systems in terms of the total execution time and the quality of load balancing.

High Performance Computing - HiPC 2008 O'Reilly Media
A User Facilitated Autonomous Load Balancing Framework for UCNA Framework for Supporting Application Level Load Balancing in Distributing Systems A FRAMEWORK FOR SCALABLE DISTRIBUTED JOB PROCESSING WITH DYNAMIC LOAD BALANCING USING DECENTRALIZED APPROACH Lulu.com Benchmarking Framework for Performance in Load Balancing Single System Image

International Conference, San Francisco, CA, USA, May 28-30, 2001. Proceedings, Part II Springer

The ability to dynamically adapt an unstructured grid is a powerful tool for efficiently solving computational problems with evolving physical features. In this paper, we report on our experience parallelizing an edge-based adaptation scheme, called 3D_TAG, using message passing. Results show excellent speedup when a realistic helicopter rotor mesh is randomly refined. However, performance deteriorates when the mesh is refined using a solution-based error indicator since mesh adaptation for practical problems occurs in a localized region, creating a severe load imbalance. To address this problem, we have developed PLUM, a global dynamic load balancing framework for adaptive

numerical computations. Even though PLUM primarily balances processor workloads for the solution phase, it reduces the load imbalance problem within mesh adaptation by repartitioning the mesh after targeting edges for refinement but before the actual subdivision. This dramatically improves the performance of parallel 3D_TAG since refinement occurs in a more load balanced fashion. We also present optimal and heuristic algorithms that, when applied to the default mapping of a parallel repartitioner, significantly reduce the data redistribution overhead. Finally, portability is examined by comparing performance on three state-of-the-art parallel machines.

Proceedings of International Conference on Computational Intelligence and Data Engineering Springer Nature

The emergence of the cloud and modern, fast corporate networks demands that you perform judicious balancing of computational loads. Practical Load Balancing presents an entire analytical framework to increase performance not just of one machine, but of your entire infrastructure. Practical Load Balancing starts by introducing key concepts and the tools you'll need to tackle your load-balancing issues. You'll travel through the IP layers and learn how they can create increased network traffic for you. You'll see how to account for persistence and state, and how you can judge the performance of scheduling algorithms. You'll then learn how to avoid performance degradation and any risk of the sudden disappearance of a service on a server. If you're concerned with running your load balancer for an entire network, you'll find out how to set up your network topography, and condense each topographical variety into recipes that will serve you in different situations. You'll also learn about individual servers, and load balancers that can perform cookie insertion or improve your SSL throughput. You'll also explore load balancing in the modern context of the cloud. While load balancers need to be configured for high availability once the conditions on the network have been created, modern load balancing has found its way into the cloud, where good balancing is vital for the very functioning of the cloud, and where IPv6 is becoming ever more important. You can read Practical Load Balancing from end to end or out of sequence, and indeed, if there are individual topics that interest you, you can pick up this book and work through it once you have read the first three chapters.

Middleware 2006 Springer Science & Business Media

This book covers various topics, including collective intelligence, intelligent transportation systems, fuzzy systems, Bayesian network, ant colony optimization, data privacy and security, data mining, data warehousing, big data analytics, cloud computing, natural language processing, swarm intelligence, and speech processing. This book is a collection of high-quality research work on cutting-edge technologies and the most-happening areas of computational intelligence and data engineering. It includes selected papers from the International Conference on Computational Intelligence and Data Engineering (ICCID 2021).

Proceedings CRC Press

This book illustrates various components of Distributed Computing Environment and the importance of distributed scheduling using Dynamic Load Balancing. It describes load balancing algorithms for better resource utilization, increasing throughput and improving user's response time. Various theoretical concepts, experiments, and examples enable students to understand the process of load balancing in computing cluster and server cluster. The book is suitable for students of Advance Operating Systems, High Performance Computing, Distributed Computing in B.E., M.C.A., M. Tech. and Ph.D courses.

The Architecture of Scientific Software A User Facilitated Autonomous Load Balancing Framework for UCNA Framework for Supporting Application Level Load Balancing in Distributing Systems
A FRAMEWORK FOR SCALABLE DISTRIBUTED JOB PROCESSING WITH DYNAMIC LOAD BALANCING USING DECENTRALIZED APPROACH

This book constitutes the refereed proceedings of the 31st International Conference, ISC High Performance 2016 [formerly known as the International Supercomputing Conference] held in Frankfurt, Germany, in June 2016. The 25 revised full papers presented in this book were carefully reviewed and selected from 60 submissions. The papers cover the following topics: Autotuning and Thread Mapping; Data Locality and Decomposition; Scalable Applications; Machine Learning; Datacenters and Cloud; Communication Runtime; Intel Xeon Phi; Manycore Architectures; Extreme-scale Computations; and Resilience.

11th International Conference, VECoS 2017, Montreal, QC,

Canada, August 24–25, 2017, Proceedings Springer

The book comprises selected papers presented at the International Conference on Wireless Communication (ICWiCOM), which is organized by D. J. Sanghvi College of Engineering's Department of Electronics and Telecommunication Engineering. The book focuses on specific topics of wireless communication, like signal and image processing applicable to wireless domains, networking, microwave and antenna design, and telemedicine systems. Covering three main areas – networking, antenna designs and embedded systems applicable to communication – it is a valuable resource for postgraduate and doctoral students.

Parallel Science and Engineering Applications BiblioGov

This book is intended for those developers who are keen to master the internal workings of Play Framework to effectively build and deploy web-related apps.

A Framework for Supporting Application Level Load Balancing in Distributing Systems Apress

This book constitutes the thoroughly refereed post-conference proceedings of the First JARA High-Performance Computing Symposium, JARA-HPC 2016, held in Aachen, Germany, in October 2016. The 21 full papers presented were carefully reviewed and selected from 26 submissions. They cover many diverse topics, such as coupling methods and strategies in Computational Fluid Dynamics (CFD), performance portability and applications in HPC, as well as provenance tracking for large-scale simulations.

ICCID 2021 Springer

LNCS volumes 2073 and 2074 contain the proceedings of the International Conference on Computational Science, ICCS 2001, held in San Francisco, California, May 27-31, 2001. The two volumes consist of more than 230 contributed and invited papers that reflect the aims of the conference to bring together researchers and scientists from mathematics and computer science as basic computing disciplines, researchers from various application areas who are pioneering advanced application of computational methods to sciences such as physics, chemistry, life sciences, and engineering, arts and humanitarian fields, along with software developers and vendors, to discuss problems and solutions in the area, to identify new issues, and to shape future directions for research, as well as to help industrial users apply various advanced computational techniques.