Consensus Based Dynamic Load Balancing For A Network Of

Right here, we have countless ebook **Consensus Based Dynamic Load Balancing For A Network Of** and collections to check out. We additionally manage to pay for variant types and after that type of the books to browse. The agreeable book, fiction, history, novel, scientific research, as without difficulty as various other sorts of books are readily comprehensible here.

As this Consensus Based Dynamic Load Balancing For A Network Of, it ends occurring beast one of the favored book Consensus Based Dynamic Load Balancing For A Network Of collections that we have. This is why you remain in the best website to look the amazing books to have.

Consensus Based Dynamic Load Balancing For A Network Of ***2024-07-09***

OCONNELL GUADALUPE

A Novel Agent-based Dynamic Load Balancing Model for Cloud Networks BiblioGov
A trace-driven simulation study of dynamic load balancing in homogeneous distributed systems supporting broadcasting is presented. We use information about job CPU and I/O demands collected from a production system as input to a simulation mode l that includes a representative CPU scheduling policy and considers the message exchange and job transfer costs explicitly. Seven load balancing algorithms are simulated and their performances compared. We find that load balancing is capable of significa ntly reducing the mean and standard deviation of job response times, especially under heavy system load, and for jobs with high resource demands. The performances of all hosts, even those originally with light loads, are generally improved by load balanci ng. The reduction of the mean response time increases with the number of hosts, but levels off at around 30 hosts. Algorithms based on periodic or non-periodic load information exchange provide similar performance, and, among the periodic policies, the al gorithms that use a distinguished agent to convert and distribute load information cut down the overhead and scale better. They are also the most appropriate algorithms for adaptive load balancing, which has the potential of offering near-optimal performa nce under a wide spectrum of system configurations and load conditions. System instability in the form of host overloading is possible when the load information is not up-to-date and the system is under heavy load; however, this undesirable phenomenon can be alleviated by simple measures. Load balancing is still very effective even when up to half of the eligible jobs have to be executed locally. The trace-driven simiulation approach to the study of load balancing is found to be critical and effective, an d is recommended for use before implementation efforts.
Phase-based Adaptive Dynamic Load Balancing for Parallel Tree Computation Springer
Distributed systems are often characterized by uneven loads on hosts and other resources. In this thesis, the problems concerning dynamic load balancing in loosely-coupled distributed systems are studied using trace-driven simulation, implementation, and measurement. Information about job CPU and I/O demands is collected from three production systems and used as input to a simulator that includes a representative CPU scheduling policy and considers the message exchange and job transfer costs explicitly. A prototype load balancer is implemented in the Berkeley UNIX and Sun/UNIX environments, and the results of a large number of measurement experiments performed on six workstations are presented.
Scheduling in Distributed Computing Environment Using Dynamic Load Balancing Lulu.com
A distributed system consists of many heterogeneous processors with different processing power and all processors are interconnected with a communication channel. In such a system, if some processors are less loaded or idle and others are heavily loaded, the system performance will be reduced drastically. System performance can be improved by using proper load balancing [1, 4]. The aim of load balancing is to improve the performance measures and reduce the overall completion time and cost
Dynamic Load Distribution for Parallel Applications Springer-Verlag
Abstract: "Distributed-controlled dynamic load balancing algorithms are known to have several advantages over centralised algorithms such as scalability, and fault tolerance. Distributed implies that the control is decentralised and that a copy of the algorithm (called a scheduler) is replicated on each host of the network. However, distributed control also contributes to the lack of global goals and lack of coherence. This paper presents a new algorithm called DGP (Decentralised Global Plans) that addresses the problem of coherence and coordination in distributed dynamic load balancing algorithms. The DGP algorithm is based on a strategy called Global Plans (GP), and aims at maintaining all computational loads of a distributed system within a band called [delta]. The rationale for the design of DGP is to allow each scheduler to consider the actions of its peer schedulers. With this level of coordination, the schedulers can act more as a coherent team. This new approach first explicitly specifies a global goal and then design [sic] a strategy around this global goal such that each scheduler (1) takes into account local decisions made by other schedulers; (2) takes into account the effect of its local decisions on the overall system and (3) ensures load balancing. An experimental evaluation of DGP with two other well-known dynamic load balancing algorithms published in the literature shows that DGP performs consistently better. More significantly, the results indicate that the global plan approach provides a better frameowrk for the design of distributed dynamic load balancing algorithms."
Dynamic Load Balancing Algorithms in Loosely-coupled Hard Real-time Systems Anchor Academic Publishing
In a distributed computing environment, it is important to ensure that the processor workloads are adequately balanced, Among numerous load-balancing algorithms, a unique approach due to Das and Prasad defines a symmetric broadcast network (SBN) that provides a robust communication pattern among the processors in a topology-independent manner. In this paper, we propose and analyze three efficient SBN-based dynamic load-balancing algorithms, and implement them on an SGI Origin2000. A thorough experimental study with Poisson distributed synthetic loads demonstrates that our algorithms are effective in balancing system load. By optimizing completion time and idle time, the proposed algorithms are shown to compare favorably with several existing approaches.
Loadserver - a Hybrid Dynamic Load Balancing Algorithm for Tree-based Applictions Springer Science & Business Media
An important consideration in improving the performance of a distributed computer system is the balancing of the load between the host computers. Load balancing may be either static or dynamic; static balancing strategies are generally based on information about the system's average behavior rather than its actual current state, while dynamic strategies react to the current state when making transfer decisions. Although it is often conjectured that dynamic load balancing outperforms static, careful investigation shows that this view is not always valid. Recent research on the problem of optimal static load balancing is clearly and intuitively presented, with coverage of distributed computer system models, problem formulation in load balancing, and effective algorithms for implementing optimization. Providing a thorough understanding of both static and dynamic strategies, this book will be of interest to all researchers and practitioners working to optimize performance in distributed computer systems.
Adaptive Load-Balancing Algorithms Using Symmetric Broadcast Networks
Dynamic Load Balancing (DLB) in Multi-agent Spatial Simulation (MASS) library is a thread-based load-balancing algorithm that calculates the CPU load, per thread, based on the computational time spent by CPU, using JAVA ThreadMXBean API. This thesis presents the system design, execution model of the three algorithms each based on (1) an entire history, (2) a recent time window and (3) a slope of the CPU load and our performance evaluation over two multi-threaded applications: Wave2D and SugarScape. MASS library single-node and multi-node versions are also evaluated with the objective that the Slope-based algorithm will be superior to the other two algorithms.
A Probabilistic Dynamic Load Balancing Algorithm for Homogeneous Distributed Systems (with Extension to Hypercubes)
Dynamic load balancing (DLB) is a technique for the parallel implementation of problems which generate unpredictable workloads by migrating work units to lightly loaded processors based on run-time workload measurement. Adaptive DLB is a refinement where aspects of the load balancing system itself are modified in the light of measured workloads. This thesis investigates phase-based adaptive DLB, a version of adaptive DLB in which a parallel computation moves through different load balancing phases identified on the basis of run-time workloads. The idea is explored through a case study of parallel tree computation, in which three distinct phases with intervening transitions are identified. Two major variants of phase-based adaptivity are distinguished. In parametric adaptivity, parameters of the DLB algorithm are adapted between phases; in algorithmic adaptivity, different DLB algorithms are utilised in each phase. These concepts are investigated quantitatively through a simulator for parametric adaptivity and discussed in detail for algorithmic adaptivity. The simulator permits a range of processor topologies, parameterises the performance of the underlying network, includes two different network performance models, and allows a wide range of simulated tree-structured workloads, parameterised by depth, fan-out, node granularity and imbalance. It was extensively validated in relation to the performance of two DLB algorithms on a 512-processor Cray T3D. The simulator was used to evaluate the benefit of parametric phase-based adaptivity. Preliminary experiments with non-adaptive algorithms revealed performance to be sensitive to the interval between load-balancing invocations, so this parameter was prioritised for subsequent adaptivity experiments. A performance metric called Improvement Through Adaptivity (ITA) was discussed. Two DLB algorithms were used as test cases; the well-established Generalised Dimension Exchange Method and a novel Loadserver algorithm, whose implementation is described in the thesis. Results were obtained for all combination of the transitions, and the relationships between ITA and various parameters (processor sizes, node granularity, tree imbalance and network performance) were established. Similar relationships were observed for both algorithms, though with some differences in detail. Positive values of ITA were obtained with both algorithms, for at least one transition combination, over a range of all the parameters. Thus, the potential benefits of phase-based parametric adaptivity are confirmed, justifying future work in implementing this technique.
Dynamic Load Balancing on Clusters of Heterogenous Workstations
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.
Dynamic Load Balancing
Simulation results show that the proposed algorithm operates efficiently and effectively, and it provides a significantly improved performance over existing algorithms. Comparisons with the central load balancing algorithm show that the new system provides a reduction in average wait time, a significant increase in throughput, and a dramatic reduction in CPU time consumption.
An Optimal Dynamic Load Balancing Algorithm
This book constitutes the refereed proceedings of the International Conference on Advances in Computing Communications and Control, ICAC3 2011, held in Mumbai, India, in January 2011. The 84 revised full papers presented were carefully reviewed and selected from 309 submissions. The papers address issues such as AI, artificial neural networks, computer graphics, data warehousing and mining, distributed computing, geo information and statistical computing, learning algorithms, system security, virtual reality, cloud computing, service oriented architecture, semantic web, coding techniques, modeling and simulation of communication systems, network architecture, network protocols, optical fiber/microwave communication, satellite communication, speech/image processing, wired and wireless communication, cooperative control, and nonlinear control, process control and instrumentation, industrial automation, controls in aerospace, robotics, and power systems.
Optimal Load Balancing in Distributed Computer Systems
Distributed content-based publish/subscribe systems to date suffer from performance degradation and poor scalability under load conditions typical in real-world applications. The reason for this shortcoming is due to the lack of a load balancing solution, which have rarely been studied in the context of publish/subscribe. This thesis proposes a load balancing solution specific for distributed content-based publish/subscribe systems that is distributed, dynamic, adaptive, transparent, and accommodates heterogeneity. The solution consists of three key contributions: a load balancing framework, a novel load estimation algorithm, and three offload strategies. Experimental results show that the proposed load balancing solution is efficient with less than 0.7% overhead, effective with at least 90% load estimation accuracy, and capable of load balancing with 100% of load initiated at an edge node of the entire system using real-world data sets.
Advances in Computing, Communication and Control
A Performance Study of Distributed Dynamic Load Balancing Techniques
Dynamic Load Balancing Schemes for Large-scale HLA-based Simulations
Dynamic Load Balancing in Distributed Content-based Publish/subscribe
Dynamic Load Balancing
Performance Studies of Dynamic Load Balancing in Distributed Systems
Randomised Dynamic Load Balancing
Dynamic Load Balancing on Web Server System