

The Effect of Time Delays on the Stability of Load Balancing Algorithms for Parallel Computations

John Chiasson, *Senior Member, IEEE*, Zhong Tang, *Member, IEEE*, Jean Ghanem, Chaouki T. Abdallah, *Senior Member, IEEE*, J. Douglas Birdwell, *Fellow, IEEE*, Majeed M. Hayat, *Senior Member, IEEE*, and Henry Jérez

Abstract—A deterministic dynamic nonlinear time-delay system is developed to model load balancing in a cluster of computer nodes used for parallel computations. The model is shown to be self consistent in that the queue lengths cannot go negative and the total number of tasks in all the queues and the network are conserved (i.e., load balancing can neither create nor lose tasks). Further, it is shown that using the proposed load balancing algorithms, the system is stable in the sense of Lyapunov. Experimental results are presented and compared with the predicted results from the analytical model. In particular, simulations of the models are compared with an experimental implementation of the load balancing algorithm on a distributed computing network.

Index Terms—Computer networks, load balancing, time delay systems.

I. INTRODUCTION

DISTRIBUTED computing architectures utilize a set of computational elements (CEs) to achieve performance that is not attainable on a single CE. A common architecture is the cluster of otherwise independent computers communicating through a shared network. To make use of parallel computing resources, problems must be broken down into smaller units that can be solved individually by each CE while exchanging information with CEs solving other problems. For a background on mathematical treatments of load balancing, the reader is referred to [1]–[3]. For example, The Federal Bureau of Investigation (FBI) National DNA Index System

(NDIS) and Combined DNA Index System (CODIS) software are candidates for parallelization. New methods developed by Wang *et al.* [4]–[7] lead naturally to a parallel decomposition of the DNA database search problem while providing orders of magnitude improvements in performance over the current release of the CODIS software. In this type of application, the search itself, initiated on any particular node, can initiate subsequent new searches which are added to the node's queue. Consequently, it is of great advantage to the overall system to carry out load balancing to make effective use of the overall computational resources. The projected growth of the NDIS database and the demand for searches of the database can be met by migration to a parallel computing platform.

Effective utilization of a parallel computer architecture requires the computational load to be distributed, more or less, evenly over the available CEs. The qualifier “more or less” is used because the communications required to distribute the load consume both computational resources and network bandwidth. A point of diminishing returns therefore exists.

Distribution of computational load across available resources is referred to as the *load balancing* problem in the literature. Various taxonomies of load balancing algorithms exist [8]. Direct methods examine the global distribution of computational load and assign portions of the workload to resources before processing begins. Iterative methods examine the progress of the computation and the expected utilization of resources, and adjust the workload assignments periodically as computation progresses. Assignment may be either deterministic, as with the dimension exchange/diffusion [9] and gradient methods, stochastic, or optimization based. A comparison of several deterministic methods is provided by Willebeek-LeMair and Reeves [10]. Here, a deterministic model is developed.

The present work focuses upon the effects of delays in the exchange of information among CEs, and the constraints these effects impose on the design of a load balancing strategy. Motivated by the authors' previous work in [11] and [12], a new nonlinear model is developed here (see also [13]). Specifically, a deterministic dynamic nonlinear time-delay system is developed to model load balancing. The model is shown to be self consistent in that the queue lengths cannot go negative and that the total number of tasks in all the queues and the network is conserved (i.e., load balancing can neither create nor lose tasks). Further, it is shown that the controller proposed here is asymptotically stable in the sense of Lyapunov. Simulations of the nonlinear model are compared with an *experimental implementa-*

Manuscript received April 11, 2004. Manuscript received in final form May 25, 2005. Recommended by Associate Editor J. Sarangapani. The work of C. T. Abdallah, J. Ghanem, and M. M. Hayat was supported in part by the National Science Foundation under Information Technology Grants ANI-0312611 and INT-9818312. The work of J. D. Birdwell, J. Chiasson, and Z. Tang was supported in part by the National Science Foundation under ITR Grant ANI-0312182 and in part by U.S. Department of Justice, Federal Bureau of Investigation under Contract J-FBI-98-083. J. D. Birdwell and J. Chiasson were also supported in part by a Challenge Grant Award from the Center for Information Technology Research at the University of Tennessee. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Government.

J. Chiasson, Z. Tang, and J. D. Birdwell, are with the Electrical and Computer Engineering Department, University of Tennessee, Knoxville, TN 37996 USA (e-mail: chiasson@utk.edu; ztang@utk.edu; birdwell@utk.edu).

J. Ghanem was with the Electrical and Computer Engineering Department, University of New Mexico, Albuquerque, NM 87131 USA. He is now with Verizon Communications.

C. T. Abdallah, M. M. Hayat, and H. Jérez are with the Electrical and Computer Engineering Department, University of New Mexico, Albuquerque, NM 87131 USA (e-mail: chaouki@ece.unm.edu; hayat@ece.unm.edu; hjerez@ece.unm.edu).

Digital Object Identifier 10.1109/TCST.2005.854339

tion of the load balancing algorithm performed on a distributed computing network.

Section II presents our approach to modeling the computer network and load balancing algorithm to incorporate the presence of delay in the communication between nodes and task transfers. In Section III, we show that the model captures the nonnegativity of the queue lengths as well as the fact that the totality of tasks in all the queues and in transit is conserved by the load balancing algorithm. Section IV shows that the system is asymptotically stable in the sense of Lyapunov for any choice of positive gains in the load balancing algorithm (controller). Section V presents simulations of the nonlinear models for comparison with the actual experimental data. Section VI presents experimental results from an implementation of the load balancing controller on a parallel computer consisting of a networked cluster of nodes. Section VII presents experiments conducted over a geographically-dispersed distributed environment (i.e., PlanetLab). Both the effects of the network delays and the variances in the task processing time on the behavior of the system are assessed. Finally, Section VIII is a summary and conclusion of the present work and a discussion of future work.

II. MATHEMATICAL MODEL

Continuous time models are developed in this section that model the load balancing dynamics among a network of computers. To introduce the approach, consider a computing network consisting of n computers (nodes) all of which can communicate with each other. At start up, the computers are assigned an equal number of tasks each of which has essentially the same processing time (homogenous tasks). However, in some applications when a node executes a particular task it can, in turn, generate more tasks so that very quickly the loads on various nodes become unequal. To balance the loads, each computer in the network sends (broadcasts) its queue size $q_j(t)$ to all other computers in the network. A node i receives this information from node j delayed by a finite amount of time τ_{ij} (with the convention $\tau_{ii} = 0$); that is, it receives $q_j(t - \tau_{ij})$. Each node i then uses this information to compute its estimate of the network average of the number of tasks in all n queues of the network. Based on the most recent observations, the simple (local) estimate of the network average is computed by the i th node as

$$q_i^{avg} \triangleq \frac{\sum_{j=1}^n q_j(t - \tau_{ij})}{n}.$$

Node i then compares its queue size $q_i(t)$ with its estimate of the network average by estimating its excess load, $q_i(t) - (\sum_{j=1}^n q_j(t - \tau_{ij}))/n$. If its excess load is greater than zero or some positive threshold, the node sends some of its tasks to the other nodes. If it is less than zero, no tasks are sent. Further, the tasks sent by node i are received by node j with a delay h_{ij} . The controller (load balancing algorithm) decides how often to do load balancing (transfer tasks among the nodes) and how many tasks are to be sent to each node.

The mathematical model of the task load dynamics at a given computing node is given by

$$\begin{aligned} \frac{d}{dt}x_i(t) &= \lambda_i - \mu_i + u_i(t) - \sum_{j=1}^n p_{ij} \frac{t_{p_i}}{t_{p_j}} u_j(t - h_{ij}) \\ y_i(t) &= x_i(t) - \frac{\sum_{j=1}^n x_j(t - \tau_{ij})}{n} \\ u_i(t) &= -K_i \text{sat}(y_i(t)) \end{aligned} \quad (1)$$

where $p_{ij} \geq 0$, $p_{jj} = 0$, satisfy $\sum_{i=1}^n p_{ij} = 1$, and

$$\text{sat}(y) = \begin{cases} y_{\max} & \text{if } y > y_{\max} \\ y & \text{if } 0 \leq y \leq y_{\max} \\ 0 & \text{if } y < 0. \end{cases}$$

Further, in this model, we define

- $x_i(t)$ is the expected waiting time experienced by a task inserted into the queue of the i th node. With t_{p_i} as the average time needed to process a task on the i th node, the expected (average) waiting time is given by $x_i(t) = q_i(t)t_{p_i}$. Note that $x_j/t_{p_j} = q_j$ is the number of tasks in the queue of node j . If these tasks were transferred to node i , then the waiting time transferred is $q_j t_{p_i} = x_j t_{p_i}/t_{p_j}$, so that the fraction t_{p_i}/t_{p_j} converts waiting time on node j to waiting time on node i .
- $\lambda_i \geq 0$ is the rate of generation of waiting times on the i th node caused by the addition of tasks (rate of increase in x_i).
- $\mu_i \geq 0$ is the rate of reduction in waiting time caused by the service of tasks at the i th node and is given by $\mu_i \equiv (1 \times t_{p_i})/t_{p_i} = 1$ for all i if $x_i(t) > 0$, while if $x_i(t) = 0$ then $\mu_i \triangleq 0$; that is, if there are no tasks in the queue, then the queue cannot possibly decrease.
- $u_i(t)$ is the rate of removal (transfer) of the tasks from node i at time t by the load balancing algorithm at node i . Note that $u_i(t) \leq 0$.
- p_{ij} is the fraction of the j th node's tasks to be sent out that it sends to the i th node. In more detail, $p_{ij}u_j(t)$ is the rate at which node j sends waiting time (tasks) to node i at time t where, as all the tasks must go to some node, one requires that $p_{ij} \geq 0$, $\sum_{i=1}^n p_{ij} = 1$ and $p_{jj} = 0$. That is, the transfer from node j of expected waiting time (tasks) $\int_{t_1}^{t_2} u_j(t)dt$ in the interval of time $[t_1, t_2]$ to the other nodes is carried out with the i th node receiving the fraction $p_{ij}(t_{p_i}/t_{p_j}) \int_{t_1}^{t_2} u_j(t)dt$, where the ratio t_{p_i}/t_{p_j} converts the task from waiting time on node j to waiting time on node i . As $\sum_{i=1}^n p_{ij} \int_{t_1}^{t_2} u_j(t)dt = \int_{t_1}^{t_2} u_j(t)dt$, this results in removing *all* of the excess waiting time $\int_{t_1}^{t_2} u_j(t)dt$ from node j .
- The quantity $-p_{ij}u_j(t - h_{ij})$ is the rate of increase (rate of transfer) of the expected waiting time (tasks) at time t from node j by (to) node i where h_{ij} ($h_{ii} = 0$) is the time delay for the task transfer from node j to node i .

In this model, all rates are in units of the *rate of change of expected waiting time*, or *time/time* which is dimensionless. This normalization of the queue length (i.e., $x_i = q_i/t_{p_i}$) at each node by the local average processing time (t_{p_i}) is simply a way to account for unequal task processing rates by each node. As $u_i(t) \leq 0$, node i can only send tasks to other nodes and cannot

initiate transfers from another node to itself. A delay is experienced by transmitted tasks before they are received at the other node. The control law $u_i(t) = -K_i \text{sat}(y_i(t))$ states that if the i th node output $x_i(t)$ is above its estimate of the network average $x_i^{avg} = (\sum_{j=1}^n x_j(t - \tau_{ij}))/n$, then it sends data to the other nodes, while if it is less than this average nothing is sent. The j th node receives the fraction $\int_{t_1}^{t_2} p_{ji}(t_{p_i}/t_{p_j})u_i(t)dt$ of transferred waiting time $\int_{t_1}^{t_2} u_i(t)dt$ delayed by the time h_{ij} .

A. Specification of the Factors p_{ij}

The model described in (1) is the basic model, but an important detail remains unspecified, namely the exact form of the p_{ij} for each sending node i . One approach is to choose them as constant and equal, that is, $p_{ij} = 1/(n-1)$ for $j \neq i$. Another approach is to use the local information of the waiting times $x_i(t)$, $i = 1, \dots, n$ to set their values. The quantity $x_i(t - \tau_{ji}) - x_j^{avg}$ is node j 's estimate of the excess (or deficit) waiting time in the queue of node i with respect to the local average of node j . If node i 's queue is above the local average, then node j does not send tasks to it. Therefore, $\text{sat}(x_j^{avg} - x_i(t - \tau_{ji}))$ is a measure by node j as to how much node i is *below* the local average. Node j performs this computation for all the other nodes and then portions out its tasks among the other nodes according to the amounts they are below the local average, that is

$$p_{ij} \triangleq \frac{\text{sat}(x_j^{avg} - x_i(t - \tau_{ji}))}{\sum_{i \neq j} \text{sat}(x_j^{avg} - x_i(t - \tau_{ji}))}. \quad (2)$$

If the denominator $\sum_{i \neq j} \text{sat}(x_j^{avg} - x_i(t - \tau_{ji})) = 0$, then the p_{ij} are defined to be zero and no load is transferred. This is illustrated in Fig. 1 for node 1.

Remark: If the denominator

$$\sum_{i \neq j} \text{sat}(x_j^{avg} - x_i(t - \tau_{ji}))$$

is zero, then $x_j^{avg} - x_i(t - \tau_{ji}) \leq 0$ for all $i \neq j$. However, by definition of the average

$$\begin{aligned} \sum_{i \neq j} (x_j^{avg} - x_i(t - \tau_{ji})) + x_j^{avg} - x_j(t) \\ = \sum_i (x_j^{avg} - x_i(t - \tau_{ji})) = 0 \end{aligned}$$

which implies

$$x_j^{avg} - x_j(t) = - \sum_{i \neq j} (x_j^{avg} - x_i(t - \tau_{ji})) > 0.$$

That is, if the denominator is zero, the node j is not greater than the local average, so $u_j(t) = -K_j \text{sat}(y_j(t)) = 0$ and is therefore not sending out any tasks.

III. MODEL CONSISTENCY

It is now shown that the model is consistent with actual working systems in that the queue lengths cannot go negative, and the load balancing algorithm cannot create or lose tasks; it can only move tasks between nodes [13], [14].




Fig. 1. Illustration of a hypothetical distribution p_{i1} of the load at some time t from node 1's point of view. Node 1 will send data out to node i in proportion p_{i1} it estimates node i is below the average where $\sum_{i=1}^n p_{i1} = 1$ and $p_{11} = 0$.

A. Nonnegativity of the Queue Lengths

To show the nonnegativity of the queue lengths, recall that the queue length of each node is given by $q_i(t) = x_i(t)/t_{p_i}$. The model is rewritten in terms of these quantities as

$$\frac{d}{dt} \left(\frac{x_i(t)}{t_{p_i}} \right) = \frac{\lambda_i - \mu_i}{t_{p_i}} + \frac{1}{t_{p_i}} u_i(t) - \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t - h_{ij}). \quad (3)$$

Given that $x_i(0) > 0$ for all i , it follows from the right-hand side of (3) that $q_i(t) = x_i(t)/t_{p_i} \geq 0$ for all $t \geq 0$ and all i . To see this, suppose without loss of generality that $q_i(t) = x_i(t)/t_{p_i}$ is the first queue to go to zero, and let t_1 be the time when $x_i(t_1) = 0$. At the time t_1 , $\lambda_i - \mu_i = \lambda_i \geq 0$ by the definition of μ_i and $-\sum_{j=1}^n (p_{ij}/t_{p_j})u_j(t - h_{ij}) \geq 0$ for all time by the definition of the u_j . Further, the term $u_i(t_1)$ is negative only if

$$x_i(t_1) > \frac{\left(\sum_{j=1}^n x_j(t_1 - \tau_{ij}) \right)}{n}. \quad (4)$$

By supposition (up to time t_1) all the $x_j(t_1 - \tau_{ij}) > 0$ for $j \neq i$ and $x_i(t_1) = 0$ so that $u_i(t_1) = 0$ as the right side of (4) is positive at time t_1 . Consequently, at time t_1 all terms on the right-hand side of (3) are nonnegative. Further, $x_i(t)$ cannot go negative in a neighborhood of t_1 . For if it did, as the right-hand side of (4) is continuous, it follows that:

$$x_i(t) < 0 < \frac{\left(\sum_{j=1}^n x_j(t - \tau_{ij}) \right)}{n} \quad (5)$$

for some $t \in (t_1, t_1 + \delta)$ with $\delta > 0$. Therefore, $u_i(t) = 0$ for all $t \in [t_1, t_1 + \delta]$ and the right-hand side of (3) is nonnegative for all $t \in [t_1, t_1 + \delta]$ which contradicts $x_i(t) < 0$. Note that $t_1 + \delta$ can be taken to be at least as large as the time at which some x_k goes to zero, that is, $q_k(t_1 + \delta) = 0$ as the right-hand side of (5) must remain positive for $t \in [t_1, t_1 + \delta]$.

If $x_i(t)$ goes positive after t_1 , then the previous argument is repeated at the next time a queue goes to zero. If $x_i(t)$ remains identically zero in the interval $(t_1, t_1 + \delta)$, then the argument is also similar in that at time $t_1 + \delta$, both $x_i(t_1 + \delta)$, $x_k(t_1 + \delta)$ are then zero. As the remaining $n - 2$ nodes are still positive, the right-hand side of (5) continues to hold with both x_i and x_k zero at time $t_1 + \delta$ and one again gets a contradiction if either x_i or x_k goes negative in an interval $(t_1 + \delta, t_2)$. Continuing in this manner, it follows that $q_i(t) = x_i(t)/t_{p_i}$ cannot go negative for all i .

B. Conservation of Queue Lengths

It is now shown that the total number of tasks in all the queues and the network are conserved. To do so, sum up (3) from $i = 1, \dots, n$ to obtain

$$\frac{d}{dt} \left(\sum_{i=1}^n q_i(t) \right) = \sum_{i=1}^n \left(\frac{\lambda_i - \mu_i}{t_{p_i}} \right) + \sum_{i=1}^n u_i(t)/t_{p_i} - \sum_{i=1}^n \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t - h_{ij}) \quad (6)$$

which is the rate of change of the total queue lengths on all the nodes. However, the network itself also contains tasks in transit between nodes. The dynamic model of the queue lengths in the network is given by

$$\frac{d}{dt} q_{\text{net}_i}(t) = \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t - h_{ij}) - \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t). \quad (7)$$

Here q_{net_i} is the number of tasks put on the network that are being sent to node i . This equation simply says that the j th node is putting tasks on the network to be sent to node i at the rate $(p_{ij}/t_{p_j})u_j(t)$ while the i th node is taking these tasks from node j off the network at the rate $(p_{ij}/t_{p_j})u_j(t - h_{ij})$. Summing (7) over all the nodes, one obtains

$$\begin{aligned} \frac{d}{dt} \left(\sum_{i=1}^n q_{\text{net}_i}(t) \right) &= \sum_{i=1}^n \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t - h_{ij}) - \sum_{i=1}^n \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t) \\ &= \sum_{i=1}^n \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} u_j(t - h_{ij}) - \sum_{j=1}^n \frac{u_j(t)}{t_{p_j}}. \end{aligned} \quad (8)$$

Adding (6) and (8), one obtains the conservation of queue lengths given by

$$\frac{d}{dt} \sum_{i=1}^n (q_i(t) + q_{\text{net}_i}(t)) = \sum_{i=1}^n \left(\frac{\lambda_i - \mu_i}{t_{p_i}} \right). \quad (9)$$

In words, the total number of tasks which are in the system (i.e., in the nodes and/or in the network) can increase only by the rate of arrival of tasks $\sum_{i=1}^n \lambda_i/t_{p_i}$ at all the nodes, or similarly, decrease by the rate of processing of tasks $\sum_{i=1}^n \mu_i/t_{p_i}$ at all the nodes. The load balancing itself cannot increase or decrease the total number of tasks in all the queues.

IV. STABILITY OF THE CONTROLLER

The controller in the model (1) is

$$u_i(t) = -K_i \text{sat}(y_i(t))$$

where the gains $K_i > 0$ are to be specified. Physically, these gains are limited by the bandwidth constraints in the network. One can also view the p_{ij} as controller parameters to be specified subject to the constraints given previously.

Interestingly, it turns out that the system (1) is asymptotically stable in the sense of Lyapunov for any set of gains $K_i > 0$ and any set of $p_{ij} \geq 0$ with $\sum_{i=1}^n p_{ij} = 1$. Specifically, we have the following theorem.

Theorem: Given the system described by (1) and (7) with $\lambda_i = 0$ for $i = 1, \dots, n$ and initial conditions $x_i(0) \geq 0$, then $(q_i(t), q_{\text{net}_i}(t)) \rightarrow 0$ as $t \rightarrow \infty$.

Proof: First note that the q_{net_i} are nonnegative since by (7), it follows that

$$q_{\text{net}_i}(t) = - \sum_{j=1}^n \frac{p_{ij}}{t_{p_j}} \left(\int_{t-h_{ij}}^t u_j(\tau) d\tau \right) \geq 0. \quad (10)$$

Under the conditions of the theorem, (9) becomes

$$\frac{d}{dt} \sum_{i=1}^n (q_i(t) + q_{\text{net}_i}(t)) = - \sum_{i=1}^n \frac{\mu_i(q_i)}{t_{p_i}}. \quad (11)$$

Let $V(t) \triangleq \sum_{i=1}^n (q_i(t) + q_{\text{net}_i}(t))$ and, as the $q_i(t)$, $q_{\text{net}_i}(t)$ are nonnegative, $V(t) \geq 0$ and is equal to zero if and only if $q_i(t) = q_{\text{net}_i}(t) = 0$ for every i . Further, as $\mu_i(q_i(t)) = 1$ for $q_i(t) > 0$ and $\mu_i(q_i(t)) = 0$ if only if $q_i(t) = 0$, it follows that $dV/dt = - \sum_{i=1}^n \mu_i(q_i(t))/t_{p_i} \leq 0$. This then implies that

$$V(t) = V(0) - \int_0^t \sum_{i=1}^n \frac{\mu_i(q_i(t))}{t_{p_i}} dt \geq 0 \quad (12)$$

is monotonically decreasing. As $V(t)$ is bounded below, we have $V(t) \downarrow V_f \geq 0$, or

$$\lim_{t \rightarrow \infty} \int_0^t \sum_{i=1}^n \frac{\mu_i}{t_{p_i}} dt = V(0) - V_f \geq 0. \quad (13)$$

The quantity $\mu_i(q_i(t))$ is either 1 or 0 depending on whether $q_i(t)$ is positive or zero, so $\mu_i(q_i(t))$ can be viewed as a set of pulses of unit height and varying width. The integral $\int_0^\infty \mu_i(q_i(t)) dt$ is finite by (13) which implies that the widths of the unit-height pulses making up $\mu_i(q_i(t))$ must go to zero as $t \rightarrow \infty$. So, even if a $q_i(t)$ ($= x_i(t)/t_{p_i}$) continues to switch between zero and positive values, the time intervals for which it is nonzero must go to zero as $t \rightarrow \infty$. Summarizing, the $q_i(t)$ are nonnegative, continuous functions, bounded by the nonnegative monotonically decreasing function $V(t)$, and the intervals for which the $q_i(t)$ are nonzero goes to zero as $t \rightarrow \infty$. More precisely, let $I_{t,h} = [t - h, t]$ and if we define $E_{t,h} = \{s \in I_{t,h} : q_i(s) > 0\}$, then the Lebesgue measure of

$E_{t,h}$, denoted by $m(E_{t,h})$, converges to 0 as $t \rightarrow \infty$ for every h . Further, as $u_i(t) \leq 0$ is always true and

$$\begin{aligned} u_i(t) &= -K_i \text{sat} \left(x_i(t) - \frac{\sum_{j=1}^n x_j(t - \tau_{ij})}{n} \right) \\ &= -K_i \text{sat} \left(t_{p_i} q_i(t) - \frac{\sum_{j=1}^n \left(\frac{t_{p_i}}{t_{p_j}} \right) q_j(t - \tau_{ij})}{n} \right) \end{aligned}$$

it follows that the time intervals for which the bounded functions $u_i(t)$ are nonzero must go to zero as $t \rightarrow \infty$. This follows from the observation that $u_i(t) < 0$ necessarily implies $q_i(t) > 0$. Thus, the integral in (10) can be upper bounded by $\int_{[t-h_{ij}, t] \cap E_{t,h_{ij}}} |K_i| y_{\max} d\tau = |K_i| y_{\max} m(E_{t,h_{ij}})$ (recall that $E_{t,h} \subset I_{t,h}$, which converges to zero as $t \rightarrow \infty$). Consequently, by (10), $q_{\text{net}_i}(t) \rightarrow 0$ as $t \rightarrow \infty$.

We now show that the monotonically decreasing function $V(t)$ must go to zero, that is, $\lim_{t \rightarrow \infty} V(t) = V_f = 0$. Suppose not, so that $V_f > 0$. As $q_{\text{net}_i}(t) \rightarrow 0$, choose t_1 large enough so that $0 \leq \sum_{i=1}^n q_{\text{net}_i}(t) < \epsilon V_f$ for $t > t_1$ where $0 < \epsilon < 1$. Since

$$V(t) = \sum_{i=1}^n (q_i(t) + q_{\text{net}_i}(t)) \geq V_f \text{ for all } t$$

we have

$$\sum_{i=1}^n q_i(t) \geq (1 - \epsilon)V_f > 0 \text{ for } t > t_1.$$

For every $t > t_1$, there exists at least one i (which depends on t) for which $q_i(t) > 0$. Therefore, $\sum_{i=1}^n \mu_i(q_i(t))/t_{p_i} dt \geq \min\{1/t_{p_i}\}$ for all $t > t_1$. By (11), we then have

$$V(t) = V(t_1) - \int_{t_1}^t \sum_{i=1}^n \frac{\mu_i(q_i(t))}{t_{p_i}} dt \leq V(t_1) - \int_{t_1}^t \min\left\{\frac{1}{t_{p_i}}\right\} dt. \quad (14)$$

As the right side of (14) eventually becomes negative, we have a contradiction and therefore $V_f = 0$. As it has already been shown that $q_{\text{net}_i} \rightarrow 0$ for all i , $V(t) \rightarrow 0$ then implies that $q_i(t) \rightarrow 0$ for all i . This completes the proof of the theorem.

V. SIMULATION RESULTS

Experimental procedures to determine the delay values are given in [15] and summarized in [16]. These give representative values for a Fast Ethernet network with three nodes of $\tau_{ij} = \tau = 200 \mu\text{s}$ for $i \neq j$, $\tau_{ii} = 0$, and $h_{ij} = 2\tau = 400 \mu\text{s}$ for $i \neq j$, $h_{ii} = 0$. The initial conditions were $x_1(0) = 0.6$, $x_2(0) = 0.4$ and $x_3(0) = 0.2$. The inputs were set as $\lambda_1 = 3\mu_1$, $\lambda_2 = 0$, $\lambda_3 = 0$, $\mu_1 = \mu_2 = \mu_3 = 1$. The t_{p_i} 's were taken to be equal to $10 \mu\text{s}$.

In this set of simulations, the model (1) is used. Figs. 2 and 3 show the responses with the gains set as $K_1 = K_2 = K_3 = K = 1000$ and as $K_1 = 6667$, $K_2 = 4167$, $K_3 = 5000$, respectively.

VI. EXPERIMENTAL RESULTS

A parallel machine has been built to implement an experimental facility for evaluation of load balancing strategies and




Fig. 2. Constant p_{ij} nonlinear output responses with $K = 1000$.




Fig. 3. Nonlinear simulation with constant p_{ij} and $K_1 = 6666.7$; $K_2 = 4166.7$; $K_3 = 5000$.

parallel databases. A root node communicates with k groups of computer networks. Each of these groups is composed of n nodes (hosts) holding identical copies of a portion of the database. Any pair of groups correspond to different databases, which are not necessarily disjoint. A specific record (or DNA profile in our specific case) is in general stored in two groups for redundancy to protect against failure of a node. Within each node, there are either one or two processors. In the experimental facility, the dual processor machines use 1.4 GHz Athlon MP processors, and the single processor machines use 1.3 GHz Athlon processors. All run the Linux operating system. Our interest here is in the load balancing in any one group of n nodes.

The database is implemented as a set of queues with associated search engine threads, typically assigned one per node of the parallel machine. The search requests are created not only by the database clients; the search process also creates search requests as the index tree is descended by any search thread. This creates the opportunity for parallelism; search requests that

await processing may be placed in any queue associated with a search engine, and the contents of these queues may be moved arbitrarily among the processing nodes of a group to achieve a balance of the load.

An important point is that the actual delays experienced by the network traffic in the parallel machine are *random*. Work has been performed to characterize the bandwidth and delay on unloaded and loaded network switches, in order to identify the delay parameters of the analytic models and is reported in [15] and [16]. The value $\tau = 200 \mu\text{s}$ used for simulations represents an average value for the delay and was found using the procedure described in [16]. The interest here is to compare the experimental data with that from the three models previously developed.

To explain the connection between the control gain K_i and the actual implementation, recall that the waiting time is related to the number of tasks by $x_i(t) = q_i(t)t_{p_i}$ where t_{p_i} is the average time to carry out a task. The continuous time control law is

$$u_i(t) = -K_i \text{sat}(y_i(t))$$

where $u_i(t)$ is the rate of decrease of waiting time $x_i(t)$ per unit time. Consequently, the gain K_i represents the rate of reduction of waiting time per second in the continuous time model. Also, $y_i(t) = (q_i(t) - n^{-1} \sum_{j=1}^n q_j(t - \tau_{ij}))t_{p_i} = r_i(t)t_{p_i}$, where $r_i(t)$ is simply the number of tasks above the estimated (local) average number of tasks and, as the interest here is the case $y_i(t) > 0$, consider $u_i(t) = -K_i y_i(t)$. The implementation must execute the load balancing control law repetitively with a (possibly random) time interval between balancing actions. With Δt the time interval between successive executions of the load balancing algorithm, a discrete time control law is defined that removes a fraction of the queue $K_z r_i(t) (0 < K_z < 1)$ in the time Δt . The rate of reduction of *waiting time* is $-K_z r_i(t)t_{p_i}/\Delta t = -K_z y_i(t)/\Delta t$ so that an equivalent continuous time control law, given the discrete time gain K_z and control interval Δt , is

$$u(t) = -\frac{K_z y_i(t)}{\Delta t} \implies K_i = \frac{K_z}{\Delta t}. \quad (15)$$

This shows that the gain K_i is related to the actual implementation by how fast the load balancing can be carried out and how much (fraction) of the load is transferred. In the experimental work reported here, Δt actually varies each time the load is balanced. As a consequence, the value of Δt used in (15) is an average value for that run. The average time t_{p_i} to process a task is the same on all nodes used for the experiments (identical processors) and is equal $10 \mu\text{s}$, while the time it takes to ready a load for transfer is about $5 \mu\text{s}$. The initial conditions were taken as $q_1(0) = 6000$, $q_2(0) = 4000$, $q_3(0) = 2000$ (corresponding to $x_1(0) = q_1(0)t_{p_i} = 0.06$, $x_2(0) = 0.04$, $x_3(0) = 0.02$). All of the experimental responses were carried out with constant $p_{ij} = 1/2$ for $i \neq j$.

Fig. 4 is a plot of the responses $r_i(t) = q_i(t) - (\sum_{j=1}^n q_j(t - \tau_{ij}))/n$ for $i = 1, 2, 3$ (recall that $y_i(t) = r_i(t)t_{p_i}$). The (average) value of the gains were ($K_z = 0.5$) $K_1 = 0.5/75 \mu\text{s} = 6667$, $K_2 = 0.5/120 \mu\text{s} = 4167$, $K_3 = 0.5/100 \mu\text{s} = 5000$. This figure compares favorably with Fig. 3 except for the time scale being off; that is, the experimental responses are

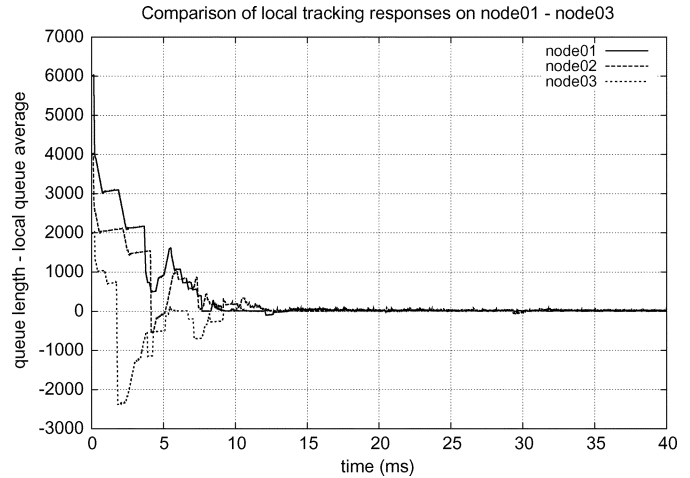


Fig. 4. Experimental response of the load balancing algorithm. The average value of the gains are ($K_z = 0.5$) $K_1 = 6667$, $K_2 = 4167$, $K_3 = 5000$ with constant p_{ij} .

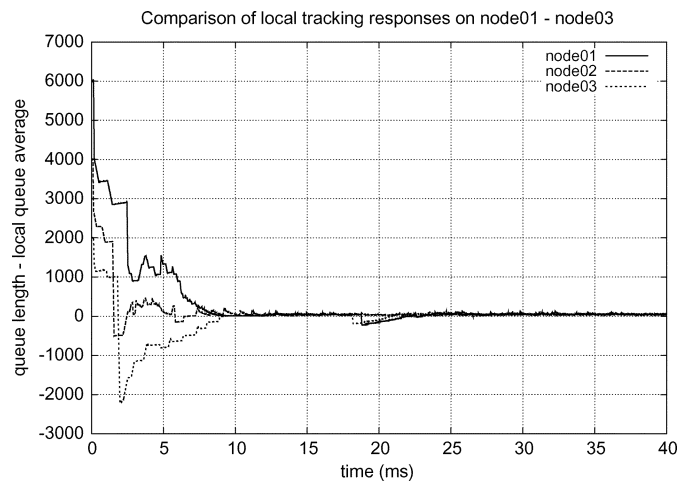


Fig. 5. Experimental response of the load balancing algorithm. The average value of the gains are ($K_z = 0.3$) $K_1 = 2400$, $K_2 = 7273$, $K_3 = 2500$ with constant p_{ij} .

slower. The explanation for this is that the discrete load balancing implementation is *not* accurately modeled in the continuous time simulations, only its average effect is represented in the gains K_i . That is, the continuous time model does not stop processing jobs (at the average rate t_{p_i}) while it is transferring tasks to do the load balancing. Fig. 5 shows the plots of the response for the (average) value of the gains given by ($K_z = 0.3$) $K_1 = 0.3/125 \mu\text{s} = 2400$, $K_2 = 0.3/110 \mu\text{s} = 7273$, $K_3 = 0.3/120 \mu\text{s} = 2500$.

Fig. 6 shows the plots of the response for the (average) value of the gains given by ($K_z = 0.2$) $K_1 = 0.2/125 \mu\text{s} = 1600$, $K_2 = 0.2/80 \mu\text{s} = 2500$, $K_3 = 0.2/70 \mu\text{s} = 2857$. The initial conditions were $q_1(0) = 6000$, $q_2(0) = 4000$, $q_3(0) = 2000$ ($x_1(0) = q_1(0)t_{p_i} = 0.06$, $x_2(0) = 0.04$, $x_3(0) = 0.02$).

Fig. 7 summarizes the data from several experimental runs of the type shown in Figs. 4–6. For $K_z = 0.1, 0.2, 0.3, 0.4, 0.5$, ten runs were made and the settling time (time to load balance) were determined. These are marked as small horizontal ticks on Fig. 7. (For all such runs, the initial queues were the same and equal to $q_1(0) = 600$, $q_2(0) = 400$, $q_3(0) = 200$. For




Fig. 6. Experimental response of the load balancing algorithm. The average value of the gains are ($K_z = 0.2$) $K_1 = 1600$, $K_2 = 2500$, $K_3 = 2857$ with constant p_{ij} .

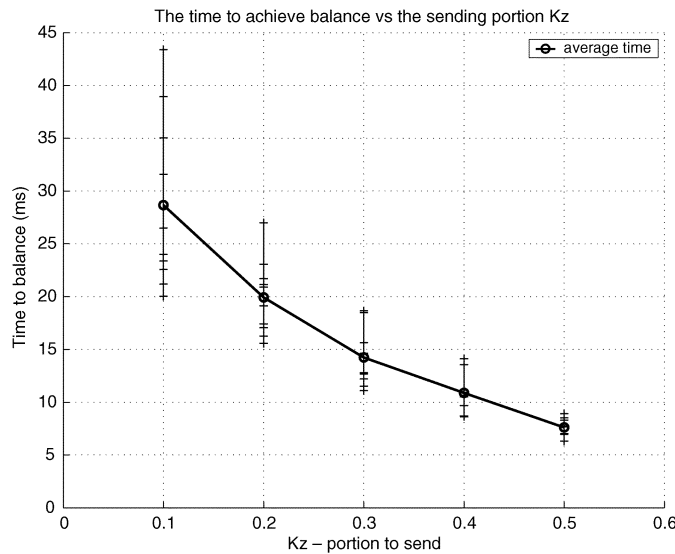


Fig. 7. Summary of the load balance time as a function of the feedback gain K_z .

each value of K_z , the average settling time for these ten runs was computed and is marked as a dot on given on Fig. 7. For values of $K_z = 0.6$ and higher (with increments of 0.1 in K_z), consistent results could not be obtained. In many cases, ringing extended throughout the experiment's time interval (200 ms).

For example, Fig. 8 shows the plots of the queue length less the local queue average for an experimental run with $K_z = 0.6$ where the settling time is approximately 7 ms. In contrast, Fig. 9 shows the experimental results under the same conditions where persistent ringing regenerates for 40 ms. It was found the response was so oscillatory that a settling time was not possible to determine accurately. However, Fig. 7 shows that one desires to choose the gain to be close to 0.5 to achieve a faster response time without breaking into oscillatory behavior.

VII. EXPERIMENTS OVER PLANETLAB

A geographically distributed system was developed to validate the theoretical work for large delays and to assess different




Fig. 8. $K_z = 0.6$ —settling time is approximately 7 ms.




Fig. 9. $K_z = 0.6$ These are the same conditions as Fig. 8, but now the ringing persists.

load balancing policies in a real environment. The system consists of several nodes running the same code. The nodes are part of PlanetLab, a planetary-scale network involving more than 350 nodes positioned around the globe and connected via the Internet (www.planetlab.org). The application used to illustrate

| | node 1 | node 2 | node 3 |
|--|-------------------------------|-----------------|---------------------|
| Location | University of New Mexico (US) | Taipei - Taiwan | Frankfurt - Germany |
| Initial Distribution | 6000 tasks | 4000 tasks | 2000 tasks |
| Average Task Processing Time t_{pi} | | | 10.2 ms |
| Standard Deviation for t_{pi} | | | 2.5 ms |
| Interval between load balancing instances Δt | | | 150 ms |
| Interval between 2 comm. transmissions | | | 50 ms |

Fig. 10. Parameters and settings of the experiment.

| | Roundtrip delay τ_{ij} | Data transmission rate | Average Transmission of one Task |
|---------|-----------------------------|------------------------|----------------------------------|
| n1 - n2 | 215 ms | 1.34 KB/s | 14 ms |
| n1 - n3 | 200 ms | 1.42 KB/s | 16 ms |
| n2 - n3 | 307 ms | 1.03 KB/s | 20 ms |

Fig. 11. Average network delays and transmission rates.

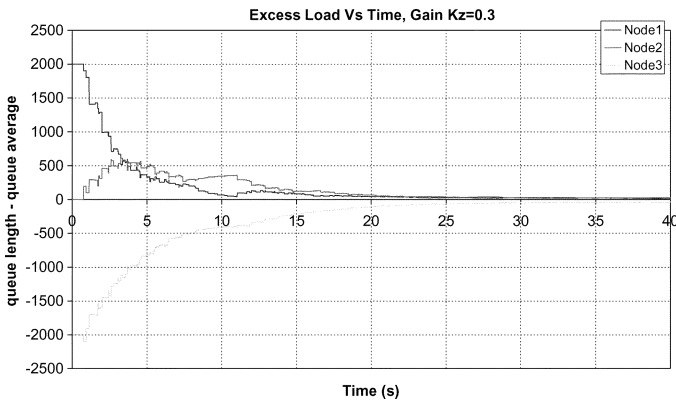


Fig. 12. Experimental response of the load balancing algorithm under large delays. Gain $K_z = 0.3$ and $p_{ij} = 0.5$.

the load balancing process was matrix multiplication, where one task is defined as the multiplication of one row by a static matrix duplicated on all nodes (3 nodes in our experiment). The number of elements in each row were generated randomly from a specified range, which made the execution time of a task variable. The network protocol UDP was used to exchange queue size information among the nodes, and TCP (connection-based) was used to transfer task data from one machine to another.

To match the experimental settings of the previous sections, 3 nodes were used; node1 at the University of New Mexico, node2 in Taipei, Taiwan and node3 in Frankfurt, Germany. The p_{ij} were set to $1/2$ for $i \neq j$. The initial parameters and settings for the experiment are summarized in the Table given in Fig. 10.

Throughout the experiments, network statistics related to transmission rates and delays were collected. The averages of these statistics are shown in the Table given in Fig. 11. Large delays were observed in the network due to the dispersed geographical location of the nodes. Moreover, the transmission rates detected between the nodes were very low mainly because the amount of data exchanged (in bytes) is small. Indeed, the average size of data needed to transmit a single task was 20 bytes, which caused variability in the observed transmission




Fig. 13. Experimental response of the load balancing algorithm under large delays. Gain $K_z = 0.5$ and $p_{ij} = 0.5$.




Fig. 14. Summary of the load balancing time as function of the gain K_z .

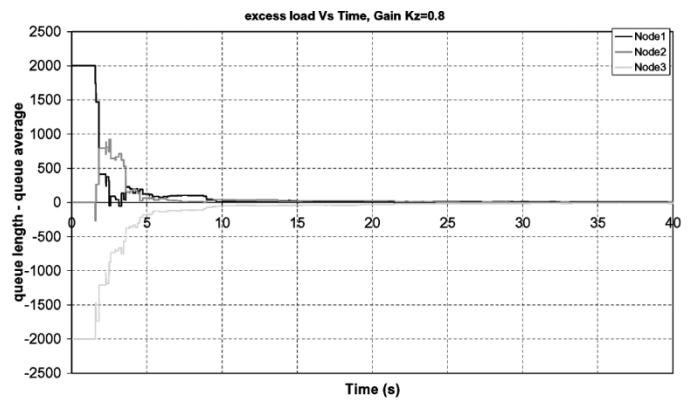


Fig. 15. Experimental response of the load balancing algorithm under large delays. Gain $K_z = 0.8$ and $p_{ij} = 0.5$.

rates due to the large communication delays and their variation. In order to observe the behavior of the system under various gains, several experiments were conducted for different gain values K_z ranging from 0.1 to 1.0. Fig. 12 is a plot of the responses $r_i(t)$ corresponding to each node i where the gain K_z was set to 0.3. Similarly, Fig. 13 shows the system response for gain K_z equal to 0.5. Fig. 14 summarizes several runs corresponding to different gain values. For each $K_z = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7$, ten runs were made, and the settling times (time to load balance) were determined. For gain values higher than 0.8, consistent results could not be obtained. For instance, in most of the runs no settling time could be achieved.

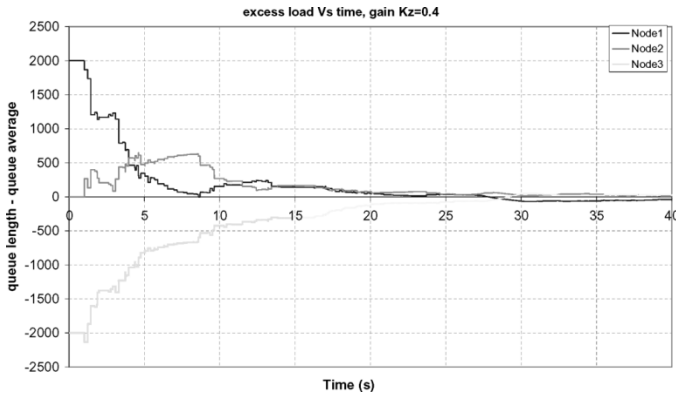


Fig. 16. Experimental response of the load balancing algorithm under large delays. Gain $K_z = 0.4$ and $p_{ij} = 0.5$.

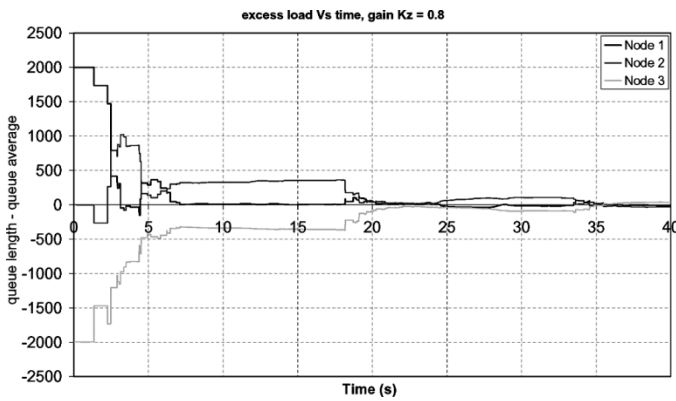


Fig. 17. Experimental response of the load balancing algorithm under large delays. Gain $K_z = 0.8$ and $p_{ij} = 0.5$.

However, when the observed network delays were stable, the system response was steady and converged quickly to a balanced state when K_z was equal to 0.8 (Fig. 15). As previously mentioned, this scenario was not frequently seen. The system behaviors in these experiments do not exactly match, and the results obtained in the previous sections. This is due to the difference in network topology and delays (and very likely due to the random nature of the delay [17], [18]). For instance, the ratio between the average delay and the task process time is 20 ($200 \mu\text{s}/10 \mu\text{s}$) for the local area network (LAN) setting and 12 ($120 \text{ms}/10 \text{ms}$) for the distributed setting. This fact is one of the reasons ringing is observed earlier (for $K_z = 0.6$) in the LAN experiment whereas under PlanetLab the ringing responses were observed starting at $K_z = 0.8$. Interestingly, we have previously observed a similar behavior, using a Monte Carlo simulation of a 3-node distributed system with *random* delays in [17] as well as experiments on a wireless LAN in [18].

The previous experiments have been conducted under normal network conditions stated in the Table given in Fig. 11. However, another set of experiments were carried out where the network condition worsens and larger delays were observed. In particular, the data transmission rate between node 2 (Taiwan) and node 3 (Germany) dropped from 1.03 KB/s to 407 KB/s. Figs. 16 and 17 show the system responses for gains $K_z = 0.4$ and $K_z = 0.8$, respectively. These experiments clearly show the negative effect of the delay on the stability of the system.




Fig. 18. Experimental response of the load balancing algorithm under large variance in the tasks processing time. Gain $K_z = 0.3$ and $p_{ij} = 0.5$.




Fig. 19. Experimental response of the load balancing algorithm under large variance in the tasks processing time. Gain $K_z = 0.8$ and $p_{ij} = 0.5$.

Nevertheless, we see that with a low gain ($K_z = 0.4$), the settling time is approximately 22 ms. On the other hand, when the gain was set to 0.8, the system did not reach an equilibrium as shown by the nodes' responses $r_i(t)$ in Fig. 17. At this point, only the effect of delay on the stability of the system was tested. In order to study the effect of the variability of the task processing time on system behavior, the matrix multiplication application was adjusted in a way to obtain the following results; the average task processing time was kept at 10.2 ms, but the standard deviation became 7.15 ms instead of 2.5 ms. Figs. 18 and 19 show the respective system responses for gains $K_z = 0.3$ and $K_z = 0.8$. Comparing Figs. 12 and 18, we can see that in the latter case, some ringing persists and the system response did not completely settle out. On the other hand, setting the gain K_z to 0.8 led the system to accommodate the variances in the task processing time.

The experiments presented in this section support the ones reported in the previous section where a local area network was used. In particular, high gains were shown to lead to persistent ringing. Conversely, systems with low gain values lead to slow responses in the load balancing.

VIII. SUMMARY AND CONCLUSION

A load balancing algorithm was modeled as a nonlinear time-delay system. The model was shown to be consistent in that the total number of tasks was conserved and the queues were always

nonnegative. Further, the system was shown to be always stable, but the delays do create a limit on the size of the controller gains in order to ensure performance (fast enough response without oscillatory behavior). Experiments demonstrated a correlation of the continuous time model with the actual implementation. Future work will consider the fact that the load balancing operation involves processor time which is not being used to process tasks. There is a tradeoff between using processor time/network bandwidth and the advantage of distributing the load evenly between the nodes to reduce overall processing time, which has not been fully captured in the present work.

REFERENCES

[1] E. Altman and H. Kameda, "Equilibria for multiclass routing in multi-agent networks," in *Proc. 40th IEEE Conf. Decision and Control*, Orlando, FL, Dec. 2001, pp. 604–609.

[2] C. K. Hisao Kameda, J. Li, and Y. Zhang, *Optimal Load Balancing in Distributed Computer Systems*. New York: Springer-Verlag, 1997.

[3] H. Kameda, I. R. El-Zoghdy Said Fathy, and J. Li, "A performance comparison of dynamic versus static load balancing policies in a mainframe," in *Proc. 39th IEEE Conf. Decision and Control*, Sydney, Australia, Dec. 2000, pp. 1415–1420.

[4] J. D. Birdwell, R. D. Horn, D. J. Icove, T. W. Wang, P. Yadav, and S. Niezgod, "A hierarchical database design and search method for CODIS," in *Proc. 10th Int. Symp. Human Identification*, Orlando, FL, Sep. 1999.

[5] J. D. Birdwell, T. W. Wang, R. D. Horn, P. Yadav, and D. J. Icove, "Method of indexed storage and retrieval of multidimensional information," in *Proc. 10th SIAM Conf. Parallel Processing for Scientific Computation*, Sep. 2000.

[6] J. D. Birdwell, T. W. Wang, and M. Rader, "The university of Tennessee's new search engine for CODIS," in *Proc. 6th CODIS Users Conf.*, Arlington, VA, Feb. 2001.

[7] T. W. Wang, J. D. Birdwell, P. Yadav, D. J. Icove, S. Niezgod, and S. Jones, "Natural clustering of DNA/STR profiles," in *Proc. 10th Int. Symp. Human Identification*, Orlando, FL, Sep. 1999.

[8] H. G. Rotithor, "Taxonomy of dynamic task scheduling schemes in distributed computing systems," *Inst. Elect. Eng. Proc. Comput. Dig. Techniques*, vol. 141, no. 1, pp. 1–10, 1994.

[9] A. Corradi, L. Leonardi, and F. Zambonelli, "Diffusive load-balancing policies for dynamic applications," *IEEE Concurrency*, vol. 22, no. 1, pp. 979–993, Jan.–Feb. 1999.

[10] M. H. Willebeek-LeMair and A. P. Reeves, "Strategies for dynamic load balancing on highly parallel computers," *IEEE Trans. Parallel Distrib. Syst.*, vol. 4, no. 9, pp. 979–993, Sep. 1993.

[11] C. T. Abdallah, J. D. Birdwell, J. Chiasson, V. Chupryna, Z. Tang, and T. W. Wang, "Load balancing instabilities due to time delays in parallel computation," in *Proc. 3rd IFAC Conf. Time Delay Systems*, Sante Fe, NM, Dec. 2001, pp. 198–202.

[12] J. D. Birdwell, J. Chiasson, Z. Tang, C. T. Abdallah, M. Hayat, and T. Wang, "Dynamic time delay models for load balancing part I: deterministic models," in *Proc. CNRS-NSF Workshop: Advances in Control of Time-Delay Systems*, Paris, France, Jan. 2003, pp. 355–370.

[13] J. D. Birdwell, J. Chiasson, Z. Tang, C. T. Abdallah, and M. M. Hayat, "The effect of feedback gains on the performance of a load balancing network with time delays," in *Proc. IFAC Workshop on Time-Delay Systems (TDS'03)*, Rocquencourt, France, Sep. 2003, pp. 371–385.

[14] J. D. Birdwell, J. Chiasson, C. T. Abdallah, Z. Tang, N. Alluri, and T. Wang, "The effect of time delays in the stability of load balancing algorithms for parallel computations," in *Proc. 42nd IEEE Conf. Decision and Control*, Maui, HI, Dec. 2003, pp. 582–587.

[15] P. Dasgupta, "Performance evaluation of fast ethernet, ATM and myrinet under PVM," MS thesis, Univ. Tennessee, Knoxville, 2001.

[16] P. Dasgupta, J. D. Birdwell, and T. W. Wang, "Timing and congestion studies under PVM," in *Proc. 10th SIAM Conf. Parallel Processing for Scientific Computation*, Portsmouth, VA, Mar. 2001.

[17] S. Dhakal, B. Paskaleva, M. Hayat, E. Schamiloglu, and C. Abdallah, "Dynamical discrete-time load balancing in distributed systems in the presence of time delays," in *Proc. IEEE Conf. Decision Control*, Maui, HI, Dec. 2003, pp. 5128–5134.

[18] J. Ghanem, S. Dhakal, M. M. Hayat, H. Jerez, C. T. Abdallah, and J. Chiasson, "On load balancing in distributed systems with large time delays: theory and experiments," in *Proc. IEEE Mediterranean Control Conf. (MED'04)*, Kusadasi, Turkey, Jun. 6–9, 2004.



John Chiasson (S'81–M'85–SM'03) received the B.S. degree in mathematics from the University of Arizona, Tucson, the M.S. degree in electrical engineering from Washington State University, Pullman, and the Ph.D. degree in controls from the University of Minnesota, Minneapolis.

He has worked in industry at Boeing Aerospace, Control Data, and ABB Daimler-Benz Transportation. Since 1999, has been on the faculty of Electrical and Computer Engineering at the University of Tennessee, Knoxville.



Zhong Tang (S'99–M'05) received the B.S. and M.S. degrees in automatic control engineering from Huazhong University of Science and Technology, Wuhan, China, in 1994 and 1997, respectively, and the Ph.D. degree in electrical engineering from the University of Tennessee, Knoxville, in 2005.

His research interests include control systems, parallel databases, and distributed computing.



Jean Ghanem received the B.S. degree in computer and communication engineering from the American University of Beirut and the M.S. degree in computer engineering from the University of New Mexico, Albuquerque.

His research focused on the analysis and implementation of load balancing policies in large time delay networks mainly the Internet and wireless network. He has also performed work related to securing 802.11 wireless networks. He has software development experience in network, multithreading,

database and parallel programming. He is currently a solution architect consultant for Verizon Communications.



Chaouki T. Abdallah (M'81–SM'85) received the M.S. and Ph.D. degrees in electrical engineering from the Georgia Institute of Technology, Atlanta, in 1982 and 1988, respectively.

He is currently a Professor, Associate Chair, and the Director of the Graduate Program in the Electrical and Computer Engineering Department, University of New Mexico, Albuquerque. He conducts research and teaches courses in the general area of systems theory with focus on control and communications systems. His research has been

funded by national funding agencies (NSF, AFOSR, NRL), national laboratories (SNL, LANL), and by various companies (Boeing, HP). He has also been active in designing and implementing various international graduate programs with Latin American and European countries. He was a cofounder in 1990 of the ISTECON consortium, which currently includes more than 150 universities in the US, Spain, and Latin America. He has coauthored four books, and more than 150 peer-reviewed papers.

Dr. Abdallah's IEEE professional service credits include serving as an IEEE CSS BOG Appointed Member (2004–2005), as a member of the IEEE CSS Long Range Planning Committee (2004–2005), and as the Program Chair of the IEEE Conference on Decision and Control, Hawaii, 2003. He will also be serving as the General Chair of CDC'08. He is a recipient of the IEEE Millennium medal.



J. Douglas Birdwell (S'73–M'78–SM'85–F'99) received the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978.

He is currently a Professor of Electrical and Computer Engineering at the University of Tennessee (UT), Knoxville. He joined the faculty at UT in 1978, and is Director of the Laboratory for Information Technologies, which develops secure distributed information systems and analysis tools.

His experience includes computer hardware and software applications development, high performance data base design with applications in bioinformatics, parallel computation and load balancing, and artificial intelligence.

Dr. Birdwell has served in a number of positions in the IEEE Control Systems Society, including President (2004), member of the Board of Governors (1990–2001), General Chair of IEEE Conference on Decision and Control (1998), and Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL.



Majeed M. Hayat (S'89–M'92–SM'00) was born in Kuwait in 1963. He received the B.S. degree (*summa cum laude*) in 1985 in electrical engineering from the University of the Pacific, Stockton, CA, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Wisconsin-Madison, in 1988 and 1992, respectively.

From 1993 to 1996, he worked at the University of Wisconsin-Madison as a Research Associate and co-principal investigator of a project on statistical minefield modeling and detection, which was

funded by the Office of Naval Research. In 1996, he joined the faculty of the Electro-Optics Graduate Program and the Department of Electrical and Computer Engineering, University of Dayton, OH. He is currently an Associate Professor in the Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque. His research contributions cover a broad range of topics in statistical communication theory, optoelectronics, signal processing, and applied probability theory including avalanche photodiodes, optical communication systems, image processing, models for distributed computing, as well as infrared and spectral imaging.

Dr. Hayat was a recipient of a 1998 National Science Foundation Early Faculty Career Award. He is a member of SPIE and OSA.



Henry Jérez received the B.Sc. degree in systems engineering from the Private University of Bolivia and the M.Sc. and Ph.D. degrees in computer engineering from the University of New Mexico, Albuquerque.

He is a Research Scientist with the Corporation for National Research Initiatives, Reston, VA. His research in the area of distributed digital object repositories, load balancing, and wireless communications has been funded by AISTI, Hewlett Packard, Microsoft and Sun Microsystems Latin America. His primary work has been in the area of digital libraries

and distributed digital object repositories as a member of the Prototyping team for the Research library at Los Alamos National Laboratory working on LANL's new Digital Object Storage Architecture. He has served as faculty member, conference reviewer, enterprise consultant and international advisor to several companies and universities across Latin and North America; and is part of the advisory board for the Digital Library Linkages Initiative from the Ibero American Science and Technology Education Consortium ISTEAC.