Dynamic Time Delay Models for Load Balancing Part I: Deterministic Models

J. Douglas Birdwell[1], John Chiasson[1], Zhong Tang[1], Chaouki Abdallah[2], Majeed M. Hayat[2], and Tsewei Wang[3]

[1] ECE Dept, University of Tennessee, Knoxville TN 37996, USA
   {birdwell,chiasson,ztang}@utk.edu
[2] ECE Dept, University of New Mexico, Albuquerque NM 87131-1356, USA
   {chaouki,hayat}@eece.unm.edu
[3] ChE Dept, University of Tennessee, Knoxville TN 37996, USA
   twang@utk.edu

Summary. Parallel computer architectures utilize a set of computational elements (CE) to achieve performance that is not attainable on a single processor, or CE, computer. 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.

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 exists.

In this work, a nonlinear deterministic dynamic time-delay systems is developed to model load balancing in a cluster of computer nodes used for parallel computations. This model is then compared with an experimental implementation of the load balancing algorithm on a parallel computer network.

1 Introduction

Parallel computer architectures utilize a set of computational elements (CE) to achieve performance that is not attainable on a single processor, or CE, computer. 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.

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. [3][4][5][11] 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. The projected growth of the NDIS database and in the demand for searches of the database necessitates 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 exists. The 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. 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 [6] and gradient methods, stochastic, or optimization based. A comparison of several deterministic methods is provided by Willeback-LeMain and Reeves [12].

To adequately model load balancing problems, several features of the parallel computation environment should be captured (1) The workload awaiting processing at each CE; (2) the relative performances of the CEs; (3) the computational requirements of each workload component; (4) the delays and bandwidth constraints of CEs and network components involved in the exchange of workloads, and (5) the delays imposed by CEs and the network on the exchange of measurements. A queuing theory [9] approach is well-suited to the modelling requirements and has been used in the literature by Spies [10] and others. However, whereas Spies assumes a homogeneous network of CEs and models the queues in detail, the present work generalizes queue length to an expected waiting time, normalizing to account for differences among CEs, and aggregates the behavior of each queue using a continuous state model. 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. Preliminary results by the authors appear in [1] with a stability analysis for a proposed *linear* model given in [2]. Here, a nonlinear model is developed to obtain better fidelity and experimental results are presented and compared to that given by the model.

Section 2 presents our approach to modelling the computer network and load balancing algorithms to incorporate the presence of delay in communicating between nodes and transferring tasks. Section 3 presents simulations of the nonlinear model. Section 4 presents experimental data from an actual implementation of a load balancing algorithm which is compared with the

simulations. Finally, Section 5 is a summary and conclusion of the present work and a discussion of future work.

2 Models of Load Balancing Algorithms

In this section, a continuous time model in the form of a nonlinear delay-differential system of equations is developed to model load balancing among a network of computers. A modification to the model is presented so that the number of tasks a node distributes to the other nodes is based on their relative load levels.

To introduce the basic approach to load balancing, 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. However, 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 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}, that is, it receives $q_j(t-\tau_{ij})$. Each node i then uses this information to compute its local estimate[4] of the average number of tasks in the queues of the n computers in the network. In this work, the simple estimator $\left(\sum_{j=1}^{n} q_j(t - \tau_{ij})\right)/n$ $(\tau_{ii} = 0)$ is used which is based on the most recent observations is used. Node i then compares its queue size $q_i(t)$ with its estimate of the network average as $q_i(t) - \left(\sum_{j=1}^{n} q_j(t - \tau_{ij})\right)/n$ and, if this is greater than zero, the node sends some of its tasks to the other nodes while if it is less than zero, no tasks are sent (see Figure 1). 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 and fast to do load balancing (transfer tasks among the nodes) and how many tasks are to be sent to each node. As just explained, each node controller (load balancing algorithm) has only *delayed* values of the queue lengths of the other nodes, and each transfer of data from one node to another is received only after a finite time delay. An important issue considered here is to study the effect of these delays on system performance. Specifically, the continuous time model developed here represents our effort to capture the effect of the delays in load balancing techniques and were developed so that system theoretic methods could be used to analyze them.

2.1 Basic Model

The basic mathematical model of a given computing node for load balancing is given by

---

[4] It is an estimate because at any time, each node only has the delayed value of the number of tasks in the other nodes.
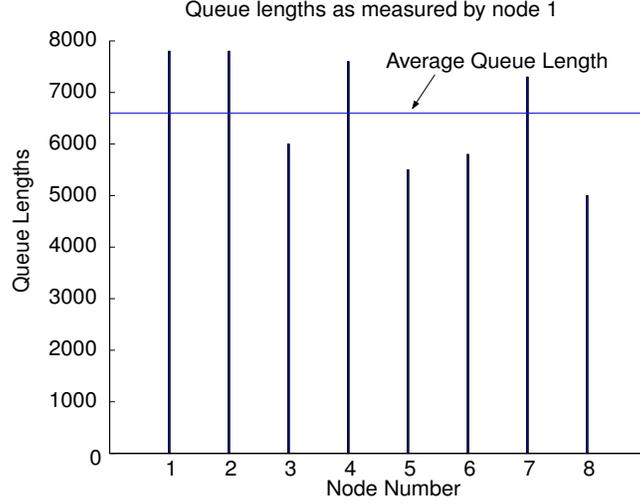
Fig. 1. Graphical description of load balancing. This bar graph shows the load for each computer vs. node of the network. The thin horizontal line is the average load as estimated by node 1. Node 1 will transfer (part of) its load only if it is above its estimate of the average. Also, it only transfers to nodes that it estimates are below the node average.

$$
\begin{aligned}
\frac{dx_i(t)}{dt} &= \lambda_i - \mu_i + u_i(t) - \sum_{j=1}^{n} p_{ij} \frac{t_{p_i}}{t_{p_j}} u_j(t - h_{ij}) \\
y_i(t) &= x_i(t) - \frac{\sum_{j=1}^{n} x_j(t - \tau_{ij})}{n} \\
u_i(t) &= -K_i \mathrm{sat}\left(y_i(t)\right) \\
p_{ij} &\geqslant 0, p_{jj} = 0, \sum_{i=1}^{n} p_{ij} = 1
\end{aligned}
\tag{1}
$$

where

$$
\begin{aligned}
\mathrm{sat}\left(y\right) &= y \text{ if } y \geqslant 0 \\
&= 0 \text{ if } y < 0.
\end{aligned}
$$

In this model we have

- n is the number of nodes.
- $x_i(t)$ is the *expected waiting time* experienced by a task inserted into the queue of the i^{th} node. With $q_i(t)$ the number of *tasks* in the i^{th} node and t_{p_i} the average time needed to process a task on the i^{th} node, the

expected (average) waiting time is then 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 node 1 queue. 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.

- λ_i is the rate of generation of waiting time on the i^{th} node caused by the addition of tasks (rate of increase in x_i)

- μ_i 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.

- $u_i(t)$ 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 $u_i(t) \leq 0$.

- $p_{ij}u_j(t)$ is the rate that node j sends waiting time (tasks) to node i at time t where $p_{ij} \geqslant 0$, $\sum_{i=1}^n p_{ij} = 1$ 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 being sent the fraction $p_{ij}\frac{t_{p_i}}{t_{p_j}} \int_{t_1}^{t_2} u_j(t)dt$ where the fraction t_{p_i}/t_{p_j} converts the task from waiting time on node j to waiting time on node i. As $\sum_{i=1}^n \left(p_{ij} \int_{t_1}^{t_2} u_j(t)dt \right) = \int_{t_1}^{t_2} u_j(t)dt$, this results in a removing *all* the 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.

- The quantities τ_{ij} ($\tau_{ii} = 0$) denote the time delay for communicating the expected waiting time x_j from node j to node i.

- The quantity $x_i^{avg} = \left(\sum_{j=1}^n x_j(t-\tau_{ij}) \right)/n$ is the estimate[5] by the i^{th} node of the average waiting time of the network and is referred to as the *local average* (local estimate of the average).

In this model, all rates are in units of the *rate of change of expected waiting time*, or *time/time* which is dimensionless. 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 $u_i(t) = -K_i \mathrm{sat}(y_i(t))$ states that if the i^{th} node output $x_i(t)$ is above the local average $\left(\sum_{j=1}^n x_j(t-\tau_{ij}) \right)/n$, then it sends data to the other nodes, while if it is less than the local average nothing is sent. The j^{th} node receives the fraction $\int_{t_1}^{t_2} p_{ji}u_i(t)dt$ of transferred waiting time $\int_{t_1}^{t_2} u_i(t)dt$ delayed by the time h_{ij}.

2.2 Constant p_{ij}

The model (1) is the basic model but one important detail remains unspecified, namely the exact form p_{ji} for each sending node i. One approach is to choose

---

[5] This is an only an estimate due to the delays.

them as constant and equal

$$p_{ji} = \frac{1}{n-1}\delta_{ji} \qquad (2)$$

where δ_{ji} is the standard Kronecker delta function. It is clear that $p_{ji} \geqslant 0$, $\sum_{j=1}^{n} p_{ji} = 1$.

Remark If the p_{ij} are specified by (2) and the saturation functions in (1) are removed, the following *linear time invariant* model results

$$\frac{dx_i(t)}{dt} = \lambda_i - \mu_i + u_i(t) - \sum_{j \neq i} pu_j(t - h_{ij})$$

$$y_i(t) = x_i(t) - \frac{\sum_{j=1}^{n} x_j(t - \tau_{ij})}{n} \qquad (3)$$

$$u_i(t) = -K_i y_i(t), \ \ p = \frac{1}{n-1}.$$

When $u_i(t) = -K_i y_i(t) < 0$, this operates as in (1) in that the tasks are immediately removed and sent to the other nodes where each of those nodes experiences a delay (h_{ij}) in getting these tasks. However, a fundamental problem with this linear model is that when $y_i(t) < 0$ the controller (load balancing algorithm) $u_i(t) = -K_i y_i(t) > 0$ so that the node is *instantaneously* taking on waiting time (tasks) from the other nodes before those tasks are removed from the other nodes' queues. That is, it is accepting the waiting times (tasks) $pu_j(t)$ from each of the other nodes. There is a finite time delay associated with this transfer of tasks, and this model ignores this fact. In spite of this fact, it is still of value to consider the system (3) because it can be completely analyzed with regards to stability, and it does capture the oscillatory behavior of the $y_i(t)$. A stability analysis of this linear model is presented in [2].

2.3 Non Constant p_{ij}

It could be useful to use the local information of the waiting times $x_i(t), i = 1, .., n$ to set the values of the p_{ij}. Recall that p_{ij} is the fraction of $u_j(t)$ that node j allocates (transfers) to node i at time t, and conservation of the tasks requires $p_{ij} \geqslant 0, \sum_{i=1}^{n} p_{ij} = 1$ and $p_{jj} = 0$. The quantity $x_i(t - \tau_{ji}) - x_j^{avg}$ represents what node j estimates[6] the waiting time in the queue of node i is with respect to the local average of node j. If queue of node i is above the local average, then node j does not send tasks to it. Therefore $\text{sat}\left(x_j^{avg} - x_i(t - \tau_{ji})\right)$ is an appropriate measure by node j as to how much node i is *below* the local average. Node j then repeats this computation for

---

[6] Again, the term "estimates" is used because node j does not know the current value of $x_i(t)$, but only its earlier value $x_i(t - \tau_{ij})$.

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,

$$p_{ij} = \frac{\text{sat}\left(x_j^{avg} - x_i(t - \tau_{ji})\right)}{\displaystyle\sum_{i \, \ni \, i \neq j} \text{sat}\left(x_j^{avg} - x_i(t - \tau_{ji})\right)}. \tag{4}$$

A p_{ij} is defined to be zero if the denominator $\displaystyle\sum_{i \, \ni \, i \neq j} \text{sat}\left(x_j^{avg} - x_i(t - \tau_{ji})\right) = 0$.
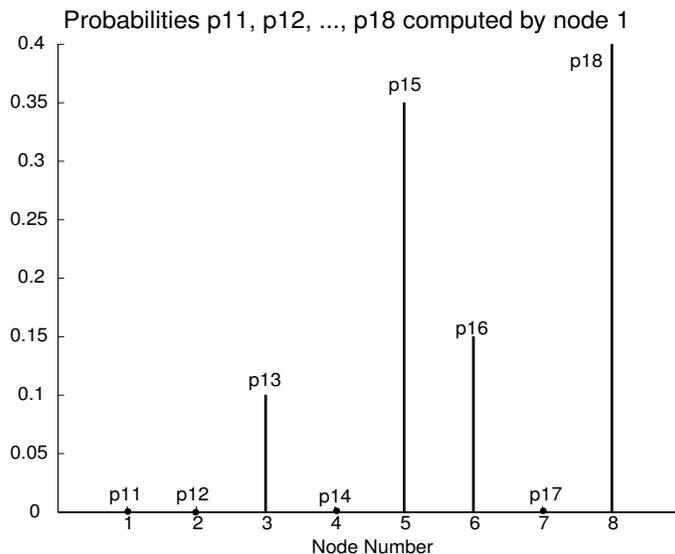
Fig. 2. Illustration of a hypothetical distribution p_{i1} of the load at some time t from node 1's point of view. Node 1 will send data out to node i in proportion p_{i1} it estimates node i is below the average where $\sum_{i=1}^n p_{i1} = 1$ and $p_{11} = 0$

Remark If the denominator $\displaystyle\sum_{i \, \ni \, i \neq j} \text{sat}\left(x_j^{avg} - x_i(t - \tau_{ji})\right)$ is zero, then $x_j^{avg} - x_i(t - \tau_{ji}) < 0$ for all $i \neq j$. However, by definition of the average,

$$\sum_{i \, \ni \, i \neq j} \left(x_j^{avg} - x_i(t - \tau_{ji})\right) + x_j^{avg} - x_j(t) = \sum_i \left(x_j^{avg} - x_i(t - \tau_{ji})\right) = 0$$

which implies $x_j^{avg} - x_j(t) = -\displaystyle\sum_{i \, \ni \, i \neq j} \left(x_j^{avg} - x_i(t - \tau_{ji})\right) > 0$. That is, if the denominator is zero, the node j is below the local average so that $u_j(t) = -K_j\text{sat}(y_j(t)) = 0$ and is therefore not sending out any tasks.

With the definition of the p_{ij} given by (4), a load balancing algorithm which portions out the tasks in proportion to the amounts they are below the local average, is given by the following nonlinear differential-delay system

$$\frac{dx_i(t)}{dt} = \lambda_i - \mu_i + u_i(t) - \sum_{j \neq i} p_{ij} u_j(t - h_{ij})$$

$$x_i^{avg} = \frac{\sum_{j=1}^n x_j(t - \tau_{ij})}{n}$$

$$y_i(t) = x_i(t) - x_i^{avg}$$

$$u_i(t) = -K_i \text{sat}\,(y_i(t)) \tag{5}$$

$$p_{ij} = \frac{\text{sat}\,\left(x_j^{avg} - x_i(t - \tau_{ji})\right)}{\displaystyle\sum_{i \ni i \neq j} \text{sat}\,\left(x_j^{avg} - x_i(t - \tau_{ji})\right)} \delta_{ij}$$

3 Simulations

The simulations here were performed using the model (1) in order to compare with the actual experimental data in the next section. Experimental procedures to determine the delay values are given in [7] and summarized in [8]. These give representative values for a Fast Ethernet network with three nodes of $\tau_{ij} = \tau = 200\ \mu$sec for $i \neq j, \tau_{ii} = 0$, and $h_{ij} = 2\tau = 400\ \mu$sec for $i \neq j, h_{ii} = 0$. The initial conditions for the waiting times were chosen as $x_1(0) = 0.6$, $x_2(0) = 0.4$ 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 and $p_{ij} = (1/2)\delta_{ij}$ for all i, j.

Figures 3 and 4 show the responses with the gains set as $K = 1000$ and $K = 5000$, respectively. These figures indicate that the value of the gain K has a significant effect on the response of the system. Many simulations were performed that are not presented here, and it was found that the system did not go unstable. However, for low values of the gains, the response was sluggish as in Figure 3 while for high values of the gains, the response was quite oscillatory.

To compare with the experimental results given in Figure 8 of the next section, Figure 5 shows the output responses with the gains set as $K_1 = 6667, K_2 = 4167, K_3 = 5000$, respectively.

It is important to note that these plots are of the quantities y_i in equation (1) which are the amount of waiting time *relative* to the local average and therefore can and do go negative. The actual waiting times x_i do *not* go negative.
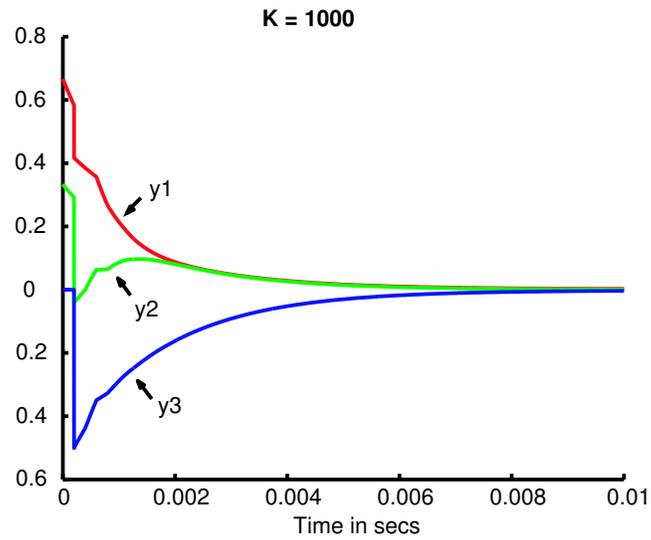
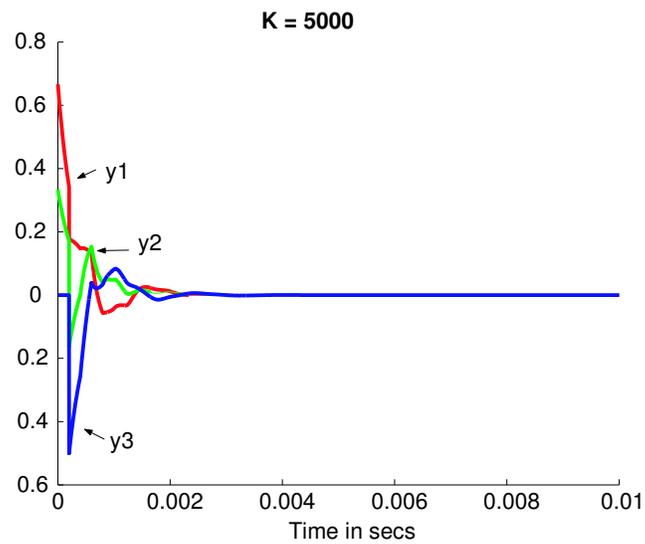**Fig. 3.** Output responses with $K = 1000$.

Fig. 4. Output responses with $K = 5000$.

4 Experimental Results

A parallel machine has been built to implement an experimental facility for evaluation of load balancing strategies. To date, this work has been performed
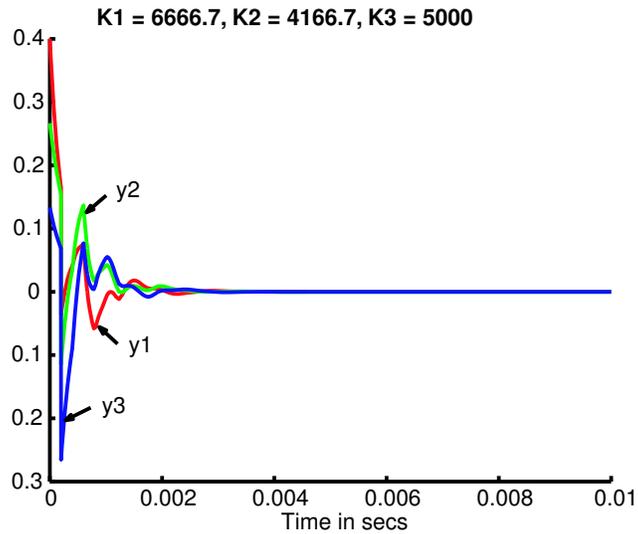
Fig. 5. Output responses with $K1 = 6666.7; K2 = 4166.7; K3 = 5000$

for the FBI Laboratory to evaluate candidate designs of the parallel CODIS database. The design layout of the parallel database is shown in Figure 6. 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, 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 1.6 GHz Athlon MP processors, and the single processor machines use 1.33 GHz Athlon processors. All run the Linux operating system. Our interest here is in the load balancing in any one group of n nodes/hosts.

The database is implemented as a set of queues with associated search engine threads, typically assigned one per node of the parallel machine. Due to the structure of the search process, search requests can be formulated for any target DNA profile and associated with any node of the index tree.

These 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. This structure is shown in Figure 7.
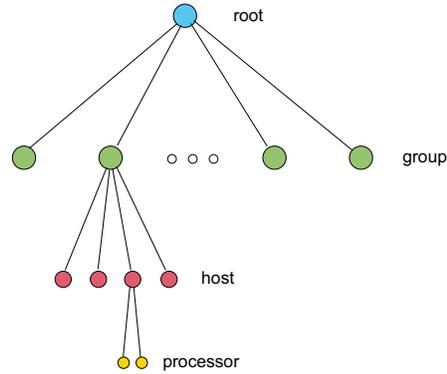
Fig. 6. Hardware structure of the parallel database. Each of the host computers in any given group have the same database. Load balancing is carried out only within a given group as they all perform the same task of searching a particular database.
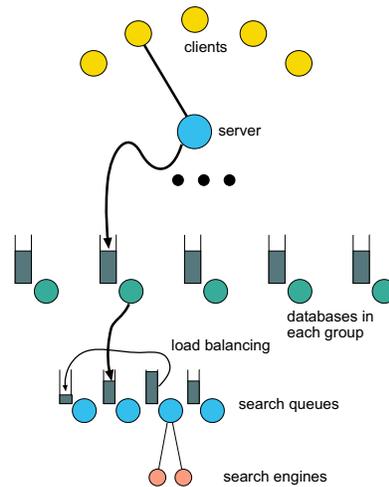

Fig. 7. A depiction of multiple search threads in the database index tree. Here the server corresponds to the "root" in Figure 6. To even out the search queues in each database group, load balancing is done between the nodes (hosts) of a group. If a node has a dual processor, then it can be considered to have two search engines for its queue.

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 [7][8]. The value $\tau = 200 \, \mu \sec$ used for simulations represents an

average value for the delay and was found using the procedure described in [8]. The interest here is to compare the experimental data with that from the simulations given presented in the previous section.

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

$$u(t) = -K\text{sat}\left(y_i(t)\right)$$

where $u(t)$ is the rate of decrease of waiting time $x_i(t)$ per unit time. Consequently, the gain K represents the rate of reduction of waiting time per second in the continuous time model. Also, $y_i(t) = \left(q_i(t) - \left(\sum_{j=1}^{n} q_j(t - \tau_{ij})\right)/n\right)t_{p_i} = r_i(t)t_{p_i}$ where $r_i(t)$ is simply the number of tasks above the estimated (local) average number of tasks. As the interest here is the case $y_i(t) > 0$, consider $u(t) = -Ky_i(t)$. With Δt the time interval between successive executions of the load balancing algorithm, the control law says that a fraction of the queue $K_z r_i(t)$ $(0 < K_z < 1)$ is removed in the time Δt so 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

$$u(t) = -\frac{K_z y_i(t)}{\Delta t} \text{ or } K = \frac{K_z}{\Delta t}. \tag{6}$$

This shows that the gain K 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 (6) is an average value for that run. The average time t_{p_i} to process a task is the same on all nodes (identical processors) and is equal 10μ sec while the time it takes to ready a load for transfer is about 5μ sec. The initial conditions were taken as $q_1(0) = 60000, q_2(0) = 40000, q_3(0) = 20000$ (corresponding to $x_1(0) = q_1(0)t_{pi} = 0.06, x_2(0) = 0.04, x_3(0) = 0.02$). All of the experimental responses were carried out with constant $p_{ij} = 1/2$ for $i \neq j$.

Figure 8 is a plot of the responses $r_i(t) = q_i(t) - \left(\sum_{j=1}^{n} q_j(t - \tau_{ij})\right)/n$ for $i = 1, 2, 3$ (recall that $y_i(t) = r_i(t)t_{p_i}$). The (average) value of the gains were $(K_z = 0.5)$ $K_1 = 0.5/75\mu$ sec $= 6667, K_2 = 0.5/120\mu$ sec $= 4167, K_3 = 0.5/100\mu$ sec $= 5000$. This figure compares favorably with Figure 5 except for the time scale being off, that is, the experimental responses are slower. The explanation for this it that the gains here vary during the run because Δt (the time interval between successive executions of the load balancing algorithm) varies during the run. Further, this time Δt is *not* modelled in the continuous time simulations, only its average effect in the gains K_i. That is, unlike the actual computer network, 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.
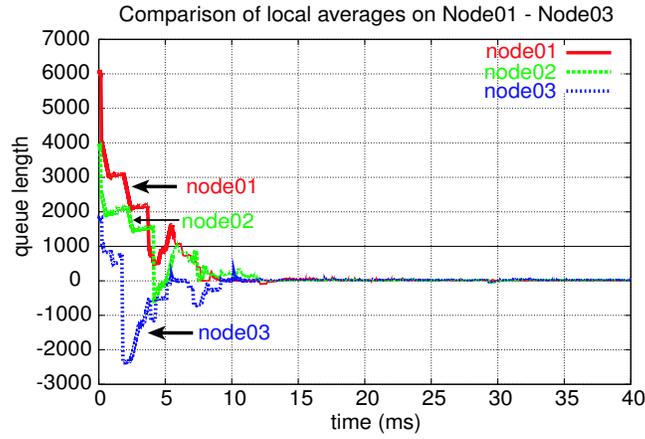
Fig. 8. Experimental response of the load balancing algorithm. The average value of the gains are $(K_z = 0.5)$ $K_1 = 6667, K_2 = 4167, K_3 = 5000$ with constant p_{ij}.
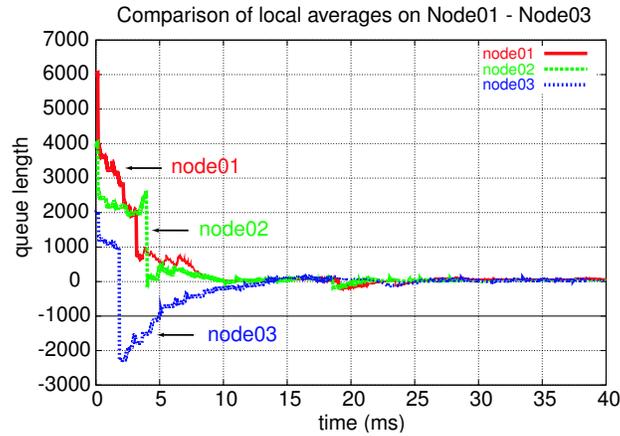

Fig. 9. Experimental response of the load balancing algorithm. The average value of the gains are $(K_z = 0.2)$ $K_1 = 1600, K_2 = 2500, K_3 = 2857$.

Figure 9 shows the plots of the response for the (average) value of the gains given by $(K_z = 0.2)$ $K_1 = 0.2/125\mu\,\sec = 1600, K_2 = 0.2/80\mu\,\sec = 2500, K_3 = 0.2/70\mu\,\sec = 2857$. The initial conditions were $q_1(0) = 60000, q_2(0) = 40000, q_3(0) = 20000$ $(x_1(0) = q_1(0)t_{pi} = 0.06, x_2(0) = 0.04, x_3(0) = 0.02)$.

Figure 10 shows the plots of the response for the (average) value of the gains given by $(K_z = 0.3)$ $K_1 = 0.3/125\mu\,\sec = 2400, K_2 = 0.3/110\mu\,\sec = 7273, K_3 = 0.3/120\mu\,\sec = 2500$.
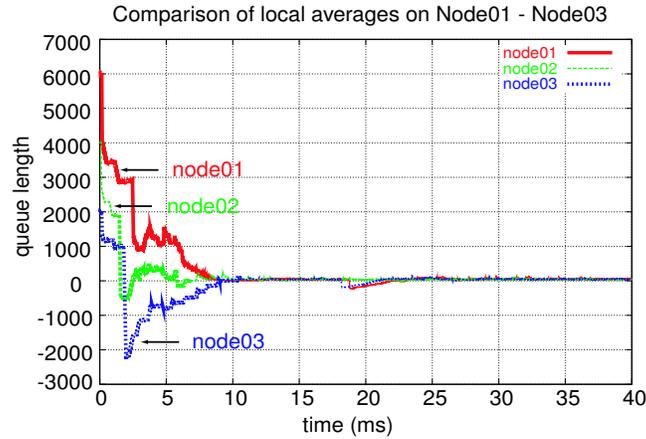
Fig. 10. Experimental response of the load balancing algorithm. The average value of the gains are $(K_z = 0.3)$ $K_1 = 2400, K_2 = 7273, K_3 = 2500$.

5 Summary and Conclusions

In this work, a load balancing algorithm was modelled by a system of non-linear delay-differential equations. Simulations were preformed and compared with actual experimental data. The comparison indicates that the model does indeed capture dynamic behavior of the load balancing network.

A consideration for future work is the fact that the load balancing operation involves processor time which is not being used to process tasks. Consequently, there is a trade-off between using processor time/network bandwidth and the advantage of distributing the load evenly between the nodes to reduce overall processing time. Another issue is that the delays in actuality are not constant and depend on such factors as network availability, the execution of the software, etc. An approach to modelling using a discrete-event / hybrid state formulation that accounts for block transfers that occur after random intervals may also be advantageous in analyzing the network.

6 Acknowledgements

and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Government.

References

1. C. ABDALLAH, J. BIRDWELL, J. CHIASSON, V. CHURPRYNA, Z. TANG, AND T. WANG. Load balancing instabilities due to time delays in parallel computation. In "Proceedings of the 3rd IFAC Conference on Time Delay Systems" (December 2001). Sante Fe NM.

2. C. T. ABDALLAH, N. ALLURI, J. D. BIRDWELL, J. CHIASSON, V. CHUPRYNA, Z. TANG, AND T. WANG. A linear time delay model for studying load balancing instabilities in parallel computations. *The International Journal of System Science* (2003). To appear.

3. J. BIRDWELL, R. HORN, D. ICOVE, T. WANG, P. YADAV, AND S. NIEZGODA. A hierarchical database design and search method for codis. In "Tenth International Symposium on Human Identification" (September 1999). Orlando, FL.

4. J. BIRDWELL, T. WANG, R. HORN, P. YADAV, AND D. ICOVE. Method of indexed storage and retrieval of multidimensional information. In "Tenth SIAM Conference on Parallel Processing for Scientific Computation" (September 2000). U. S. Patent Application 09/671,304.

5. J. BIRDWELL, T.-W. WANG, AND M. RADER. The university of tennessee's new search engine for codis. In "6th CODIS Users Conference" (February 2001). Arlington, VA.

6. A. CORRADI, L. LEONARDI, AND F. ZAMBONELLI. Diffusive load-balancing polices for dynamic applications. *IEEE Concurrency* **22**(31), 979–993 (Jan-Feb 1999).

7. P. DASGUPTA. "Performance Evaluation of Fast Ethernet, ATM and Myrinet under PVM, MS Thesis". University of Tennesse (2001).

8. P. DASGUPTA, J. D. BIRDWELL, AND T. W. WANG. Timing and congestion studies under PVM. In "Tenth SIAM Conference on Parallel Processing for Scientific Computation" (March 2001). Portsmouth, VA.

9. L. KLEINROCK. "Queuing Systems Vol I : Theory". John Wiley & Sons (1975). New York.

10. F. SPIES. Modeling of optimal load balancing strategy using queuing theory. *Microprocessors and Microprogramming* **41**, 555–570 (1996).

11. T. WANG, J. BIRDWELL, P. YADAV, D. ICOVE, S. NIEZGODA, AND S. JONES. Natural clustering of DNA/STR profiles. In "Tenth International Symposium on Human Identification" (September 1999). Orlando, FL.

12. M. WILLEBEEK-LEMAIR AND A. REEVES. Strategies for dynamic load balancing on highly parallel computers. *IEEE Transactions on Parallel and Distributed Systems* **4**(9), 979–993 (1993).