


23577
UCRL-JC-125614
PREPRINT

Improved Utilization and Responsiveness with Gang Scheduling

D.G. Feitelson
M.A. Jette

This paper was prepared for submittal to the
Job Scheduling Strategies for Parallel Processing
Geneva, Switzerland
April 5, 1997

October 1996



This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Improved Utilization and Responsiveness with Gang Scheduling

Dror G. Feitelson

Institute of Computer Science
The Hebrew University of Jerusalem
91904 Jerusalem, Israel

Morris A. Jette

Livermore Computing
Lawrence Livermore National Laboratory
Livermore, CA 94550

Abstract

Most commercial multicomputers use space-slicing schemes in which each scheduling decision has an unknown impact on the future: should a job be scheduled, risking that it will block other larger jobs later, or should the processors be left idle for now in anticipation of future arrivals? This dilemma is solved by using gang scheduling, because then the impact of each decision is limited to its time slice, and future arrivals can be accommodated in other time slices. This added flexibility is shown to improve overall system utilization and responsiveness. Empirical evidence from using gang scheduling on a Cray T3D installed at Lawrence Livermore National Lab corroborates these results, and shows conclusively that gang scheduling can be very effective with current technology.

1 Introduction

As parallel computers become more popular, there is a growing need for good schedulers that will manage these expensive shared resources. And indeed, many scheduling schemes have been designed, evaluated, and implemented in recent years [5, 10].

Many papers investigate scheduling schemes from a system point of view, asking what the system can do to improve utilization and response time, but disregarding the effect on the user. As a result they sometimes advocate solutions that require users to depart from common practice, e.g. to write applications in a style that supports dynamic partitioning (i.e. the allocation may change at runtime) [26, 17], rather than the prevalent SPMD style.

We take different approach, and ask what the system can do given the constraint that users require jobs to execute on a fixed number of processors (as in SPMD). Within this framework, we compare variable partitioning, possibly with reordering of the jobs in the queue, with gang scheduling. We show that although gang scheduling suffers from more overhead than variable partitioning, it can lead to significant improvements due to its added

flexibility. Indeed, gang scheduling can actually give better service (reduced response time) *and* improved utilization, so using it leads to a win-win situation relative to variable partitioning.

The results agree with actual experience on the LLNL Cray T3D, which employs a home-grown gang scheduler [11] (the original system software uses variable partitioning). When this scheduler was ported to the new Cray machine, utilization nearly doubled from 33.4% to 60.9% on average. Additional tuning has led to weekly utilizations that top 90%.

2 Approaches to Scheduling Jobs of Given Size

The schedulers of most commercial parallel systems use variable partitioning. The user specifies the number of processors to use at the time of submitting the job. The scheduler then carves out a partition of the required size, and dedicates it to the job for the duration of its execution. If the required number of processors is not available, the job is either rejected or queued. In most systems a time limit is also imposed, and if the job exceeds it it is killed.

The problem with this scheme is that scheduling decisions have a potentially large, persistent, and unpredictable impact on the future. Specifically, when a new job arrives, the system is faced with the following dilemma:

- if the new job can be accommodated, then scheduling it immediately will utilize unused resources, so it is good.
- however, if this job runs for a long time, and will block other jobs in the future, it may lead to more future loss than current gain. so maybe it should be left aside.

Consider the following simple case as an example: a 128-node system is currently running a 64-node job, and there are a 32-node job and a 128-node job in the queue. The question is, should the 32-node job be scheduled to run concurrently with the 64-node job? Two outcomes are possible. If the 32-node job is scheduled and it terminates before the 64-node job, resource utilization is improved from 50% possibly up to 75%. But if the 64-node job terminates soon after the 32-node job is scheduled, and the 32-node job runs for a long time, the utilization drops from 50% to 25%. And, in order not to starve the 128-node job, it might be necessary to just let the 64-node job run to completion, and settle for 50% utilization.

As the future is usually unknown, there is no solution to this dilemma, and any decision may lead to fragmentation. Thus using variable partitioning may lead to significant loss of computing power [15, 29], either because jobs do not fit together, or because processors are intentionally left idle in anticipation of future arrivals [23].

The most common solution is to reorder jobs in queues so as to pack them more tightly [14]. One promising approach is to allow small jobs to move forward in the queue if they can be scheduled immediately. However, this may cause starvation of large jobs, so it is typically combined with allowing large jobs to make reservations of processors for some future time. Only *short* jobs are then allowed to move ahead in the queue (Fig. 2) [3, 16].




Figure 1: *Example of the problems faced by variable partitioning.*

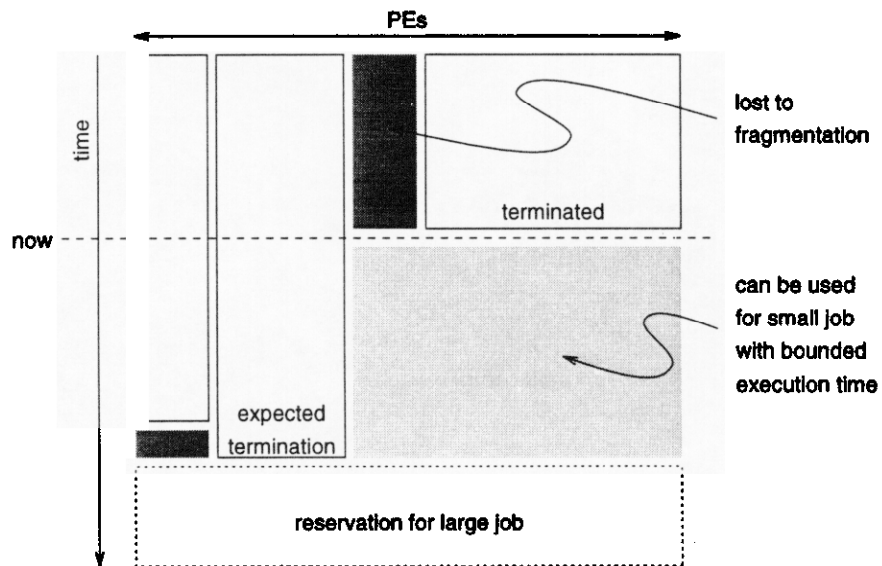


Figure 2: *Runtime bounds on executing jobs allow reservations to be made for large jobs and then backfilling with smaller jobs to reduce fragmentation.*

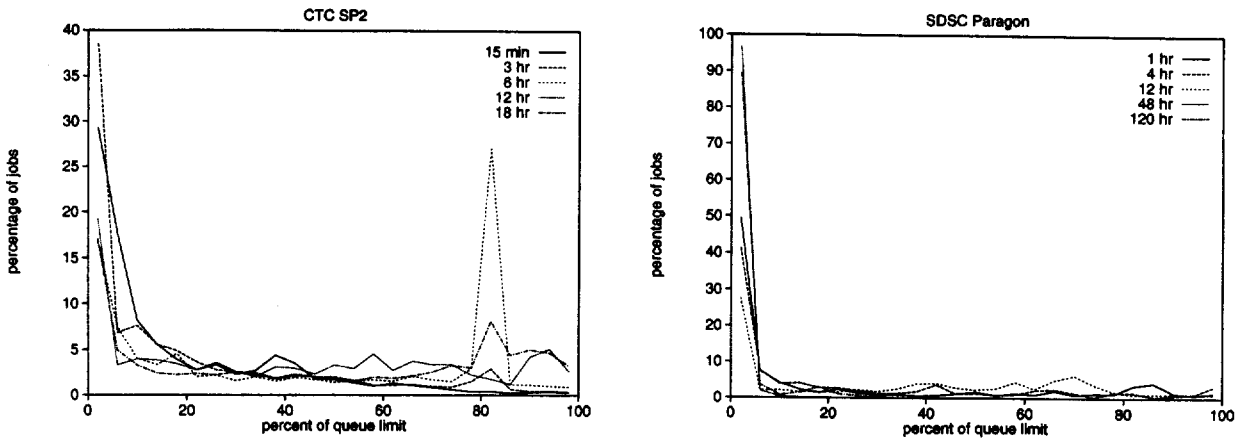


Figure 3: *Job runtimes as a fraction of the batch queue time limit, showing that most jobs use only a fraction of the time limit, even for queues with very long limits. The plot for each queue limit is normalized independently.*

The problem with this idea is that it requires information about job runtimes. A rough approximation may be obtained from the queue time limit (in most systems users may choose which queue to use, the difference being that each queue has a distinct set of resource limits associated with it). The idea is that the user would choose the queue that best represents the application's needs, and the system would then be able to select jobs from the different queues to create a job mix that uses the system's resources effectively [27]. However, experience indicates that this information is unreliable, as shown by the distributions of queue-time utilization in Fig. 3. The graphs show that users tend to be extremely sloppy in selecting the queue, thus undermining the whole scheme. (the graphs show the distributions in buckets of 4 percentage points. Thus the top left data point in the left graph shows that about 38% of the jobs submitted to all the 6-hour queues on the Cornell SP2 only used between 0 and 4% of their time limit, i.e. they were actually shorter than 15 minutes)

Another solution to the fragmentation problem is to use adaptive partitioning rather than variable partitioning [25, 24, 22]. The idea here is that the number of processors used is a compromise between the user's request and what the system can provide. Thus the system can take the current load into account, and reduce the partition sizes under high load conditions. However, this scheme also requires change of user interfaces, albeit much less disruptive than dynamic partitioning.

The preferred solution is to use gang scheduling [19, 7, 8]. With gang scheduling jobs receive the number of processors requested, but only for a limited time quantum. Then a "multi-context-switch" is performed on all the processors at once, and another job (or set of jobs) is scheduled instead. Thus all jobs can execute concurrently using time slicing, as in conventional uniprocessors. As a result, a scheduling decision only impacts the scheduling slot to which it pertains; other slots are available to handle other jobs and future arrivals. this adds flexibility and boosts performance.




Figure 4: Example of how the flexibility afforded by time slicing can increase system utilization; compare with Fig. 1.

Returning to the example considered earlier, the situation with gang scheduling is illustrated in Fig. 4. The 32-node job can safely run in the same time-slot with the 64-node job, while the 128-node job gets a separate time-slot. There is no danger of starvation. As long as all three jobs are active, the utilization is 87.5%. Even if the 64-node job terminates, leaving the 32-node job to run alone in its time-slot, the utilization is 62.5%. Naturally, a few percentage points should be shaved off these figures to account for context-switching overhead. Nevertheless, this is a unique case where time-slicing, despite its added overhead, can lead to better resource utilization than batch scheduling.

Using gang scheduling not only improves utilization — it also reduces mean response time. It is well known that mean response time is reduced by the shortest-job-first discipline. In workloads with high variability this is approximated by time slicing, because chances are that a new job will have a short runtime [21, 20]. As production workloads do indeed exhibit a high variability [6], it follows that gang scheduling will reduce mean response time. Indeed, gang scheduling has even been advocated in conjunction with dynamic partitioning [18].

3 Simulation Results

3.1 The Compared Scheduling Schemes

In order to demonstrate the ideas described above, we simulate the performance of a multicomputer subjected to a realistic workload and using one of a set of different scheduling schemes. these are:

FCFS: the base case we use for comparison is variable partitioning with first-come-first-serve queueing. This scheme is expected to suffer from significant fragmentation.

Backfill: backfilling was developed for the Argonne National Lab SP1 machine [16], and has recently also been installed on the Cornell SP2 and other machines. It allows short jobs to move forward in the queue provided they do not cause delays for any other job. Only jobs that do not cause delay are moved forward. We assume the scheduler has perfect information when making such decisions, i.e. it knows the exact runtimes of all the jobs in the queue.

Prime: this policy is a simplified version of a policy used on the SP2 machine at NASA Ames [12]. The idea is to distinguish between prime time and non-prime time¹: during prime time, large jobs (more than 32 nodes) are restricted to 10 minutes, while small jobs are allowed up to 4 hours provided at least 32 nodes are available. Thus, if only a few nodes are available, all jobs are restricted to 10 minutes, and responsiveness for short jobs is improved. This achieves a similar effect to setting aside a pool of nodes for interactive jobs [27]. During non-prime time these restrictions are removed. Again, we assume the scheduler knows the runtimes of all jobs.

Gang: gang scheduling with no information regarding runtimes. The jobs are packed into slots using the buddy scheme, including alternate scheduling [4]. Two versions with different scheduling time quanta are compared: one has relatively small time quantum of 10 seconds, so most jobs effectively run immediately, and the other has a time quantum of 10 minutes (600 seconds), so jobs may be queued for a certain time before getting to run.

In all cases, exactly the same workload is used.

3.2 Simulation Methodology

The workload model is an improved version of the model used in [4]. It is based on workload analysis from a number of production systems [6, 13, 28], and is characterized as follows (Fig. 5)

- The distribution of job sizes emphasizes small jobs and powers of two.
- The distribution of runtimes is a three-stage hyperexponential, where the relative weights of the three stages depends on the job size. This dependence is used to create a correlation between the job size and the runtime.
- The arrivals are poisson, except for jobs that are re-run a number of times, in which case they are re-submitted immediately upon completion.

¹Our workload model does not include a daily cycle of job submittals — it is a continuous stream of jobs with the same statistical properties. Thus in our simulations the distinction is only in the scheduling policy, which is switched every 12 hours.

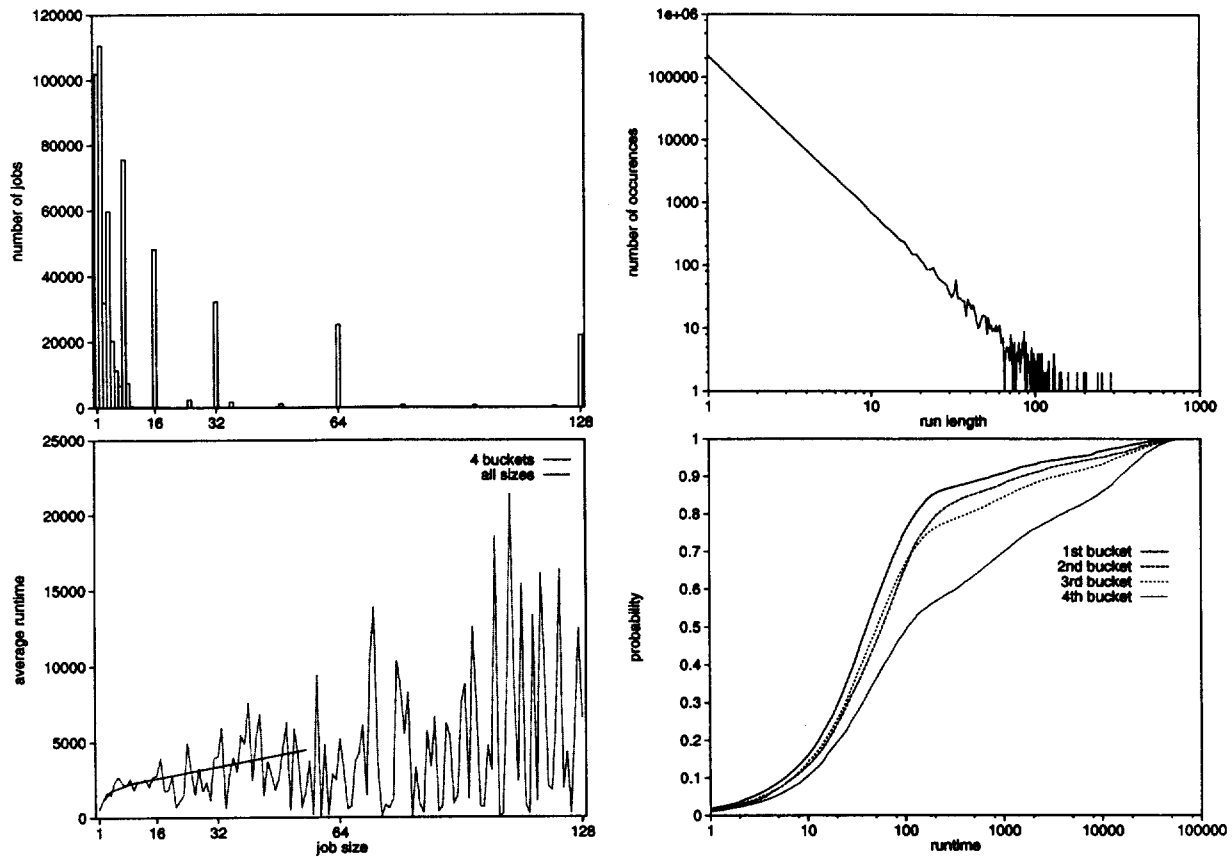


Figure 5: *Statistical properties of the workload model: (a) distribution of job sizes (b) distribution of runlengths (c) correlation of runtime with job size, for all sizes and when jobs are grouped into four buckets according to size (d) distributions of runtimes for the jobs in the four buckets.*

The simulation uses the batch means method to evaluate confidence intervals. Each batch includes 3333 job terminations. The first batch is discarded to account for simulation warmup. The length of each experiment (i.e. the simulation for each data point in the results) is at least 3 batches, or more as required so that the 90% confidence interval is no larger than 10% of the data point value, up to a maximum of 100 batches. Interestingly, simulations of all scheduling schemes except gang scheduling with short time quanta used all 100 batches, without a significant reduction in the confidence interval: it was typically in the range of 20-40% of the data point value. This reflects the very high variance present in the workload.

The sequence of job arrivals is generated once and reused for each data point and each scheme. Only the mean interarrival time is changed to create different load conditions.




Figure 6: *Simulation results.*

3.3 Experimental Results

The results are shown in Fig. 6. As expected, FCFS saturates at extremely low loads, and even before saturation it tends to create very high slowdowns. Backfilling and delaying large jobs to non-prime time are both much better, but backfilling can sustain a higher load and produces lower slowdowns. Attempts to improve the performance of the prime/non-prime policy by fiddling with its parameters (the threshold between small and large jobs, and the length of the prime shift) showed that it is relatively insensitive to the exact values of these parameters. However, it should be remembered that our workload is not optimal for checking the prime/non-prime policy, because it does not display a daily cycle.

Gang scheduling, even with relatively long quanta of 10 minutes, takes the cake: the slowdowns are very low, and saturation is delayed until system load approaches 1. This agrees with the informal arguments presented in Section 2. While our simulations may be criticized for not modeling the overheads involved in gang scheduling — for example, the overhead of context switching and the effect of corrupting cache state — we feel that with long enough time quanta these overheads can be kept relatively low, so the main results remain valid.

4 Experience with Gang Scheduling on the Cray T3D

The default mode of operation for the Cray T3D is that of variable partitioning. Generally, jobs are allocated a partition of processors as soon as a suitable set becomes available. In the case where a job has waited for longer than some configurable period of time, the initiation of all other jobs is deferred until it begins execution. The partitions are held until the job completes and relinquishes them, effectively locking out any other use for those processors.

With such a processor allocation mechanism, the computational requirements of long-running production jobs directly conflict with those of interactive code development work.

Our goals in the development of the Gang Scheduler for the Cray T3D were:

- Provide better response for interactive jobs
- Provide better throughput for batch jobs
- Permit larger jobs to be executed
- Provide optimum throughput for specific jobs, as designated by management

While achieving all of these objectives would seem impossible, our initial utilization rate of 33.4% provided us a great deal of room for improvement. By creating a larger virtual machine through gang scheduling, we believed these objectives could be met by:

- Time sharing processors for interactive jobs when experiencing an extremely heavy interactive workload
- Keep processors fully utilized with batch jobs until preempted by an interactive or other high priority job
- Make processors available in a timely fashion for large jobs
- Make processors available in a timely for specific jobs and make those jobs non-preemptable

Some might argue that interactive computing with a massively parallel computer is unreasonable, but interactive computing accounts for a substantial portion of our workload and is used for code development and rapid throughput. Interactive jobs currently account for 79% of all jobs executed and 11% of all CPU cycles used in our environment. A single interactive job can be allocated up to 25% of all processors and memory. The aggregate of all interactive work will normally consume between zero and 150% of all processors and memory in our environment. While the Cray T3D is well suited for addressing the execution of grand challenge problems, we wanted to expand its range of functionality into general purpose support of these challenge problems.

4.1 Design of the Cray T3D

The Cray T3D is a massively parallel computer incorporating DEC alpha 21064 microprocessors, capable of 150 MFLOPS peak performance. Each processor has its own local memory. The system is configured into nodes, consisting of two processors with their local memory and a network interconnect. The nodes are connected by a bidirectional three-dimensional torus communications network. There are also four synchronization circuits (barrier wires) connected to all processors in a tree shaped structure. The system at Lawrence Livermore

National Laboratory (LLNL) has 256 processors, each with 64 megabytes of DRAM. Disk storage is required to store the job state information for preempted jobs. This can be either shared or private storage space. We have created a shared 48 gigabyte file system for this purpose. The storage requirements will depend upon the T3D configuration and the amount of resource oversubscription permitted.

Without getting into great detail, the T3D severely constrains processor and barrier wire assignments to jobs. Jobs must be allocated a processor count which is a power of two, with a minimum of two processors or one node. The processors allocated to a job must have a specific shape with specific dimensions for a given problem size. For example, an allocation of 32 processors must be made with a contiguous block with 8 processors in the X direction, 2 processors in the Y direction and 2 processors in the Z direction. Furthermore, the possible location of the processors assignments is restricted. These very specific shapes and locations for processor assignments are the result of the barrier wire structure. Jobs must be allocated one of the four barrier wires when initiated. The barrier wire assigned to a job can not change if the job is relocated and, under some circumstances, two jobs sharing a single barrier wire may not be located adjacent to each other. The number of processors assigned to a job can not change during execution [2].

There are two fundamentally different ways of providing for timesharing of processors. The entire state of a job, including memory contents, register contents and switch state information can be written to disk. Alternately, the register and switch state information can be saved and the memory shared through paging. Saving the entire job state clearly makes context switches very time consuming, however, it can provide a means of relocating jobs to different processors and provide a means of preserving executing jobs over computer restarts. Sharing memory through paging can make for much faster context switches.

The Cray T3D provides timesharing by saving the entire state of a job to disk. Given the T3D's constraints on processor assignment, the ability to relocate jobs with this mechanism clearly make it preferable. While the ability to preserve executing jobs over computer restarts has proven to be of some use, most programs complete in a few hours and can be restarted without substantial impact upon the system. Unfortunately, the high context switch time provides lower interactivity than would be desirable. It should also be noted that the Cray T3D only supports the movement of a job's state in it's entirety. It is not possible to initiate state transfers on a processor by processor basis, although that would capability would improve the context switch time.

The original version of this Gang Scheduler was developed for the BBN TC2000 computer. The BBN computer permitted programs to be assigned processors without locality constraints. Its timesharing through shared memory and paging was successful at providing both excellent interactivity and utilization [?, 11].

4.2 Implementation Overview

The T3D Gang Scheduler allocates processors and barrier circuits for all programs. In order to satisfy the diverse computational requirements of our clients, the programs are classified

by access requirements:

- Interactive class jobs require responsive service
- Debug class jobs require responsive service and can not be preempted
- Production class jobs require good throughput
- Benchmark class jobs can not be preempted
- Standby class jobs have low priority and are suitable for absorbing otherwise idle compute resources

There are several class dependent scheduling parameters to achieve the desired performance characteristics.

- Priority: Job classes are prioritized for service
- Wait time: The maximum time that a job should wait before (or between) processor access
- Do-not-disturb time multiplier: This parameter is multiplied by the number of processors to arrive at the do-not-disturb time, the minimum processor allocation time before preemption
- Processor limit: The maximum number of processors which can be allocated to jobs of this class

The priority assigns a relative importance to loading jobs of a given class. We make interactive jobs higher priority than production jobs during the daytime and assign them equal priority at night.

The wait time is designed to insure timely responsiveness, especially for interactive and debug class jobs. After a job has waited to be loaded for the maximum wait time, an attempt will be made to reserve a block of processors for it. This processor reservation mechanism frequently preempts multiple small jobs to prevent starvation of large jobs.

The desire for timely response needs to be balanced against the cost of moving a job's state onto disk and back to memory. In order to prevent job thrashing, a job is assigned processors for a minimum of its do-not-disturb time before preemption. After a job's do-not-disturb time has been exceeded, the job may be preempted by any other job, although preference is given to jobs of higher priority classes and larger jobs. Some jobs might not be preempted until long after their do-not-disturb time has been exceeded. This scheme has been shown to keep the processors efficiently utilized by high priority class jobs. A job's do-not-disturb time is computed by multiplying the do-not-disturb time multiplier for the job's class by the number of processors used. The do-not-disturb time multiplier should be set to a value substantially larger than the time required to move a job's state in one processor

from memory to disk and back to memory. This time will vary with the disk configuration. On the LLNL T3D with 256 processors and 64 megabytes of memory each, the entire torus or processors can be repacked in about eight minutes or one second per processor.

We use the class processor limit to restrict the number of processors allocated to non-preemptable jobs during the daytime.

Several non-class dependent scheduling parameters also exist to regulate computer-wide resource use.

- Large job size: The minimum number of processors requested by a job for it to be considered “large”
- Large processor limit: The maximum number of processors which can be allocated to “large” jobs at any time
- Job processor limit: The maximum number of processors which can be allocated to any single job
- System processor limit: The maximum number of processors used by jobs either running or swapped to disk

Since “large” jobs can take a significant period of time to have their state moved between memory and disk, interactivity can be improved by restricting the number of processors allocated to them.

The system processor limit is used to avoid filling the file system used for job state information. We are conservative in our allocation of this storage area because it is shared. Jobs will be queued, but not initiated to avoid exceeding this parameter. If an attempt is made to preempt a job when insufficient storage is available, that job will continue execution and no further attempts will be made to preempt it.

Several sets of the jobs scheduling parameters can be defined for different hours of the day and job scheduling parameters may be altered in real time. For example, it may be desirable to provide a lower level of interactivity at night. The reduced level of job preemption (and reduced time for such idled processors) can result in improved system throughput. It might also be desirable to severely restrict processor availability to benchmark class jobs except at night or on weekends.

These parameters can result in conflicting job scheduling rules if not set appropriately or if the workload is irregular. Performance will degrade under such circumstances, but “reasonable” behavior can be expected with the highest priority job classes receiving the best performance.

The scheduling parameters currently being used during the daytime on weekdays are:

<i>Job Class</i>	<i>Priority</i>	<i>Wait Time</i>	<i>Do-not-disturb Time per Processor</i>	<i>Processor Limit</i>
Interactive	4	0 Sec	10 Sec	256
Debug	4	300 Sec	1 Year	96
Production	3	1 Hour	10 Sec	256
Benchmark	2	1 Year	1 Year	64
Standby	1	1 Year	3 Seconds	256

Large Job Size	64 processors
Large Processor Limit	192 processors
Job Processor Limit	256 processors
System Processor Limit	576 processors

The time of one year is used in several cases to insure no preemption or an indefinite wait for some job classes.

4.3 Job Scheduling Algorithm

We have implemented a two pass scheduling algorithm. The first pass checks for jobs which have waited for loading longer than their job class' maximum wait time. These jobs are viewed as having a high priority for loading and special measures are taken for loading them. If there is more than one such job, a list of these jobs is constructed then sorted by job class priority and within each priority value by the time waiting for loading. Each of these jobs is considered for loading in the sorted order. The processor requirement for the job will be compared against the scheduler's job processor limit. If the job's processor request can not be satisfied, that job will no longer be considered a candidate for loading. Possible processor assignments for the job are considered. For each possible processor assignment, a cost is computed. The cost considers the number of nodes occupied by the potentially preempted jobs, their relative priority, and how much time remains in their do-not-disturb time. In no case will a job be preempted for another job of a lower priority class. Jobs of benchmark and debug class will never be preempted. If no possible processor assignment for loading the waiting job is located, its loading will be deferred. If a possible processor assignment is located, the lowest cost set of processors will be reserved for the exclusive use of this waiting job and jobs occupying those processors will be preempted when their do-not-disturb times have been exhausted. Only one job will have processors reserved for it at any point in time. Once a set of processors have been reserved for a waiting job, the reservation of processors for other waiting jobs will be deferred until the selected job has been loaded. An exception is made only in the case that a higher priority class job exceeds its maximum wait time. For example, an interactive class job could preempt the reservation of processors for a production class job. The job with reserved processors can be loaded into other processors if another compatible set of processors becomes available at an earlier time. As soon as that job is loaded, the reserved processors are made generally available. This mechanism insures timely interactivity and prevents the starvation of large jobs.

In the second scheduler pass, other executable jobs are recorded in a list sorted by job class priority and within each priority by the time waiting for loading. Each job in the sorted list is considered for processor assignment. First the limits (job processor limit, large job limit, and job class limit) are checked to determine if the job should be allocated processors. Any job satisfying these limits will have its barrier wire circuit and processor requirements considered. If the job can have its requirements met either with unallocated resources or resources which can be made available by preempting jobs which have exceeded their do-not-disturb time, it will have a barrier wire circuit and processors assigned. If a specific barrier wire is not requested, one of those available will be assigned. All four barrier wire circuits are considered for use and selected on the basis of lowest contention. More efficient relocation of jobs can be achieved by using all four barrier wire circuits.

The time required to save the state of a job on disk can be up to four minutes. Given this delay, it is not ideal to queue the loading of a job until the processors assigned to it are actually available. Whenever processors are actually made available, the job scheduler is executed again. This insures that when processors become available, they are assigned to the most appropriate jobs then available.

When a newly started job can immediately begin execution in a variety of possible sets of processor, a best-fit algorithm is used to make the selection. We also try to locate debug and benchmark class jobs, which can not be preempted, together in order to avoid blocking large jobs.

4.4 Client Interface

The default mode of operation for the Cray T3D requires all jobs, batch and interactive, to be initiated through a program called `mppexec`, which will accept as arguments the number of processors required, specific processor requirements, specific barrier wire requirements, etc. The Gang Scheduler takes advantage of this feature by creating a wrapper for `mppexec` which is upwardly compatible with it. The interface registers the job with the Gang Scheduler and waits for an assignment of processors and barrier circuit before continuing. On a heavily utilized computer, this typically takes a matter of seconds for small numbers of processors and possibly much longer for large jobs. The only additional argument to the Gang Scheduler interface is the job class, which is optional. By default, interactive jobs are assigned to the interactive job class, the Totalview debugger jobs are assigned to the debug class, and batch jobs are assigned to the production job class.

4.5 The Gangster Tool

We provide users with an interactive tool, `gangster`, for observing the state of the system and controlling some aspects of their jobs. `Gangster` communicates with the Gang Scheduler to determine the state of the machine's processors and individual jobs. `Gangster's` three-dimensional node map displays the status of each node (each node consists of two processing elements on the T3D). `Gangster's` job summary reports the state of each job, including jobs


```

      h h c c a a a a CLAS JOB-USER      PID COMMAND #PE BASE W ST MM:SS
      h h c c a a a a Int d - colombo    2976 icl1      8 100 1 R 42:44
      h h c c a a a a Int h - mshaw      9529 icf3d     32 020 2 R 00:33
h h c c a a a a
                                     Bmrk g - grote    9264 warpslav  32 200 1 R 00:20
      h h c c a a a a
      h h c c a a a a Prod a - caturla  95396 moldy    128 400 2 R 107:35
      h h c c a a a a Prod b - colombo  98057 vdif      8 000 3 R 91:54
h h c c a a a a Prod c - wenski   98484 pproto6.  32 220 3 R 85:12
                                     Prod e - colombo   8712 icl3      8 004 1 R 6:23
      b d g g a a a a Prod f - colombo  8873 icl2      8 104 2 R 4:54
      b d g g a a a a Prod i - dan      5684 camille   32 020 0 0 32:08
      e f g g a a a a Prod j - vickie  99393 kiten    64 400 3 0 132:59
e f g g a a a a
                                     b d g g a a a a
      b d g g a a a a
      e f g g a a a a
      e f g g a a a a
gangster:

```

Figure 7: *Sample gangster display.*

moving between processors and disk (see the sample display below). Users can use gangster to change the class of their own jobs or to explicitly move their job's state to disk (suspending execution) or make it available for execution (resume).

Gangster communicates with the Gang Scheduler via sockets in the /tmp directory. A socket with a file name of the form `gsched#` is used to establish communications, where # represents the Gang Scheduler daemon version number. A second socket with a file name of the form `user#` is used for communications with a specific gangster process, where # represents a process ID number. Authentication is provided by checking the owner of the socket.

A sample gangster display is shown in Fig. 7. This display identifies jobs in the system and assigned processors. The node map is on the left. A dot or letter denotes each node (two processing elements on the T3D): a dot indicates the node is not in use, a letter designates the job currently occupying that node. On the right is a summary of all jobs. The W field shows the barrier wire used. The MM:SS field shows the total execution time. The ST field shows the job's state:

i = swapping in
 N = new job, not yet assigned nodes or barrier wire
 o = swapping out
 O = swapped out
 R = running
 S = suspended
 W = waiting job, assigned nodes and barrier wire

Node number 000 is in the upper left corner of the lowest plane. The *X* axis extends downward within a plane. The *Y* axis extends up, with one *Y* value in each plane. The *Z* axis extends to the right. This orientation was selected for ease of display for a 256 processor T3D configuration. The diagram below shows a typical weekday gangster output.

4.6 Configuration

Prior to installation of the Gang Scheduler, our NQS batch system was configured to leave an adequate number of processors available for interactive computing. The following table summarizes the original global NQS processor limits at various times of the day:

<i>Start Time</i>	<i>End Time</i>	<i>User Limit</i>	<i>Run Limit</i>	<i>Aggregate mpp-pe-limit</i>
00:00	04:00	2	8	256
04:00	18:00	2	8	96
18:00	24:00	2	8	192

The *User Limit* specifies the maximum number of batch jobs a single user may have executing. The *Run Limit* specifies the maximum number of batch jobs the system may have executing at one time. The *mpp-pe-limit* is the maximum number of processors which the batch jobs may have allocate at one time. These limits apply to the entire batch system, an aggregate of all queues. The limits on individual NQS queues were configured by processor limits and time limit as shown below:

<i>Queue Name</i>	<i>User Limit</i>	<i>Run Limit</i>	<i>Time Limit</i>	<i>Job Processor Limit</i>	<i>Single Job Aggregate mpp-pe-limit</i>
pe32	1	4	4 Hr	32	128
pe64	1	3	4 Hr	64	192
pe64_long	2	2	19 Hr	64	96
pe128_short	1	4	15 Min	128	128
pe128	1	1	4 Hr	128	128
pe256_short	1	1	15 Min	256	256
pe256	1	1	4 Hr	256	256

Note that jobs requiring more than a four hour time limit were limited to 64 processors. Also note that substantial compute resources were sacrificed in order to insure processors for

interactive computing. This was particularly noticeable in the early morning hours as the mpp-pe-limit dropped to 96 at 04:00 in order to insure the availability of 160 processors for interactive use at 08:00. Frequently this left many processors idle. Under the occasionally heavy interactive workload, all processors would be allocated and interactive jobs experienced lengthy initiation delays.

Significant changes in the NQS configuration have been made possible by the Gang Scheduler. Instead of leaving a significant number of processors available for interactive use, NQS fully subscribes the computer during times of peak use and over allocates the processors by as much as 100% at night to process jobs needing all of the processors. The overallocation of processors permits the execution of larger jobs and makes more jobs available for fully packing the T3D's torus of processors. The global limits and those of most queues have been changed as shown below.

<i>Start Time</i>	<i>End Time</i>	<i>User Limit</i>	<i>Run Limit</i>	<i>Aggregate mpp-pe-limit</i>
00:00	04:00	5	20	512
04:00	07:00	5	20	320
07:00	18:00	5	20	256-320*
18:00	24:00	5	20	320

<i>Queue Name</i>	<i>User Limit</i>	<i>Run Limit</i>	<i>Time Limit</i>	<i>Job Processor Limit</i>	<i>Job Aggregate mpp-pe-limit</i>
pe32	3	8	4 Hr	32	128
pe64	2	2-3*	4 Hr	64	128-192*
pe64_long	2	2	40 Hr	128	96-128*
pe128_short	1	4	15 Min	96	128
pe128	0-2*	1	4 Hr	128	0-256*
pe256_short	0-1*	1	15 Min	256	0-256*
pe256	0-1*	1	4 Hr	256	0-256*

*Varies by time of day and/or day of week

NQS jobs now relinquish their processors only as needed, not in anticipation of interactive work. During periods of heavy use, this improves our realized throughput substantially while preserving good interactivity. While average interactivity has decreased slightly due to interference from NQS jobs, the worse case startup time has dropped from tens of minutes to about one minute. This is still quite acceptable to our user community, especially when accompanied by a substantial increase in realized batch throughput. We also see a few jobs relocated to better utilize the available processors during periods of heavy use, especially when jobs requiring 64 or 128 processors exist.

<i>Scheduler</i>	<i>Batch Only</i>	<i>Batch and Interactive</i>	
		<i>Interactive Part</i>	<i>Batch Part</i>
UNICOS MAX	3327 Sec	3553 Sec	3347 Sec
DJM	2848 Sec	2339 Sec	3248 Sec
Gang Scheduler	2950 Sec	2797 Sec	3565 Sec

Table 1: *Benchmark results for the three schedulers.*

4.7 Performance Results

In order to quantify the effect upon system throughput and interactivity under heavy load, we have tested the Gang Scheduler against the standard UNICOS MAX scheduler and the Distributed Job Manager (DJM). DJM is a gang scheduler developed by the Minnesota Supercomputer Center. DJM has undergone substantial modification for performance enhancements by Cray analysts at LLNL. All of the DJM code to accomplish job swapping is new. The enhanced version of DJM was used for testing purposes.

All three schedulers were tested under a very heavy load of jobs typical of those executed at LLNL. One test consisted of a single stream of batch jobs initiated at one time. A second test consisted of the same batch workload with the addition of an interactive workload. The initiation of each interactive job was fixed in relationship to the completion of the preceding interactive job. The results are summarized in the below tables. The additional interactivity of DJM and the Gang Scheduler come at the cost of processors idled during job swapping, although movement of some jobs can result in more efficient packing of the processors and a net increase in throughput. Note that the interactive execution time is not an absolute measure of responsiveness, but the time required to execute the benchmark workload. The results are shown in Table 1.

It should be noted that the Gang Scheduler results are based upon an early version which has undergone substantial enhancement and tuning. Current Gang Scheduler results should be within a percent or two of DJM results. The Gang Scheduler and DJM do give some indication as to the additional throughput and interactivity achievable through gang scheduling on a massively parallel computer.

Fig. 8 demonstrates the Gang Scheduler's ability to execute large jobs, note the dramatic improvement in throughput of 128 and 256 processor jobs. This charts the distribution of resources allocated to each job size as percentage of CPU resources actually delivered to customers. The percentage of gross CPU resources which are delivered to large jobs has increased by an even wider margin. The January 1996 period is the last full month of operation with the standard UNICOS MAX operating system. The later two dates show periods of Gang Scheduler use.

The best measure of success is probably actual throughput achieved. While utilization is quite low on weekend, the improvement in throughput at other times has dramatically improved with preemptive schedulers. Current utilization on weekends is typically about ten




Figure 8: Changes in workload distribution due to use of the Gang Scheduler.




Figure 9: Weekly utilization with the three schedulers.

percent, while utilization during normal work hours with a preemptive scheduler normally exceeds 90 percent. Fig. 9 summaries utilization of processor resources over the course of several entire weeks.

For the period recorded above, the percent utilization realized by the three schedulers tested were:

UNICOS MAX	28.6
DJM	38.2
Gang Scheduler	50.2

Over the longer term, utilization has improved even more dramatically while providing




Figure 10: Monthly utilization with the three schedulers.

good interactivity, as shown in Fig. 10. CPU utilization reported is the percentage of all CPU cycles available which are delivered to customer computation. Weekly utilization rates have reached over 93 percent.

While DJM would have provided for good interactivity and throughput, it became available at the same time as our Gang Scheduler was completed and we felt that continued development of our Gang Scheduler was worthwhile. In addition, our Gang Scheduler provided the means to arbitrarily lock jobs into processors. This was important for us to be able to insure optimal throughput for jobs specified by our management.

4.8 Future Development

While the Gang Scheduler manages the currently active jobs well, the NQS batch system selects the jobs to be started. It would be desirable to integrate the Gang Scheduler with NQS in order to more efficiently schedule all available jobs. Work is also planned for the gang scheduling of jobs across a heterogenous collection of computers.

5 Conclusions

Gang scheduling has often been advocated based on its advantages of

- presenting jobs with an environment similar to that of a dedicated machine, thus allowing fine grain interactions based on user-level communication and busy waiting [9],
- support for interactive work by using time slicing, which guarantees a reasonable response time for short jobs, and

- not placing any restrictions or requirements on the model of computation and programming style.

However, many researchers have expressed the fear that using gang scheduling would lead to unacceptable system performance due to the overheads involved in context switching and the loss of resources to fragmentation.

In contrast, we have shown that gang scheduling can *improve* system performance significantly relative to static space slicing policies often used in practice on parallel supercomputers. Gang scheduling adds flexibility to resource allocations, and reduces the impact of bad decisions. This contributes directly to a reduction in fragmentation, and more than offsets the cost of overheads. Indeed, experience with using gang scheduling for a production workload on the Cray T3D at Lawrence Livermore National Lab has shown a dramatic raise in system utilization.

The main obstacle to widespread use of gang scheduling is memory pressure. If gang scheduling is performed at a fine granularity, all jobs need to be memory resident at the same time, so each has less memory available. The alternative is to swap jobs to disk when they are de-scheduled, and swap them in again when scheduled. This is a viable approach, but it requires sufficient resources to be invested in adequate I/O facilities. The combination of demand paging and prefetching with gang scheduling remains an interesting topic for future research.

References

- [1] G. Alverson, S. Kahan, R. Korry, C. McCann, and B. Smith, “*Scheduling on the Tera MTA*”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 19–44, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.
- [2] Cray Research, Inc., *Cray T3D System Architecture Overview*. Order number HR-04033.
- [3] D. Das Sharma and D. K. Pradhan, “*Job scheduling in mesh multicomputers*”. In *Intl. Conf. Parallel Processing*, vol. II, pp. 251–258, Aug 1994.
- [4] D. G. Feitelson, “*Packing schemes for gang scheduling*”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), Springer-Verlag, 1996. Lecture Notes in Computer Science Vol. 1162.
- [5] D. G. Feitelson, *A Survey of Scheduling in Multiprogrammed Parallel Systems*. Research Report RC 19790 (87657), IBM T. J. Watson Research Center, Oct 1994.
- [6] D. G. Feitelson and B. Nitzberg, “*Job characteristics of a production parallel scientific workload on the NASA Ames iPSC/860*”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 337–360, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.

- [7] D. G. Feitelson and L. Rudolph, "Distributed hierarchical control for parallel processing". *Computer* **23(5)**, pp. 65–77, May 1990.
- [8] D. G. Feitelson and L. Rudolph, "Evaluation of design choices for gang scheduling using distributed hierarchical control". *J. Parallel & Distributed Comput.* **35(1)**, pp. 18–34, May 1996.
- [9] D. G. Feitelson and L. Rudolph, "Gang scheduling performance benefits for fine-grain synchronization". *J. Parallel & Distributed Comput.* **16(4)**, pp. 306–318, Dec 1992.
- [10] D. G. Feitelson and L. Rudolph, "Parallel job scheduling: issues and approaches". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 1–18, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.
- [11] B. Gorda and R. Wolski, "Time sharing massively parallel machines". In *Intl. Conf. Parallel Processing*, vol. II, pp. 214–217, Aug 1995.
- [12] R. L. Henderson, "Job scheduling under the portable batch system". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 279–294, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.
- [13] S. Hotovy, "Workload evolution on the Cornell Theory Center IBM SP2". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), Springer-Verlag, 1996. Lecture Notes in Computer Science Vol. 1162.
- [14] Intel Corp., *iPSC/860 Multi-User Accounting, Control, and Scheduling Utilities Manual*. Order number 312261-002, May 1992.
- [15] K. Li and K-H. Cheng, "A two-dimensional buddy system for dynamic resource allocation in a partitionable mesh connected system". *J. Parallel & Distributed Comput.* **12(1)**, pp. 79–83, May 1991.
- [16] D. Lifka, "The ANL/IBM SP scheduling system". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 295–303, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.
- [17] C. McCann, R. Vaswani, and J. Zahorjan, "A dynamic processor allocation policy for multiprogrammed shared-memory multiprocessors". *ACM Trans. Comput. Syst.* **11(2)**, pp. 146–178, May 1993.
- [18] C. McCann and J. Zahorjan, "Scheduling memory constrained jobs on distributed memory parallel computers". In *SIGMETRICS Conf. Measurement & Modeling of Comput. Syst.*, pp. 208–219, May 1995.
- [19] J. K. Ousterhout, "Scheduling techniques for concurrent systems". In *3rd Intl. Conf. Distributed Comput. Syst.*, pp. 22–30, Oct 1982.

- [20] E. W. Parsons and K. C. Sevcik, "Multiprocessor scheduling for high-variability service time distributions". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 127-145, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.
- [21] R. C. Regis, "Multiserver queueing models of multiprocessing systems". *IEEE Trans. Comput.* **C-22(8)**, pp. 736-745, Aug 1973.
- [22] E. Rosti, E. Smirni, L. W. Dowdy, G. Serazzi, and B. M. Carlson, "Robust partitioning schemes of multiprocessor systems". *Performance Evaluation* **19(2-3)**, pp. 141-165, Mar 1994.
- [23] E. Rosti, E. Smirni, G. Serazzi, and L. W. Dowdy, "Analysis of non-work-conserving processor partitioning policies". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 165-181, Springer-Verlag, 1995. Lecture Notes in Computer Science Vol. 949.
- [24] K. C. Sevcik, "Application scheduling and processor allocation in multiprogrammed parallel processing systems". *Performance Evaluation* **19(2-3)**, pp. 107-140, Mar 1994.
- [25] K. C. Sevcik, "Characterization of parallelism in applications and their use in scheduling". In *SIGMETRICS Conf. Measurement & Modeling of Comput. Syst.*, pp. 171-180, May 1989.
- [26] A. Tucker and A. Gupta, "Process control and scheduling issues for multiprogrammed shared-memory multiprocessors". In *12th Symp. Operating Systems Principles*, pp. 159-166, Dec 1989.
- [27] M. Wan, R. Moore, G. Kremenek, and K. Steube, "A batch scheduler for the Intel Paragon MPP system with a non-contiguous node allocation algorithm". In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), Springer-Verlag, 1996. Lecture Notes in Computer Science Vol. 1162.
- [28] K. Windisch, V. Lo, D. Feitelson, B. Nitzberg, and R. Moore, "A comparison of workload traces from two production parallel machines". In *Frontiers '96*, Oct 1996.
- [29] Q. Yang and H. Wang, "A new graph approach to minimizing processor fragmentation in hypercube multiprocessors". *IEEE Trans. Parallel & Distributed Syst.* **4(10)**, pp. 1165-1171, Oct 1993.

*This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

Technical Information Department • Lawrence Livermore National Laboratory
University of California • Livermore, California 94551

