Job Scheduling for the BlueGene/L System

Elie Krevat[1], José G. Castaños[2], and José E. Moreira[2]

[1] krevat@mit.edu
Massachusetts Institute of Technology
Cambridge, MA 02139-4307
[2] {castanos,jmoreira}@us.ibm.com
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598-0218

Abstract. BlueGene/L is a massively parallel cellular architecture system with a toroidal interconnect. Cellular architectures with a toroidal interconnect are effective at producing highly scalable computing systems, but typically require job partitions to be both rectangular and contiguous. These restrictions introduce fragmentation issues that affect the utilization of the system and the wait time and slowdown of queued jobs. We propose to solve these problems for the BlueGene/L system through scheduling algorithms that augment a baseline first come first serve (FCFS) scheduler. Restricting ourselves to space-sharing techniques, which constitute a simpler solution to the requirements of cellular computing, we present simulation results for migration and backfilling techniques on BlueGene/L. These techniques are explored individually and jointly to determine their impact on the system. Our results demonstrate that migration can be effective for a pure FCFS scheduler but that backfilling produces even more benefits. We also show that migration can be combined with backfilling to produce more opportunities to better utilize a parallel machine.

1 Introduction

BlueGene/L (BG/L) is a massively parallel cellular architecture system. 65,536 self-contained computing nodes, or *cells*, are interconnected in a three-dimensional toroidal pattern [19]. In that pattern, each cell is directly connected to its six nearest neighbors, two each along the x, y, and z axes. Three-dimensional torus interconnects are simple, modular, and scalable, particularly when compared with systems that have a separate, typically multistage, interconnection network [13]. Examples of successful toroidal-interconnected parallel systems include the Cray T3D and T3E machines [11].

There is, however, a price to pay with toroidal interconnects. We cannot view the system as a simple fully-connected interconnection network of nodes that are equidistant to each other (*i.e.*, a flat network). In particular, we lose an important feature of systems like the IBM RS/6000 SP, which lets us pick any set of nodes for execution of a parallel job, irrespective of their physical location in the machine [1]. In a toroidal-interconnected system, the spatial allocation of nodes to jobs is of critical importance. In most toroidal systems, including BG/L, job partitions must be both rectangular (in a multidimensional sense) and contiguous. It has been shown by Feitelson and Jette [7] that, because of these restrictions, significant machine fragmentation occurs in a toroidal system. Fragmentation results in low system utilization and high wait time for queued jobs.

In this paper, we analyze a set of strictly space-sharing scheduling techniques to improve system utilization and reduce the wait time of jobs for the BG/L system. Time-sharing techniques such as gang-scheduling are not explored since these types of schedulers demand more hardware resources than are practically available in a cellular computing environment. We analyze the two techniques of backfilling [8, 14, 17] and migration [3, 20] in the context of a toroidal-interconnected system. Backfilling is a technique that moves lower priority jobs ahead of other higher priority jobs, as long as execution of the higher priority jobs is not delayed. Migration moves jobs around the toroidal machine, performing on-the-fly defragmentation to create larger contiguous free space for waiting jobs.

We conduct a simulation-based study of the impact of our scheduling algorithms on the system performance of BG/L. Using actual job logs of supercomputing centers, we measure the impact of migration and backfilling as enhancements to a first-come first-serve (FCFS) job scheduling policy. Migration is shown to be effective in improving maximum system utilization while enforcing a strict FCFS policy. We also find that backfilling, which bypasses the FCFS order, can lead to even higher utilization and lower wait times. Finally, we show that there is a small benefit from combining backfilling and migration.

The rest of this paper is organized as follows. Section 2 discusses the scheduling algorithms used to improve job scheduling on a toroidal-interconnected parallel system. Section 3 describes the simulation procedure to evaluate these algorithms and presents our simulation results. Section 4 describes related work and suggests future work opportunities. Finally, Section 5 presents the conclusions.

2 Scheduling Algorithms

System utilization and average job wait time in a parallel system can be improved through better job scheduling algorithms [4, 5, 7, 9, 10, 12, 14–17, 21, 22, 26]. The opportunity for improvement over a simple first-come first-serve (FCFS) scheduler is much greater for toroidal interconnected systems because of the fragmentation issues discussed in Section 1. The following section describes four job scheduling algorithms that we evaluate in the context of BG/L. In all algorithms, arriving jobs are first placed in a queue of waiting jobs, prioritized according to the order of arrival. The scheduler is invoked for every job arrival and job termination event in order to schedule new jobs for execution.

Scheduler 1: First Come First Serve (FCFS). For FCFS, we adopt the heuristic of traversing the waiting queue in order and scheduling each job in a way that maximizes the largest free rectangular partition remaining in the torus. For each job of size p, we try all the possible rectangular shapes of size p that fit in the torus. For each shape, we try all the legal allocations in the torus that do not conflict with running jobs. Finally, we select the shape and allocation that results in the maximal largest free rectangular partition remaining after allocation of this job. We stop when we find the first job in the queue that cannot be scheduled.

A valid rectangular partition does not always exist for a job. There are job sizes which are always impossible for the torus, such as prime numbers greater than the largest dimension size. Because job sizes are known at job arrival time, before execution, jobs with impossible sizes are modified to request the next largest possible size. Additionally, there are legal job sizes that cannot be scheduled because of the current state of the torus. Therefore, if a particular job of size p cannot be scheduled, but some free partition of size $q > p$ exists, the job will be increased in size by the minimum amount required to schedule it. For example, consider a 4×4 (two-dimensional) torus with a single free partition of size 2×2. If a user submits a job requesting 3 nodes, that job cannot be run. The scheduler increases the job size by one, to 4, and successfully schedules the job.

Determining the size of the largest rectangular partition in a given three-dimensional torus is the most time-intensive operation required to implement the maximal partition heuristic. When considering a torus of shape $M \times M \times M$, a straightforward exhaustive search of all possible partitions takes $O(M^9)$ time. We have developed a more efficient algorithm that computes incremental projections of planes and uses dynamic programming techniques. This projection algorithm has complexity $O(M^5)$ and is described in Appendix A.

An FCFS scheduler that searches the torus in a predictable incremental fashion, implements the maximal partition heuristic, and modifies job sizes when necessary is the simplest algorithm considered, against which more sophisticated algorithms are compared.

Scheduler 2: FCFS With Backfilling. Backfilling is a space-sharing optimization technique. With backfilling, we can bypass the priority order imposed by the job queuing policy. This allows a lower priority job j to be

scheduled before a higher priority job i as long as this reschedule does not delay the estimated start time of job i.

The effect of backfilling on a particular schedule for a one-dimensional machine can be visualized in Figure 1. Suppose we have to schedule five jobs, numbered from 1 to 5 in order of arrival. Figure 1(a) shows the schedule that would be produced by a FCFS policy without backfilling. Note the empty space between times T_1 and T_2, while job 3 waits for job 2 to finish. Figure 1(b) shows the schedule that would be produced by a FCFS policy with backfilling. The empty space was filled with job 5, which can be executed before job 3 without delaying it.

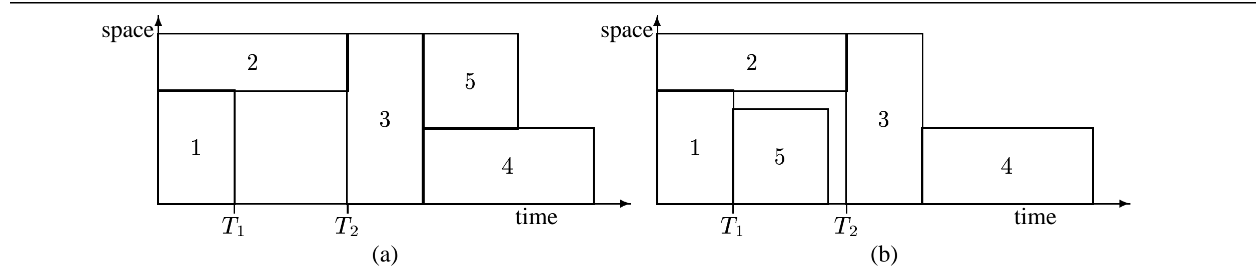

Fig. 1. FCFS policy without (a) and with (b) backfilling. Job numbers correspond to their position in the priority queue.

The backfilling algorithm seeks to increase system utilization without job starvation. It requires an estimation of job execution time, which is usually not very accurate. However, previous work [8, 18, 23] has shown that overestimating execution time does not significantly affect backfilling results. Backfilling has been shown to increase system utilization in a fair manner on an IBM RS/6000 SP [8, 23].

Backfilling is used in conjunction with the FCFS scheduler and is only invoked when there are jobs in the waiting queue and FCFS halts because a job does not fit in the torus. A reservation time for the highest-priority job is then calculated, based on the worst case execution time of jobs currently running in the torus. The reservation guarantees that the job will be scheduled no later than that time, and if jobs end earlier than expected the reservation time may improve. Then, if there are additional jobs in the waiting queue, a job is scheduled out of order so long as it does not prevent the first job in the queue from being scheduled at the reservation time. (Jobs behind the first one may be delayed.)

Just as the FCFS scheduler dynamically increases the size of jobs that cannot be scheduled with their current size, similar situations may arise during backfilling. Unlike FCFS, however, the size increase is performed more conservatively during backfilling, because there are other jobs in the queue which might better utilize the free nodes of the torus. Therefore, a parameter I specifies the maximum size by which the scheduler will increase a job. For example, by setting $I = 1$ (our default value), backfilling increases a job size by at most one node. This parameter is used only during the backfilling phase of scheduling; the FCFS phase will always increase the first job in the queue as much as is required to schedule it.

Scheduler 3: FCFS With Migration. The migration algorithm rearranges the running jobs in the torus in order to increase the size of the maximal contiguous rectangular free partition. Migration in a toroidal-interconnected system compacts the running jobs and counteracts the effects of fragmentation.

While migration does not require any more information than FCFS, it may require additional hardware and software functionality. This paper does not attempt to quantify the overhead of that functionality. However, accepting that this overhead exists, migration is only undertaken when the expected benefits are deemed substantial. The decision to migrate is therefore based on two parameters: FN_{tor}, the ratio of free nodes in the system compared to the size of the torus, and FN_{max}, the fraction of free nodes contained in the

maximal free partition. In order for migration to establish a significant larger maximal free partition, FN_{tor} must be sufficiently high and FN_{max} must be sufficiently low. Section 3.4 contains further analysis of these parameters.

The migration process is undertaken immediately after the FCFS phase fails to schedule a job in the waiting queue. Jobs already running in the torus are organized in a queue of migrating jobs sorted by size, from largest to smallest. Each job is then reassigned a new partition, using the same algorithm as FCFS and starting with an empty torus. After migration, FCFS is performed again in an attempt to start more jobs in the rearranged torus.

In order to ensure that all jobs fit in the torus after migration, job sizes are not increased if a reassignment requires a larger size to fit in the torus. Instead, the job is removed from the queue of migrating jobs, remaining in its original partition, and reassignment begins again for all remaining jobs in the queue. If the maximal free partition size after migration is worse than the original assignment, which is possible but generally infrequent under the current scheduling heuristics, migration is not performed.

Scheduler 4: FCFS with Backfilling and Migration. Backfilling and migration are independent scheduling concepts, and an FCFS scheduler may implement both of these functions simultaneously. First, we schedule as many jobs as possible via FCFS. Next, we rearrange the torus through migration to minimize fragmentation, and then repeat FCFS. Finally, the backfilling algorithm from Scheduler 2 is performed to make a reservation for the highest-priority job and attempt to schedule jobs with lower priority so long as they do not conflict with the reservation. The combination of these policies should lead to an even more efficient utilization of the torus. For simplicity, we call this scheduling technique, that combines backfilling and migration, B+M.

3 Experiments

We use a simulation-based approach to perform quantitative measurements of the efficiency of the proposed scheduling algorithms. An event-driven simulator was developed to process actual job logs of supercomputing centers. The results of simulations for all four schedulers were then studied to determine the impact of their respective algorithms. We begin this section with a short overview of the BG/L system. We then describe our simulation environment. We proceed with a discussion of the workload characteristics for the two job logs we consider. Finally, we present the experimental results from the simulations.

3.1 The BlueGene/L System

The BG/L system is organized as a $32 \times 32 \times 64$ three-dimensional torus of nodes (cells). Each node contains processors, memory, and links for interconnecting to its six neighbors. The unit of allocation for job execution in BG/L is a 512-node ensemble organized in an $8 \times 8 \times 8$ configuration. Therefore, BG/L behaves as a $4 \times 4 \times 8$ torus of these *supernodes*. We use this supernode abstraction when performing job scheduling for BG/L.

3.2 The Simulation Environment

The simulation environment models a torus of 128 (super)nodes in a three-dimensional $4 \times 4 \times 8$ configuration. The event-driven simulator receives as input a job log and the type of scheduler (FCFS, Backfill, Migration, or B+M) to simulate. There are four primary events in the simulator: (1) an *arrival event* occurs when a job is first submitted for execution and placed in the scheduler's waiting queue; (2) a *schedule event* occurs when a job is allocated onto the torus, (3) a *start event* occurs after a standard delay of one second following a schedule event, at which time a job begins to run, and (4) a *finish event* occurs upon completion

of a job, at which point the job is deallocated from the torus. The scheduler is invoked at the conclusion of every event that affects the states of the torus or the waiting queue (*i.e.*, the arrival and finish events).

A job log contains information on the arrival time, execution time, and size of all jobs. Given a torus of size N, and for each job j the arrival time t_j^a, execution time t_j^e and size s_j, the simulation produces values for the start time t_j^s and finish time t_j^f of each job. These results are analyzed to determine the following parameters for each job: (1) wait time $t_j^w = t_j^s - t_j^a$, (2) response time $t_j^r = t_j^f - t_j^a$, and (3) bounded slowdown $t_j^{bs} = \frac{\max(t_j^r, \Gamma)}{max(t_j^e, \Gamma)}$ for $\Gamma = 10$ seconds. The Γ term appears according to recommendations in [8], because some jobs have very short execution time, which may distort the slowdown.

Global system statistics are also determined. Let the simulation time span be $T = \max_{\forall j}(t_j^f) - \min_{\forall k}(t_k^a)$. We then define system utilization (also called *capacity utilized*) as

$$w_{\text{util}} = \sum_{\forall j} \frac{s_j t_j^e}{TN}. \tag{1}$$

Similarly, let $f(t)$ denote the number of free nodes in the torus at time t and $q(t)$ denote the total number of nodes requested by jobs in the waiting queue at time t. Then, the total amount of unused capacity in the system, w_{unused}, is defined as:

$$w_{\text{unused}} = \int_{\min(t_j^a)}^{\max(t_j^f)} \frac{\max(0, f(t) - q(t))}{TN} dt. \tag{2}$$

This parameter is a measure of the work unused by the system because there is a lack of jobs requesting free nodes. The $\max$ term is included because the amount of unused work cannot be less than zero. The balance of the system capacity is lost despite the presence of jobs that could have used it. The measure of lost capacity in the system, which includes capacity lost because of the inability to schedule jobs and the delay before a scheduled job begins, is then derived as:

$$w_{\text{lost}} = 1 - w_{\text{util}} - w_{\text{unused}} \tag{3}$$

3.3 Workload characteristics

We performed experiments on a 10,000-job span of two job logs obtained from the *Parallel Workloads Archive* [6]. The first log is from NASA Ames's 128-node iPSC/860 machine (from the year 1993). The second log is from the San Diego Supercomputer Center's (SDSC) 128-node IBM RS/6000 SP (from the years 1998-2000). For our purposes, we will treat each node in those two systems as representing one supernode (512-node unit) of BG/L. This is equivalent to scaling all job sizes in the log by 512, which is the ratio of the number of nodes in BG/L to the number of nodes in these 128-node machines. Table 1 presents the workload statistics and Figure 2 summarizes the distribution of job sizes and the contribution of each job size to the total workload of the system. Using these two logs as a basis, we generate logs of varying workloads by multiplying the execution time of each job by a coefficient c, mostly varying c from 0.7 to 1.4 in increments of 0.05. Simulations are performed for all scheduler types on each of the logs. With these modified logs, we plot wait time and bounded slowdown as a function of system utilization.

3.4 Simulation Results

Figures 3 and 4 present plots of average job wait time (t_j^w) and average job bounded slowdown (t_j^{bs}), respectively, *vs* system utilization (w_{util}) for each of the four schedulers considered and each of the two job logs. We observe that the overall shapes of the curves for wait time and bounded slowdown are similar.

Table 1. Statistics for 10,000-job NASA and SDSC logs.

	NASA Ames iPSC/860 log	*SDSC IBM RS/6000 SP log*
Number of nodes:	128	128
Job size restrictions:	powers of 2	none
Job size (nodes)		
Mean:	6.3	9.7
Standard deviation:	14.4	14.8
Workload(node-seconds)		
Mean:	0.881×10^6	7.1×10^6
Standard deviation:	5.41×10^6	25.5×10^6

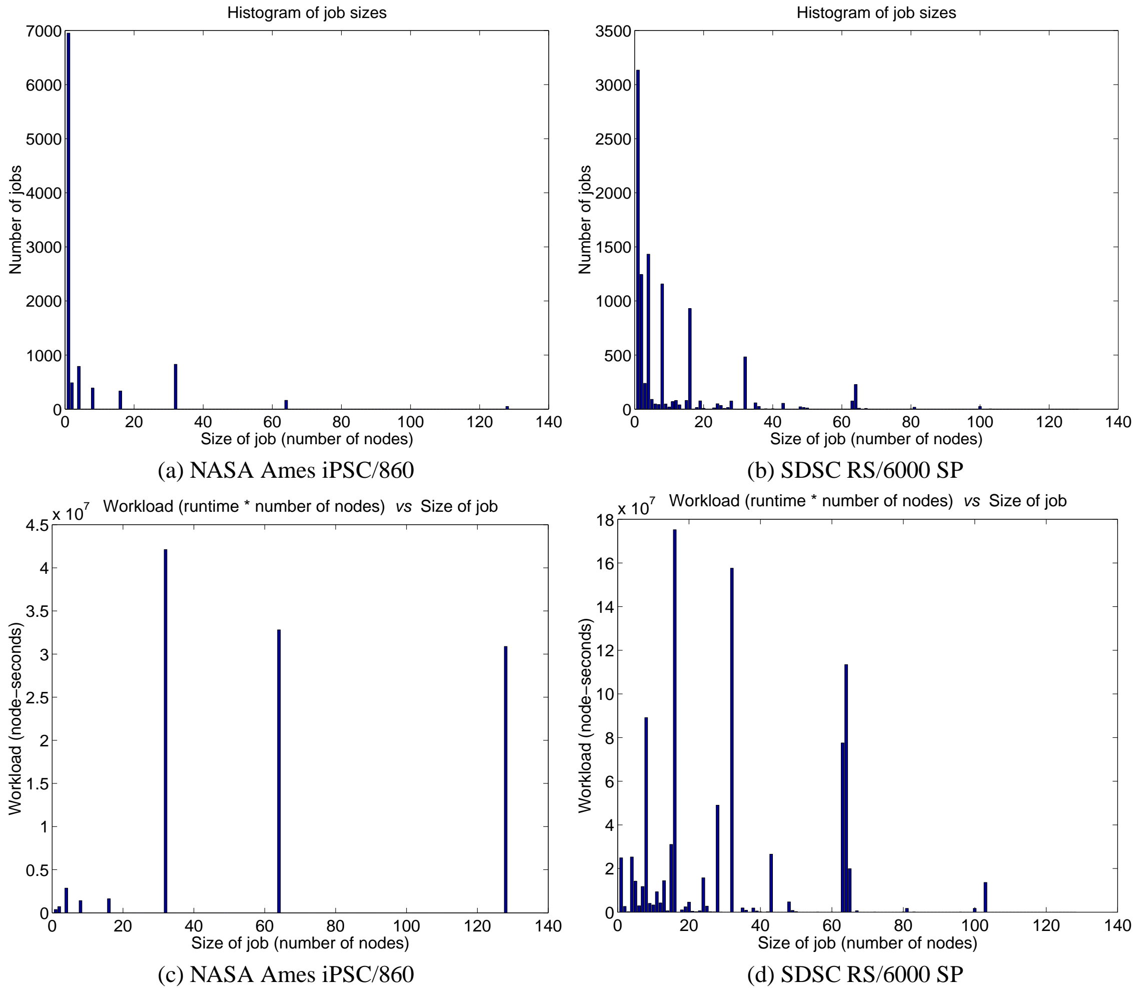

Fig. 2. Job sizes and total workload for NASA Ames iPSC/860((a) and (c)) and San Diego Supercomputer Center (SDSC) IBM RS/6000 SP((b) and (d)).

The most significant performance improvement is attained through backfilling, for both the NASA and SDSC logs. Also, for both logs, there is a certain benefit from migration, whether combined with backfilling or not. We analyze these results from each log separately.

NASA log: All four schedulers provide similar average job wait time and average job bounded slowdown for utilizations up to 65%. The FCFS scheduler saturates at about 77% utilization, whereas the Migration

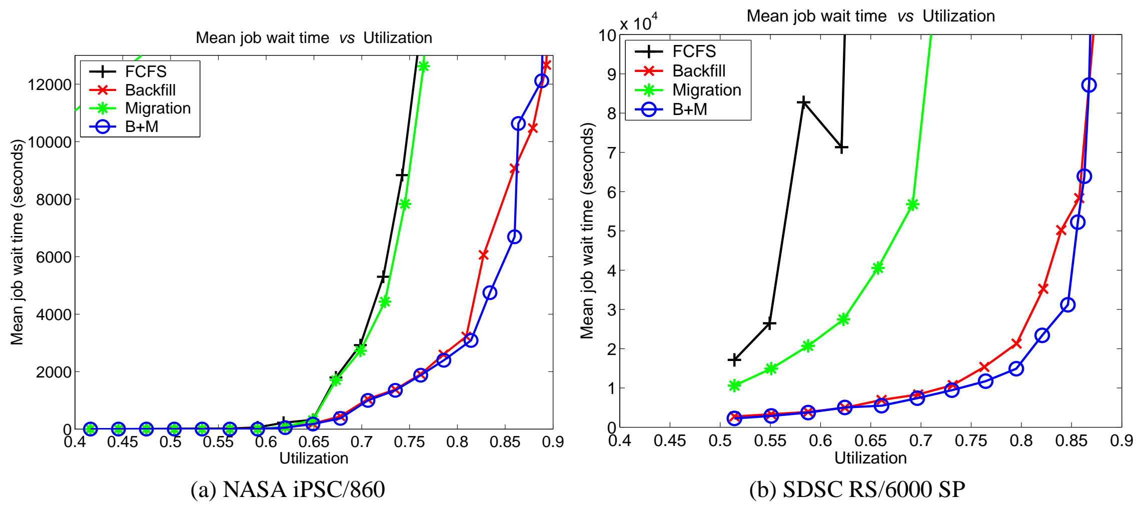

(a) NASA iPSC/860 (b) SDSC RS/6000 SP

Fig. 3. Mean job wait time *vs* utilization for (a) NASA and (b) SDSC logs.

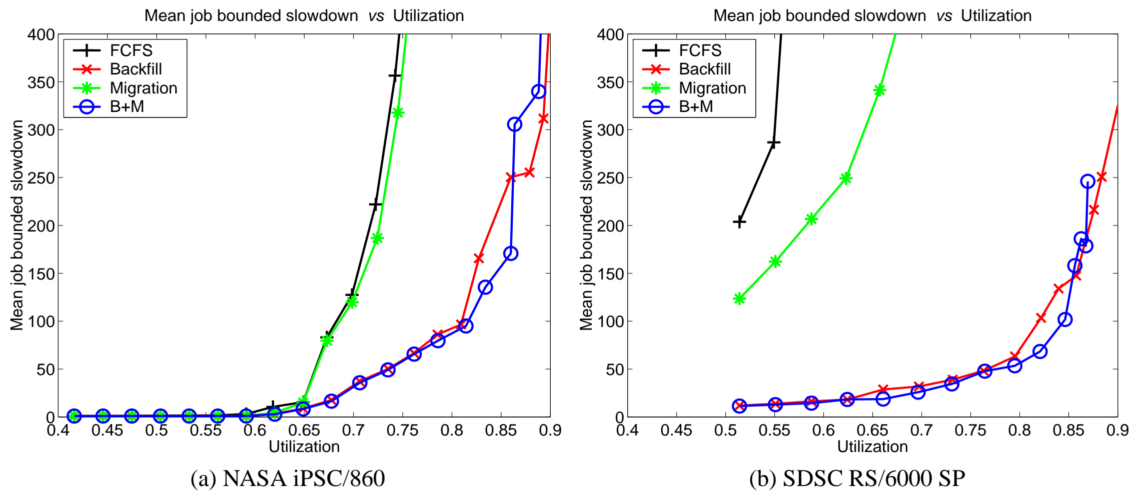

(a) NASA iPSC/860 (b) SDSC RS/6000 SP

Fig. 4. Mean job bounded slowdown *vs* utilization for (a) NASA and (b) SDSC logs.

scheduler saturates at about 80% utilization. Backfilling (with or without migration) allows utilizations above 80% and saturates closer to 90% (the saturation region for these schedulers is shown here by plotting values of $c > 1.4$). We note that migration provides only a small improvement in wait time and bounded slowdown for most of the utilization range, and the additional benefits of migration with backfilling becomes unpredictable for utilization values close to the saturation region. In the NASA log, all jobs are of sizes that are powers of two, which results in a good packing of the torus. Therefore, the benefits of migration are limited.

SDSC log: With the SDSC log, the FCFS scheduler saturates at 63%, while the stand-alone Migration scheduler saturates at 73%. In this log, with jobs of more varied sizes, fragmentation occurs more frequently. Therefore, migration has a much bigger impact on FCFS, significantly improving the range of utilizations at which the system can operate. However, we note that when backfilling is used there is again only a small additional benefit from migration, more noticeable for utilizations between 75 and 85%. Utilization above 85% can be achieved, but only with exponentially growing wait time and bounded slowdown, independent of performing migration.

Figure 5 presents a plot of average job bounded slowdown (t_j^{bs}) *vs* system utilization (w_{util}) for each of the four schedulers considered and each of the two job logs. We also include results from the simulation of a fully-connected (*flat*) machine, with and without backfilling. (A fully-connected machine does not suffer from fragmentation.) This allows us to assess the effectiveness of our schedulers in overcoming the difficulties imposed by a toroidal interconnect. The overall shapes of the curves for wait time are similar to those for bounded slowdown.

Migration by itself cannot make the results for a toroidal machine as good as those for a fully connected machine. For the SDSC log, in particular, a fully connected machine saturates at about 80% utilization with just the FCFS scheduler. For the NASA log, results for backfilling with or without migration in the toroidal machine are just as good as the backfilling results in the fully connected machine. For utilizations above 85% in the SDSC log, not even a combination of backfilling and migration will perform as well as backfilling on a fully connected machine.

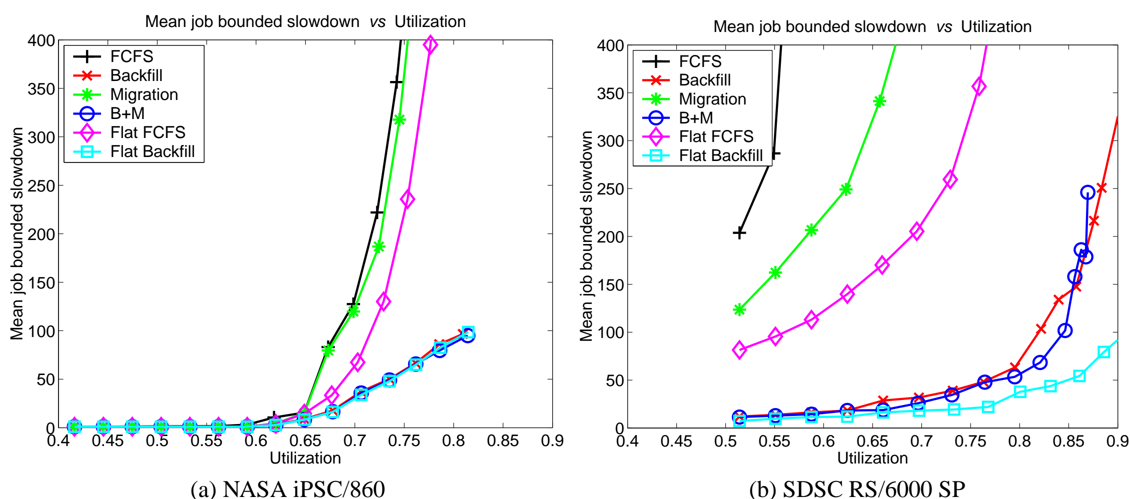

(a) NASA iPSC/860 (b) SDSC RS/6000 SP

Fig. 5. Mean job bounded slowdown *vs* utilization for the NASA and SDSC logs, comparing toroidal and flat machines.

Figure 6 plots the number of migrations performed and the average time between migrations *vs* system utilization for both workloads. We show results for the number of *total* migrations attempted, the number of *successful* migrations, and the maximum possible number of successful migrations (*max successful*). As described in Section 2, the parameters which determine if a migration should be attempted are FN_{tor}, the ratio of free nodes in the system compared to the size of the torus, and FN_{max}, the fraction of free nodes contained in the maximal free partition. According to our standard migration policy, a migration is only attempted when $FN_{tor} \geq 0.1$ and $FN_{max} \leq 0.7$. A *successful* migration is defined as a migration attempt that improves the maximal free partition size. The *max successful* value is the number of migrations that are successful when a migration is always attempted (*i.e.*, $FN_{tor} \geq 0.0$ and $FN_{max} \leq 1.0$).

Almost all migration attempts were successful for the NASA log. This property of the NASA log is a reflection of the better packing caused by having jobs that are exclusively power of two in size. For the SDSC log, we notice that many more total attempts are made while about 80% of them are successful. If we always try to migrate every time the state of the torus is modified, no more than 20% of these migrations are successful, and usually much less.

For the NASA log, the number of migrations increases linearly while the average time between these migrations varies from about 90 to 30 minutes, depending on the utilization level and its effect on the amount of fragmentation in the torus. In contrast to the NASA log, the number of migrations in the SDSC log do not

increase linearly as utilization levels increase. Instead, the relationship is closer to an elongated bell curve. As utilization levels increase, at first migration attempts and successes also increase slightly to a fairly steady level. Around the first signs of saturation the migrations tend to decrease (*i.e.*, at around 70% utilization for the Migration scheduler and 77% for B+M). Even though the number of successful migrations is greater for the SDSC log, the average time between migrations is still longer as a result of the larger average job execution time.

Most of the benefit of migration is achieved when we only perform migration according to our parameters. Applying these parameters has two main advantages: we reduce the frequency of migration attempts so as not to always suffer the required overhead of migration, and we increase the average benefits of a successful migration. This second advantage is apparent when we compare the mean job wait time results for our standard FN_{tor} and FN_{max} settings to that of the scheduler that always attempts to migrate. Even though the maximum possible number of successful migrations is sometimes twice as many as our actual number of successes, Figure 7 reveals that the additional benefit of these successful migrations is very small.

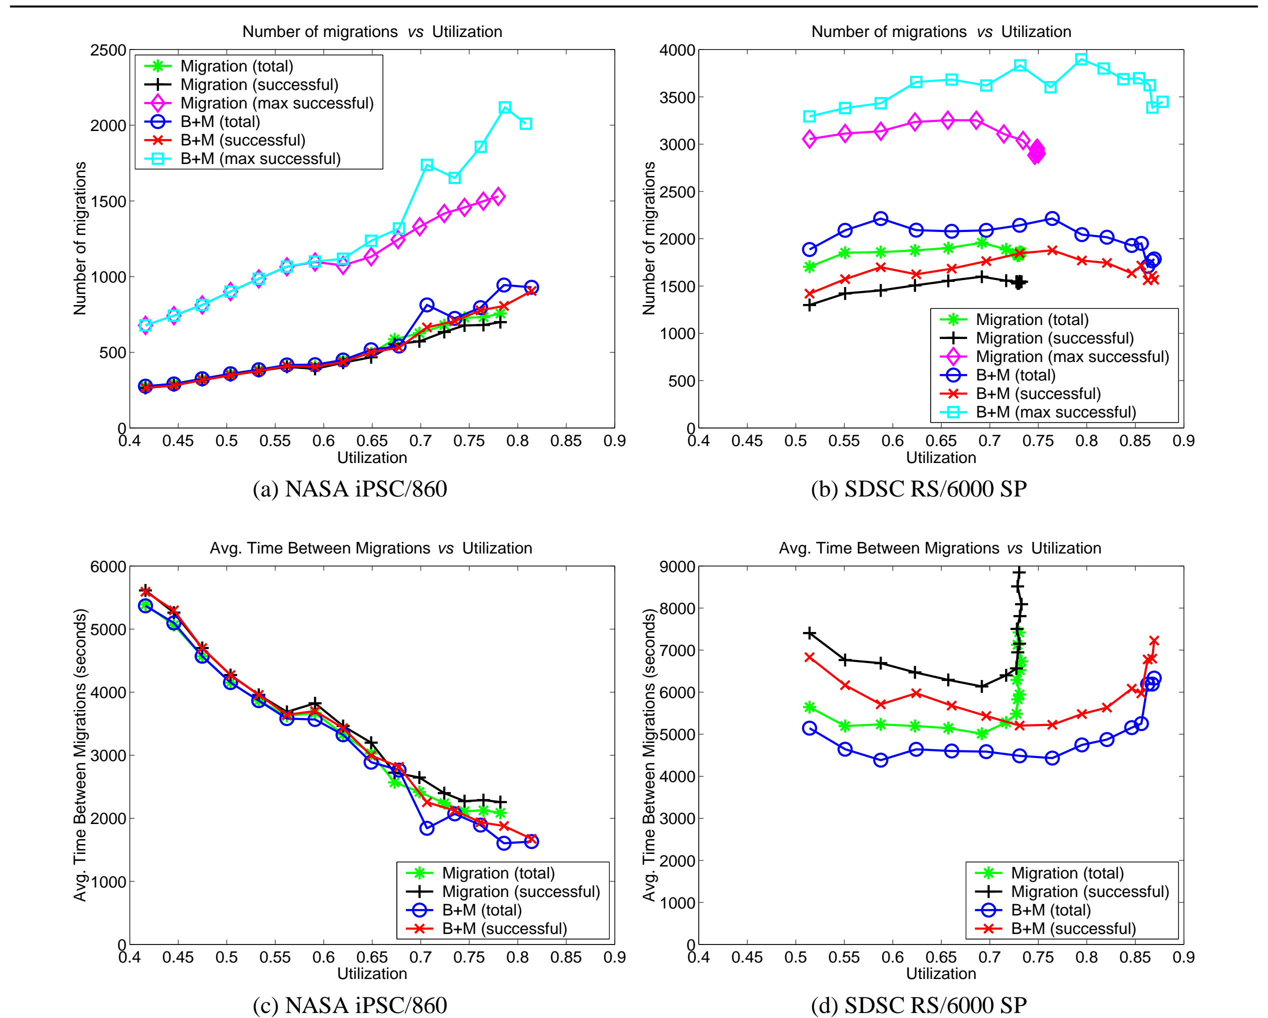

(a) NASA iPSC/860 (b) SDSC RS/6000 SP

(c) NASA iPSC/860 (d) SDSC RS/6000 SP

Fig. 6. Number of total, successful, and maximum possible successful migrations *vs* utilization ((a) and (b)), and average time between migrations *vs* utilization ((c) and (d)).

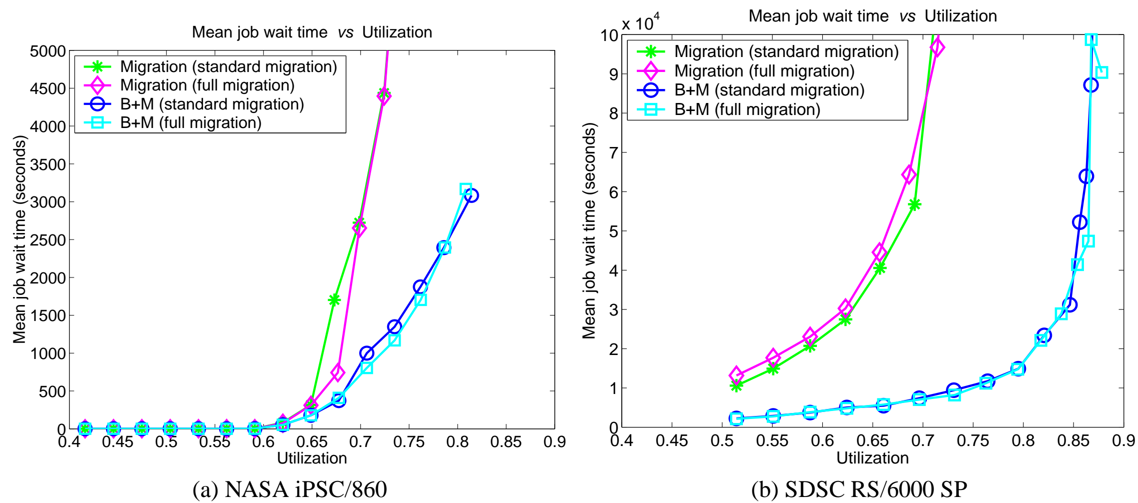

(a) NASA iPSC/860 (b) SDSC RS/6000 SP

Fig. 7. Mean job wait time *vs* utilization for the NASA and SDSC logs, comparing the standard migration policy to a full migration policy that always attempts to migrate.

We complete this section with an analysis of results for system capacity utilized, unused capacity, and lost capacity. The results for each scheduler type and both standard job logs ($c = 1.0$) are plotted in Figure 8. The utilization improvements for the NASA log are barely noticeable – again, because its jobs fill the torus more compactly. The SDSC log, however, shows the greatest improvement when using B+M over FCFS, with a 15% increase in capacity utilized and a 54% decrease in the amount of capacity lost. By themselves, the Backfill and Migration schedulers each increase capacity utilization by 15% and 13%, respectively, while decreasing capacity loss by 44% and 32%, respectively. These results show that B+M is significantly more effective at transforming lost capacity into unused capacity. Under the right circumstances, it should be possible to utilize this unused capacity more effectively.

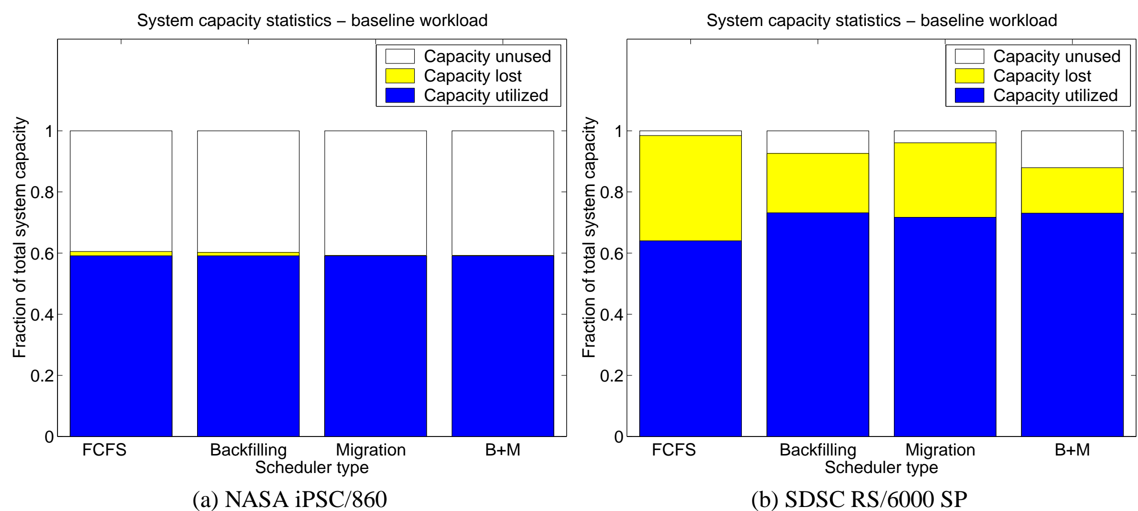

(a) NASA iPSC/860 (b) SDSC RS/6000 SP

Fig. 8. Capacity utilized, lost, and unused as a fraction of the total system capacity.

4 Related and Future Work

The topics of our work have been the subject of extensive previous research. In particular, [8, 14, 17] have shown that backfilling on a flat machine like the IBM RS/6000 SP is an effective means of improving quality of service. The benefits of combining migration and gang-scheduling have been demonstrated both for flat machines [24, 25] and toroidal machines like the Cray T3D [7]. The results in [7] are particularly remarkable, as system utilization was improved from 33%, with a pure space-sharing approach, to 96% with a combination of migration and gang-scheduling. The work in [21] discusses techniques to optimize spatial allocation of jobs in mesh-connected multicomputers, including changing the job size, and how to combine spatial- and time-sharing scheduling algorithms. An efficient job scheduling technique for a three-dimensional torus is described in [2]. This paper, therefore, builds on this previous research by applying a combination of backfilling and migration algorithms, exclusively through space-sharing techniques, to improve system performance on a toroidal-interconnected system.

Future work opportunities can further build on the results of this paper. The impact of different FCFS scheduling heuristics for a torus, besides the largest free partition heuristic currently used, can be studied. It is also important to identify how the current heuristic relates to the optimal solution in different cases. Additional study of the parameters I, FN_{tor}, and FN_{max} may determine further tradeoffs associated with partition size increases and more or less frequent migration attempts. Finally, while we do not attempt to implement complex time-sharing schedulers such as those used in gang-scheduling, a more limited time-sharing feature may be beneficial. Preemption, for example, allows for the suspension of a job until it is resumed at a later time. These time-sharing techniques may provide the means to further enhance the B+M scheduler and make the system performance of a toroidal-interconnected machine more similar to that of a flat machine.

5 Conclusions

We have investigated the behavior of various scheduling algorithms to determine their ability to increase processor utilization and decrease job wait time in the BG/L system. We have shown that a scheduler which uses only a backfilling algorithm performs better than a scheduler which uses only a migration algorithm, and that migration is particularly effective under a workload that produces a large amount of fragmentation (*i.e.*, when many small to mid-sized jobs of varied sizes represent much of the workload). Migration has a significant implementation overhead but it does not require any additional information besides what is required by the FCFS scheduler. Backfilling, on the other hand, does not have a significant implementation overhead but requires additional information pertaining to the execution time of jobs.

Simulations of FCFS, backfilling, and migration space-sharing scheduling algorithms have shown that B+M, a scheduler which implements all of these algorithms, shows a small performance improvement over just FCFS and backfilling. However, B+M does convert significantly more lost capacity into unused capacity than just backfilling. Additional enhancements to the B+M scheduler may harness this unused capacity to provide further system improvements. Even with the performance enhancements of backfilling and migration techniques, a toroidal-interconnected machine such as BG/L can only approximate the job scheduling efficiency of a fully connected machine in which all nodes are equidistant.

References

1. T. Agerwala, J. L. Martin, J. H. Mirza, D. C. Sadler, D. M. Dias, and M. Snir. **SP2 system architecture**. *IBM Systems Journal*, 34(2):152–184, 1995.
2. H. Choo, S.-M. Yoo, and H. Y. Youn. **Processor Scheduling and Allocation for 3D Torus Multicomputer Systems**. *IEEE Transactions on Parallel and Distributed Systems*, 11(5):475–484, May 2000.
3. D. H. J. Epema, M. Livny, R. van Dantzig, X. Evers, and J. Pruyne. **A worldwide flock of Condors: Load sharing among workstation clusters**. *Future Generation Computer Systems*, 12(1):53–65, May 1996.
4. D. G. Feitelson. **A Survey of Scheduling in Multiprogrammed Parallel Systems**. Technical Report RC 19790 (87657), IBM T. J. Watson Research Center, October 1994.
5. D. G. Feitelson. Packing schemes for gang scheduling. In *Job Scheduling Strategies for Parallel Processing, IPPS'96 Workshop*, volume 1162 of *Lecture Notes in Computer Science*, pages 89–110, Berlin, March 1996. Springer-Verlag.
6. D. G. Feitelson. **Parallel Workloads Archive**. URL: `http://www.cs.huji.ac.il/labs/parallel/workload/index.html`, 2001.
7. D. G. Feitelson and M. A. Jette. **Improved Utilization and Responsiveness with Gang Scheduling**. In *IPPS'97 Workshop on Job Scheduling Strategies for Parallel Processing*, volume 1291 of *Lecture Notes in Computer Science*, pages 238–261. Springer-Verlag, April 1997.
8. D. G. Feitelson and A. M. Weil. **Utilization and predictability in scheduling the IBM SP2 with backfilling**. In *12th International Parallel Processing Symposium*, pages 542–546, April 1998.
9. H. Franke, J. Jann, J. E. Moreira, and P. Pattnaik. **An Evaluation of Parallel Job Scheduling for ASCI Blue-Pacific**. In *Proceedings of SC99, Portland, OR*, November 1999. IBM Research Report RC21559.
10. B. Gorda and R. Wolski. **Time Sharing Massively Parallel Machines**. In *International Conference on Parallel Processing*, volume II, pages 214–217, August 1995.
11. D. Hyatt. **A Beginner's Guide to the Cray T3D/T3E**. URL: `http://www.jics.utk.edu/SUPER_COMPS/T3D/T3D_guide/T3D_guideJul97.html`, July 1997.
12. H. D. Karatza. **A Simulation-Based Performance Analysis of Gang Scheduling in a Distributed System**. In *Proceedings 32nd Annual Simulation Symposium*, pages 26–33, San Diego, CA, April 11-15 1999.
13. D. H. Lawrie. **Access and Alignment of Data in an Array Processor**. *IEEE Transactions on Computers*, 24(12):1145–1155, December 1975.
14. D. Lifka. **The ANL/IBM SP scheduling system**. In *IPPS'95 Workshop on Job Scheduling Strategies for Parallel Processing*, volume 949 of *Lecture Notes in Computer Science*, pages 295–303. Springer-Verlag, April 1995.
15. J. E. Moreira, W. Chan, L. L. Fong, H. Franke, and M. A. Jette. **An Infrastructure for Efficient Parallel Job Execution in Terascale Computing Environments**. In *Proceedings of SC98, Orlando, FL*, November 1998.
16. U. Schwiegelshohn and R. Yahyapour. **Improving First-Come-First-Serve Job Scheduling by Gang Scheduling**. In *IPPS'98 Workshop on Job Scheduling Strategies for Parallel Processing*, March 1998.
17. J. Skovira, W. Chan, H. Zhou, and D. Lifka. **The EASY-LoadLeveler API project**. In *IPPS'96 Workshop on Job Scheduling Strategies for Parallel Processing*, volume 1162 of *Lecture Notes in Computer Science*, pages 41–47. Springer-Verlag, April 1996.
18. W. Smith, V. Taylor, and I. Foster. **Using Run-Time Predictions to Estimate Queue Wait Times and Improve Scheduler Performance**. In *Proceedings of the 5th Annual Workshop on Job Scheduling Strategies for Parallel Processing*, April 1999. In conjunction with IPPS/SPDP'99, Condado Plaza Hotel & Casino, San Juan, Puerto Rico.
19. H. S. Stone. **High-Performance Computer Architecture**. Addison-Wesley, 1993.
20. C. Z. Xu and F. C. M. Lau. **Load Balancing in Parallel Computers: Theory and Practice**. Kluwer Academic Publishers, Boston, MA, 1996.
21. B. S. Yoo and C. R. Das. **Processor Management Techniques for Mesh-Connected Multiprocessors**. In *Proceedings of the International Conference on Parallel Processing (ICPP'95)*, volume 2, pages 105–112, August 1995.
22. K. K. Yue and D. J. Lilja. **Comparing Processor Allocation Strategies in Multiprogrammed Shared-Memory Multiprocessors**. *Journal of Parallel and Distributed Computing*, 49(2):245–258, March 1998.
23. Y. Zhang, H. Franke, J. E. Moreira, and A. Sivasubramaniam. **Improving Parallel Job Scheduling by Combining Gang Scheduling and Backfilling Techniques**. In *Proceedings of IPDPS 2000*, Cancun, Mexico, May 2000.
24. Y. Zhang, H. Franke, J. E. Moreira, and A. Sivasubramaniam. **The Impact of Migration on Parallel Job Scheduling for Distributed Systems**. In *Proceedings of the 6th International Euro-Par Conference*, pages 242–251, August 29 - September 1 2000.
25. Y. Zhang, H. Franke, J. E. Moreira, and A. Sivasubramaniam. **An Analysis of Space- and Time-Sharing Techniques for Parallel Job Scheduling**. In *Job Scheduling Strategies for Parallel Processing, Sigmetrics'01 Workshop*, June 2001.
26. B. B. Zhou, R. P. Brent, C. W. Jonhson, and D. Walsh. **Job Re-packing for Enhancing the Performance of Gang Scheduling**. In *Job Scheduling Strategies for Parallel Processing, IPPS'99 Workshop*, pages 129–143, April 1999. LNCS 1659.

A Projection Of Partitions (POP) Algorithm

In a given three-dimensional torus of shape $M \times M \times M$ where some nodes have been allocated for jobs, the POP algorithm provides a $O(M^5)$ time algorithm for determining the size of the largest free rectangular partition. This algorithm is a substantial improvement over an exhaustive search algorithm that takes $O(M^9)$ time.

Let FREEPART $= \{\langle B, S\rangle \mid B$ is a base location (i, j, k) and S is a partition size (a, b, c) such that $\forall\, x, y, z,\ i \le x < (i + a),\ j \le y < (j + b),\ k \le z < (k + c)$, node $(x \bmod M,\ y \bmod M,\ z \bmod M)$ is free$\}$. POP narrows the scope of the problem by determining the largest rectangular partition $P \in$ FREEPART rooted at each of the M^3 possible base locations and then deriving a global maximum. Given a base location, POP works by finding the largest partition first in one dimension, then by projecting adjacent one-dimensional columns onto each other to find the largest partition in two dimensions, and iteratively projecting adjacent two-dimensional planes onto each other to find the largest partition in three dimensions.

First, a partition table of the largest one-dimensional partitions $P \in$ FREEPART is pre-computed for all three dimensions and at every possible base location in $O(M^4)$ time. This is done by iterating through each partition and whenever an allocated node is reached, all entries for the current "row" may be filled in from a counter value, where the counter is incremented for each adjacent free node and reset to zero whenever an additional allocated node is reached.

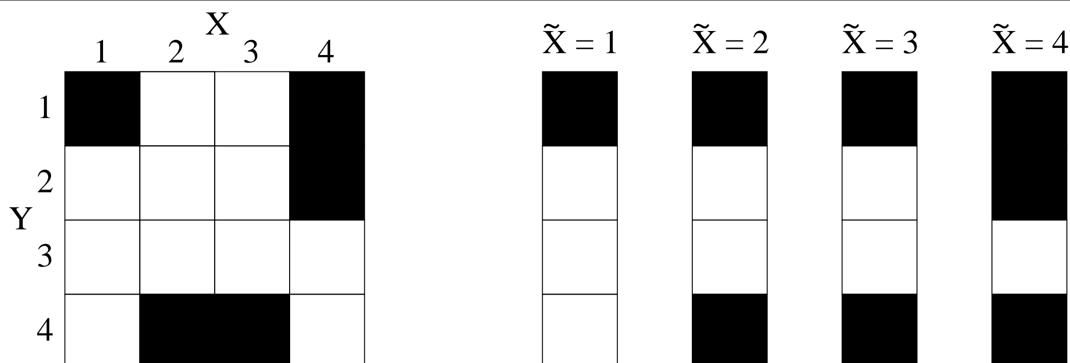

Fig. 9. 2-dimensional POP Algorithm applied to Base Location (1,2): Adjacent 1-dimensional columns are projected onto each other as $\tilde{X}$ is incremented.

For a given base location (i, j, k), we fix one dimension (e.g., k), start a counter $\tilde{X} = i$ in the next dimension, and multiply $\tilde{X}$ by the minimum partition table entry of the third dimension for $(x \bmod M, j, k)$, where x varies as $i \le x \le \tilde{X}$ and $\tilde{X}$ varies as $i \le \tilde{X} \le (i + M)$. As the example in Figure 9 shows, when $\tilde{X} = 1$ for some fixed k at base location $(1, 2, k)$ the partition table entry in the Y dimension will equal 3 since there are 3 consecutive free nodes, and our largest possible partition size is initially set to 3. When $\tilde{X}$ increases to 2, the minimum table entry becomes 2 because of the allocated node at location $(2, 4, k)$ and the largest possible partition size is increased to 4. When $\tilde{X} = 3$, we calculate a new largest possible partition size of 6. Finally, when we come across a partition table entry in the Y dimension of 0 because of the allocated node at location $(4, 2, k)$, we stop increasing $\tilde{X}$. We would also have to repeat a similar calculation along the Y dimension, by starting a counter $\tilde{Y}$.

Finally, this same idea is extended to work for 3 dimensions. Given a similar base location (i, j, k), we start a counter $\tilde{Z}$ in the Z dimension and calculate the maximum two-dimensional partition given the current value of $\tilde{Z}$. Then we project the adjacent two-dimensional planes by incrementing $\tilde{Z}$ and calculating the

largest two-dimensional partition while using the minimum partition table entry of the X and Y dimensions for $(i, j, z \bmod M)$, where z varies as $k \leq z \leq \tilde{Z}$.

Using the initial partition table, it takes $O(M)$ time to calculate a projection for two adjacent planes and to determine the largest two-dimensional partition. Since there are $O(M)$ projections required for $O(M^3)$ base locations, our final algorithm runs in $O(M^5)$ time.

When we implemented this algorithm in our scheduling simulator, we achieved a significant speed improvement. For the original NASA log, scheduling time improved from an average of 0.51 seconds for every successfully scheduled job to 0.16 seconds, while the SDSC log improved from an average of 0.125 seconds to 0.063 seconds. The longest time to successfully schedule a job also improved from 38 seconds to 8.3 seconds in the NASA log, and from 50 seconds to 8.5 seconds in the SDSC log.