

# Experimental Analysis of the Root Causes of Performance Evaluation Results: A Backfilling Case Study

Dror G. Feitelson

School of Computer Science and Engineering

The Hebrew University of Jerusalem

91904 Jerusalem, Israel

## Abstract

Modern computer systems are very complex, and many factors influence their performance. It may therefore happen that seemingly minor variations or assumptions in the performance evaluation procedures actually determine the outcome. We provide a detailed case study, in which two parallel job schedulers based on backfilling are compared, using four different workloads and two different metrics. Simulations show that sometimes the same workload gives different results when using different metrics; this is explained by the fact that the metrics are sensitive to different job classes, and do not measure the performance of the whole workload in an impartial manner. Likewise, the same metric sometimes gives different results when using different workloads; this turned out to depend on subtle details, such as the use of accurate runtime estimates when real estimates are not available, or using a machine size that is not a power of two. More worrisome is the fact that such seemingly minor details had more effect than the basic features of the different workloads, such as the distributions of job sizes and runtimes. This motivates the use of experimental methodology to investigate the effect of different factors, and uncover those that lead to the observed results.

## Keywords:

C.1.4.d Scheduling and task partitioning

C.4.g Measurement, evaluation, modeling, simulation of multiple-processor systems

D.4.8.f Simulation

K.6.2.d Performance and usage measurement

## 1 Introduction

The goal of performance evaluation is often to compare different system designs or implementations. The evaluation is expected to bring out performance differences that will allow for an educated decision regarding what design to employ or what system to buy. Thus it is implicitly assumed that observed performance differences indeed reflect important differences between the systems being studied.

However, performance differences may also be an artifact of the evaluation methodology. The performance of a system is not only a function of the system design and implementation. It may also be affected by the workload to which the system is subjected, and even by the metric being used to gauge the performance. Thus it is important to understand the effect of the workload and metrics on the evaluation. To complicate matters, in some cases the effect is not due to the metric or workload alone, but rather to an interaction between them and the system [6].

We present a case study of a detailed analysis of such a situation. We start by providing the required background: first a description of the systems being compared, which are two rather similar versions of parallel job schedulers (Section 2), and then a description of the methodology, including the workloads used to drive the schedulers, and the performance metrics that were measured (Section 3). Next we report the performance results obtained by running simulations as described above (Section 4). While each set of results seems to answer the question of the relative merit of the two schedulers, the set as a whole includes some discrepancies when using different workloads and metrics. The main part of the paper is the analysis of these results (Sections 5 and 6). It shows that they do not stem from inherent differences between the schedulers, or from significant statistical differences between the workloads (even though such differences do in fact exist). Rather, the different results stem from intricate interactions between the schedulers, the workloads, and the metrics. Moreover, these interactions involve minor details and assumptions that are easy to overlook. We conclude that a detailed analysis as to the source of performance differences is required in order to assess whether the results are credible, and that experimental methodology is needed to verify that suspected factors are indeed the cause of observed differences.

## 2 Parallel Job Scheduling with Backfilling

The domain we will use is the scheduling of parallel jobs for execution on a parallel supercomputer. Such scheduling is typically done in two dimensions: *partitioning*, in which disjoint subsets of processors are allocated to different jobs, and *time slicing*, in which processors serve processes belonging to different jobs. We will focus on partitioning, and specifically, on an optimization called backfilling.

Scheduling jobs using partitioning is akin to packing in 2D. Regard one dimension (say the vertical) as representing processors, and the other (the horizontal) as representing time. A parallel job is a rectangle, representing the use of a certain number of processors for a certain duration of time. The scheduler has to pack these rectangles as tightly as possible, within the space provided by the available resources. In the context of backfilling, the sizes of the rectangles are known: each submitted job comes with a specification of how many processors to use, and an estimate of how long it will run.

Given that jobs come in various sizes, they typically do not pack perfectly. Therefore holes are left in the schedule. Backfilling is the process of trying to fill in these holes using newly submitted small jobs. Two main versions of backfilling have been studied in the literature: EASY backfilling and conservative backfilling.

EASY backfilling is part of the Extensible Argonne Scheduling sYstem developed for the IBM SP1 machine [16]. This version of backfilling is rather aggressive in trying to maximize system utilization. The scheduler is called whenever there are idle processors (e.g. due to the termination




Figure 1: An example of EASY backfilling.

of a running job) and jobs are waiting (because sufficient processors were not available earlier). Its operations can be summarized as follows:

1. Scan the queue of waiting jobs in the order they had arrived (FCFS), and allocate processors as requested.
2. When you reach a job that requires more processors than remain available, calculate when the required processors are expected to become free (based on the running jobs' runtime estimates). Make a reservation to run this job at that time.
3. Continue to scan the queue of waiting jobs. Allocate processors to any job that is small enough (i.e. sufficient processors are available) and will not interfere with the commitment to run the first queued job.

A simple example is given in Figure 1.

EASY backfilling makes a reservation for the first queued job at each stage, but its backfilling may cause delays to other queued jobs [19]. Conservative backfilling aims to reduce this risk. To do so, it makes a reservation for each newly submitted job at some future time, based on the current knowledge of the system state. This commitment may be based on backfilling the new job, but is subject to the condition that such backfilling will not violate any previous commitments.

The two versions of backfilling therefore make different tradeoffs in their backfilling policy. EASY performs as much backfilling as possible, thus reducing the response time of backfilled jobs, at the possible expense of delaying other jobs. Conservative prevents the delaying of jobs, so their response time will not grow in an uncontrolled manner, but this comes at the price of limiting the reduction of response time of backfill jobs. The question is then which of the two approaches is better. Our case study is a set of simulations designed to answer this question.

It should be noted that other variants of backfilling have also been suggested. The main ideas are to allow the scheduler more flexibility, by allowing it some slack in fulfilling reservations [24, 23], or by considering different sets of jobs for backfilling [21]. Backfilling has also been integrated with multiple-queue scheduling [14] and with gang scheduling [26]. But here we focus on the two simpler variants.

### 3 Evaluation Methodology

There are many ways to evaluate the performance of schedulers. A major distinction is between the off-line scheduling of a given set of jobs, and the online scheduling of a very long sequence of jobs that arrive at unknown times. We prefer the latter, which is more realistic.

#### 3.1 Simulation

We perform the evaluation by discrete event simulation. The events of interest are the arrival and termination of jobs. Each job requires a certain number of processors, runs for a certain time, and also has a user estimate of its runtime. As each job terminates, the quality of the service it had received is tabulated. One of the most commonly used metrics for scheduling is the average response time (the time from when a job was submitted until it terminated). The problem with this metric is that it is typically dominated by long jobs: a job that runs for several hours has much more weight than a job that only runs for a few seconds.

An attempt to solve this problem is the use of slowdown as a metric. Slowdown is the response time normalized by the job’s actual running time. Thus it measures how much slower the job ran due to scheduling conflicts with competing jobs, and seems to better capture users’ expectations that a job’s response time will be proportional to its runtime. While this seems to put all jobs on an equal footing, it turns out to be dominated by short jobs, because the job’s runtime appears in the denominator. This is especially problematic when using workload models which may contain jobs that are much shorter than a second. We therefore use a version known as “bounded slowdown” [8], defined as

$$\text{bounded slowdown} = \max \left\{ \frac{T_{resp}}{\max\{T_{run}, 10\}}, 1 \right\}$$

The threshold of 10 seconds is chosen to represent the maximal interactive time — that is, short jobs are allowed this much delay without penalty. In addition, the metric is set to 1 if the response time is shorter than this threshold, to avoid values smaller than 1 (which would imply a false sense of speedup).

The simulation uses the batch means approach to evaluate confidence intervals for the average response time and slowdown [9]. The batch size is set to be 5000 jobs, as recommended by MacDougall for open systems under high load [18]. Depending on the length of the workload log, 10 to 17 batches were completed, and the first discarded to bring the system to a steady state. The number of queued jobs at the end of each batch is also recorded, in order to enable the identification of situations in which the system is actually saturated and jobs accumulate in the queue.

Due to the large variability of job durations, the metrics are also highly variable. As a consequence the results often seem not to be statistically significant, as the 90% confidence intervals overlap. We therefore performed a more sophisticated analysis of these results, using the common random numbers<sup>1</sup> variance reduction technique [13]. In this analysis, we first compute the difference in response times or slowdowns between the two schedulers on a per-job basis, which is possible because we are using the same workload log. We then compute confidence intervals on

---

<sup>1</sup>The name is somewhat of a misnomer in this case, as we are using a logged workload rather than generating it using a random number generator.

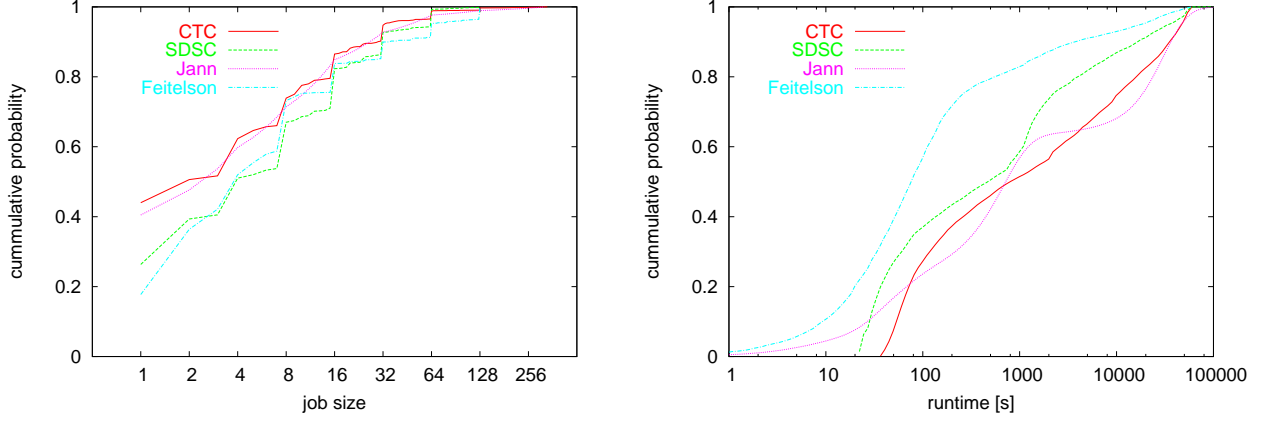


Figure 2: *Comparison of workloads.*

these differences using the batch means approach. If the confidence intervals do not include zero, the difference is significant.

### 3.2 Workloads

Workloads for parallel job scheduling are interesting due to the combination of being relatively small and at the same time relatively complex. The size of typical workloads is tens of thousands of jobs, as opposed to millions of packets in communication workloads. These workloads are characterized by a large number of factors, including the job sizes, runtimes, runtime estimates, and arrival patterns. The complexity derives not only from the multiple factors themselves, but from various correlations between them. Research on these issues is facilitated by the availability of data and models in the Parallel Workloads Archive [1]. In addition, these workloads are often used for parallel system evaluations (e.g. [5, 22, 26, 19, 3, 12]).

The main workloads we used are a log of jobs run on the Cornell Theory Center (CTC) IBM SP2 machine from July 1996 through May 1997<sup>2</sup>, and a model based on the first  $2\frac{1}{2}$  months of this workload which was created by Jann et al. [10]. Despite the close relationship between the two workloads, there are several differences (Figure 2):

- In the CTC workload, most jobs request a power-of-two number of nodes. In the Jann workload it was decided not to favor powers of two, based on the assumption that this is not a real attribute of the workload but rather a consequence of administrative policies. This notion was later borne out by a user survey conducted by Cirne and Berman [4].
- The CTC workload had a sharp administrative limit on runtime at 18 hours. In the Jann model there is no such limit, as a continuous distribution is used to represent the runtimes. But given that the chosen distribution is supposed to mimic the CTC workload statistics,

<sup>2</sup>For the simulations and analysis reported here we removed a flurry of 1999 one-node jobs lasting about 42–48 seconds each that were submitted by user 135 in one day; it is assumed that this was the result of a misbehaving script. This large burst of small jobs is very sensitive to details of the scheduling and tends to cause large fluctuations in simulation results. It represents 2.55% of the total workload of 78500 jobs. The justification for removing such flurries is elaborated in [25].

relatively few jobs have runtimes longer than 18 hours. At the other extreme, the CTC log hardly has any jobs shorter than 30 seconds, whereas the Jann model has many such jobs.

- The Jann model uses a hyper-Erlang distribution to model runtimes and interarrival times. While this matches the first three moments of the original distribution, it also leads to a somewhat bimodal structure that does not exist in the original.
- The Jann model does not include the modeling of user estimates of runtime. We therefore used the actual runtime in the simulations. This is equivalent to assuming that the estimates are perfect.

In addition, we also use a log of jobs run on the San-Diego Supercomputer Center (SDSC) IBM SP2 from May 1998 through April 2000<sup>3</sup>, and a model proposed by Feitelson [7]. As can be seen from the graphs, these two workloads have a similar distribution of job sizes which has fewer small jobs than the CTC and Jann workloads. The Feitelson model is also unique in having many more short jobs. The SDSC log also has shorter jobs than CTC and Jann, but not as short as in the Feitelson model.

Each workload is represented by a single data file specifying the arrival time, runtime, and number of processors used by each job, and in the case of the CTC and SDSC workloads, also the estimated runtime. In order to create different load conditions, all the arrival times are multiplied by a suitable constant. For example, if the original workload file leads to a load of 0.7 of capacity, multiplying all interarrival times by a factor of  $7/8 = 0.875$  will cause the jobs to arrive faster, and increase the load to 0.8 of capacity.

## 4 Simulation Results

The typical methodology for evaluating systems (in our case, schedulers) calls for simulating their behavior, and tabulating the resulting performance. This is usually repeated numerous times under different conditions, including different load levels and different parameter settings for the system. This enables a study of how the system responds to load, and what parameter settings lead to optimal performance.



In this paper we emphasize two other dimensions instead: the workload used to drive the simulation, and the metric used to measure the results. Figure 3 shows results of comparing our two schedulers using four different workloads, and two different metrics. In all cases, simulations are run for different load levels, and the behavior of the system as a function of load is plotted. These results can be divided into four groups:

- Results which show that EASY is better than conservative. These include the results obtained with the CTC, SDSC, and Jann workloads, when using the response time metric.
- Results which show that conservative is better than EASY. Such results were obtained from the Jann workload when using the bounded slowdown metric.



---

<sup>3</sup>In this log we also removed two flurries, one composed of 30-second one-node jobs by user 365, and the other composed of 1-minute 32-node jobs by user 319 [25]; they are known to affect simulation results, but in the specific simulations used here they seem to have only a marginal effect.



### CTC workload



### SDSC workload



### Jann workload



### Feitelson workload





Figure 3: Simulation results with 90% confidence intervals.

- Results showing that both schedulers are essentially the same. The results from the Feitelson model fall into this class.
- A combination: results that show the schedulers to be equivalent under most loads, and show an advantage for EASY only under high load. Such results were observed for CTC and SDSC when using the bounded slowdown metric. It should be noted that such high loads are seldom achieved in practice on this class of machines [11, 20].

In short, if only a single set of results is used, the conclusions would depend on which specific set was chosen. Our goal in the next sections is to uncover the root causes for these discrepant results.

## 5 Analysis of Jann vs. CTC Results

As noted above, the Jann model is specifically meant to mimic the CTC workload. Nevertheless, the simulation results indicate that the two produce discrepant predictions. For the Jann workload the response time metric favors EASY backfilling, whereas the slowdown metric favored conservative backfilling. The CTC workload, at the same time, favors EASY for both metrics (albeit for slowdown only under very high load). This is therefore actually a triple interaction of scheduler, workload, and metric. In this section we focus on a detailed analysis of how this triple interaction comes about.

### 5.1 Producing Discrepant Results

Both the slowdown metric and the backfilling policy are sensitive to job duration. We therefore start by partitioning the jobs into five classes according to their duration, and tabulating the results for each class separately. The classes used were very short ( $< 30\text{sec}$ ), short ( $< 5\text{min}$ ), medium ( $< 1\text{hr}$ ), long ( $< 10\text{hr}$ ), and very long ( $> 10\text{hr}$ ). The results are shown in Table 1. For the CTC workload, both metrics favor EASY backfilling for each class individually, and also for all of them together. But in the Jann workload we indeed see a difference that depends on job class. For jobs that are longer than 1 hour, both metrics favor EASY. But for shorter jobs, both metrics favor conservative backfilling.

Given that for each job class both metrics agree, how does this turn into discrepant results when the whole workload is considered? The answer is due to simple arithmetic. The average response time of all jobs is a weighted average of response times of the different classes, with the weights proportional to the number of jobs in each class. The same goes for slowdown. But when response times are used, the high values that dominate the average come from the long jobs, whereas when we calculate the average slowdown the high values come from the short jobs. Thus the average response time is similar to the response time for long jobs, which favors EASY, whereas the average slowdown is similar to the slowdown of short jobs, which favors conservative.

### 5.2 Source of Performance Differences

The results in Table 1 indicate that the difference between the CTC and Jann results is due to the short jobs, which fare better under conservative in the Jann workload, but not in the CTC workload.



job class	number of jobs*	response time		bounded slowdown	
		EASY	vs. cons	EASY	vs. cons
Jann model					
very short	33728	10737.46	> 7699.22	837.22	> 488.89
short	74642	11384.41	> 8368.95	144.14	> 79.72
medium	106331	10484.76	> 8788.11	15.44	> 9.38
long	76974	48652.08	< 66560.68	2.56	< 2.59
very long	41624	91151.88	< 112610.74	1.79	< 1.82
all	333300	29600.57	< 34892.05	122.75	> 90.53
CTC workload					
very short	380	6640.05	< 11020.90	627.92	< 1062.75
short	28746	6378.98	< 6564.45	64.38	< 75.91
medium	16906	9689.79	< 10580.02	11.76	< 12.65
long	20334	28829.00	< 36771.25	2.36	< 2.87
very long	10133	67616.53	< 84720.34	1.29	< 1.63
all	76500	21190.66	< 25855.42	30.71	< 37.58

\* the numbers may differ slightly for the two schedulers as different jobs may remain in the queue at the end of the simulation.

Table 1: *Simulation results for different job classes. EASY backfilling is better in all cases for the CTC workload, and for long jobs in the Jann workload. Results shown are for load of 0.85.*

To try and understand why this happens, we need to understand how the scheduler interacts with the workload.

As the difference between our two schedulers is in their backfilling policy, a good place to start is investigating their backfilling patterns. Figure 4 shows the absolute numbers of backfilled jobs as a function of job length (top), as well as the additional backfilling achieved by EASY, as a fraction of the backfilling achieved by conservative (bottom); thus a value of 0.1 means that EASY did 10% more backfilling. Surprisingly, for CTC conservative backfilling actually achieves somewhat higher backfilling levels! But the main difference between the workloads is that under the Jann workload, EASY achieved much more backfilling of long jobs.

After checking many other possibilities (see below), the root cause for this was found to be the use of accurate runtime estimates with the Jann workload. Accurate runtime estimates provide full information for backfilling decisions. The conservative algorithm has to take multiple reservations into account, and this is especially problematic for long jobs, that have the potential to interact with many other jobs. EASY, on the other hand, only has to consider one reservation. Therefore conservative achieves much less backfilling. But when estimates are inaccurate, jobs tend to terminate before the time expected by the scheduler. This creates holes in the schedule that provide new backfilling opportunities, that can be exploited by both schedulers.

To confirm this hypothesis, we re-ran the CTC simulations but using the actual runtimes rather than the original user estimates to control the backfilling. The results, shown in Figure 5, largely confirm the conjecture. When using accurate estimates, conservative performed much less backfilling of long jobs than before. And as expected, this led to an inversion of the results, with




Figure 4: *Increased backfilling with EASY relative to conservative scheduler. Data for load of 0.85.*

conservative achieving better average slowdown scores. A similar but weaker effect also occurred when running the SDSC workload with accurate runtime estimates. The results were inverted, but the differences was much closer to zero.

We are now left with one last question: how does the reduced backfilling under the conservative policy lead to better performance when measured using the slowdown metric? The answer is in the details. The reduced backfilling applies to *long* jobs. On the other hand, slowdown is sensitive mainly to *short* jobs. So the performance change results from less backfilling of long jobs, which leads to better performance of short jobs. This is explained as follows. Under the EASY policy, backfilling is allowed provided it does not delay the first queued job. But it might delay subsequent queued jobs, including short ones, that will then suffer from a very large slowdown. The conservative policy prohibits such backfilling, with the specific goal of preventing such delays for subsequent jobs. In the simulations with the Jann workload, this turned out to be the decisive factor. Similar observations have been made independently by Srinivasan et al. [23].

To summarize, our analysis exposed the following triple interaction:

- The Jann and CTC workloads differ (among other things) in that the CTC workload is a real trace including user estimates of runtime, whereas the Jann model does not include this detail.
- Due to using accurate estimates for the Jann model, the conservative scheduler achieved less backfilling of long jobs that use few processors. This is obviously detrimental to the performance of these long jobs, but turned out to be beneficial for short jobs that don't get delayed by these long jobs.

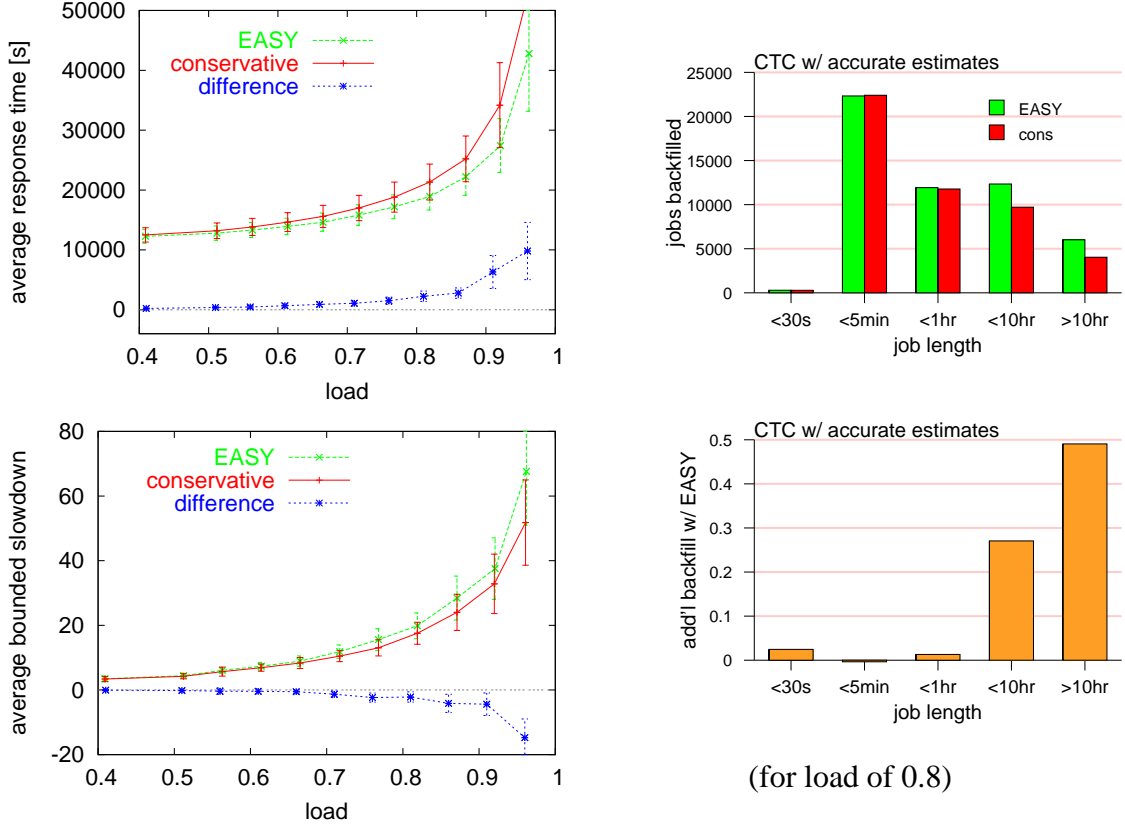


Figure 5: Results for the CTC workload when using actual runtimes as estimates, to verify that this is the cause of the Jann results. Compare with Figures 3 and 4.

- As response time is dominated by long jobs, the response time metric showed that EASY is better than conservative for the Jann workload. The slowdown metric, on the other hand, is dominated by short jobs, so it showed conservative to be better.

As real workloads have inaccurate runtime estimates [19], it seems that in this particular case the CTC results should be favored over the Jann results.

It should be noted that other works have also investigated the effect of user runtime estimates on performance [19, 27, 2]. However, these works did not elucidate the mechanisms by which the inaccurate estimates cause their effect.

### 5.3 Non-Issues

Finding that the difference between the CTC and Jann results hinges on the runtime estimates was surprising not only because this it routinely brushed aside as unimportant, but also because there is no lack of other candidates, which seem much more significant. In this section we demonstrate that they are indeed unimportant in our case.

The usual suspects regarding performance variations are the statistical differences between the workloads. Figure 6 shows histograms of the distribution of job sizes for jobs with different runtimes, comparing the two workloads. The most striking difference is that the Jann workload

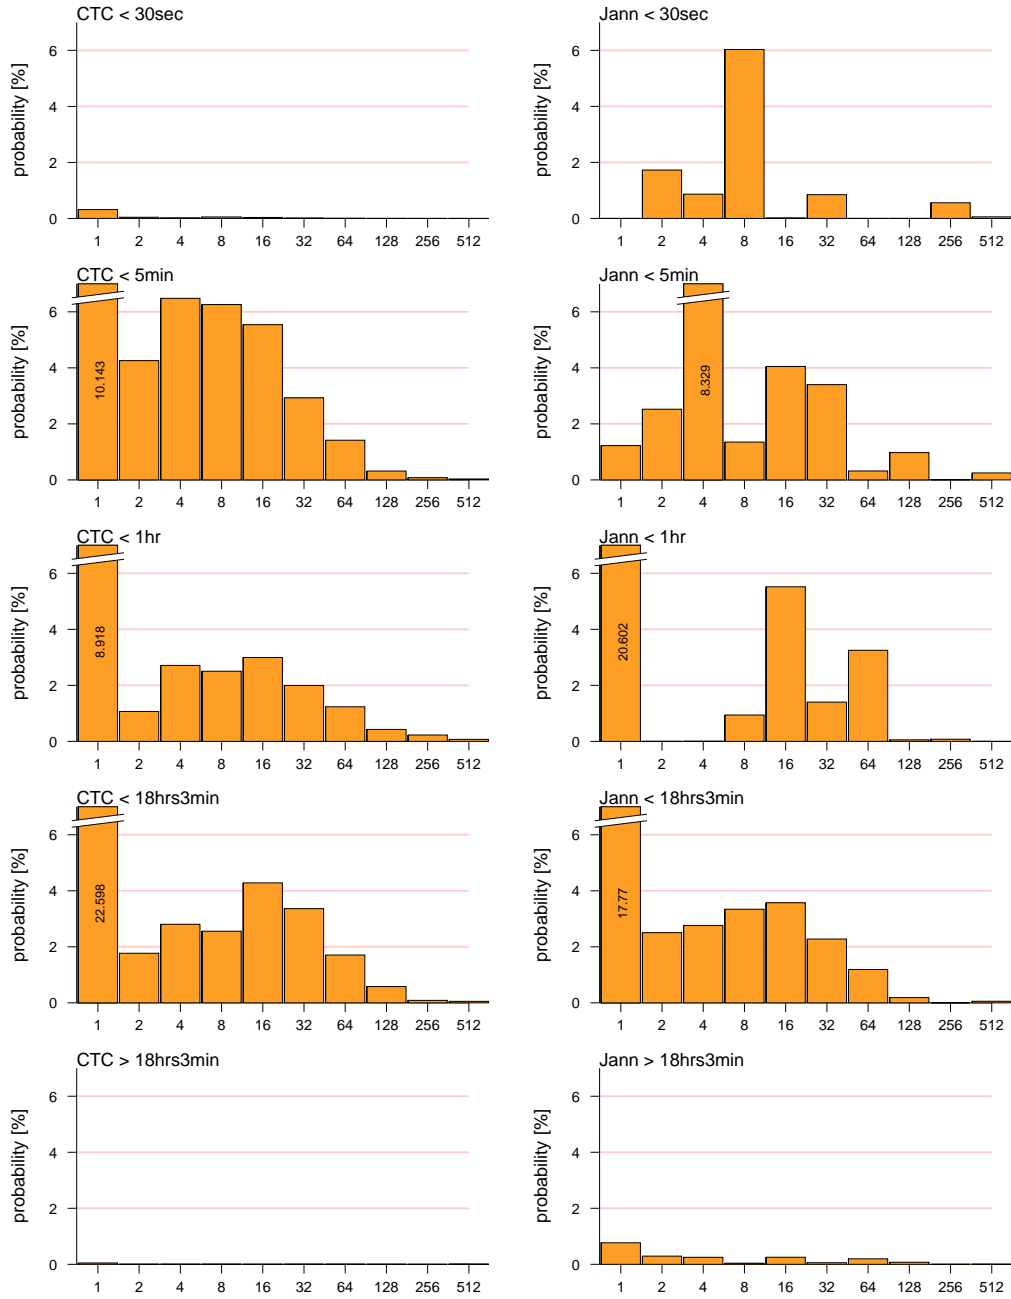


Figure 6: Histograms of job sizes for jobs of different durations. Each bar represents a sequence of sizes between powers of two. Note that some of the higher bars are truncated.

has tails at both ends, which the CTC workload does not. The CTC workload is bounded at 18 hours<sup>4</sup> (an administrative issue), whereas the Jann workload has a tail that extends beyond 30 hours. The CTC workload hardly has any jobs shorter than 30 seconds, probably due to the fact that the measurement includes the time to start up the required processes on all the nodes, and to report their termination. In the Jann model, by contradistinction, over 10% of the jobs are shorter than 30 seconds, and many jobs only run for a fraction of a second.

The long jobs in the tail could affect the results by causing longer delays to other jobs that wait for their termination because they need their processors. To check this, we re-ran the simulations with a modified version of the Jann workload, in which all jobs longer than 18 hours were deleted. The results were essentially the same as for the original workload.

The short jobs could affect the results by contributing very high values to the average slowdown metric. This argument is not very strong, as we use the bounded slowdown metric, which is designed to reduce the impact of very short jobs. Nevertheless we re-ran the simulations with a modified Jann workload, this time removing all the jobs shorter than 30 seconds. Again, the results were not significantly different from those of the original workload.

Another major difference between the workloads is that in the original CTC workload most jobs use power-of-two nodes, whereas in the Jann model jobs are spread evenly between each two consecutive powers of two. Previous work has shown that the fraction of jobs that are powers of two is important for performance, as it is easier to pack power-of-two jobs [17]. However, in our case this seemed not to make a qualitative difference. It was checked by running the simulations on a modified version of the Jann workload in which the sizes of 80% of the jobs were rounded up to the next power of two.

Finally, the size of the system is also important. The Jann model specifies the use of 322 processors, which is the number used in the simulations. For CTC, we used 430 processors, which is the size of the batch partition on the CTC machine. Repeating the simulations using 512 nodes led to somewhat lower backfilling rates, because 512 is a power of two, and therefore jobs pack much better, leaving less holes in the schedule. However, the relationship between the two schedulers did not change significantly.

## 6 Comparison with Feitelson Results

Having explained the results obtained using the CTC and Jann workloads, and the reasons for the differences between them, we now turn to the Feitelson workload (the behavior of the SDSC workload is less interesting as it is similar to that of CTC). Comparing this workload with the previous two, the question is why no major differences are observed between the performance of the two schedulers in this case.

### 6.1 Long Serial Jobs

A special feature of the CTC and Jann workloads is that they have very many serial jobs, and moreover, that most long jobs are serial. This is a result of the fact that when the CTC SP2 machine

---

<sup>4</sup>About 2.4% of the workload exceed this bound and are killed by the system. Of these, 90% are killed within one minute, and another 5% within 5 minutes. The highest runtime observed is 20 hours.

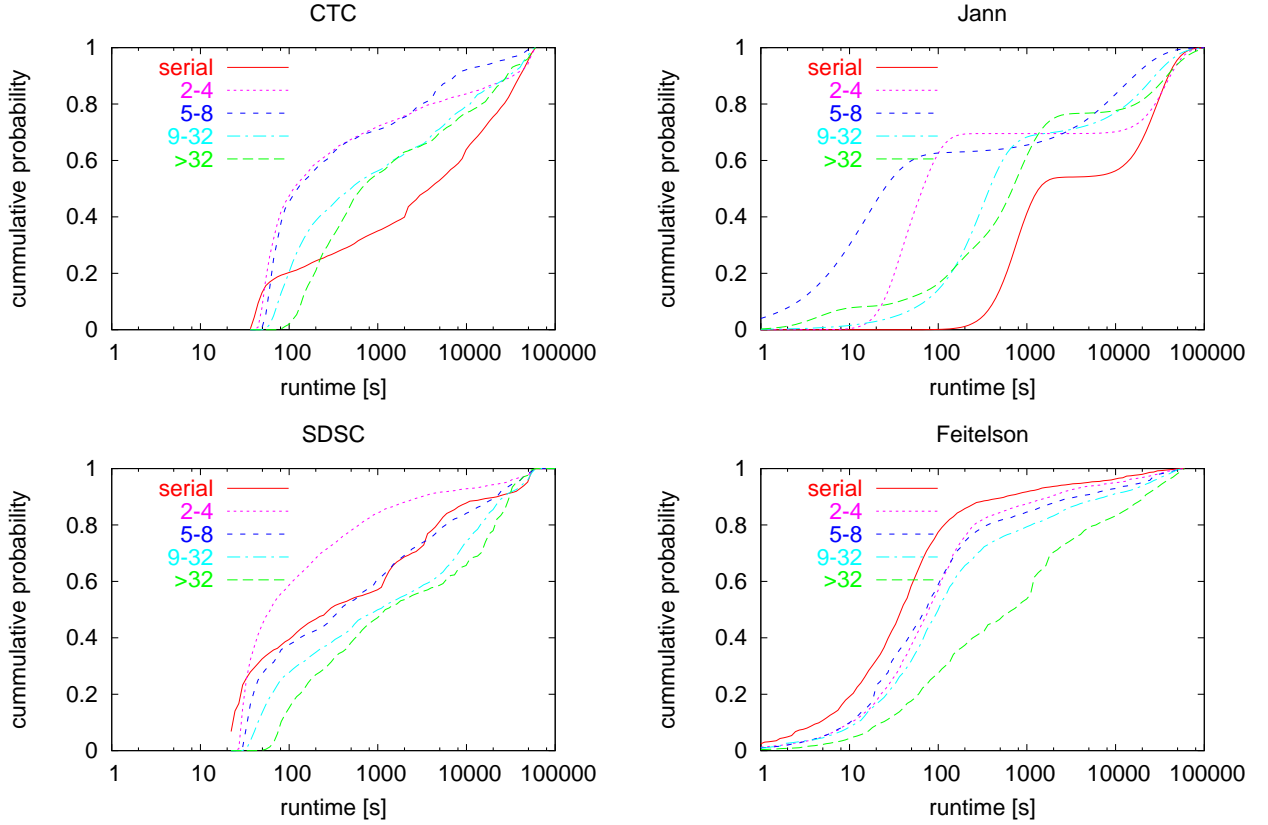


Figure 7: *Distributions of runtimes for jobs of different sizes in the various workloads.*

was installed, it replaced a large ES/9000 mainframe, and inherited the mainframe’s workload (the Jann model, in an attempt to model the CTC workload, copied this feature). Indeed, the large number of serial jobs and the large proportion of long serial jobs are seen clearly in Figures 2 and 6. We may therefore conjecture that the performance differences we saw emanate specifically from the backfilling of long serial jobs, and the subsequent delaying of larger short jobs. In the Feitelson workload, in contrast, serial jobs tend to be shorter than larger ones.

To check that this explanation is viable, we first compare the distributions of runtimes in the various workloads, and how it correlates with job size. The graphs are shown in Figure 7, with a separate line for jobs in different ranges of size. Apart from many interesting variations, one feature does indeed stand out: in the CTC and Jann workloads, serial jobs are generally very long in comparison with larger job sizes. In the Feitelson workload, by contradistinction, serial jobs are generally very short in comparison with larger job sizes. In the SDSC workload, they are in the middle.

## 6.2 Verification Using the Feitelson Model

To verify this conjecture, we modified the Feitelson workload and re-ran the simulations. Part of this has been done in the past, as reported in [19]. That paper suggested the following two modifications to the Feitelson model (Figure 8 “Feitelson→CTC”):

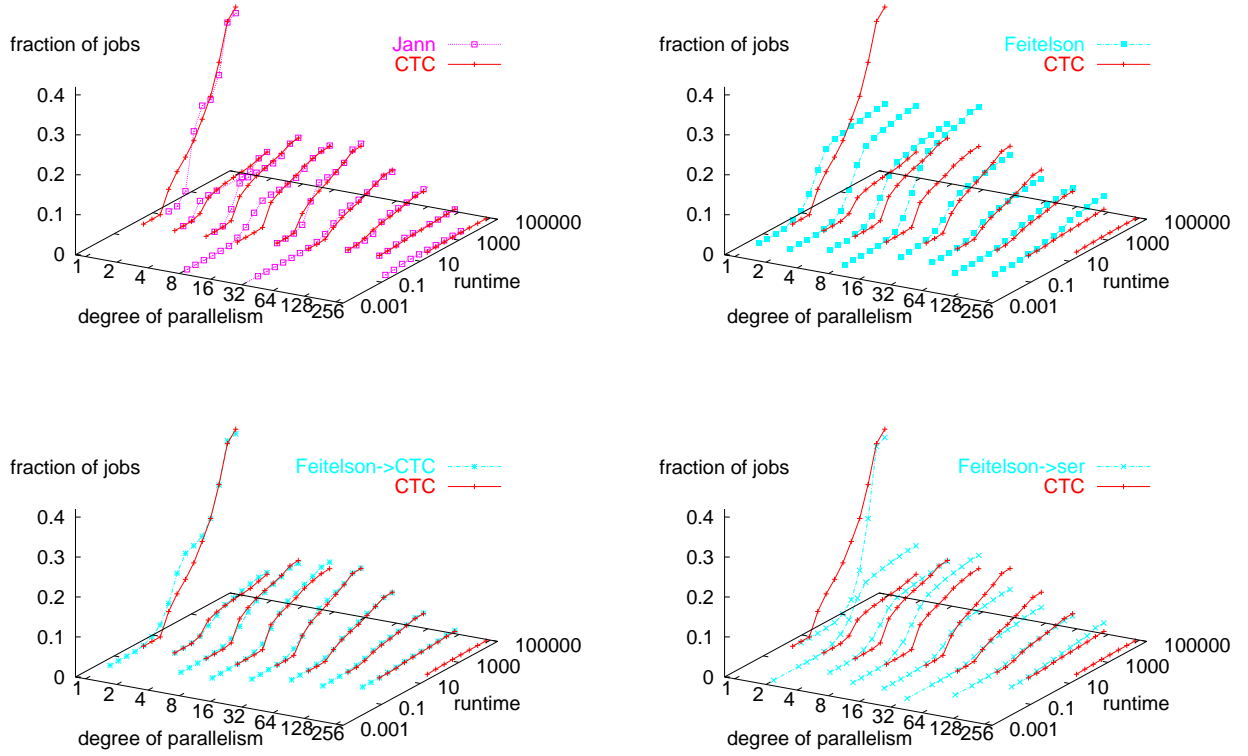


Figure 8: *Detailed comparison of different workloads, using the CTC workload as a reference. The second modification of the Feitelson model only affects serial jobs, and does not in general resemble the CTC or Jann workloads.*

- Modify the distribution of job sizes to emphasize small jobs, and
- Modify the distribution of runtimes to emphasize longer jobs.

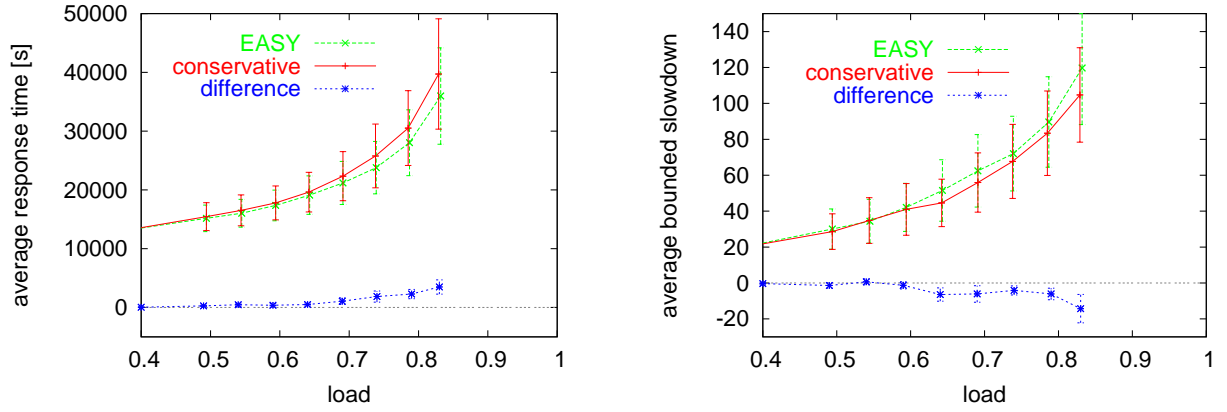
The details of implementing these modifications were hand tailored to create a workload that closely resembled both the CTC and Jann workloads. In particular, the distribution of sizes emphasized serial jobs and jobs in the range of sizes from 8 to 31, and serial jobs also received special treatment in terms of making them longer. Note that this is at odds with the original Feitelson model, in which a weak correlation exists between job size and runtime [7]. Simulating the behavior of the EASY and conservative schedulers on this modified Feitelson workload indeed showed behavior similar to that of the Jann workload (Figure 9 top). However, the magnitude of the effect was smaller.

Based on our current understanding of the matter, we can suggest a simpler modification:

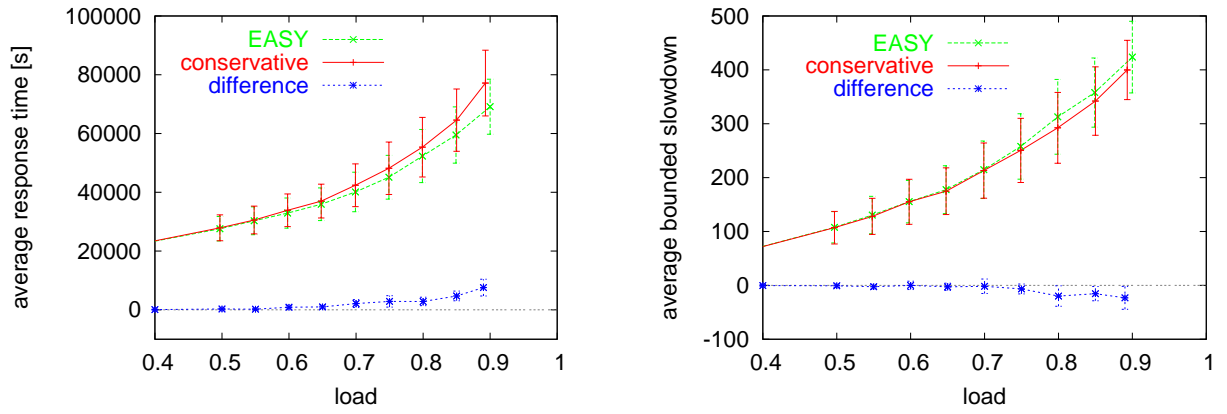
- Specifically create many long serial jobs.

This leads to a model that is quite different from the Jann and CTC workloads, because only the serial jobs have been changed (Figure 8 “Feitelson→ser”). However, making this change and re-

Feitelson workload modified to resemble Jann/CTC



Feitelson workload modified to include numerous long serial jobs



Feitelson workload modified to include numerous long serial jobs, using 139 nodes

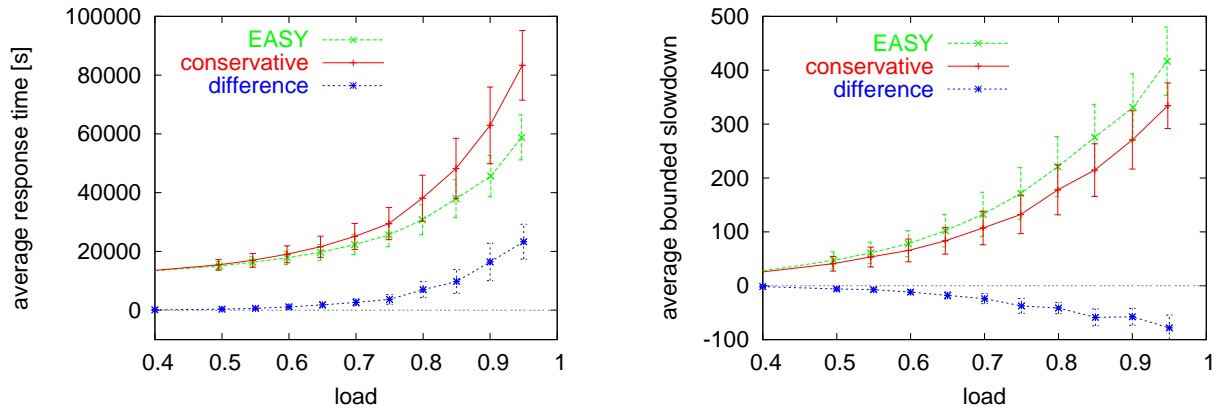


Figure 9: *Simulation results for modified Feitelson workloads. Compare with results shown in Figure 3.*



running the simulations<sup>5</sup> produced results that are quite similar to the previous, more elaborate modified model (Figure 9 middle). Thus we can indeed claim to have identified the crucial aspect of the workload that is needed in order to create the distinctive results seen with the Jann and CTC workloads.

However, the results are still much less pronounced than those obtained from the Jann workload. Given that the original modification of the Feitelson workload (Feitelson→CTC) is extremely similar to the CTC and Jann workloads, the explanation for this is not the workload statistics. Rather, it is the size of the machine used in the simulation: using 139 nodes instead of 128 significantly increases the gap in performance observed between EASY and conservative (Figure 9 bottom). The reason for this is that most job sizes are powers of two, so using a machine size that is not a power of two leads to more idle processors and more options for backfilling. Both the CTC and Jann workloads specify non-power-of-2 sizes.

## 7 Conclusions

Simulations of the relative performance of EASY and conservative backfilling, using different workloads and metrics, exhibit discrepant results. However, detailed analysis of the specific circumstances of each simulation allows us to elucidate the following general conclusions:

- If the workload does not contain many long serial jobs, both backfilling policies lead to similar performance results.
- If the workload does indeed contain many long serial jobs, as was the case at CTC, and to a lesser degree at SDSC, the relative performance depends on the accuracy of user runtime estimates and on the number of nodes in the system.
  - If user runtime estimates are inaccurate, the EASY policy leads to better results. This is expected to be the more common case.
  - But if user runtime estimates are highly accurate, conservative backfilling degrades the performance of the long serial jobs and enhances the performance of larger short jobs. This leads to better overall slowdown results.
  - This effect is intensified when the number of nodes is not a power of two.

This can be formulated as a decision tree for use in predicting performance.

Apart for settling the issue of the relative merits of EASY and conservative backfilling, this work also underscores the importance of workload models and their details. All the workloads we investigated have been accepted as reasonable and representative by researchers who have used them in the past. But using different workloads can lead to completely different results, none of which are universally correct. If this happens, one should find the workload features that lead

---

<sup>5</sup>The Feitelson model is somewhat problematic in the sense that it sometimes produces workloads that are not well-behaved when the load is increased by multiplying all arrival times by a constant smaller than unity. The causes for this behavior, and indeed how to model different load conditions in general, is the subject of a separate research effort. In our case, the suggested modification seemed to accentuate the problem. The solution was to generate multiple workloads using different random number streams, and use the one in which the load could be increased in the most predictable manner.

to the performance differences. This then enables the prediction of performance results for new workload conditions.

In addition to different workloads, seemingly benign assumptions can also have a decisive impact on evaluation results. In our case, the assumption of accurate runtime estimates, shown to be unlikely to be true in [19, 15], was now shown to be instrumental in the results obtained using the Jann workload. This underscores the need to collect real data that can serve as the basis for performance evaluation, reducing the risk of using unbiased assumptions.

## Acknowledgments

This research was supported in part by the Israel Science Foundation (grant no. 219/99). The workload models and logs on which it is based are available on-line from the Parallel Workloads Archive [1]. The workload log from the CTC SP2 was graciously provided by the Cornell Theory Center, a high-performance computing center at Cornell University, Ithaca, New York, USA. The workload log from the SDSC SP2 was graciously provided by the HPC Systems group of the San Diego Supercomputer Center (SDSC), which is the leading-edge site of the National Partnership for Advanced Computational Infrastructure (NPACI). The code for the Jann model was graciously provided by Joe Jann of IBM Research. Many thanks to them for making the data available and for their help with background information and interpretation.

## References

- [1] “Parallel workloads archive”. URL <http://www.cs.huji.ac.il/labs/parallel/workload/>.
- [2] S-H. Chiang, A. Arpaci-Dusseau, and M. K. Vernon, “*The impact of more accurate requested runtimes on production job scheduling performance*”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson, L. Rudolph, and U. Schwiegelshohn (eds.), pp. 103–127, Springer Verlag, 2002. Lect. Notes Comput. Sci. vol. 2537.
- [3] W. Cirne and F. Berman, “*A comprehensive model of the supercomputer workload*”. In *4th Workshop on Workload Characterization*, Dec 2001.
- [4] W. Cirne and F. Berman, “*A model for moldable supercomputer jobs*”. In *15th Intl. Parallel & Distributed Processing Symp.*, Apr 2001.
- [5] A. B. Downey, “*Predicting queue times on space-sharing parallel computers*”. In *11th Intl. Parallel Processing Symp.*, pp. 209–218, Apr 1997.
- [6] D. G. Feitelson, “*Metric and workload effects on computer systems evaluation*”. *Computer* **36(9)**, pp. 18–25, Sep 2003.
- [7] D. G. Feitelson, “*Packing schemes for gang scheduling*”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 89–110, Springer-Verlag, 1996. Lect. Notes Comput. Sci. vol. 1162.
- [8] D. G. Feitelson, L. Rudolph, U. Schwiegelshohn, K. C. Sevcik, and P. Wong, “*Theory and practice in parallel job scheduling*”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 1–34, Springer Verlag, 1997. Lect. Notes Comput. Sci. vol. 1291.
- [9] R. Jain, *The Art of Computer Systems Performance Analysis*. John Wiley & Sons, 1991.

- [10] J. Jann, P. Pattnaik, H. Franke, F. Wang, J. Skovira, and J. Riordan, “Modeling of workload in MPPs”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 95–116, Springer Verlag, 1997. Lect. Notes Comput. Sci. vol. 1291.
- [11] J. P. Jones and B. Nitzberg, “Scheduling for parallel supercomputing: a historical perspective of achievable utilization”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 1–16, Springer-Verlag, 1999. Lect. Notes Comput. Sci. vol. 1659.
- [12] E. Krevat, J. G. Castaños, and J. E. Moreira, “Job scheduling for the BlueGene/L system”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson, L. Rudolph, and U. Schwiegelshohn (eds.), pp. 38–54, Springer Verlag, 2002. Lect. Notes Comput. Sci. vol. 2537.
- [13] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*. McGraw Hill, 3rd ed., 2000.
- [14] B. G. Lawson and E. Smirni, “Multiple-queue backfilling scheduling with priorities and reservations for parallel systems”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson, L. Rudolph, and U. Schwiegelshohn (eds.), pp. 72–87, Springer Verlag, 2002. Lect. Notes Comput. Sci. vol. 2537.
- [15] C. B. Lee, Y. Schwartzman, J. Hardy, and A. Snaveley, “Are user runtime estimates inherently inaccurate?”. In *Job Scheduling Strategies for Parallel Processing*, Springer-Verlag, 2004. (to appear).
- [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. Lect. Notes Comput. Sci. vol. 949.
- [17] V. Lo, J. Mache, and K. Windisch, “A comparative study of real workload traces and synthetic workload models for parallel job scheduling”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 25–46, Springer Verlag, 1998. Lect. Notes Comput. Sci. vol. 1459.
- [18] M. H. MacDougall, *Simulating Computer Systems: Techniques and Tools*. MIT Press, 1987.
- [19] A. W. Mu’alem and D. G. Feitelson, “Utilization, predictability, workloads, and user runtime estimates in scheduling the IBM SP2 with backfilling”. *IEEE Trans. Parallel & Distributed Syst.* **12(6)**, pp. 529–543, Jun 2001.
- [20] L. Rudolph and P. Smith, “Valuation of ultra-scale computing systems”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph (eds.), pp. 39–55, Springer Verlag, 2000. Lect. Notes Comput. Sci. vol. 1911.
- [21] E. Shmueli and D. G. Feitelson, “Backfilling with lookahead to optimize the performance of parallel job scheduling”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson, L. Rudolph, and U. Schwiegelshohn (eds.), pp. 228–251, Springer-Verlag, 2003. Lect. Notes Comput. Sci. vol. 2862.
- [22] M. S. Squillante, D. D. Yao, and L. Zhang, “The impact of job arrival patterns on parallel scheduling”. *Performance Evaluation Rev.* **26(4)**, pp. 52–59, Mar 1999.
- [23] S. Srinivasan, R. Kettimuthu, V. Subramani, and P. Sadayappan, “Selective reservation strategies for backfill job scheduling”. In *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson, L. Rudolph, and U. Schwiegelshohn (eds.), pp. 55–71, Springer-Verlag, 2002. Lect. Notes Comput. Sci. vol. 2537.

- [24] D. Talby and D. G. Feitelson, “Supporting priorities and improving utilization of the IBM SP scheduler using slack-based backfilling”. In *13th Intl. Parallel Processing Symp.*, pp. 513–517, Apr 1999.
- [25] D. Tsafrir and D. G. Feitelson, *Workload Flurries*. Technical Report 2003-85, Hebrew University, Nov 2003.
- [26] Y. Zhang, H. Franke, J. E. Moreira, and A. Sivasubramaniam, “Improving parallel job scheduling by combining gang scheduling and backfilling techniques”. In *14th Intl. Parallel & Distributed Processing Symp.*, pp. 133–142, May 2000.
- [27] D. Zotkin and P. J. Keleher, “Job-length estimation and performance in backfilling schedulers”. In *8th Intl. Symp. High Performance Distributed Comput.*, Aug 1999.