

The Impact of More Accurate Requested Runtimes on Production Job Scheduling Performance

Su-Hui Chiang, Andrea Arpaci-Dusseau, and Mary K. Vernon

Computer Sciences Department
University of Wisconsin
1210 W. Dayton Street, Madison, Wisconsin
{suhui, dusseau, vernon}@cs.wisc.edu

Abstract. The question of whether more accurate requested runtimes can significantly improve production parallel system performance has previously been studied for the FCFS-backfill scheduler, using a limited set of system performance measures. This paper examines the question for higher performance backfill policies, heavier system loads as are observed in current leading edge production systems such as the large Origin 2000 system at NCSA, and a broader range of system performance measures. The new results show that more accurate requested runtimes can improve system performance much more significantly than suggested in previous results. For example, average slowdown decreases by a factor of two to six, depending on system load and the fraction of jobs that have the more accurate requests. The new results also show that (a) nearly all of the performance improvement is realized even if the more accurate runtime requests are a factor of two higher than the actual runtimes, (b) most of the performance improvement is achieved when test runs are used to obtain more accurate runtime requests, and (c) in systems where only a fraction (e.g., 60%) of the jobs provide approximately accurate runtime requests, the users that provide the approximately accurate requests achieve even greater improvements in performance, such as an order of magnitude improvement in average slowdown for jobs that have runtime up to fifty hours.

1 Introduction

Many state-of-the-art production parallel job schedulers are non-preemptive and use a requested runtime for each job to make scheduling decisions. For example, the EASY Scheduler for the SP2 [3, 4] implements the First-Come First-Served (FCFS)-backfill policy, in which the requested runtime is used to determine whether a job is short enough to be backfilled on a subset of the nodes during a period when those nodes would otherwise be idle. The more recent Maui Scheduler ported to the NCSA Origin 2000 (O2K) [1] and the large NCSA Linux Cluster [2] implements a parameterized priority-backfill policy that uses the requested runtime to determine job priority as well as whether it can be backfilled.

Recent work [5] has shown that the priority-backfill policy on the O2K has similar performance to FCFS-backfill, but that modifying the priority function to favor jobs with short requested runtimes provides superior average wait, 95th-percentile wait, and average slowdown, as well as similar maximum wait time as for FCFS-backfill, for the large production workloads that run on the NCSA O2K. Thus, requested runtimes are needed not only for backfill decisions but also to enable favoring short jobs in a way that improves service for nearly all jobs. For example, in the LXF&W-backfill priority function derived from high-performance uniprocessor scheduling policies, a job's priority increases linearly with its *expansion factor*, where the expansion factor is the ratio of the job's wait time plus requested runtime to requested runtime.

The key advantage of nonpreemptive scheduling policies is that they are significantly easier to implement than preemptive policies, particularly for systems with many processors. Furthermore, simulation results for the O2K job traces and system show that the non-preemptive LXF&W-backfill policy has performance that is reasonably competitive with high performance (but more difficult to implement) preemptive policies such as gang scheduling or spatial equi-partitioning [5]. This relatively high performance is achieved in spite of the fact that user requested runtimes are often highly inaccurate [6–9]. For example, analysis of the NCSA O2K logs shows that 30% of the jobs that request 200 or more hours of runtime terminate in under ten hours [9].

The key open question addressed in this paper is whether the high performance backfill policies can be further improved with more accurate requested runtimes. Several previous simulation studies of FCFS-backfill show that more accurate requested runtimes have only minimal impact on the average wait time or average slowdown [6, 10–12, 7]. We briefly revisit the question for FCFS-backfill, using workloads from recent months on the O2K that have significantly heavier system load (e.g., up to 100% cpu demand), and using a more complete set of performance measures. More importantly, we investigate the question of whether more accurate requested runtimes can significantly improve the high performance backfill policies such as LXF&W-backfill that use requested runtimes to favor short jobs. We evaluate this question using complete workload traces from the NCSA O2K and consider not only average wait time and average slowdown as in previous studies, but also the maximum and 95th-percentile wait time. Each of these measures is obtained as a function of actual job runtime and as a function of the number of requested processors, to determine how performance varies with these job parameters.

To study the above key question, two design issues that relate to preventing starvation in backfill policies that favor short jobs require further investigation. As discussed further in Sections 2.3 and 3.1, preventing starvation was not fully addressed in previous policy evaluations. In particular, the problem is more significant for the heavier system load in recent months on the O2K. The first policy design issues relate to reservations; that is, how many jobs are given reservations and, in the case of dynamic priority functions, whether the reservations are *fixed* or *dynamic*. The second design issue is the relative weight in the priority function

for requested job runtime and current job wait time. A more complete analysis of these issues is needed in order to set these policy parameters properly for studying the potential improvement of more accurate requested runtimes.

The key results in the paper are as follows:

- For a set of high performance backfill policies that favor short jobs (i.e., LXF&W-, SJF&W-, $L\sqrt{X}F&W$ -, and $S\sqrt{T}F&W$ -backfill), accurate requested runtimes dramatically improve the average slowdown, greatly improve the average and maximum wait for short jobs without increasing the average wait for long jobs, and greatly improve the 95th-percentile wait for all jobs. More precisely, the 95th-percentile wait for all jobs is reduced by up to a factor of two, while the average slowdown is reduced by up to a factor of six, depending on the system load. Policy performance with more accurate requested runtimes is thus even more similar to the high performance pre-emptive policies such as gang scheduling or spatial equipartitioning.
- Nearly all of the improvement is realized even if requested runtimes are only approximately accurate; that is, if all requested runtimes are up to a factor of two times the actual runtime. Furthermore, most of the improvement can be achieved (a) even if only 60% - 80% of the jobs provide the approximately accurate runtime requests, and (b) test runs are used to more accurately estimate requested runtime.
- In systems where only a fraction (e.g., 60%) of the jobs provide approximately accurate requested runtimes, the jobs with improved runtime requests have even greater improvements in performance, such as more than an order of magnitude improvement in average slowdown of the jobs that have actual runtime up to fifty hours. Thus, there is a significant incentive for individual users to improve the accuracy of their runtime requests.

Additional contributions of the paper include:

- A summary of the very recent workloads (October 2000 - July 2001) on the O2K, including several months with heavier processor and memory demand than workloads used previously to design scheduling policies. Note that heavier system load can have a significant impact on the magnitude of the performance differences among alternative scheduling policies. For example, more accurate requested runtimes improves the average slowdown of FCFS-backfill more significantly for the recent heavy loads on the O2K.
- For the NCSA O2K architecture and workload, using a small number of reservations (2 to 4) outperforms a single reservation, but a larger number of reservations results in poor performance during months with exceptionally heavy load.
- Compared to the highest performance previous backfill policy, namely LXF&W-backfill with single reservation, LXF&W-backfill with two to four reservations, as well as two proposed new priority backfill policies ($L\sqrt{X}F&W$ and $S\sqrt{T}F&W$ -backfill) with two reservations, significantly improve the maximum wait time.

The remainder of this paper is organized as follows. Section 2 provides background on the system and workloads used in this study, and on related previous

work. Section 3 evaluates the impact of reservation policies and the relative priority weight between job requested runtime and current job wait time for backfill policies. Section 4 evaluates the potential benefit of using more accurate requested runtimes in priority backfill policies. Section 5 shows whether the performance benefit of more accurate requested runtimes can be achieved if test runs are used to estimate the more accurate requested runtimes. Section 6 provides the conclusions of this work.

2 Background

2.1 The NCSA Origin 2000 System

The NCSA O2K is a large production parallel system that provides 960 processors and 336 GB of memory for processing batch jobs that do not request a dedicated host. The processors are partitioned into eight hosts, each of which has 64 or 128 processors and 32 or 64 GB of memory. The jobs are scheduled using a ported version of the Maui Scheduler that implements a backfill policy with a parameterized priority function, and evicts a job if it has run one hour longer than its requested runtime. More detail about the system and scheduler configuration can be found in [1, 5].

2.2 Workloads

In this study, we have evaluated scheduling policy performance using simulations with ten different one-month job traces obtained during October 2000 - July 2001 from the O2K. Three of these months (October - December 2000) were fully characterized in [9]. The load during each month is summarized in Table 1. The overall processing demand ("proc demand") per month is the actual runtime of a job times the requested processors for the job, summed over all jobs submitted that month, expressed as a percentage of the total available processor-minutes for the month. The memory demand ("mem demand") is the equivalent measure for the job memory requests. Processor and memory demand are also given for each job class, where job class is defined by the requested runtime and requested processor and memory resources, as defined below the table.

There are two key differences in the traces summarized in the table compared to those considered previously [5, 9]. First, the actual job runtime in these traces includes the initial data *setup time*, during which the job occupies its requested resources (i.e., processors and memory) but it has not yet started its computation. The data setup time adds negligible ($\leq 1\%$) total cpu and memory load each month, but it is significant (e.g., 10 hours) for some jobs. Second, the traces include four months (January - March and May 2001) that have exceptionally high demand for processor resources (i.e., very close to 100%), and three of those months (February, March, and May 2001) also have exceptionally high memory demand ($> 90\%$). The other three months in 2001 (April, June, and July) have cpu demand (80 - 90%) and memory demand (70 - 80%) that is typical in earlier O2K workloads [5]. Results will be shown in the paper for three of the heavy load months (January - March 2001), one of the months that follows a heavy load month (June 2001) and one typical month (July 2001).

Table 1. Summary of Monthly NCSA O2K Workloads

Month	Total	Job Class											
		sj				mj				lj			
		vst	st	mt	lt	vst	st	mt	lt	vst	st	mt	lt
Oct00													
#jobs	6552	1342	2491	576	276	248	624	240	50	57	362	208	78
proc demand	82%	1%	11%	9%	7%	0%	10%	11%	2%	0%	14%	13%	4%
mem demand	81%	0%	6%	7%	9%	0%	6%	6%	2%	0%	6%	18%	20%
Nov00													
#jobs	6257	1719	2279	417	60	287	499	186	16	146	513	110	25
proc demand	85%	1%	10%	8%	3%	1%	9%	12%	3%	1%	21%	13%	3%
mem demand	61%	1%	5%	5%	2%	0%	5%	6%	1%	0%	11%	11%	14%
Dec00													
#jobs	4782	1114	2056	563	164	100	203	215	59	45	135	113	15
proc demand	89%	0%	10%	10%	9%	0%	4%	18%	4%	0%	8%	13%	12%
mem demand	63%	0%	6%	8%	5%	0%	2%	10%	6%	0%	3%	13%	9%
Jan01													
#jobs	4837	945	2000	649	164	185	267	158	151	37	170	97	14
proc demand	*102%	1%	9%	13%	7%	0%	4%	18%	10%	0%	9%	15%	14%
mem demand	76%	0%	6%	8%	5%	0%	3%	10%	14%	0%	6%	9%	14%
Feb01													
#jobs	6784	2328	2264	479	180	357	333	119	63	281	219	91	70
proc demand	*97%	1%	9%	9%	8%	0%	6%	13%	7%	0%	11%	12%	22%
mem demand	*87%	1%	6%	5%	5%	0%	4%	8%	8%	0%	8%	14%	28%
Mar01													
#jobs	5929	1915	1869	644	221	372	290	140	50	78	224	87	39
proc demand	*100%	1%	12%	11%	10%	1%	4%	10%	5%	0%	11%	18%	17%
mem demand	*92%	1%	7%	8%	9%	0%	3%	6%	8%	0%	9%	11%	30%
Apr01													
#jobs	6206	2106	2304	643	202	235	238	70	78	47	159	90	34
proc demand	78%	1%	13%	12%	9%	0%	5%	8%	5%	0%	8%	8%	8%
mem demand	77%	1%	6%	7%	7%	0%	3%	3%	9%	0%	9%	8%	25%
May01													
#jobs	6573	2220	2012	611	191	364	355	115	96	214	246	104	45
proc demand	*99%	2%	12%	10%	10%	1%	6%	10%	12%	1%	8%	19%	9%
mem demand	*92%	1%	5%	9%	6%	0%	3%	4%	14%	1%	10%	18%	20%
Jun01													
#jobs	6364	2076	2317	690	82	271	346	113	86	91	189	84	19
proc demand	86%	2%	12%	15%	6%	1%	8%	10%	9%	1%	9%	10%	4%
mem demand	75%	1%	7%	11%	4%	0%	4%	4%	12%	0%	9%	17%	6%
Jul01													
#jobs	5705	1363	2070	664	136	243	415	177	111	102	263	131	30
proc demand	89%	1%	12%	15%	5%	1%	7%	14%	6%	1%	12%	10%	5%
mem demand	81%	1%	8%	9%	5%	0%	4%	7%	13%	0%	9%	18%	8%

* indicates exceptionally high load.

Job Class Definition						
Requested Run Time Class				Space Class		
vst	st	mt	lt	sj	mj	lj
≤5hrs	[5, 50)	[50, 200)	[200, 400)	P ≤ 8	P ≤ 16	P ≤ 64
	hrs	hrs	hrs	M ≤ 2GB	M ≤ 4GB	M ≤ 25GB

P = requested processors; M = requested memory

Other characteristics of the workloads during 2001 are similar to the previous months in 2000. In particular, there is an approximately similar mix of job classes (i.e., sizes) from month to month (as shown in Table 1), and there is a large discrepancy between requested runtime (R) and actual runtime (T), as illustrated in Figure 1. Figure 1(a) plots the distribution of actual runtime for the jobs submitted during January - July 2001 that requested 50 hours (R = 50h) or 400 hours (R = 400h). These results show that almost 30% of the jobs

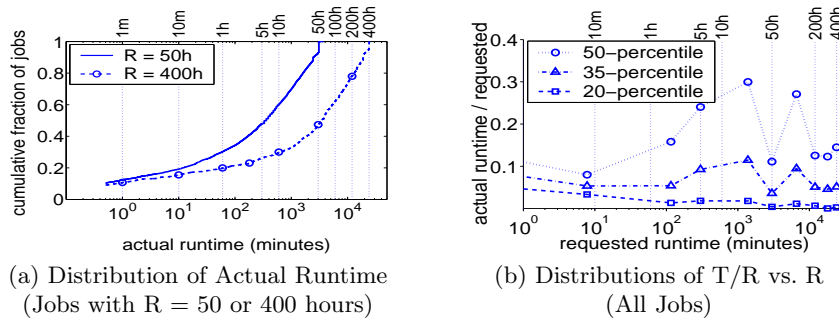


Fig. 1. Requested Runtime (R) Compared with Actual Runtime (T)
(O2K Workloads, January 2001 - July 2001)

Table 2. Priority Functions of Previous Backfill Policies

Priority Weight				Job Measure
FCFS	SJF	LXF	LXF&W(w)	
1	0	0	$w = 0.02$	current wait time, J_w , in hours
0	1	0	0	inverse of requested runtime ($\frac{1}{R}$)
0	0	1	1	current job expansion factor ($\frac{J_w + R \text{ in hours}}{R \text{ in hours}}$)

with $R = 400$ hours terminate in under 10 hours, another 10% have actual runtime between 10 and 50 hours, and approximately 10% of the jobs with $R = 50$ hours or $R = 400$ hours terminate in under one minute. Figure 1(b) plots points in the distribution (i.e., the 20th-, 35th-, and 50th-percentile) of T/R as a function of the requested runtime of the job. This figure shows that for any requested runtime greater than one minute, 35% of the jobs use less than 10% of their requested runtime (i.e., $R \geq 10 T$), and another 15% of the jobs have actual runtime between 10% and 30% of the requested runtime. Similarly large discrepancies between requested and actual runtimes have also recently been reported for many SP2 traces [7, 8]. In particular, the results by Cirne and Berman [8] show that for four SP2 traces, 50-60% of the jobs use under 20% of the requested runtime, which is very similar to the results for the NCSA O2K workloads examined here.

2.3 Previous Work

In this section, we review previous work on three topics: alternative priority functions for backfill policies, the impact of reservation policies, and the impact of using more accurate requested runtimes on backfill policies.

The most comprehensive previous comparison of alternative priority backfill policies [5] shows that, among the priority functions defined in Table 2, the LXF&W-backfill policy that gives priority to short jobs while taking current job waiting time into account outperforms FCFS-backfill, whereas SJF-backfill has the problem of starvation (i.e., large maximum wait) under high load. This

previous paper also provides a review of earlier papers [13, 11, 14] that compare the SJF-backfill and FCFS-backfill policies.

Reservation policies concern (a) the number of jobs waiting in the queue that are given (earliest possible) reservations for processor and memory resources, and (b) whether the reservations are dynamic or fixed. Previous results by Feitelson and Weil [6] show that, for FCFS-backfill and a set of SP workloads, average slowdown is similar when only one (i.e., the oldest) waiting job has a reservation or when all jobs have a reservation. In more recent work [7] they find similar results for further SP2 workloads, for workloads from other systems, and for many synthetic workloads, but they find that for many other SP2 monthly workloads, a single reservation significantly improves the average slowdown (by $> 40\%$) and average response time (by $> 30\%$). Several papers evaluate backfill policies that have reservations for all waiting jobs [10, 14, 11], while still other papers evaluate backfill policies that give reservations to only one waiting job [3, 4, 15, 13, 5].

With dynamic reservations, job reservations and the ordering of job reservations can change when a new job arrives, or if the relative priorities of the waiting jobs change with time. For example, in SJF-backfill with a single dynamic reservation, an arriving job will preempt the reservation held by a longer job. With *fixed* reservations, in contrast, once a job is given a reservation, it may be given an earlier reservation when another job terminates earlier than its requested runtime, but recomputed job reservations will have the same order as the existing reservations, even if a job that has no reservation or a later reservation attains a higher priority. A single fixed reservation is used to reduce starvation in SJF-backfill in [5]. In [14], each job is given a reservation when it arrives. They compare a form of dynamic ("no guarantee") reservations, in which reservations are only recomputed if and when a job finishes early but the recomputed reservations are done in priority (i.e., FCFS or SJF) order, against "guaranteed reservations", in which job reservations are recomputed only in the same order as the existing reservations. They find that the dynamic reservations have lower average slowdown and average wait than guaranteed reservations for the priority backfill policies studied, including SJF-backfill. This paper includes the *maximum wait* measure and concludes that fixed reservations significantly improve the performance of SJF-backfill; otherwise the results in this paper are consistent with their results.

Two previous papers show that perfectly accurate requested runtimes for FCFS-backfill improve the average slowdown by no more than 30% [7] and the average wait time by only 10 - 20% [10], compared to using the highly inaccurate requested runtimes given in SP traces. Several papers [6, 13, 11, 12, 7] compare the performance of various models of requested runtimes against perfectly accurate runtime requests. For a given actual runtime, they model the requested runtime overestimation (i.e., requested runtime - actual runtime) as a factor times the actual runtime, where the factor is drawn from a uniform distribution between 0 and a fixed parameter C. The paper [13] also includes a model where the factor is deterministic. The results in those papers show that even for C as large as 300 [6, 7] (or 50 [13] or 10 [11, 12]), the average slowdown or average wait

is similar to, or even slightly better than that of $C = 0$. Additional results in [7] show that multiplying the user requested runtimes by two slightly improves on average slowdown and response time for SP workloads and FCFS-backfill. These papers conclude that there is no benefit of using accurate requested runtimes for FCFS-backfill and SJF-backfill. We note that for large C (or when multiplying requested runtime by two), jobs with long runtimes can have very large runtime overestimation, which leaves larger holes for backfilling shorter jobs. As a result, average slowdown and average wait may be lower, as reported in these previous papers. On the other hand, these systems may have poorer maximum wait, which was not studied in any of these previous papers.

3 Reducing Starvation in Systems that Favor Short Jobs

Backfill policies that favor short jobs have the potential problem of poor maximum wait for long jobs. Mechanisms for reducing the maximum wait include using a larger number of reservations, and increasing the priority weight on the current job wait time. On the other hand, either of these mechanisms may increase the average and 95th-percentile wait for all jobs. The goal of this section is to provide a more comprehensive evaluation of the trade-offs in the wait time measures for different reservation policies and for alternative priority functions that give different relative weight to the current job waiting time. In evaluating the tradeoffs for each policy, we seek to achieve a maximum wait that is no greater than the maximum wait in FCFS-backfill, while reducing the average and 95th-percentile wait time as much as possible.

In this section, and in the remainder of the paper, policy comparisons will be shown for five representative workloads. These workloads include (a) three of the four new *exceptionally heavy load* months (i.e., January - March 2001), which are the most important months for policy optimization, (b) June 2001, which has similar policy performance as in April 2001 since both of these months follow an exceptionally heavy load month, and (c) July 2001 which has a typical load and policy performance similar to October - December 2000 and other previously studied workloads. The other new exceptionally heavy load month (May 2001) has somewhat lower wait time statistics for each policy than the other three exceptionally heavy months, due to a larger number of short jobs submitted that month.

Section 3.1 re-evaluates previous backfill policies, showing that starvation is a more significant issue for the new exceptionally heavy load months on the NCSA O2K. Section 3.2 evaluates several alternative reservation policies. Section 3.3 evaluates several new priority functions with different relative weights on the current job waiting time and compares the best new priority backfill policies against FCFS-backfill.

3.1 Re-evaluation of Previous Policies

In this section, we use the recent O2K workloads to re-evaluate the FCFS-backfill, SJF-backfill, and LXF&W-backfill policies that are defined in Table 2.

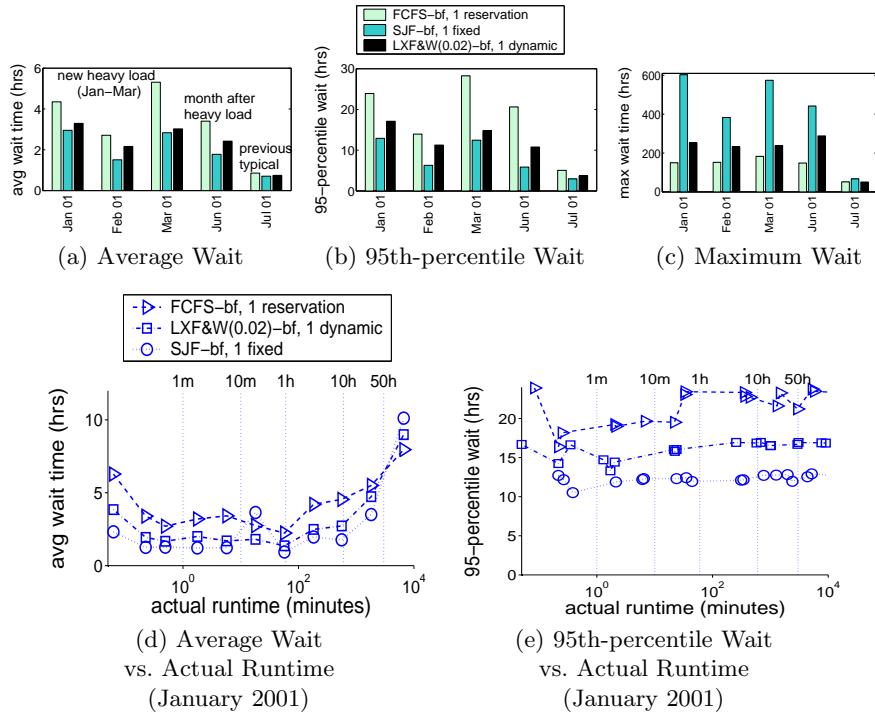


Fig. 2. Performance of Previous Backfill Policies

Note that both SJF-backfill and LXF&W-backfill favor short jobs, but LXF&W-backfill also has a priority weight for current job wait time. The reservation policies in these previously defined schedulers are: FCFS-backfill uses one reservation, LXF&W-backfill uses one dynamic reservation, and SJF-backfill uses one fixed reservation (to reduce the maximum wait).

Figure 2 compares the three policies, showing (a) overall average wait, (b) 95th-percentile wait, (c) maximum wait, and (d)-(e) average and 95th-percentile wait, respectively, as a function of actual runtime, during a representative heavy load month. Results in previous work [5] are similar to the results for the July 2001 workload in figures (a) - (c). Conclusions for the new heavy load months that are *similar* to previous work are that (1) SJF-backfill and LXF&W-backfill have significantly lower 95th-percentile wait (for all ranges of actual runtime) than that of FCFS-backfill, and (2) SJF-backfill has the problem of poor maximum wait for many of the workloads, as shown in figure (c). Conclusions for the new heavy load months that *differ* from the results in previous work (and also differ from the results for July 2001), are that (1) LXF&W-backfill and SJF-backfill have even greater improvement in *average wait* compared with FCFS-backfill (for most ranges of actual runtimes), and (2) LXF&W-backfill has *higher maximum wait* than FCFS-backfill.

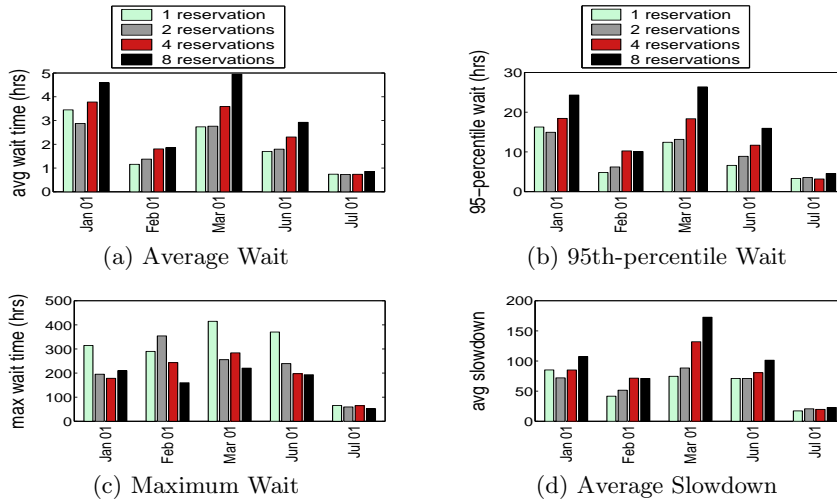


Fig. 3. Impact of Number of Reservations on LXF-backfill (Dynamic Reservations)

The starvation problems that lead to high maximum wait in LXF&W-backfill and SJF-backfill systems are addressed in the next two sections. The questions are (1) whether multiple reservations can improve the performance, particularly the maximum wait, of SJF-backfill and LXF&W-backfill, (2) whether fixed reservations can improve the maximum wait for LXF&W-backfill, and (3) whether new priority functions, such as adding a priority weight for current waiting time to the SJF-backfill priority function, or more generally whether new relative priority weights between requested job runtime and current job wait time, can improve on the previous policy priority functions. Section 3.2 addresses the first two questions. Section 3.3 studies the third question.

3.2 New Reservation Policy Comparisons

This section studies the impact of reservation policies, i.e., the number of reservations and dynamic versus fixed reservations, on backfill policies. We use three simple priority backfill policies to evaluate the reservation policies, namely: FCFS-backfill, SJF-backfill, and LXF-backfill (all with weight on current waiting time equal to zero). Adding weights for current waiting time will be studied in the next section.

For each of the three policies, we evaluated the performance for the following numbers of reservations: 1, 2, 4, 6, 8, 12, and 16. For the LXF-backfill and SJF-backfill policies that have dynamic priority functions, we evaluate the performance of both dynamic and fixed reservations, each over the entire range of number of reservations.

Figure 3 shows the performance of LXF-backfill with up to eight dynamic reservations. Twelve and sixteen reservations have similar or worse performance

Table 3. Weights for New Priority Backfill Policies

Priority Weight			Job Measure
SJF&W(w)	S \sqrt{T} F&W(w)	L \sqrt{X} F&W(w)	
$w = 0.05-0.2$	$w = 0.01-0.05$	$w = 0.01-0.02$	current wait time, J_w , in hours
1	0	0	$J_r = \frac{400^*}{R \text{ in hours}}$
0	1	0	$\sqrt{J_r}$
0	0	1	$\sqrt{J_x}$, where $J_x = \frac{J_w + R \text{ in hours}}{R \text{ in hours}}$

(* R = requested runtime. The maximum value of R is 400 hours.)

as that of eight reservations. The impact of the number of reservations is similar for FCFS-backfill and SJF-backfill (not shown to conserve space), except that four reservations performs slightly better for SJF-backfill. For months with a typical O2K load (e.g., July 2001), the impact of reservation policies on backfill policies is minimal, which agrees with previous results for the average slowdown of FCFS-backfill in [6]. However, for most of the new heavy load months, as shown in Figure 3(c), the key new result is that using a small number of reservations (i.e., 2-4) reduces the maximum wait time (by about 30%) compared to using a single reservation. Furthermore, as shown in Figure 3(a) - (c), using more than four reservations usually makes minimal further improvement for the maximum wait, yet significantly increases the average and 95th-percentile wait, for the new heavy load workloads or immediately following the heavy load months.

Other results omitted to conserve space show that fixed and dynamic reservations (with 2-4 reservations) have similar performance for LXF-backfill and the policies developed in the next section. However, for SJF-backfill, dynamic reservations has higher maximum wait than fixed reservations because (particularly under heavy load) dynamic reservations for jobs with long requested runtimes are often usurped by newly arriving jobs that have short requested runtimes.

3.3 New Priority Functions

In this section, we propose three alternative new priority functions, evaluate the impact of the alternative priority weights for current job wait time together with the impact of reservation policies, and compare the best new priority functions against the previous backfill policies.

The new (dynamic) priority functions are defined in Table 3. Note that in the J_r metric for the SJF&W and S \sqrt{T} F&W priority functions, the inverse of requested runtime ($1/R$) is normalized to the maximum allowed requested runtime (i.e., 400 hours). The SJF&W priority function extends the previous SJF function with a weight for the current job wait time. The S \sqrt{T} F&W and the L \sqrt{X} F&W priority functions are designed to increase the relative weight on the current job waiting time (J_w), by applying a square root to the job metric that includes requested runtime. Results below show that S \sqrt{T} F&W-backfill and L \sqrt{X} F&W-backfill only very slightly outperform SJF&W-backfill and LXF&W-




Fig. 4. Alternative Job Wait Priority Values (w) for $L\sqrt{XF}\&W$ -backfill (One Dynamic Reservation)

backfill. Thus, further alternative relative weights for these job measures are not likely to lead to any significant improvement.

Let w denote the priority weight on current job wait time in the new priority functions. Figure 4 shows the performance of alternative values of w in the $L\sqrt{XF}\&W$ -backfill scheduler with one dynamic reservation. Results are similar for 2-4 reservations, and for each of the other two new priority functions, as well as for LXF&W (not shown). The Figure shows that average and 95th-percentile wait are not highly sensitive to w in the range of 0.005 - 0.05, and that during heavy load months, this range of w values significantly reduces the maximum wait (by 30-50%) compared to $w = 0$. Larger values of w , such as $w = 1$, significantly increase the average and 95th-percentile wait time, with only small improvements in the maximum wait (not shown).

Similar to the previous section, we find that using a small number of reservations (i.e., two or three) outperforms a single reservation for each of the alternative new priority functions.

Figure 5 compares the performance of FCFS-backfill, LXF&W(0.02)-backfill, and the two best alternative new priority backfill policies (i.e., $L\sqrt{XF}\&W(0.01)$ and $S\sqrt{TF}\&W(0.05)$ -backfill, which slightly outperform SJF&W-backfill), each with 2 - 3 reservations. One key result is that using 2-4 reservations instead of one reservation has improved the overall performance of all four policies. For example, compared to Figure 2, the maximum wait for FCFS-backfill and LXF&W(0.02)-backfill is reduced by up to 30% while the average or 95th-percentile wait is increased by on the order of 10% or less. Another key result is that $L\sqrt{XF}\&W(0.01)$ -backfill with 2-4 reservations has maximum wait that is reasonably competitive with FCFS-backfill, yet significantly outperforms FCFS-backfill with respect to the other wait time statistics, particularly average slowdown. $L\sqrt{XF}\&W$ -backfill has slightly better overall performance than LXF&W-backfill. Finally, $S\sqrt{TF}\&W$ -backfill has better average and 95th-percentile wait than $L\sqrt{XF}\&W$ -backfill, but more often has significantly poorer maximum wait than FCFS-backfill (e.g., in February and June 2001).

The overall conclusion is that, similar to results in [5], giving priority to short jobs but also using an appropriate weight for current job wait can significantly

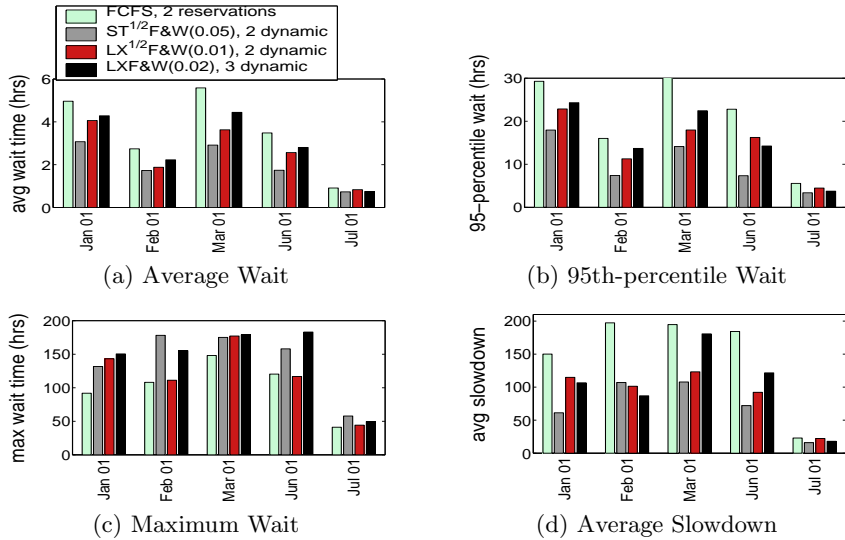


Fig. 5. Performance of New Priority Backfill Policies that Favor Short Jobs

outperform FCFS-backfill, particularly with respect to the 95th percentile wait time and the average slowdown measures. In the remainder of this paper, we study the impact of more accurate requested runtimes on such high performance backfill policies.

4 More Accurate Requested Runtimes

There is reason to believe that runtimes can be more accurately estimated for the jobs that run on the O2K. In particular, a majority of the jobs use one of the default requested runtimes, which are 5, 50, 200, or 400 hours. This indicates that users specify highly approximate requested runtimes due to the course-grain defaults that are available. Furthermore, since the current priority-backfill policy provides similar 95th-percentile waiting time for the entire range of job runtimes (see Figure 2(e) and results in [5]), there isn't currently any incentive for an individual user to provide a more accurate requested runtime. These factors explain why, for example, many of the jobs that have actual runtime of 10 hours have requested runtime of 50, 200, or 400 hours.

Previous results suggest that using more accurate requested runtimes has only minimal impact on the average slowdown and average wait time for FCFS-backfill [6, 10, 13, 11, 12]. This section investigates whether the benefit of more accurate requested runtimes is more significant for the high performance priority backfill policies that use requested runtimes to favor short jobs. We consider various scenarios more accurate runtime requests, the O2K workloads that include exceptionally heavy load in recent months, and more extensive performance measures than in the previous evaluations of FCFS-backfill. Section 4.1 describes the

Table 4. Notation

Symbol	Definition
T	Actual job runtime
R	User requested runtime from the O2K logs
R*	Simulated requested runtime
P	Number of requested processors

scenarios of more accurate requested runtime that will be evaluated. Section 4.2 reassesses the impact of more accurate requested runtimes on FCFS-backfill, whereas Section 4.3 evaluates the impact of more accurate requested runtimes on the policies that favor short jobs.

4.1 More Accurate Requested Runtime Scenarios

Using the notation defined in Table 4, we consider three different scenarios of more accurate runtime requests for the O2K workloads. In the first (implausible) case, each requested runtime is perfect (i.e., the simulated requested runtime, $R^* = T$). In the second case, all requested runtimes are imperfect but are *approximately* accurate (i.e., $R^* = \min\{R, kT\}$, $1 < k \leq 2$). In the third case, only a fraction (e.g., 80% or 60%) of the jobs have the approximately accurate requested runtimes, while the rest of the jobs, selected randomly, have requested runtimes as given in the job log, which are generally highly inaccurate. In the third case, the fraction of jobs that have the inaccurate requested runtimes from the trace represent carelessly specified runtime requests or runtimes that can't (easily) be accurately estimated. This fraction is varied in the experiments. The first case is used to provide a bound on the maximum benefit of more accurate requested runtimes, while the second and third cases are used to assess performance gains that are more likely to be achievable. Section 5 will explore the performance impact of using short test runs to achieve the more accurate runtime requests.

We present results for $k = 2$. We also considered smaller values of k , in particular $k = 1.2$, which results in slightly better performance, but we omit those results below to conserve space. As noted in Section 2.3, several previous papers [6, 13, 11, 12] have used a uniform distribution of requested runtime overestimations, with a large upper bound factor (e.g., 10, 50, or 300). In contrast, our scenarios assume that requested runtime is not larger than the actual requested runtime in the workload trace.

4.2 FCFS-backfill Results

Figure 6 compares the performance of perfectly accurate requested runtimes (i.e., $R^* = T$) against user requested runtimes from the trace (i.e., $R^* = R$) for FCFS-backfill with two reservations. The results for previous typical O2K workloads (e.g., July 2001) agree with previous results in [6]; that is, using more accurate runtimes has only very slight impact on system performance. Moreover,

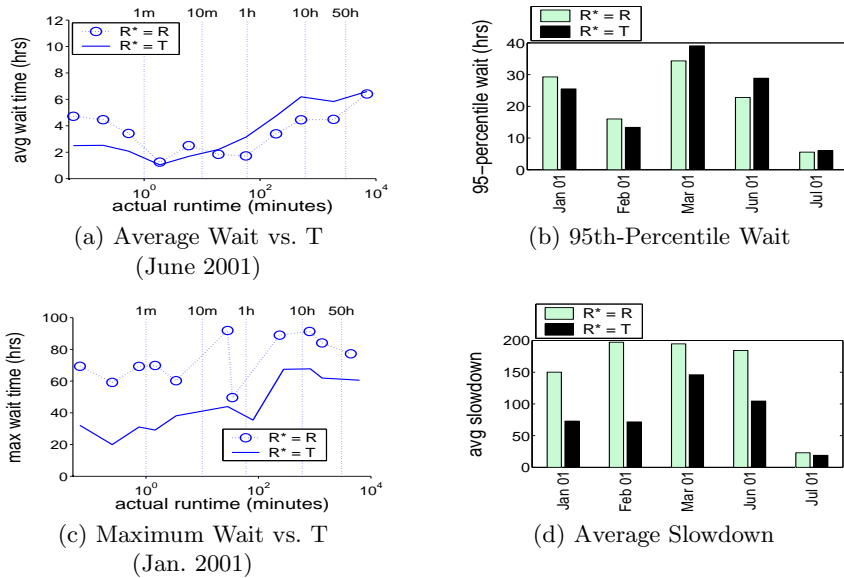


Fig. 6. Impact of Perfect Requested Runtimes for FCFS-backfill

perfect requested runtimes have minimal impact on the overall average waiting time for each month (not shown), and on the 95th-percentile wait each month, shown in Figure 6(b). On the other hand, as shown in Figure 6(a) for June 2001, accurate runtime requests improve the average wait of very short jobs ($T < 30$ minutes) during and immediately following the new exceptionally heavy load months. More significantly, Figure 6(c) shows that accurate requested runtimes significantly improve maximum wait time for most actual job runtimes, for many of the exceptionally heavy load months and immediately following new heavy load months. Figure 6(d) shows that actual runtimes significantly reduce average slowdown under and immediately following new heavy load months (by up to 60% in Feb. 2001).

We note that perfect requested runtimes generally improves the wait time for short jobs because these jobs can be backfilled more easily. Accurate requested runtimes also improve the maximum wait for long jobs due to shorter backfill windows. Using approximately accurate requested runtimes (i.e., $R^* = kT$) has a somewhat lower impact on system performance than using perfect runtime requests (not shown to conserve space).

4.3 Results for High Performance Backfill Policies

This section evaluates the impact of more accurate requested runtimes on the performance of high performance backfill policies that favor short jobs. We present the results for $L\sqrt{X}F&W$ -backfill. Results are similar for the other high performance backfill policies such as $S\sqrt{T}F&W$ -backfill and $LXF&W$ -backfill.

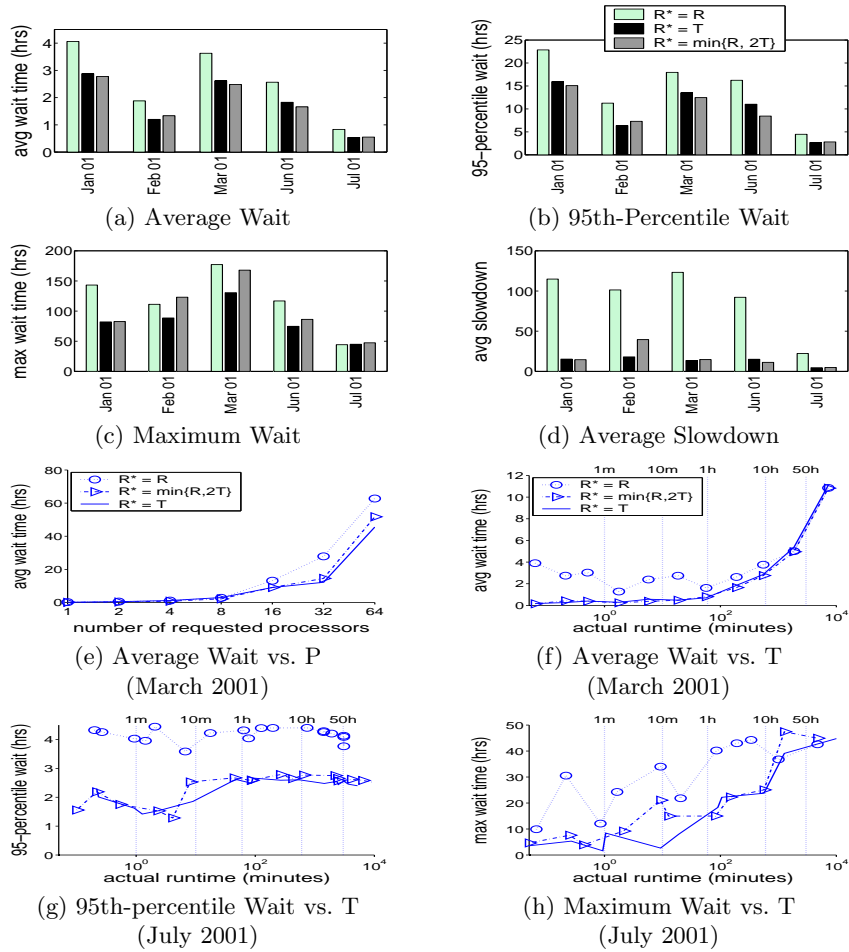


Fig. 7. Impact of Accurate Runtime Requests for $L\sqrt{XF}$ &W-backfill

Initially, we consider the case where *all* jobs have requested runtimes within a small factor of their actual runtimes. Then, we consider the case where only 80% or 60% of the jobs have approximately accurate requested runtimes.

Figure 7 compares the performance of perfectly accurate runtime requests (i.e., $R^* = T$) and approximately accurate runtime requests (i.e., $R^* = \text{Min}\{R, 2T\}$) against requested runtimes from the trace (i.e., $R^* = R$). Graphs (a)-(d) contain the overall performance measures each month, whereas graphs (e)-(h) show performance versus requested number of processors or actual runtime for a given representative monthly workload. Results for other months (not shown) are similar.

In contrast to the FCFS-backfill results shown in Figure 6, there is an even larger benefit of using more accurate requested runtimes for $L\sqrt{XF}$ &W-backfill, because accurate runtime requests enable the scheduler to give priority to jobs

that are actually shorter. In particular, Figures 7(a) - (d) show that perfectly accurate runtime requests improve not only the maximum wait and average slowdown, but also the average and 95th-percentile wait time over all jobs. Furthermore, the average slowdown is dramatically improved in *every* month, including the months with typical O2K load (e.g., July 2001). These four graphs also show that even if the the requested runtimes are only approximately accurate (i.e., $R^* = \text{Min}\{R, 2T\}$), similar improvements are achieved.

Figure 7(e) shows that accurate or approximately accurate requested runtimes improve the average wait time for jobs that require a large number of processors (i.e., greater than 16 processors).

Figures 7(f)-(h) show that more accurate requested runtimes improve the average wait for short jobs (up to 10 hours), 95th-percentile wait for all jobs, and the maximum wait for all but the very largest jobs. Note that these improvements occur for typical system loads as well as exceptionally heavy loads that during months following the exceptionally heavy load. Note also that the improvement in the average wait for short jobs is significantly larger than the improvement for FCFS-backfill systems, *and* the improvement is achieved without increasing the average wait time for longer jobs. Furthermore, when requested runtimes are accurate or approximately accurate, the average wait under $L\sqrt{X}F\&W$ -backfill decreases (monotonically) as the actual runtime decreases; this is a desirable property that, to our knowledge, has not been demonstrated for any previous backfill scheduler.

Next, we consider scenarios in which not all jobs have approximately accurate requested runtimes. Two systems are evaluated: hybrid(4:1) and hybrid(3:2). In the hybrid(4:1) system, 4 out of 5 jobs (i.e., 80% of jobs), selected randomly, have approximately accurate requested runtime (i.e., $R^* = \text{min}R, 2T$). The hybrid(3:2) system is similar to the hybrid(4:1) system, except that only three out of five (or 60%) of the jobs have the approximately accurate runtime requests. Results will be shown for $L\sqrt{X}F\&W$ -backfill; the results are similar for the other priority backfill policies that favor short jobs.

Figure 8 compares hybrid(4:1) and hybrid(3:2) against the case where all jobs have perfectly accurate runtime requests (i.e, $R^* = T$), and the case where all jobs use requested runtimes from the trace (i.e, $R^* = R$). The key conclusion is that much of the benefit of accurate requested runtimes can be achieved even if only 60% or 80% of the jobs have approximately accurate requested runtimes. Specifically, Figures 8(a) and (b) show that hybrid(4:1) has similar average and 95th-percentile wait time as the system where all jobs have $R^* = T$. Figure 8(c) shows that hybrid(4:1) has somewhat higher maximum wait than when requested runtimes are perfectly accurate, but has lower maximum wait than for user requested runtimes in the trace. Figure 8(d) shows that hybrid(4:1) has much lower average slowdown than the system with user requested runtimes from the trace. If only 60% of the jobs have improved requested runtimes, i.e., hybrid(3:2), the performance improvement is smaller than that in hybrid(4:1), but hybrid(3:2) still has lower average and 95th-percentile wait time and significantly lower average slowdown than that of using very inaccurate requested




Fig. 8. Hybrid(x:y) Approximately Accurate:Inaccurate Requested Runtimes ($L\sqrt{XF}$ &W-backfill; Approximately Accurate $R^* = \text{Min}\{R, 2T\}$)

runtimes from the trace. Further reducing the fraction of the jobs to have improved requested runtimes results in a system increasingly more similar to the system where all jobs have requested runtimes from the trace.

The next results show that the jobs in the hybrid systems that have more accurate requested runtimes experience substantial performance benefit. In particular, Figure 9 compares the wait time statistics for ‘*approx. accurate jobs*’ (i.e., $R^* \leq 2T$) in the hybrid system against the wait time statistics for ‘*inaccurate jobs*’ (i.e., $R^* = R > 2T$) in the hybrid system. The figure also includes the performance when all jobs have requested runtimes as in the workload trace (i.e., all jobs have $R^* = R$). The results are shown for hybrid(3:2), in which only 60% of the jobs have approximately accurate requested runtimes. As noted in the figure captions, only the jobs with under 50 hours of actual runtime are considered in the first four graphs because requested runtime accuracy has little impact on jobs with larger actual runtime (as can be seen in Figure 9(e)). The key results are:

- Figures 9(a) - (c) show that during and immediately following the extremely heavy load months, for actual runtime up to 50 hours, jobs with approximately accurate runtime requests have 20% lower average and 95th-percentile wait time and up to 50% lower maximum wait time than the jobs with inaccurate runtime requests.
- Furthermore, the jobs with approximately accurate runtime requests improve the average and 95th-percentile wait time of inaccurate jobs, compared to when all jobs have the requested runtimes from the trace.
- Figure 9(d) shows that for any month, the average slowdown of jobs with approximately accurate runtime requests is dramatically (i.e., more than an

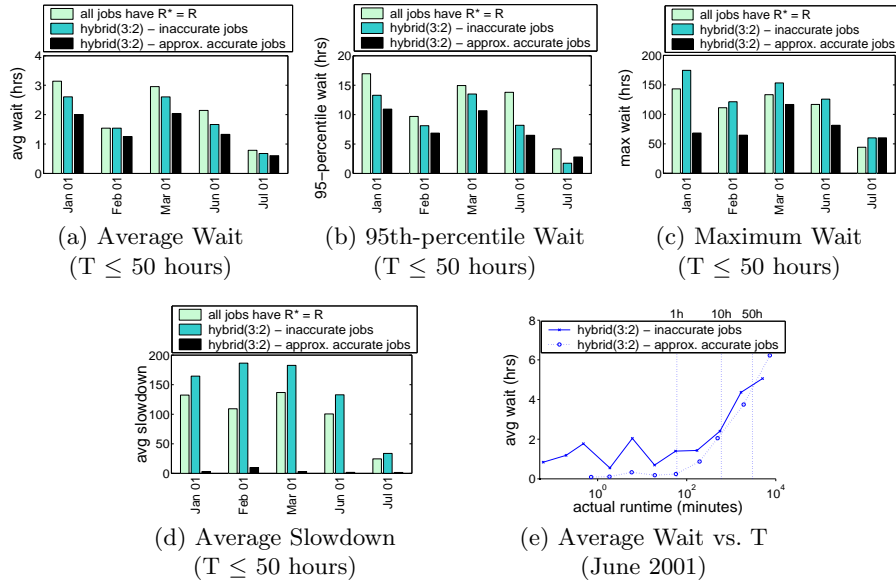


Fig. 9. Performance for Jobs with More Accurate Requested Runtimes ($L\sqrt{XF}$ &W-backfill; Approximately Accurate $R^* = \text{Min}\{R, 2T\}$)

order of magnitude) lower than either the average slowdown of jobs with inaccurate requests, or the overall average slowdown when all jobs use the requested runtime from the trace (i.e., $R^* = R$).

- Figure 9(e) further shows that for actual runtime of up to 10 hours, jobs with approximately accurate requests achieve significantly lower average wait time than that of inaccurate jobs, and average wait decreases monotonically as actual runtime decreases for the jobs with approximately accurate requests.

5 Test Runs for Improving Requested Runtimes

Results in Section 4 show that if a majority of the jobs (e.g., 60% or more) have estimated runtimes within a factor of two of their actual runtime, system performance improves greatly, particularly for the jobs that have such approximately accurate runtime requests. Thus, if users are provided with incentives and tools to provide more accurate requested runtimes, the users will reap significant performance benefit.

We hypothesize that approximately accurate requested runtimes are feasible in at least three cases. First, many large scale simulations are run with similar input parameters to previous runs, or with changes in the input parameters that will affect run time in an approximately predictable way (e.g., runtime can be estimated within a factor of two). Second, for other applications, the runtime request can be more accurate if a short test run is made before the full run. Example applications that can estimate requested runtime after a test

run include those that involve iterative computation in which the number of iterations and/or the time per iteration are dependent on the input data, but can be estimated reasonably well after having run the first few iterations. In the third case, many applications such as stochastic optimization have a number of iterations that is dependent on how quickly the solution converges, which generally can't be predicted ahead of time. In this case approximately accurate requested runtimes could be provided if the computation is performed in several runs, each except the last run having requested runtime that is shorter than needed to reach final convergence, and with the solution from one run being input to the next run.

The remainder of this section investigates whether most of the benefit of more accurate requested runtimes shown in the previous section can still be realized if some jobs perform a short test run before providing an approximately accurate requested runtime.

To address this question, the following assumptions are made regarding the test runs. If the user requested runtime is already within a factor of two of the actual runtime (i.e., $R \leq 2T$), we assume that the user is aware that a test run is not needed, and the job is simply submitted with the requested runtime supplied by the user. For the remaining jobs, a specified fraction (i.e., 100% in section 5.1 or 50% - 80% in section 5.2) are assigned more accurate requested runtimes than specified by the user. The jobs that do not have more accurate requested runtimes represent jobs for which the user is either not able or not interested in estimating runtime more accurately. Of the jobs that are given more accurate requested runtimes, some fraction (e.g., 25%) require a test run before the more accurate request is given. If the test run is not used, the more accurate runtime request is assumed to be estimated from previous runs of the application. The requested runtime for a test run is equal to: (a) 10% of the user requested runtime if the user requested runtime is under 10 hours, or (b) one hour if the user requested runtime is greater than 10 hours. That is, the requested runtime for the test run is equal to the minimum of 1 hour and 10% of the user requested runtime (R). The requested runtime for the test run represents the runtime the user believes is needed to estimate the full job runtime within a small factor. Note that because the user requested runtimes can be highly inaccurate, the actual job runtime may be shorter than the requested runtime for the test run. In such cases only the test run is needed. This represents the case in the actual system, in which jobs complete during the test run either due to the user's lack of experience in how long the test run should be, or due to an unexpected error in the execution. If the actual job runtime is longer than the test run or a test run is not needed, the job is submitted with an approximately accurate requested runtime (i.e., a requested runtime equal to twice the actual runtime, $2T$).

Section 5.1 considers the scenario in which all full runs have requested runtime within a factor of two of the actual runtime, but two different fractions of jobs (i.e., 100% or 25%) (randomly selected) make test runs before submitting with the approximately accurate requested runtime. Section 5.2 considers the scenario in which only 50% or 80% of the jobs provide approximately accurate

requested runtimes, whereas the other 50% or 20% of the jobs provide the same requested runtimes as in the job trace. Of the jobs that provide approximately accurate runtime requests, 25% make the test run before submitting with the approximately accurate request.

5.1 Improved Requested Runtimes for All Jobs

This section studies the impact of test runs for the optimistic ("best case performance") scenario in which all jobs provide approximately accurate requested runtimes. In one case ("25% testrun"), 25% of jobs that do not have approximately accurate user requested runtime from the trace are randomly selected to have a test run. In the other case ("100% testrun"), every job with improved runtime request requires a test run. Note that "100% testrun" is a pessimistic assumption that is not likely to occur in practice, since many applications are run a large number of times, and in many cases previous executions can be used to provide approximately accurate runtime requests. Thus, we consider the "25% testrun" experiments to be more representative of the practical impact of using test runs to improve runtime estimate accuracy.

During each month, 35-45% of the jobs have inaccurate requested runtimes (i.e., $R > 2T$) and actual runtime T greater than the minimum of one hour and 10% of the user runtime request. For such jobs, if a test run is used to improve requested runtime, the job is resubmitted after the test run. The total extra load due to the test runs is very small (only 1-3% increase in processor and memory demand each month), even for 100% testrun. However, the additional waiting time for the test run, and the test run, must be included in the measures of total job waiting time. The results below address the impact of this extra waiting time on the overall system performance.

Figure 10 compares 100% testrun and 25% testrun against the optimal case where all jobs use actual runtimes (i.e., $R^* = T$) without test runs and the case where all jobs use the requested runtimes from the trace (i.e., $R^* = R$). The average total wait, 95th-percentile total wait, maximum total wait, and average slowdown, are shown for representative recent O2K workloads. For each of these measures except average slowdown during February 2001, the performance of the 25% testrun case is very similar to the case where $R^* = T$. Overall the results show that a significant fraction of test runs can be made to improve requested runtimes, and if the improved requested runtimes are within a factor of two of the actual runtime, then nearly the maximum possible benefit of accurate requested runtimes can be achieved.

The test run overhead becomes prominent if all jobs with $R > 2T$ require a test run (i.e., 100% testrun). Even so, Figure 10(a) shows that during and immediately following the heavy load months, "100% testrun" has lower average and 95th-percentile wait, and especially lower average slowdown, than using requested runtimes from the trace.

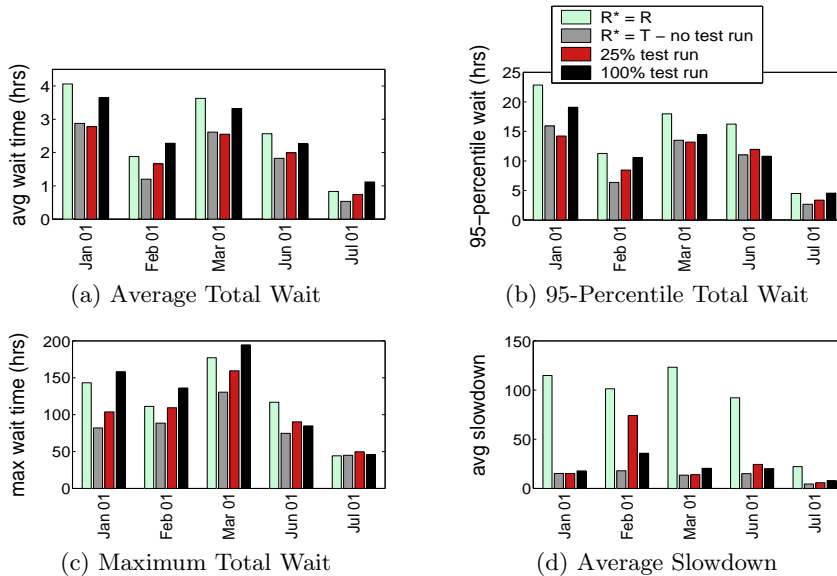


Fig. 10. Impact of Test-Runs to Determine Requested Runtimes ($L\sqrt{XF\&W}$ -backfill; Wait includes testrun wait and overhead; $R^* = \text{Min}\{R, 2T\}$)

5.2 Improved Requested Runtimes for a Majority of the Jobs

This section evaluates scenarios where only 50% or 80% of the jobs have improved requested runtime accuracy, and test runs are needed for 25% of the jobs that have improved requested runtimes. The two scenarios are named hybrid(1:1) - 25% testrun and hybrid(4:1) - 25% testrun, respectively. Note that hybrid(1:1) with 25% testrun represents a reasonably pessimistic, but possibly realistic scenario, in which only 50% of the jobs have approximately accurate requested runtimes and one out of four jobs requires a test run to improve requested runtime accuracy. Again, we use $R^* = 2T$ for approximately accurate runtime requests.

Figure 11 compares the above two scenarios with 25% test run against that of using requested runtimes from the trace (i.e., " $R^* = R$ "). The performance for hybrid(4:1) without test run is also included in Figure 11(b) - (d) for comparison with hybrid(4:1) with 25% testrun. The results show that for both hybrid systems with testruns, the average wait for short jobs, the 95th-percentile wait, and the average slowdown is significantly better than for the requested runtimes in the O2K traces. The results also show that test runs do not introduce significant overhead in the hybrid (4:1) system.

6 Conclusions

In this paper, we used ten one-month recent traces from the NCSA O2K to evaluate whether high-performance backfill policies can be significantly improved

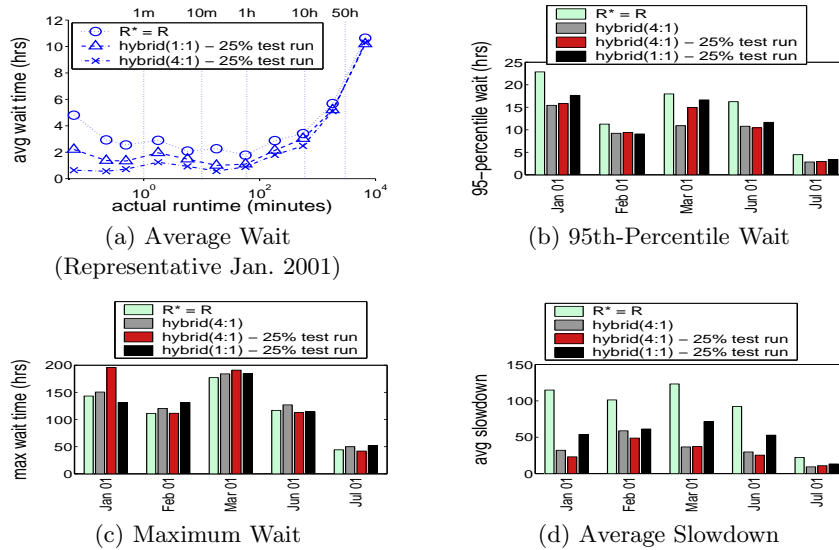


Fig. 11. Performance of Hybrid(x:y) with Test Runs
($L\sqrt{XF}$ &W-backfill; $R^* = \text{Min}\{R, 2T\}$)

if the requested runtimes can be more accurate. Several of these months have exceptionally heavy load, which tends to result in larger policy performance differentials than for the lower loads used in previous work.

To select the best backfill policies for studying this key question, we more fully evaluated the issues related to starvation in backfill policies that favor short jobs. The results show (1) that a few reservations (2 - 4) can significantly reduce the maximum wait time but a larger number of reservations result in poor performance, and (2) fixed reservations have similar performance to dynamic reservations in most cases, except for SJF-backfill which requires fixed reservation to reduce starvation. The results also show that two new priority backfill policies, namely $L\sqrt{XF}$ &W-backfill and $S\sqrt{TF}$ &W-backfill, achieve a high performance trade-off between favoring short jobs and preventing starvation.

The results for the high-performance backfill policies, heavier system load, and a more complete set of performance measures show that the potential benefit of more accurate requested runtimes is significantly larger than suggested in previous results for FCFS-backfill. Furthermore, the results show that most of the benefit of more accurate requested runtimes can be achieved by using test runs to improve requested runtime accuracy, in spite of the time needed to perform the test run. Another key result is that users who provide more accurate requested runtimes can expect improved performance, even if other jobs do not provide more accurate requested runtimes.

Topics for future work include developing approaches for achieving more accurate requested runtimes in actual systems, improving job placement during backfill to reduce system fragmentation, and extending the high performance

policies for use in more complex distributed Grid architectures such as the TeraGrid.

References

1. National Computational Science Alliance Scientific Computing: Silicon Graphics Origin2000. (<http://www.ncsa.uiuc.edu/SCD/Hardware/Origin2000>)
2. NCSA Scientific Computing: IA-32 Linux Cluster. (<http://www.ncsa.uiuc.edu/UserInfo/Resources/Hardware/IA32LinuxCluster>)
3. Lifka, D.: The ANL/IBM SP scheduling system. In: Proc. 1st Workshop on Job Scheduling Strategies for Parallel Processing, Santa Barbara, Lecture Notes in Comp. Sci. Vol. 949, Springer-Verlag (1995) 295–303
4. Skovira, J., Chan, W., Zhou, H., Lifka, K.: The EASY-Loadleveler API Project. In: Proc. 2nd Workshop on Job Scheduling Strategies for Parallel Processing, Honolulu, Lecture Notes in Comp. Sci. Vol. 1162, Springer-Verlag (1996) 41–47
5. Chiang, S.H., Vernon, M.K.: Production job scheduling for parallel shared memory systems. In: Proc. Int'l. Parallel and Distributed Processing Symp. (IPDPS) 2001, San Francisco (2001)
6. Feitelson, D.G., Weil, A.M.: Utilization and predictability in scheduling the IBM SP2 with backfilling. In: Proc. 12th Int'l. Parallel Processing Symp., Orlando (1998) 542–546
7. Mu'alem, A.W., Feitelson, D.G.: Utilization, predictability, workloads, and user runtime estimates in scheduling the IBM SP2 with backfilling. *IEEE Trans. Parallel and Distributed Syst.* **12** (2001) 529–543
8. Cirne, W., Berman, F.: A comprehensive model of the supercomputer workload. In: Proc. IEEE 4th Annual Workshop on Workload Characterization, Austin, TX. (2001)
9. Chiang, S.H., Vernon, M.K.: Characteristics of a large shared memory production workload. In: Proc. 7th Workshop on Job Scheduling Strategies for Parallel Processing, Cambridge, MA. (2001)
10. Smith, W., Taylor, V., Foster, I.: Using run-time predictions to estimate queue wait times and improve scheduler performance. In: Proc. 5th Workshop on Job Scheduling Strategies for Parallel Processing, San Juan, Lecture Notes in Comp. Sci. Vol. 1659, Springer-Verlag (1999) 202–219
11. Zhang, Y., Franke, H., Moreira, J.E., Sivasubramaniam, A.: Improving parallel job scheduling by combining gang scheduling and backfilling techniques. In: Proc. Int'l. Parallel and Distributed Processing Symp. (IPDPS) 2000, Cancun (2000)
12. Zhang, Y., Franke, H., Moreira, J.E., Sivasubramaniam, A.: An analysis of space- and time-sharing techniques for parallel job scheduling. In: Proc. 7th Workshop on Job Scheduling Strategies for Parallel Processing, Cambridge, MA. (2001)
13. Zotkin, D., Keleher, P.J.: Job-length estimation and performance in backfilling schedulers. In: 8th IEEE Int'l Symp. on High Performance Distributed Computing, Redondo Beach (1999) 236–243
14. Perkovic, D., Keleher, P.J.: Randomization, speculation, and adaptation in batch schedulers. In: Proc. 2000 ACM/IEEE Supercomputing Conf., Dallas (2000)
15. Gibbons, R.: A historical application profiler for use by parallel schedulers. In: Proc. 3rd Workshop on Job Scheduling Strategies for Parallel Processing, Geneva, Lecture Notes in Comp. Sci. Vol. 1291, Springer-Verlag (1997)