Workload Characteristics of the DAS-2 Supercomputer

Hui Li* Lex Wolters* David Groep†

*Leiden Institute of Advanced Computer Science (LIACS), Leiden University
2333 CA, Leiden, The Netherlands

†National Institute for Nuclear and High Energy Physics (NIKHEF)
1098 SJ, Amsterdam, The Netherlands

hli@liacs.nl llexx@liacs.nl davidg@nikhef.nl

Keywords: Workload Characterization, Performance Metric, DAS-2 Supercomputer

Abstract

This paper presents a comprehensive characterization of the DAS-2[1] workloads using twelve-month scientific traces. Metrics that we characterize include system utilization, job arrival rate and interarrival time, job size (degree of parallelism), job run time, memory usage, and job queue wait time. Differences with previous reported workloads are recognized and statistical distributions are fitted for generating synthetic workloads with the same characteristics. This study provides a realistic basis for experiments in resource management and evaluations of different scheduling strategies in a multi-cluster environment.

1 Introduction

Workload characterization of parallel supercomputers is important to understand the system performance and develop workload models for evaluating different system designs and scheduling strategies [1, 2]. During the past ten years, lots of workload data has been collected [3], analyzed [4, 5, 6], and modeled [7, 8, 9]. Benchmarks and standards are also proposed for scheduling on parallel computers [10].

In previously studied workloads [4, 5, 6, 7], some characteristics are similar. For example, most of the workloads are collected from large custom-made machines (IBM SP2, SGI Origin, etc) in supercomputing centers. Jobs typically request "power-of-two" number of processors and have different arrival patterns in different periods (e.g. peak and none-peak hours in a daily cycle). Some characteristics, such as correlations and fitted distributions, vary across different workloads [4, 5, 11]. We compare our workload with previous reported ones on a per characteristics basis.

This paper presents a comprehensive workload characterization of the DAS-2 [12] supercomputer. The DAS-2 system is interesting in that it is built using the popular COTS (Commodity Off The Shelf) components (e.g. Intel Pentium processors and Ethernet networks) and consists of multiple distributed clusters serving the DutchGrid community [13]. We analyze twelve-month scientific workloads on DAS-2 clusters in year 2003. Characteristics that we analyzed include system utilization, job arrival rate and interarrival time, job size (degree of parallelism), job run time, memory usage, and job queue wait time.

The contributions of this paper reside in the following. Firstly, our study is based on cluster workloads. Cluster computing is a popular alternative in the HPC community and to our knowledge, not much work has been done in characterizing cluster workloads. Secondly, we present a comprehensive characterization of the DAS-2 workloads. Some new measures, such as job queue wait time, are provided. Moreover, we fit the observed data with statistical distributions to facilitate synthetic workload generation. This research serves as a realistic basis in modeling cluster workloads, which contributes to evaluations of scheduling strategies in a multi-cluster environment [14].

The rest of the paper is organized as follows. Section 2 provides an overview of the DAS-2 system and workload traces used in our study. Section 3 analyzes the overall system utilization. Section 4 describes the job arrival characteristics, including job arrival rate and interarrival time. Statistical distributions are fitted for the job interarrival time. Section 5 describes job execution characteristics. This includes job size, job actual runtime, and memory usage. Distributions are also provided for job size and runtime. Section 6 characterizes and models the job queue wait times. In section 7 conclusions are presented and future work is discussed.

[1]Distributed ASCI Supercomputer-2 (DAS-2). ASCI stands for Advanced School for Computing and Imaging in the Netherlands.

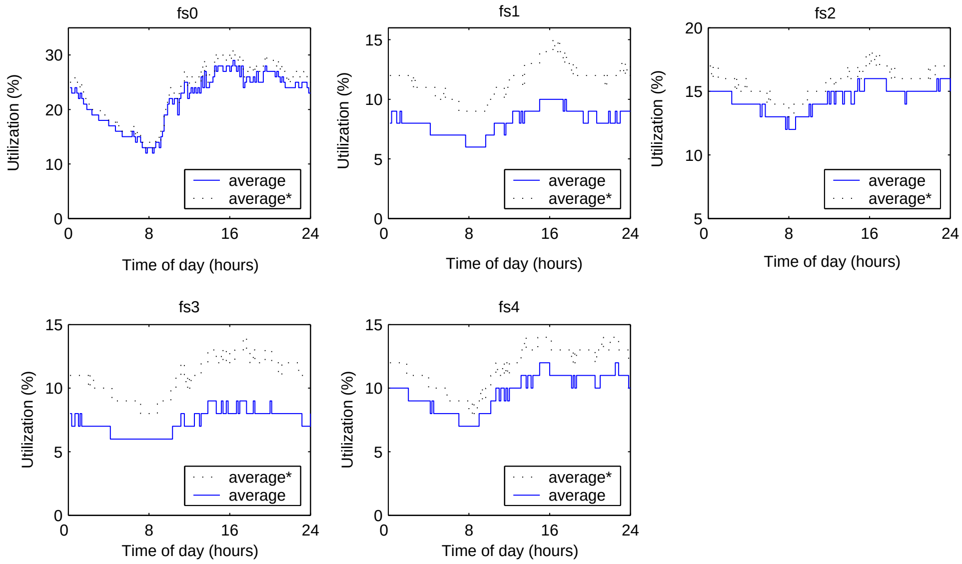

Figure 1: System utilization of DAS-2 clusters. "Average" stands for the average utilization of all days in the year. "Average*" stands for the average utilization of all active days in the year, excluding system downtime and days without job arrivals.

2 The DAS-2 supercomputer and workload traces

DAS-2 consists of five clusters located at five Dutch universities and is primarily used for parallel and distributed computing research. The largest cluster (Vije Universiteit) contains 72 nodes and the other four clusters have 32 nodes each. Every node has a dual 1GHz Pentium III processor, 1GB RAM and a 20GB local disk. The clusters are interconnected by the Dutch university internet backbone and the nodes within a local cluster are connected by high speed Myrinet as well as Fast Ethernet LANs. All clusters use openPBS [15] as local batch system and Maui (with backfilling) [16] as scheduler. Jobs that require multi-clusters can be submitted using toolkits such as Globus [17]. DAS-2 runs RedHat Linux as the operating system.

We use job traces recorded in the PBS accounting logs for twelve months in year 2003 on the five clusters[2]. All jobs in the traces are *rigid* (jobs that do not change parallelism at runtime) batch jobs. An overview of the DAS-2 system and workload traces is provided in Table 1. As we can see, fs0 (VU) is the most active cluster, with more than two hundred thousand job entries. Next we have clusters at UvA (fs2) and Delft (fs3), each with more than sixty thousand entries. Leiden (fs1) and Utrecht (fs4) are relatively less active among the DAS-2 clusters. Next section gives a more detailed analysis on the overall system utilization.

[2]Logs of January on fs4 are not available.

cluster	location	#cpus	period	#jobs
fs0	VU	144	01-12/2003	219618
fs1	Leiden	64	01-12/2003	39356
fs2	UvA	64	01-12/2003	65382
fs3	Delft	64	01-12/2003	66112
fs4	Utrecht	64	02-12/2003	32953

Table 1: DAS-2 clusters and workload traces.

3 System Utilization

Figure 1 shows the DAS-2 system utilization as function of time of day. There are two plots for every cluster. One stands for average utilization of all days and the other for average utilization of all active days in the year (excluding system down time and days without job arrivals). We can see that cluster fs0 are up and busy for most of the days in the year, while fs1 and fs3 have quite a lot of "empty" days. In average, fs0 has the highest (22%) and fs3 has the lowest system utilization (7.3%) among DAS-2 clusters. The utilization (7.3% to 22%) is substantially lower than previously reported workloads (e.g. 50% in average excluding downtime [5]). This is because DAS-2 system is mainly used for scientific research rather than production. Moreover, DAS-2 schedulers define one special policy, which forbids jobs to be scheduled on nodes (dual processor) of which one processor is already used by another job. This policy is defined for performance studies and it has certain negative impact on the overall system utilization.

We can see that the utilization approximately follows the daily job arrival rate (see Figure 2), although

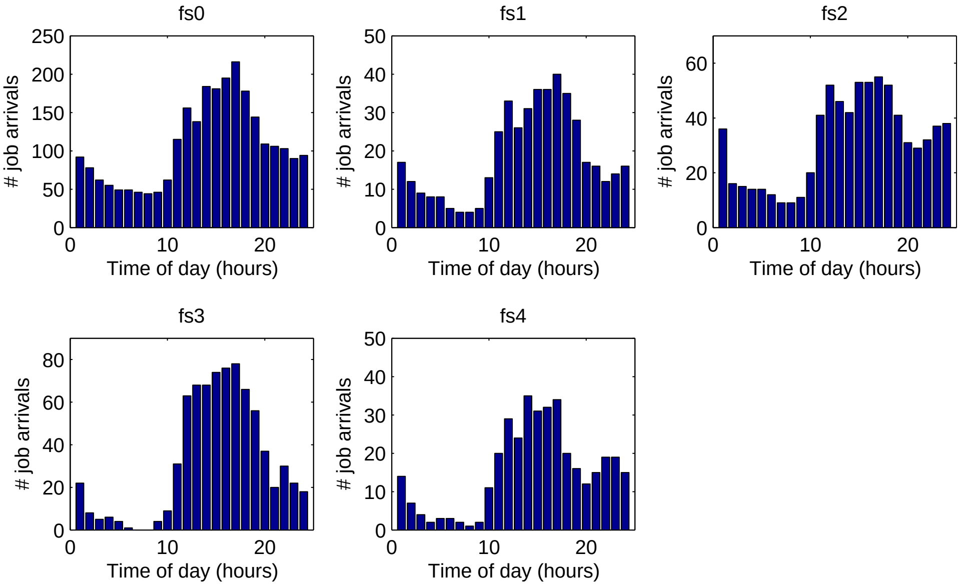

Figure 2: Daily cycle of job arrivals during weekdays.

the differences between day and night are generally smaller. It is because nightly jobs often require more processors and run longer than daily jobs, despite substantially fewer job arrivals. This is particularly evident on cluster fs3 and fs4.

4 Job arrival characteristics

In this section we analyze the job arrival characteristics. We first describe the job arrival rate, focusing mainly on daily cycles. Daily peak and non-peak hours are identified. Secondly, we characterize the job interarrival times during daily peak hours. Several statistical distributions are examined to fit the job interarrival times.

4.1 Job arrival rate

As is studied in [7], job arrivals are expected to have cycles at three levels: daily, weekly, and yearly. In a yearly cycle, we find that workloads are not distributed evenly throughout the year. Instead, workloads concentrate on specific months and job entries in these months are around two or more times above average. We call them "job-intensive" months (October, November and December on fs0, August, November on fs1, November, December on fs2, May, December on fs3, and August, November on fs4). The reason for this is that there are different active users/groups on different clusters and they are active in specific periods during the year. In a weekly cycle, all clusters share similar characteristics. Wednesday has the highest average job arrival rate and decreases alongside, with Sunday and Saturday have the lowest arrival rate. This is natural since people generally work more during weekdays (Monday - Friday) than weekends (Saturday and Sunday).

The most important cycle is the daily cycle. As is shown in Figure 2, clusters share similar daily workload distributions during weekdays. We identify the daily peak hours as from 9am to 7pm on all five clusters. This is in accordance with normal "working hours" at Dutch universities. Similar job arrival distributions are reported on other workloads with different peak hour periods (e.g. 8am to 6pm in [4], 8am to 7pm in [7]). Additionally, an intermediate period is reported from 6pm to 11pm in [4]. We observed similar characteristics on DAS-2 clusters, with an intermediate arrival period from 8pm to 0am and a low arrival period from 1am to 8am. The arrival rate per hour can be divided into three scales. The fs0 cluster has the highest one, with an average arrival rate of 108 jobs per hour and peak arrival rate exceeding 200 jobs per hour. In the middle there are fs2 and fs3, with average arrival rates of 31 and 32 jobs per hour each. Clusters fs1 and fs4 have average arrival rates of 19 and 15 jobs per hour, respectively.

4.2 Job interarrival time

Based on the observed job interarrival patterns, we choose to characterize "representative" and "high load" period of job interarrival times. The representative period is defined as the peak hours during weekdays in job-intensive months. The high load period is the peak hours of the most heavily loaded days in the year. As is shown in Table 2, during high load period the *mean* ranges from 14 to 62 seconds and the *coefficient of variation* (CV) varies from 1.3 to 3.0 on DAS-2 clusters. The mean and CV are considerably

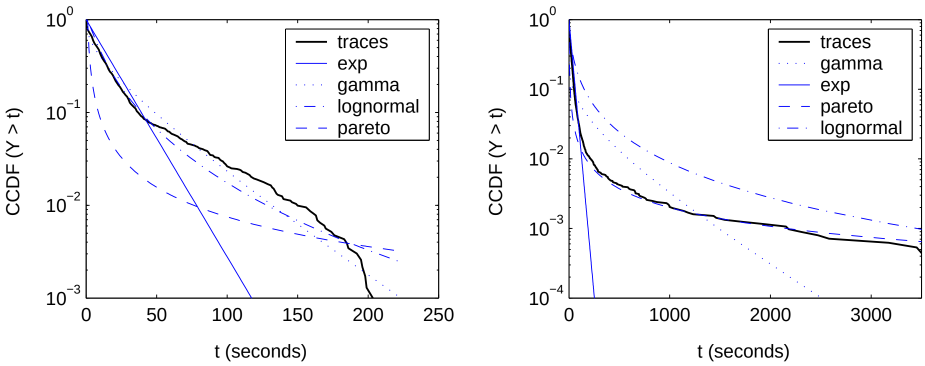

(a) High load distribution on fs0 at 12/02/2003 (mean = 17, CV = 1.6).

(b) Representative distribution on fs0 in December, 2003 (mean = 27, CV = 4.5).

Figure 3: Distributions of interarrival times during peak hours on fs0.

cluster	period	M	CV	fitted distribution
fs0	Dec 2	17	1.6	gamma (a=0.4,b=45)
fs1	Nov 25	26	2.4	logn (μ=0.4,σ=45)
fs2	Dec 29	14	1.3	gamma (a=0.6,b=23)
fs3	May 26	10	1.8	gamma (a=0.3,b=33)
fs4	Aug 13	62	3.0	logn (μ=3,σ=1.5)

Table 2: High load distribution of interarrival times during peak hours (M - mean (s), logn - lognormal).

cluster	period	M	CV	fitted distribution
fs0	Dec	27	4.5	Pareto (a=0.9)
fs1	Aug, Dec	66	3.6	logn (μ=2.9,σ=1.6)
fs2	Dec	44	5.0	Pareto (a=0.8)
fs3	May, Dec	23	6.0	Pareto (a=1.0)
fs4	Aug, Nov	86	5.1	Pareto (a=0.75)

Table 3: Representative distribution of interarrival times during peak hours (M - mean (s), logn - lognormal).

larger in the representative period (see Table 3). Both small (1-2) and large CVs (3-6) have been reported in other workloads [4, 6].

We have selected several statistical models to fit the interarrival times of representative and high load period, including exponential, gamma and heavy-tailed distributions like lognormal and Pareto [18]. We fit the above mentioned distributions by calculating their parameters from mean and CV given by the empirical CCDF (Complementary Cumulative Distribution Function) [19]. Matlab [20] is used to calculate the mean, CV and to fit the distributions.

Generally speaking, we find that the two periods (high load and representative) can be categorized into three groups based on their CVs: small (1-2), intermediate (2-4) and large ($>$ 4). Each group is best fitted by one kind of distributions. As is shown in Figure 3, for instance, the high load period on fs0 (CV=1.6, small) is best fitted by gamma distribution (a=0.38, b=45) and the representative period (CV=4.5, large) by Pareto (a=0.9). On fs1, which has intermediate CVs on both periods, lognormal fits the observed interarrival times better than other distributions. The fitted distributions during the two periods on all DAS-2 clusters are shown in Table 2 and 3. During the high load period when smaller CVs are observed, gamma and lognormal are the most suitable distributions to fit interarrival times. During the representative period with higher CVs, Pareto distribution gives the best fit when CV is larger than 4 while lognormal is more suitable for smaller CVs ($<$ 4).

5 Job execution characteristics

In this section we describe the job execution characteristics, which includes job sizes, job run times, and memory usage. For detailed analysis of correlations between these characteristics, parameters of fitted distributions we refer to [21].

5.1 Job sizes

Table 4 shows the job size characteristics of the DAS-2 clusters. The “power-of-two” phenomenon is clearly observed, as is found in many other workloads [4, 7, 9, 11]. However, the “power-of-two” sizes on cluster fs0, fs1, and fs2 are not as dominant as on fs3 and fs4. Some multiple-2 sizes also contribute to a significant portion of the total number of jobs (e.g. 6 and 14 processors on fs1, shown in Figure 4(a)). The fractions of serial (0.9-4.7%) jobs are considerably lower compared to previously reported workloads (30-40%), since the DAS-2 system is primarily used for parallel and distributed computing research.

As we all noticed in Table 4, job size of *two* processors is surprisingly popular on DAS-2 clusters and it is

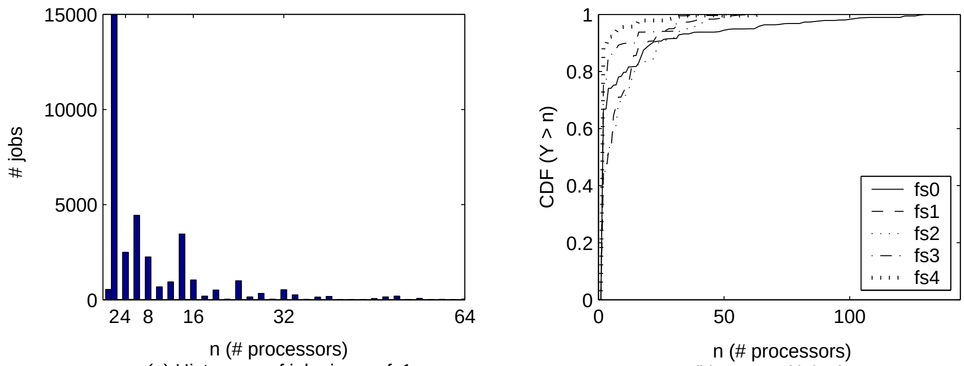

Figure 4: Distributions of job sizes on DAS-2 clusters.

cluster	serial(%)	two(%)	p-2(%)	others(%)
fs0	2.8	59.4	78.1	19.1
fs1	2.4	42.8	60.5	37.1
fs2	4.7	39.6	61.9	33.4
fs3	1.4	73.6	96.1	2.5
fs4	0.9	85.3	97.6	1.5
average	2.4	60	78.8	18.7

Table 4: Job sizes on DAS-2 clusters (“p-2” - power of two).

cluster	M	CV	Weibull parameters
fs0	374	5.3	a = 0.13, b = 0.45
fs1	648	7.9	a = 0.11, b = 0.45
fs2	531	16	a = 0.04, b = 0.65
fs3	466	12	a = 0.04, b = 0.75
fs4	2427	6.4	a = 0.04, b = 0.5
average	889	9.5	a = 0.07, b = 0.6

Table 5: Job actual runtimes on DAS-2 clusters.

chosen by a major fraction of jobs (range from 39.6% on fs2 to 85.3% on fs4). To find a proper explanation for this phenomenon, we investigate into the internal structure of the workloads. On fs0, for instance, there are ten very active users (out of 130 users in total). The most active user submitted more than 40,000 jobs (18% of the total number of jobs on fs0) in consecutive seven weeks during October and November 2003, which is his/her only active period throughout the year. All of these jobs have the same name and request two processors. For the second most active user on fs0, around 90% of his/her jobs have job sizes of two. On other DAS-2 clusters similar user behavior are observed, resulting in the popularity of job size two and power-of-two. We can see how large impact the major users have on the overall workload and therefore user behavior should be considered in workload modeling as an important structure [2, 11].

Plots of the job size distributions are shown in Figure 4(b). In [7], the best results for fitting job sizes are obtained by gamma and two-stage uniform distributions. For DAS-2 clusters (Figure 4 (b)), we find that gamma distribution provides the best fit for job sizes. Parameters of fitted gamma distributions are provided in [21].

5.2 Job Actual Runtimes

Job actual runtimes have been extensively studied in previous reported workloads. In addition, correlations between actual job runtime, job size, and requested job runtime are analyzed and established [4, 7]. We investigate these correlations in detail in [21]. Table 5 shows the characteristics of job actual runtimes on DAS-2 clusters. The average runtime is 889 seconds, which is considerably lower then previously reported workloads (e.g. 3479 seconds on SDSC SP2 [6]). However, the CV (5.3 - 16) is substantially higher than other production systems (2 - 5) [4, 5, 6]. This is in accordance with the scientific and experimental nature of the DAS-2 usage: the majority of jobs have small execution times and they vary a lot. Plots of the actual runtime distributions on a log scale are shown in Figure 5(a).

Different kinds of distributions have been examined to characterize the actual runtimes, for instance, log-uniform in [22], hypergamma in [7] and weibull in [4]. We find that weibull distribution provides the best fit for the actual runtimes on DAS-2 clusters (Figure 5 (b)). Parameters of fitted weibull distributions are listed in Table 5.

5.3 Memory Usage

The PBS [15] accounting logs record the maximum amount of physical memory used by the job. Hereafter we refer to memory usage as the maximum used physical memory. Memory usage per processor is defined as the maximum used memory divided by the number of processors requested.

Figure 6 (a) shows the distributions of memory us-

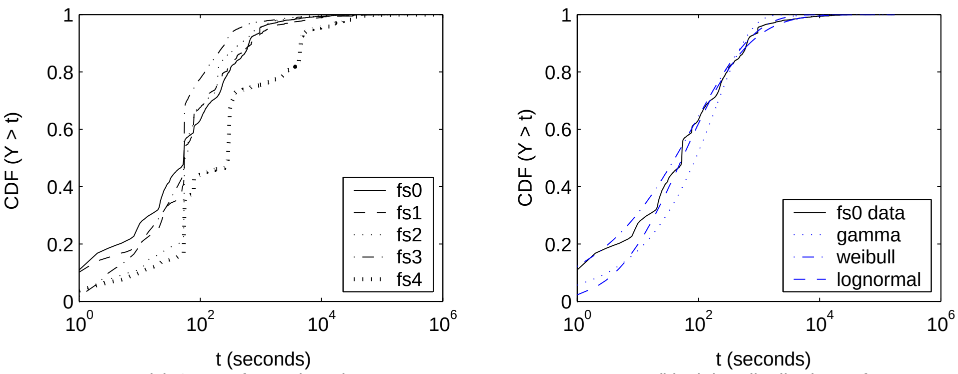

(a) CDFs of actual runtimes
(b) Fitting distributions of actual runtime on fs0

Figure 5: Distributions of job actual runtimes on DAS-2 clusters.

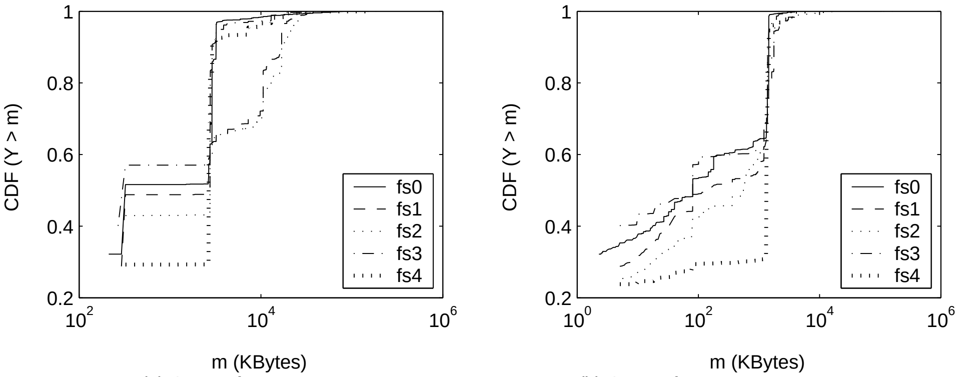

(a) CDFs of memory usage
(b) CDFs of memory usage per processor

Figure 6: Distributions of memory usage and memory usage per processor on DAS-2 clusters.

cluster	0KB (%)	324KB(%)	2.6-3MB(%)
fs0	32	19	34
fs1	29	20	16
fs2	25	18	21
fs3	40	17	34
fs4	24	6	62
Average	30	16	33

Table 6: Three special memory usage values and their corresponding job percentages.

age on DAS-2 clusters. It is clearly observed that three special values are chosen by a major fraction of jobs. These special values are 0KB, 324KB and 2600-3000KB (slightly different values in this range depending on the clusters), and their corresponding job percentages are listed in Table 6. We can see that a large fraction (30% in average) of jobs have very small memory usage[3]. 324KB and 2600-3000KB, on the other hand, contributes nearly one-sixth and one-third (in average) to the total number of jobs, respectively. The reason why memory usage concentrates on these special values might be that jobs typically have to load certain shared libraries (e.g. C, MPI, Globus), and these shared libraries normally require a fixed amount of memory. To verify this claim, we run MPI jobs (fractal computation) with different requested number of processors (4, 8, 16 and 32) on DAS-2 clusters. We found that memory usage for these jobs is almost the same (324KB, for job size 4, 8 and 16). The exception occurs for job size 32, of which memory usage jumps to 52,620KB. Other MPI programs also appears to use memory size of 324KB. Therefore, we might say that jobs which use 324KB memory most likely have to load certain libraries like MPI. Memory usage of 2600-3000KB could be other shared libraries or objects. Distributions of memory usage per processor on a log scale are shown in Figure 6 (b). In [21] correlations between memory usage and other characteristics are analyzed.

6 Job queue wait times

Queue wait times are interesting because they can be used to improve allocation strategies for malleable jobs on space-sharing parallel computers [22] or used

[3] 0KB is recorded in the PBS accounting logs. It means that the job uses very small memory (rounded to zero) instead of saying that the job does not use memory at all.

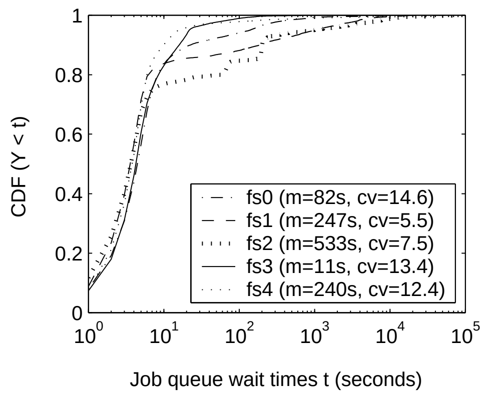

Figure 7: Job queue wait time distributions on DAS-2 clusters.

cluster	loguniform-2 parameters
fs0	l=-0.13, m=1.0, h=3.1, p=0.76
fs1	l=-0.10, m=0.7, h=4.0, p=1.00
fs2	l=-0.15, m=1.0, h=3.7, p=0.65
fs3	l=-0.14, m=1.1, h=2.2, p=0.73
fs4	l=-0.13, m=1.0, h=3.0, p=0.80

Table 7: Characteristics of job queue wait times on DAS-2 clusters.

by "superschedulers" (or resource brokers) to balance workload distribution in a Grid environment [23]. Tools and techniques have been developed to predict job queue wait times but not much characterization has been done in previous studied workloads. We analyze and model queue wait times, since a statistical model may help in improving prediction performance whereas other techniques (e.g. by simulation) are not likely to achieve [23].

The job queue wait time distributions on DAS-2 clusters are shown in Figure 7. We can see that a majority of jobs (80% in average) do not have to wait in the queue or have a short wait period (less than 10 seconds, the scheduler interval is 5 seconds). However, there are certain periods that users are very active and submit a lot of jobs in short intervals, which result in long queue wait times and contribute to the heavy tails in distributions. Generally speaking, two-stage log-uniform distributions fit the queue wait times very well on the DAS-2 clusters. Log-uniform parameters are provided in Table 7.

7 Conclusions and Future Work

In this paper, we present a comprehensive characterization of the DAS-2 workloads using the most recent traces by far (year 2003). We characterized system utilization, job arrival characteristics (arrival rate and interarrival time), job execution characteristics (job size, runtime, and memory usage), and job queue wait times. Differences of the DAS-2 workloads compared to previously reported workloads include the following:

1. A substantially lower average system utilization (from 7.3% to 22%) is observed.

2. A wide range of CVs (1-6) are obtained for the job interarrival times during peak hours, which can be well fitted using three different kinds of distributions based on the CVs.

3. Power-of-two phenomenon of job sizes is clearly observed, with an extreme popularity of job size *two*. The fraction of serial jobs (0.9%-4.7%) is much lower than other workloads (30%-40%).

4. A large portion of jobs has very small memory usage and several special values are used by a major fraction of jobs.

5. Small job queue wait times (80% of jobs wait less than 10 seconds) are observed.

To facilitate synthetic workload generation, we provide statistical distributions of the main characteristics. The distributions are summarized as follows:

1. Interarrival time — in high load period, gamma (CV 1-2) or lognormal (CV 2-4) are the most suitable distributions; in representative period, lognormal (CV 2-4) or Pareto (CV $>$ 4) give the best fit.

2. Job size — gamma gives the best fit.

3. Job actual runtime — Weibull is the best fitted distribution.

4. Queue wait time — two-stage log-uniform is the best fitted distribution.

In future work, we plan to generate workload models based on obtained results (an extended version of this paper is available in [21]). As we have partially investigated in Section 5, user behavior is very important in workload modeling since it reflects the internal structure of the workloads. We are interested in studying multi-class models [24] on DAS-2 and other cluster workloads which are more group/VO (Virtual Organization) oriented. Once good workload model(s) have been obtained, we can evaluate different scheduling strategies and tune policies on DAS-2 clusters. Moreover, a statistical model obtained from characterization can help improving performance when predicting job queue wait times. Another interesting topic in a multi-cluster environment is co-allocation. Currently multi-cluster job information is not logged on the DAS-2 clusters. We plan to instrument the Globus gatekeeper to collect the necessary traces and identify the key characteristics for multi-cluster jobs.

8 Acknowledgments

The DAS-2 supercomputer is funded by NWO (Netherlands Organization for Scientific Research) and the participating universities. We thank Dr. Dick Epema (Delft University of Technology) for his review and many valuable suggestions that improved the quality of this paper.

References

[1] M. Calzarossa and G. Serazzi. Workload characterization: A survey. *Proc. IEEE*, 81(8):1136–1150, 1993.

[2] D. G. Feitelson. Workload modeling for performance evaluation. *Lecture Notes in Computer Science*, 2459:114–141, 2002.

[3] Parallel Workload Archive. http://www.cs.huji.ac.il/labs/parallel/workload/.

[4] S.-H. Chiang and M. K. Vernon. Characteristics of a large shared memory production workload. *Lecture Notes in Computer Science*, 2221:159–187, 2001.

[5] D. Feitelson and B. Nitzberg. Job characteristics of a production parallel scientific workload on the NASA ames iPSC/860. *Job Scheduling Strategies for Parallel Processing – IPPS'95 Workshop*, volume 949, pages 337–360, Springer, 1995.

[6] K. Windisch, V. Lo, R. Moore, D. Feitelson, and B. Nitzberg. A comparison of workload traces from two production parallel machines. In *6th Symp. Frontiers Massively Parallel Comput.*, pages 319–326, 1996.

[7] U. Lublin and D. G. Feitelson. The workload on parallel supercomputers: modeling the characteristics of rigid jobs. *J. Parallel and Distributed Comput.*, 63(11):1105–1122, 2003.

[8] J. Jann, P. Pattnaik, H. Franke, F. Wang, J. Skovira, and J. Riodan. Modeling of workload in MPPs. In D. G. Feitelson and L. Rudolph, editors, *Job Scheduling Strategies for Parallel Processing*, pages 95–116. Springer Verlag, 1997.

[9] W. Cirne and F. Berman. A comprehensive model of the supercomputer workload. In *IEEE 4th Annual Workshop on Workload Characterization*, 2001.

[10] S. J. Chapin, W. Cirne, D. G. Feitelson, J. P. Jones, S. T. Leutenegger, U. Schwiegelshohn, W. Smith, and D. Talby. Benchmarks and standards for the evaluation of parallel job schedulers. In D. G. Feitelson and L. Rudolph, editors, *Job Scheduling Strategies for Parallel Processing*, pages 67–90. Springer-Verlag, 1999.

[11] A. B. Downey and D. G. Feitelson. The elusive goal of workload characterization. *Perf. Eval. Rev.*, 26(4):14–29, 1999.

[12] The DAS-2 Supercomputer. http://www.cs.vu.nl/das2.

[13] The DutchGrid community. http://www.dutchgrid.nl.

[14] S. Banen and A. I. D. Bucur and D. H. J. Epema. A Measurement-Based Simulation Study of Processor Co-Allocation in Multicluster Systems. In D. G. Feitelson and L. Rudolph, editors, *Job Scheduling Strategies for Parallel Processing*, pages 105–128. Springer-Verlag, 2003.

[15] Portable Batch System. http://www.openpbs.org.

[16] The Maui Scheduler. http://www.supercluster.org.

[17] The Globus project. http://www.globus.org.

[18] O. Allen. Probability, Statistics, and Queueing Theory with Computer Science Applications. Acdemic Press, 1978.

[19] J. Farber, S. Bodamer, and J. Charzinski. Evaluation of dial-up behaviour of internet users. citeseer.nj.nec.com/441426.html.

[20] Matlab. http://www.mathworks.com.

[21] H. Li, D. Groep, and L. Wolters. Workload Characteristics of a Multi-cluster Supercomputer (extended version). http://www.liacs.nl/˜hli.

[22] Allen B. Downey. Using Queue Time Predictions for Processor Allocation. In D. G. Feitelson and L. Rudolph, editors, *Job Scheduling Strategies for Parallel Processing*, pages 35–57. Springer-Verlag, 1997.

[23] H. Li, D. Groep, J. Templon, and L. Wolters. Predicting Job Start Times on Clusters. To appear in *proceedings of 4th IEEE/ACM International Symposium on Cluster Computing and the Grid*, Chicago, U.S.A. April, 2004.

[24] M. Calzarossa and G. Serazzi. Construction and use of multiclass workload models. *Performance Evaluation*, 19(4):341–352, 1994.