1 Outline

1.1 Supercomputer System

In SY 2014, the ISSP supercomputer center provided users with three supercomputing systems: NEC-SX9 (System A), SGI Altix ICE 8400EX (System B), and FUJITSU PRIMEHPC FX10 (System C) (Fig. 1). Systems A and B began service on July 1, 2010, and their operation has terminated on April 1, 2015. FUJITSU PRIMEHPC FX10 (System C) entered service at the beginning of SY 2013. All the systems were all installed in the main building of ISSP.

System A - NEC SX9 is a vector computer with 4 nodes (64 CPUs). Vectorization and parallelization between CPUs can automatically be done by the C/Fortran compilers. One node contains 1 TB of shared memory, and the total system achieves 6.5 TFlops theoretical peak performance. All the nodes are connected to a 13 TB storage system with high throughput.

System B - SGI Altix ICE 8400EX is a massively-parallel supercomputer with 1,920 nodes (3,840 CPUs / 15,360 cores) achieving 180.0 TFlops theoretical peak performance. Each node has 24 GB of memory (46 TB in total) and two Intel Xeon X5570 CPUs running at 2.93 GHz connected by dual QPI links (2×25.6 GB/sec). Up to 128 nodes are connected by enhanced hypercube 4×QDR InfiniBand networks with 40 GB/s bisection bandwidth. A 110 TB Lustre file system is connected to the entire system also with InfiniBand, realizing I/O throughput on the order of GB/sec.

System C - FUJITSU PRIMEHPC FX10 is highly compatible with K computer, the largest supercomputer in Japan. System C consists of 384 nodes, and each node has 1 SPARC64TM IXfx CPU (16 cores) and 32 GB of memory. The total system achieves 90.8 TFlops theoretical peak performance.

For further details, please contact ISSP Supercomputer Center (SCC-ISSP).

[Correspondence: center@issp.u-tokyo.ac.jp]

1.2 Project Proposals

The ISSP supercomputer system provides computation resources for scientists working on condensed matter sciences in Japan. All scientific staff members (including post-docs) at universities or public research institutes in Japan can submit proposals for projects related to research activities on materials and condensed matter sciences. These proposals are peer-reviewed by the Advisory Committee members (see Sec. 1.3), and then the computation resources are allocated based on the review reports. The leader of an approved project can set up user accounts for collaborators. Other types of scientists, including graduate students, may also be added. Proposal submissions, peer-review processes, and user registration are all managed via a web system.

The computation resources are distributed in a unit called "point", determined as a function of available CPU utilization time and consumed disk resources. There

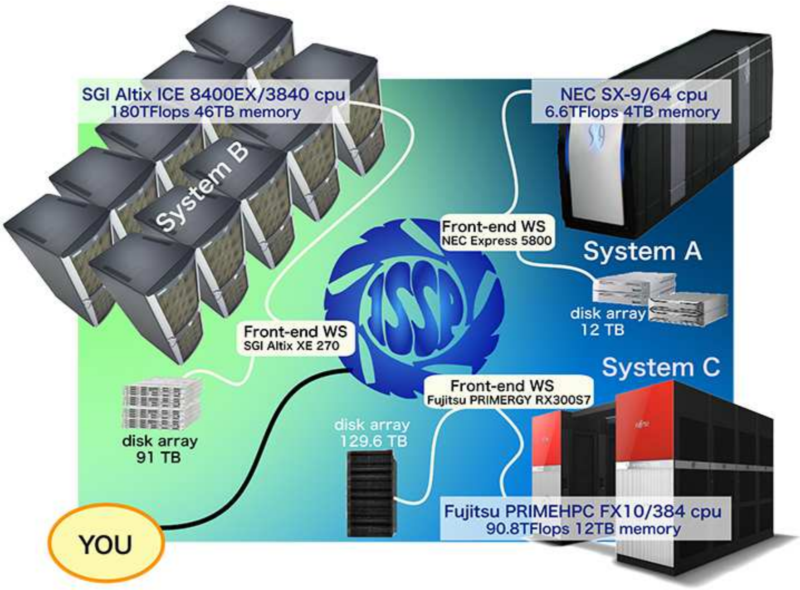

Figure 1: Supercomputer System at the SCC-ISSP

were six classes of research projects in SY 2014. The number of projects and the total number of points that were applied for and approved in this school year are listed in Table 1.

In addition, from SY 2010, ISSP Supercomputer is providing 20% of its computational resources for Computational Materials Science Initiative (CMSI), which aims at advancing parallel computations in condensed matter, molecular, and materials sciences on the 10-PFlops K Computer. The points for projects run by CMSI are distributed in accord with this policy. Computer time has also been alloted to Computational Materials Design (CMD) workshops run by CMSI.

- Proposals for projects in Classes B (small), C (mid-size), E (large-scale), and S (exceptional) can be submitted twice a year. Approved projects in Classes A, B, C, E, and S continue to the end of the school year.

- In Class D, projects can be proposed on rapidly-developing studies that need to perform urgent and relatively large calculations. An approved project continues for 6 months from its approval.

- Class S is for projects that are considered extremely important for the field of condensed matter physics and requires extremely large-scale computation. The project may be carried out either by one research group or cooperatively by several investigators at different institutions. A project of this class should be applied with at least 10,000 points; there is no maximum. We require group leaders applying for Class S to give a presentation on the proposal to the Steering Committee of the SCC-ISSP.

- Project leaders can apply for points so that the points for each system do not exceed the maximum point shown in this table.

Table 1: Classes of research projects in SY 2014

Class	Max. Point			Application
	Sys–A	Sys–B	Sys–C	
A	100	100	100	any time
B	2k	1k	500	twice a year
C	20k	10k	2.5k	twice a year
D	20k	10k	2.5k	any time
E	–	30k	2.5k	twice a year
S	(Sys-A+B)>10k		–	twice a year

Class	# of Proj.	Total points					
		Applied			Approved		
		Sys–A	Sys–B	Sys–C	Sys–A	Sys–B	Sys–C
A	9	500	600	400	500	600	400
B	53	40.8k	44.8k	6.3k	33.8k	29.4k	5.7k
C	151	637.0k	1210.2k	141.2k	429.0k	360.5k	111.4k
D	10	12.0k	71.7 k	0	12.0k	57.6k	0
E	23	–	621.0k	42.0k	–	270.0k	37.4k
S	1	0	60.0k	0	0	25.0k	0
CMSI	18	–	–	–	–	–	140.0k

1.3 Committees

In order to fairly manage the projects and to smoothly determine the system operation policies, the Materials Design and Characterization Laboratory (MDCL) of the ISSP has organized the Steering Committee of the MDCL and the Steering Committee of the SCC-ISSP, under which the Supercomputer Project Advisory Committee (SPAC) is formed to review proposals. The members of the committees in SY 2014 were as follows:

Steering Committee of the MDCL

HIROI, Zenji	ISSP (Chair person)
KATO, Takeo	ISSP
KAWASHIMA, Naoki	ISSP
MORI, Hatsumi	ISSP
NAKATSUJI, Satoru	ISSP
NOGUCHI, Hiroshi	ISSP
SUGINO, Osamu	ISSP
SUEMOTO, Toru	ISSP
TSUNEYUKI, Shinji	Univ. of Tokyo
KIMURA, Kaoru	Univ. of Tokyo
MIYASAKA, Hitoshi	Tohoku Univ.

HASEGAWA, Tadashi	Nagoya Univ.
OKAMOTO, Yuko	Nagoya Univ.
OTSUKI, Tomi	Sophia Univ.
OGUCHI, Tamio	Osaka Univ.
NOHARA, Minoru	Okayama Univ.

Steering Committee of the SCC-ISSP

NOGUCHI, Hiroshi	ISSP (Chair person)
KAWASHIMA, Naoki	ISSP
SUGINO, Osamu	ISSP
TAKADA, Yasutami	ISSP
HARADA, Yoshihisa	ISSP
TSUNETSUGU, Hirokazu	ISSP
SHIBA, Hayato	ISSP
WATANABE, Hiroshi	ISSP
KASAMATSU, Shusuke	ISSP
MORITA, Satoshi	ISSP
HATANO, Naomichi	Univ. of Tokyo
IMADA, Masatoshi	Univ. of Tokyo
NAKAJIMA, Kengo	Univ. of Tokyo
TSUNEYUKI, Shinji	Univ. of Tokyo
MOHRI, Tetsuo	Tohoku Univ.
OTSUKI, Tomi	Sophia Univ.
ODA, Tatsuki	Kanazawa Univ.
OKAMOTO, Yuko	Nagoya Univ.
MORIKAWA, Yoshitada	Osaka Univ.
SUZUKI, Takafumi	Univ. of Hyogo
YOSHIMOTO, Yoshihide	Tottori Univ.
YATA, Hiroyuki	ISSP
FUKUDA, Takaki	ISSP

Supercomputer Project Advisory Committee

NOGUCHI, Hiroshi	ISSP (Chair person)
KAWASHIMA, Naoki	ISSP
SUGINO, Osamu	ISSP
TAKADA, Yasutami	ISSP
HARADA, Yoshihisa	ISSP
TSUNETSUGU, Hirokazu	ISSP
SHIBA, Hayato	ISSP
WATANABE, Hiroshi	ISSP
KASAMATSU, Shusuke	ISSP
MORITA, Satoshi	ISSP
AOKI, Hideo	Univ. of Tokyo

HATANO, Naomichi	Univ. of Tokyo
HUKUSHIMA, Koji	Univ. of Tokyo
IKUHARA, Yuichi	Univ. of Tokyo
IMADA, Masatoshi	Univ. of Tokyo
IWATA, Jun-Ichi	Univ. of Tokyo
MIYASHITA, Seiji	Univ. of Tokyo
MOTOME, Yukitoshi	Univ. of Tokyo
NAKAJIMA, Kengo	Univ. of Tokyo
OGATA, Masao	Univ. of Tokyo
OSHIYAMA, Atsushi	Univ. of Tokyo
TSUNEYUKI, Shinji	Univ. of Tokyo
WATANABE, Satoshi	Univ. of Tokyo
NEMOTO, Koji	Hokkaido Univ.
YAKUBO, Kosuke	Hokkaido Univ.
AKAGI, Kazuto	Tohoku Univ.
KAWAKATSU, Toshihiro	Tohoku Univ.
KURAMOTO, Yoshio	Tohoku Univ.
MOHRI, Tetsuo	Tohoku Univ.
SHIBATA, Naokazu	Tohoku Univ.
YANASE, Yoichi	Niigata Univ.
ARITA, Ryotaro	RIKEN
ISHIBASHI, Shoji	AIST
MIYAMOTO, Yoshiyuki	AIST
OTANI, Minoru	AIST
KOBAYASHI, Kazuaki	NIMS
TATEYAMA, Yoshitaka	NIMS
HATSUGAI, Yasuhiro	Univ. of Tsukuba
KOBAYASHI, Nobuhiko	Univ. of Tsukuba
OKADA, Susumu	Univ. of Tsukuba
YABANA, Kazuhiro	Univ. of Tsukuba
HIDA, Kazuo	Saitama Univ.
TOMITA, Yusuke	Shibaura Inst. Tech.
NAKAYAMA, Takashi	Chiba Univ.
FURUKAWA, Nobuo	Aoyama Gakuin Univ.
MATSUKAWA, Hiroshi	Aoyama Gakuin Univ.
TAKANO, Hiroshi	Keio Univ.
YAMAUCHI, Jun	Keio Univ.
YASUOKA, Kenji	Keio Univ.
OTSUKI, Tomi	Sophia Univ.
OBATA, Shuji	Tokyo Denki Univ.
ANDO, Tsuneya	Tokyo Inst. Technology
HOTTA, Takashi	Tokyo Metropolitan Univ.
OKABE, Yutaka	Tokyo Metropolitan Univ.
TOHYAMA, Takami	Tokyo Univ. of Sci.
WATANABE, Kazuyuki	Tokyo Univ. of Sci.

HAGITA, Katsumi	National Defense Academy
INOUE, Junichiro	Nagoya Univ.
KONTANI, Hiroshi	Nagoya Univ.
OKAMOTO, Yuko	Nagoya Univ.
SHIRAISHI, Kenji	Nagoya Univ.
TANAKA, Yukio	Nagoya Univ.
ODA, Tatsuki	Kanazawa Univ.
SAITO, Mineo	Kanazawa Univ.
ARAKI, Takeaki	Kyoto Univ.
KAWAKAMI, Norio	Kyoto Univ.
MASUBUCHI, Yuichi	Kyoto Univ.
YAMAMOTO, Ryoichi	Kyoto Univ.
KASAI, Hideaki	Osaka Univ.
KAWAMURA, Hikaru	Osaka Univ.
KUROKI, Kazuhiko	Osaka Univ.
KUSAKABE, Koichi	Osaka Univ.
MORIKAWA, Yoshitada	Osaka Univ.
OGUCHI, Tamio	Osaka Univ.
SHIRAI, Koun	Osaka Univ.
YOSHIDA, Hiroshi	Osaka Univ.
YUKAWA, Satoshi	Osaka Univ.
HARIMA, Hisatomo	Kobe Univ.
SUGA, Seiichiro	Univ. of Hyogo
SUZUKI, Takafumi	Univ. of Hyogo
TATENO, Masaru	Univ. of Hyogo
SAKAI, Toru	Japan Atomic Energy Agency
HOSHINO, Kozo	Hiroshima Univ.
HOSHI, Takeo	Tottori Univ.
YOSHIMOTO, Yoshihide	Tottori Univ.
YASUDA, Chitoshi	Univ. of the Ryukyus
OZAKI, Taisuke	ISSP
KATO, Takeo	ISSP
TADA, Tomofumi	Tokyo Inst. Technology
TODO, Synge	Univ. of Tokyo

1.4 Staff

The following staff members of the SCC-ISSP usually administrate the ISSP Supercomputer.

NOGUCHI, Hiroshi	Associate Professor (Chair person)
KAWASHIMA, Naoki	Professor
SUGINO, Osamu	Associate Professor
WATANABE, Hiroshi	Research Associate
KASAMATSU, Shusuke	Research Associate

NOGUCHI, Yoshifumi	Research Associate
SHIBA, Hayato	Research Associate
MORITA, Satoshi	Research Associate
YATA, Hiroyuki	Technical Associate
FUKUDA, Takaki	Technical Associate
ARAKI, Shigeyuki	Technical Associate

2 Statistics (School Year 2014)

2.1 System and User Statistics

In the following, we present statistics for operation time taken in the period from April 2014 to March 2015 (SY 2014). In Table 2, we show general statistics of the supercomputer system in SY 2014. The total number of CPUs in System A, B, and C is 64, 3840, and 384 respectively. Consumed disk points amount to about 4%, 5%, and 1% of the total consumed points in System A, B, and C respectively.

In the left column of Fig. 2, availabilities, utilization rates, and consumed points in each system are plotted for each month. Throughout the school year, the utilization rates were high enough. Especially in System B, they were exceeding 90% throughout most of the year. In System C, roughly half of the total utilized resources were used by CMSI projects. This amounts to about 20% of the total usage of the computational resources in this school year. The user statistics are shown in the right column of Fig. 2. The horizontal axis shows the rank of the user/group arranged in the descending order of the execution time (hour×CPU). The execution time of the user/group of the first rank is the longest. The vertical axis shows the sum of the execution time up to the rank. From the saturation points of the graphs, the number of "active" users of each system is around 50, 250, and 70 for System A, B, and C respectively. The maximum ranks in the graphs correspond to the number of the users/groups that submitted at least one job.

2.2 Queue and Job Statistics

Queue structures of System A, B, and C in SY 2014 are shown in Table 3. In each system, the queues are classified by the number of CPUs the user can use and the maximum duration of each submitted job. In System A, in addition to the usual P class jobs, there is a queue "D1" for debugging, and "L1" for jobs which require only one CPU but quite a long time. Parallel jobs are executed with "P4" and "P16", 16 CPUs being available at maximum with one job using "P16".

In System B, a highly detailed classification is adopted. The biggest portion (20 racks out of 30 in total) of the resources is allotted for "F256", which mainly uses 128 or 256 CPUs at once. "F16", "F32", and "F64" are for smaller-scale jobs using 16, 32, and 64 CPUs respectively. The elapsed-time limit of the above queues is 24 hours for one job, while it is set smaller for smaller-scale queues ("F4"

Table 2: Overall statistics of SY 2014

	System-A	System-B	System-C
total service time (k hour × CPU)	544.0	32441.1	3264.4
number of executed jobs	24884	148290	25020
total consumed points (k point)	139.2	499.7	105.4
CPU points (k point)	134.9	471.1	99.9
disk points (k point)	4.3	28.6	5.5
points consumed by CMSI (k point)	–	–	62.4
total exec. time (k hour × CPU)	444.27	29096.3	2323.8
availability	97.2%	96.4%	96.8%
utilization rate	81.7%	89.5%	71.2%

and “F8”) to speed up their rotation. For time-demanding jobs, L-type queues are also introduced, whose time limit is set longer than F-type queues. “P64” queue is set up to accept jobs which require any number of CPUs more than 1 and not exceeding 64. “i32” is a queue for debugging, which corresponds to interactive mode in the previous system. In “i32”, users can execute their jobs using up to 16 nodes at once from the command line, as if they were logging into the calculation node.

In System C, the “F” and “L” queues are set up similarly to System B. In addition, a debug queue is set up for short debugging jobs utilizing 1 to 4 CPUs, and an interactive queue that can use 1 to 4 CPUs is also available.

The CPU points are set smaller for larger-scale queues for System B as shown in Table 3, while it is more uniform in System A. To prevent overuse of the storage, points are charged also for usage of disk quota in the three systems, as shown in Table 4. Disk points are revised often for optimal usage of the resources by examining usage tendencies each year.

Although we do not mention here in detail, to promote utilization of the massively parallel supercomputer, background queues (“B16”, “B32”, “B64”, and “B256”), which charge no CPU points for the jobs, have also been open in System B.

The ISSP Supercomputer also supports large-scale jobs, which use tens of thousands of cores at once by exclusively using the necessary number of CPUs. In-advance application is necessary to execute this type of job. Large-scale jobs can be executed in queues “P512”, “P1024”, “P2048”, and “P3840” just after the scheduled monthly maintenance. However, since such large-scale jobs are now covered by the K Computer, no jobs were executed in these queues since SY2013.

The number of jobs, average waiting time, and total execution time in each queue are shown in Table 5. In System A, the average waiting times of P4 and P16 are a bit long compared with the elapsed-time limit (24 hours). This is because a few active users tend to submit many jobs at once. Because fair-share scheduling is adopted, the waiting time is considered to be appropriate for fair distribution

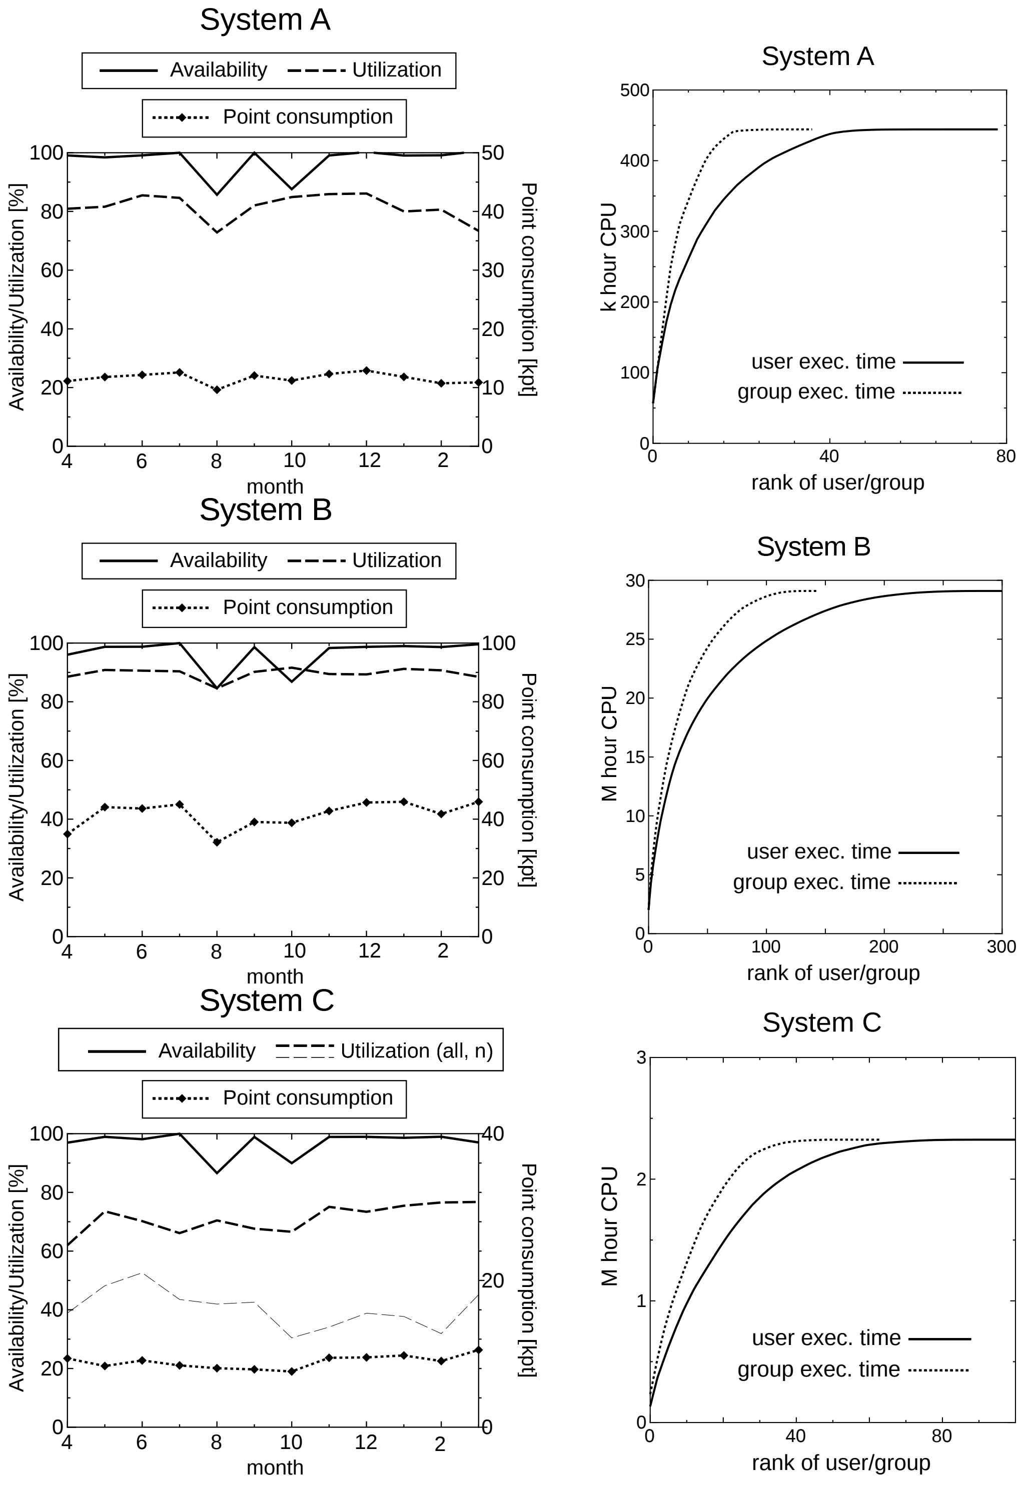

Figure 2: Left: Availabilities, utilization rates and point consumptions of each month during SY 2014. For System C, the utilization by CMSI projects (denoted by “n”) is plotted in addition to the total utilization. Right: User statistics. The horizontal axis shows the rank of the user/group arranged in the descending order of the execution time (hour×CPU). The vertical axis shows the sum of the execution time up to the rank.

Table 3: Queue structures in SY 2014

System–A

queue name	Elapsed time limit (min)	# of CPU /Job (n)	# of CPU /queue (p)	memory size	CPU points / (CPU·day)
D1	15	1	2	60GB	7.776
L1	7200	1	4	60GB	7.776
P1	1440	1	10-30	60GB	7.776
P4	1440	4	16-32	240GB	7.776
P16	1440	16	16	960GB	6.048

System–B

queue name	Elapsed time limit (min)	# of CPU /Job (n)	# of CPU /queue (p)	CPU points /(CPU·day)
P1	720	1	32	0.690
P64	720	2-64	64	0.518
F4	720	4	96	0.518
F8	720	8	96	0.518
F16	1440	16	1024	0.518
F32	1440	32	1024	0.518
F64	1440	64	1024	0.518
L16	7200	16	64	0.518
L32	7200	32	64	0.518
L64	7200	64	64	0.518
i32	20	1-32	64	0.518
F256	1440	65-256	2560	0.358
L256	7200	65-256	512	0.358
P512	–	128-512	512 or 1024	0.358
P1024	7200	384-1024	3072	0.358
P2048	–	128-2048	2048	0.358
P3840	1440	1024-3840	3840	0.358

* The available memory size is limited to 21 GB per one node.
* P queues require in-advance application (see main text). The elapsed-time limit for P512 and P2048 queues is determined on a per-application basis.

System–C

queue name	Elapsed time limit (min)	# of CPU /Job (n)	# of CPU /queue (p)	CPU points /(CPU·day)
debug	30	1-4	24	1
interactive	30	1-4	24	1
F12	1440	2-12	60	1
F96	1440	2-12	288	1
L12	7200	24-96	24	1
L96	7200	24-96	192	1

* The available memory size is limited to 28 GB per one CPU.

Table 4: Disk points of System A, B, and C

		point/day
System A	`/home`	$0.0125 \times \theta(q - 10)$
	`/work`	$0.005 \times \theta(q - 50)$
System B	`/home`	$0.05 \times \theta(q - 10)$
	`/work`	$0.005 \times \theta(q - 100)$
System C	`/home`	$0.05 \times \theta(q - 10)$
	`/work`	$0.005 \times \theta(q - 100)$

* q is denoted in unit of GB.
* $\theta(x)$ is equal to the Heaviside step function $H(x)$ multiplied by x, i.e., $xH(x)$.

of computational resources. We will continue to look for more appropriate queue settings also in the next school year to meet the user's tendency of resource usage.

In System B, a large portion of jobs have been executed in queues "F16", "F32", "F64", and "F256". As we intended, most of the execution time has been consumed in "F256" and "L256". In all of these queues, the queue settings meet the user's tendencies in that the waiting times are on the order of the elapsed-time limit.

In System C, the waiting times for the "F" queue jobs are less than twelve hours. The "L96" queue has a waiting time of nearly five days, owing to the large amount of resources the jobs occupy when run in this queue.

Acknowledgments

The staffs would like to thank Prof. Takafumi Suzuki (now at University of Hyogo) for developing WWW-based system (SCM: SuperComputer Management System) for management of project proposals, peer-review reports by the SPAC committee, and user accounts. We also thank Ms. Reiko Iwafune for creating and maintaining a new WWW page of the ISSP Supercomputer Center.

Table 5: Number of jobs, average waiting time, total execution time, and average number of used CPU's per job in each queue.

System-A

queue	# of Jobs	Waiting Time (hour)	Exec. Time (k CPU × hour)	# of CPU
D1	3886	0.01	0.20	1.00
P1	15337	14.70	122.08	1.00
L1	399	27.84	11.69	1.00
P4	4401	40.74	187.35	3.83
P16	830	108.63	122.92	15.25

System-B

queue	# of Jobs	Waiting Time (hour)	Exec. Time (k CPU × hour)	# of CPU
P1	31395	23.88	126.39	1.0
P64	22491	25.51	325.88	3.3
F4	18724	20.31	289.28	4.0
F8	9508	9.62	355.97	8.0
F16	28654	15.88	2929.68	16.0
F32	8317	16.13	1927.59	32.0
F64	4076	39.53	2503.65	64.0
L16	457	45.15	189.65	16.0
L32	46	115.47	37.48	32.0
L64	26	258.68	96.05	64.0
i32	8216	0.02	20.32	16.8
F256	8029	21.54	17437.37	193.3
L256	178	16.07	1464.29	159.0
P512	0	0	0.00	0
P1024	0	0	0.00	0
P2048	0	0	0.00	0
P3840	0	0	0.00	0

System-C

queue	# of Jobs	Waiting Time (hour)	Exec. Time (k CPU × hour)	# of CPU
F12	11890	9.27	385.52	5.2
L12	127	27.60	6.65	1.9
F96	4929	26.46	1904.18	36.2
L96	9	63.51	19.96	64.8
debug	5725	0.11	1.37	1.8
interactive	1942	0.00	0.29	1.1

Erratum

In the previous Activity Report 2013, on page 8, there is an error in the bottom right part of Fig. 2 (User statistics for System C), which should be replaced by the Fig. 3 shown below. The number of active users was about 80 in SY 2013 for System C.

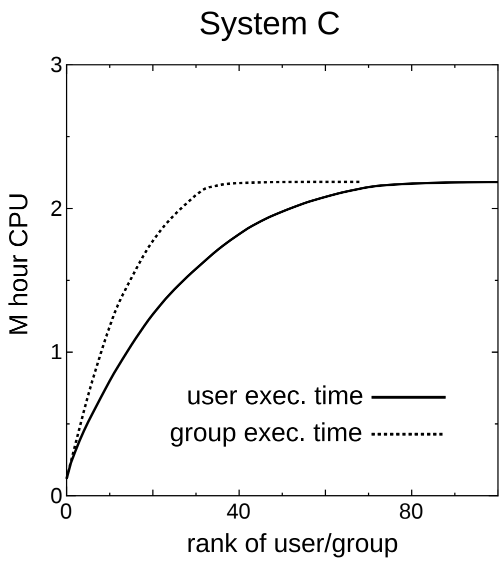

Figure 3: User statictics in SY2013 for System C.