Russian Academy of Sciences
Program Systems Institute

Research Center for Multiprocessor Systems

СКИФ

The SKIF K-1000 Supercomputer was designed and produced within the framework of the Russian Federation and Republic of Belarus Union Program

The SKIF K-1000 Supercomputer is a flagship high-performance computer built within the framework of the Program. Most technical solutions employed in the SKIF K-1000 Supercomputer are currently cutting-edge not only for the SKIF Program but also for the entire supercomputer industry; specifically:

- 64-bit AMD Opteron TM 248 (2200 MHz);
- the system network is based on the Infiniband 4x technology for interconnect.

Infiniband ensures high performance of MPI applications: bandwidth of MPI reaches 830 MB/s, latency amounts to 5 usec.

The SKIF K-1000 Supercomputer was built in record time. Assembly and setup of SKIF K-1000 were accomplished in the T-Platforms Cluster Technology Center, Moscow; then, moved and re-installed in UIIP NAS, Minsk, Belarus.

In November 2004, the SKIF K-1000 Supercomputer was ranked 98th among the world's TOP500 Supercomputers with the 2.5 Tflops peak performance. Moreover, SKIF K-1000 has one of the best price/quality ratios in the industry.

SKIF K-1000 is also ranked second-fastest — was first-fastest in November 2004 — among the 50 most powerful computer installations in the territory of the former Soviet Union and Eastern Europe.

The SKIF K-1000 Supercomputer: Structure

Computational Subsystem

Computational node (8x36=288 pcs.): 1U, 2×AMD Opteron 248 (2.2Ghz), RAM: 4GB, HDD: IDE 80GB, IB 4x Mellanox HCA MHXL-CF128 (connected to the Leaf IB Switch), 2xGbEthernet (one line connected to the Leaf Eth Switch), СКИФ-Servnet v.2

Subcontrol (1 pcs.): 2U, CPU: 2×AMD Opteron 248 (2.2Ghz), RAM: 4GB, HDD: 2× 36GB, SCSI 10K RPM, HotSwap, DVD/CD, FDD, Moxa CP104UL 4 port RS232 LP Universal card + 4 SKIF Servnet v.2 cards, 2×GbEthernet (one line connected to one of the Core Eth Switch)

The Infiniband System Network4x (MPI: ~830 MB/s, ~5 μs)

the Core IB Switch Mellanox MTS-2400 (2x6=12 pcs.): 1U, 24 ports, one line for each of the 24 Leaf IB Switches

the Leaf Switch Mellanox MTS-2400 (3x8=24 pcs.): 1U, 24 ports, 12 lines for the computational nodes of 12 lines of the Core Switches (one line for each of the nodes)

GbEthernet Managed Network

the Core Eth Switch D-Link DGS-3324SR (2 pcs.): 1U, 24 ports, 2 ports (2 Gbps trunk) per 8 Leaf Eth Switches, 40 Gbps trunk between Core Eth Switches

the Leaf Eth Switch D-Link DGS-1224T (2x8=16 pcs.): 1U (2 pcs. Are installed in 1U), 24 ports, 18 lines for computational nodes, 2 ports (2 Gbps trunk) per Core Eth Switch

The SKIF-Servnet v.2 Service Network

The Servnet boards of computational nodes are linked into the RS-485 line (two cabinets: 2x36=72 pcs. in a line) and connected to one of the four Servnet service cards in a computational node.

The SKIF K-1000 Supercomputer Data Sheet

Peak performance	**2.5 Tflops (2 534 Gflops)**
Linpack performance	**2.0 Tflops (2 032 Gflops)**
Characteristics Linpack	**Performance index=80.1%; N = 274 000; $N_{1/2}$ = 24 950**
Computational nodes/processors	**288 / 576**
CPU	**AMD Opteron™ 248 (2200 Mhz)**
RAM	**288 × (8 × 0.5 GB) = 1 152 GB**
Disk memory	**288 × 80 GB = 23 040 GB**
System network	**InfiniBand, IB 4x (MPI: ~830 MB/s, ~5 μs)**
Management network	**Gigabit Ethernet**
Service network	**СКИФ-ServNet v.2.0**
Form factor	**1U**
Racks	**8 стоек 42U**
Dimension, mass	**4.8m × 1.07m × 2.07m, 10.7 m^3, 5.4 T**
Power consumption per one rack/ the installation	**9 – 11 kW / 73 – 89 kW**
Price/peak performance	**about $700 000 per 1 Tflops**
Price/Linpack performance	**about $800 000 per 1 Tflops**
Heat emission	**about 70 kW**
Air cooling	**16 000 m^3 per hour**
Gigabit Ethernet cable system	**326 cables, about 1,500 m 652 slots**
InfiniBand cable system	**576 cables, about 1,500 m 1152 slots**
Development Stages	
- concept and design	**12/01/2003–05/31/2004 (6 months)**
- manufacturing	**07/15–10/01/2004 (2.5 months)**
- assembly	**09/04–09/17/2004 (2 weeks)**
- testing	**09/18–09/25/2004 (1 week)**

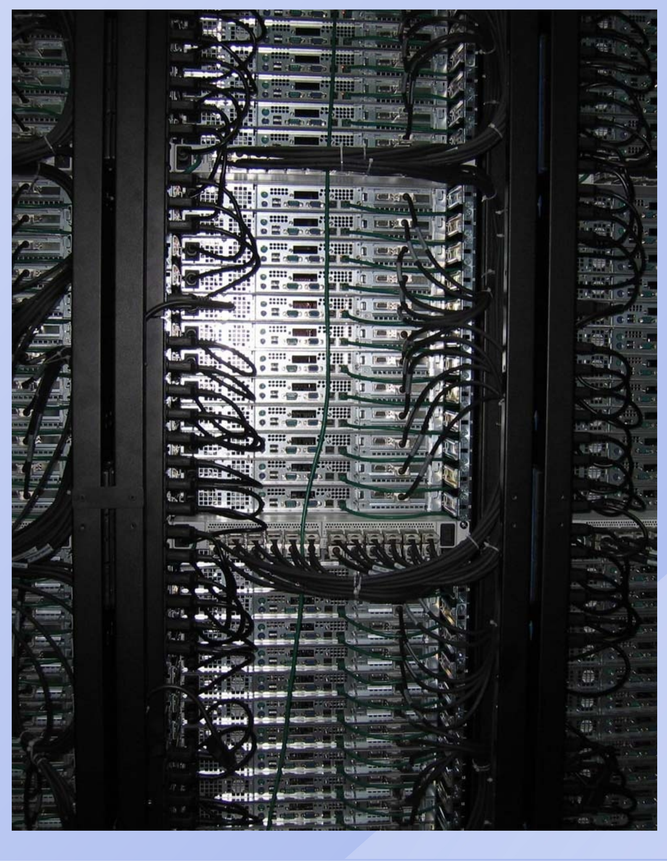

The SKIF K-1000 Supercomputer Developers

UIIP NAS, Belarus
6, Sourganov St., Minsk, 220012

Program Research Manager from Belarus:
Sergei V. Ablameyko, Doctor of Science
tel.: +375-17-284-21-75, fax: +375-17-231-84-03
e-mail: abl@newman.bas-net.by
Program Vice Research Manager from Belarus:
Vladimir V. Anischenko, Ph.D.
tel.: +375-17-284-09-85, fax: +375-17-284-03-20
e-mail: anishch@newman.bas-net.by
Program Chief Executive from Belarus:
Nicolai N. Paramonov, Ph.D.
tel.: +375-17-284-20-91
nick@newman.bas-net.by

The T-Platforms Company
113/1, Leninsky Ave.,
office E-505, Moscow, 117198

Director General:
Vsevolod Yu. Opanasenko
tel.: +7-095-956-54-90, fax: +7-095-956-54-15
info@t-platforms.ru

Cluster Technologies Center:
8, Vvedensky St., Moscow
tel.: +7-095-744-09-95, +7-095-744-09-33

Computer Research Institute
155, Bogdanovich St., Minsk, 220040

Development Manager:
Dmitry B. Javoronkov
tel.: +375-17-234-98-25, fax: +375-17-234-46-86
dmitry@ntc.niievm.minsk.by

Program Systems Institute of RAS
Pereslavl-Zalessky, 152020
Program Research Manager and Chief Executive from Russia:
Sergei M. Abramov, Doctor of Science
tel./fax: +7-08535-98-064
e-mail: abram@botik.ru

ADDRESS

Research Center for
Multiprocessor Systems
Program Systems Institute
Russian Academy of Sciences

Pereslavl-Zalessky
Yaroslavl Region
Russia, 152020
Tel/Fax: **+7 (08535) 98064**
E-mail: abram@botik.ru
Web-site: http://www.botik.ru/PSI/RCMS

RCMS ИЦМС