Supercomputers in Nagoya University

KATSUYA ISHII

The Information Technology Center (ITC) of Nagoya University was originally established as the Computing Center of Nagoya University, which all researchers of academic organization in Japan can used.

The first supercomputer of Nagoya University was Fujitsu supercomputer VP100(250Mflops) installed in the Institute of Plasma Physics in January 1984, which was the first operated supercomputer in Japan.

High Performance Computing Infrasturacture of Japan

http://www.hpci-office.jp/folders/english

HPCI

What is HPCI?

HPCI is a innovative common computational infrastructure responding to diversified needs of the users. It consists of K computer and other supercomputers installed in universities and research institutions in Japan connected by high speed back-bone networks.

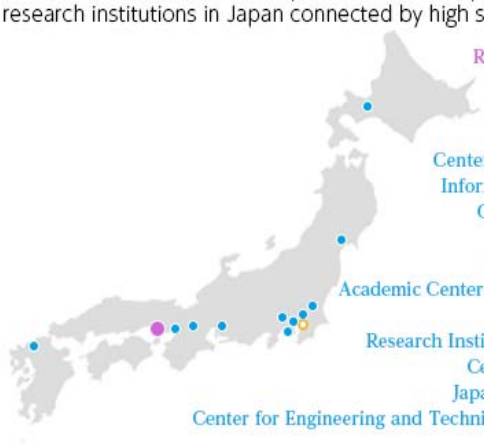

RIKEN Advanced Institute for Computational Science
National Institute of Informatics
Information Initiative Center, Hokkaido University
Cyberscience Center, Tohoku University
Center for Computational Sciences, University of Tsukuba
Information Technology Center, The University of Tokyo
Global Scientific Information and Computing Center, Tokyo Institute of Technology
Information Technology Center, Nagoya University
Academic Center for Computing and Media Studies, Kyoto University
Cybermedia Center, Osaka University
Research Institute for Information Technology, Kyushu University
Center for Earth Information Science and Technology, Japan Agency for Marine-Earth Science and Technology
Center for Engineering and Technical Support, The Institute of Statistical Mathematics

Example

1. The **K** computer（named for the Japanese word "kei" (京), meaning 10^{16}) Fujitsu 88,128 SPARC64VIIIfx's at the RIKEN Advanced Institute for Computational Science , Kobe.
11.28Petaflops (4th of Top500)

＊TIT： Tsubame 2.5: Xeon+GPU 5.375Petaflops

＊U. Tokyo ：Fujitsu FX10 SPARC64IXfx's
4,800nodes 1.135Petaflops

＊Tohoku Univ. ：NEC SX-ACE 2,560nodes
0.707Petaflops

etc


2015/4to9- Phase Ⅱ
3.9 PetaFlops
Strage 11.28PB
2014/3-2015/3
＋ High Definition Visualizaion System
0.59 PetaFlops
Strage 7.32PB
2013/10-2014/3
phase Ⅰ
0.56 PetaFlops
Strage 3.96PB
Fujitsu PRIMEHPC FX100 change
Fujitsu PRIMERGY CX400 gradeup
SGI UV2000
add strage
visualization system
Fujitsu PRIMEHPC FX10
Fujitsu PRIMERGY CX400
add new external strorage
8K High-definition Display
HMD(Mix Reality)
Dome monitor
Storage system
polarized 3D viewer

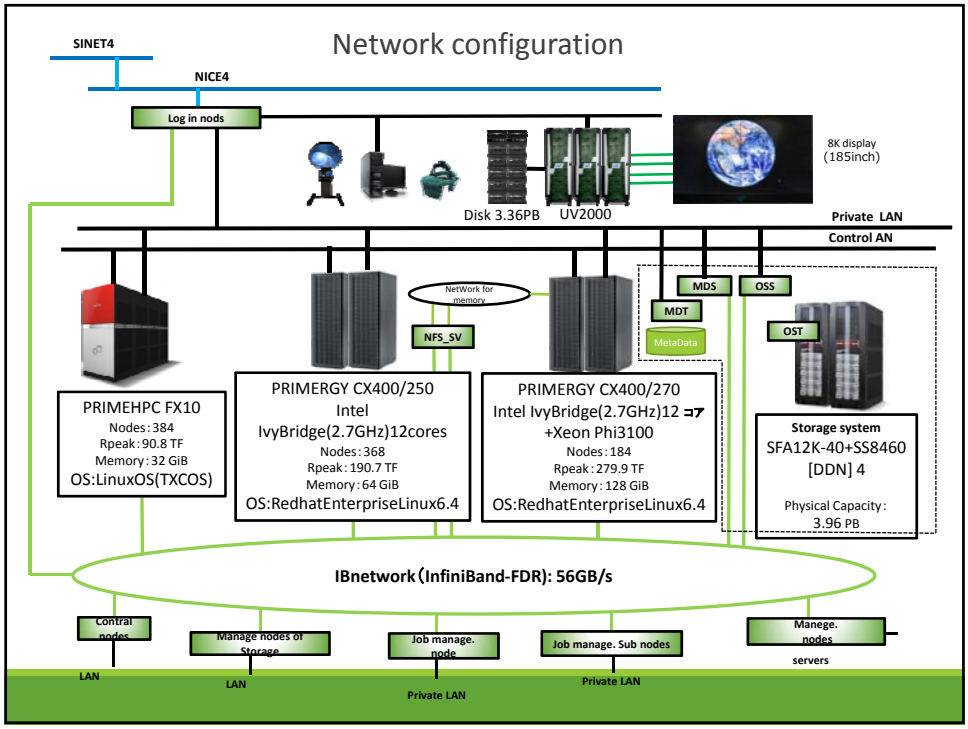
Network configuration
SINET4
NICE4
Log in nods
8K display (185inch)
Disk 3.36PB
UV2000
Private LAN
Control AN
NetWork for memory
NFS_SV
MDS
OSS
MDT
MetaData
OST
PRIMEHPC FX10
Nodes:384
Rpeak:90.8 TF
Memory:32 GiB
OS:LinuxOS(TXCOS)
PRIMERGY CX400/250
Intel IvyBridge(2.7GHz)12cores
Nodes:368
Rpeak:190.7 TF
Memory:64 GiB
OS:RedhatEnterpriseLinux6.4
PRIMERGY CX400/270
Intel IvyBridge(2.7GHz)12 コア
+Xeon Phi3100
Nodes:184
Rpeak:279.9 TF
Memory:128 GiB
OS:RedhatEnterpriseLinux6.4
Storage system
SFA12K-40+SS8460
[DDN] 4
Physical Capacity:
3.96 PB
IBnetwork（InfiniBand-FDR): 56GB/s
Contral nodes
Manage nodes of Storage
Job manage. node
Job manage. Sub nodes
Manege. nodes
servers
LAN
LAN
Private LAN
Private LAN

FX10

http://img.jp.fujitsu.com/downloads/jp/jhpc/primehpc/primehpc-fx10-catalog-en.pdf

PRIMEHPC FX10 Specifications

6,144 I/O

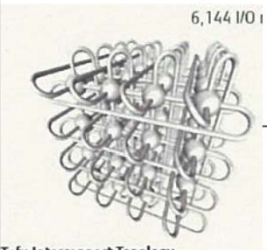

Tofu Interconnect Topology (Conceptual Diagram)

■ VISIMP

Hybrid pa

Node specifications	Theoretical peak performance	236.5 Gigaflops
	Processor	SPARC64 IXfx (1.848 GHz/16-core) x1
	Memory capacity	32GB, 64GB
	Memory bandwidth	85 GB/s
	Interconnect link bandwidth	5 GB/s x2(bi-directional)
System specifications	Number of racks	4 to 1,024
	Number of compute nodes	384 to 98,304
	Theoretical peak performance	90.8 Teraflops to 23.2 Petaflops
	Total memory capacity	12TB to 6PB
	Interconnect	Tofu Interconnect
	Cooling method	Direct water cooling + air cooling (Option: Exhaust cooling unit)

Band width/Flops = 0.36Byte/Flops

CX400 systems

	Fujitsu Primergy CX400
OS	Red Hat Enterprise Linux 6.4
Processor	Intel IvyBridge (2.7GHz) 12 core
Number of Nodes	552(13,248 cores)
Total Peak Performance	470.6Tflops
Total memory capacity	43TiB
Peak Performance of a node	518Gflops(24core)
Memory capacity of a node	64GiB(270)/128GiB(250)

coprosessor	Xeon Phi3100 family(MIC)
Processor	Accelerator(Pcle G2) 1.1GHz 57cores
Number of Nodes(cores)	184(10,488cores)
Peak performance of a node	1,003Gflops(57cores)
Memory capacity of a node	6GiB

Examples of Matrix multiplication (Block method) C=AB $10^4 \times 10^4$

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} = \begin{pmatrix} A_{11}B_{11} + A_{12}B_{21} & A_{11}B_{21} + A_{12}B_{22} \\ A_{21}B_{11} + A_{22}B_{21} & A_{21}B_{21} + A_{22}B_{22} \end{pmatrix} = C$$

Table 3.3: Result for Matrix Multiplication: Host vs Coprocessor

System	Time(sec)	Performance(GFLOPS)
Host(1 core,Non-parallelized)	404.003	4.950
Host(OMP 4 cores, 4 threads)	98.990	20.204
Host(OMP 8 cores, 8 threads)	49.827	40.139
Host(OMP 12 cores, 12 threads)	33.228	60.190
Host(OMP 16 cores, 16 threads)	24.841	80.513
Host(OMP 24 cores, 24 threads)	16.703	119.737
Xeon Phi(OMP 224 threads scatter)	10.947	170.151
Xeon Phi(OMP 224 threads compact)	9.981	186.619

Batch Job

	JOB CLass	Nodes	Cores	Max usable time		Memory	
				Default (Hours)	Max (Hours)		
FX10	fx-interactive	1	16	1	1	28GiB	For Deug
	fx-debug	4	64	1	1	28GiB × 4	
	fx-small	12	192	24	72	28GiB × 12	
	fx-middle	48	768	24	72	28GiB × 48	
	fx-large	96	1,536	24	72	28GiB × 96	
	fx-special	384	6,144	unlimited		28GiB × 384	consult
CX400/250	cx-debug	4	96	1	1	56GiB × 4	For Debug
	cx-single	1	24	24	336	56GiB × 1	
	cx-small	8	192	24	168	56GiB × 8	
	cx-middle	32	768	24	168	56GiB × 32	
	cx-large	128	3,072	24	72	56GiB × 128	
	cx-special	368	8,832	unlimited		56GiB × 368	consult
CX400/270	cx2-debug	4	96	1	1	112GiB × 4	For Debug
	cx2-small	8	192	24	72	112 GiB × 8	
	cx2-middle	32	768	24	72	112 GiB × 32	
	cx2-large	64	1,536	12	12	112 GiB × 64	
	cx2-special	184	4,416	Unlimited		112GiB × 184	consult

2nd Floor

- SuperComputer FX10
- Front end Machine（Login node）
- Storage System

3rd Floor

- Supercomputer CX400

	First Phase 2013.10～	Second Phase 2015.4～2018.3
SystemI Fujitsu Sparc	FX10: 0.1 PFLOPS	FX100: 3.2PetaFLOPS
System II Intel CPU	CX400: 0.47 PFLOPS	CX400: 0.68 PFLOPS
Storage	3.96 PB	7.92 PB

FX100 architecture (New Fujitsu SC)

An expansion to the SPARC-V9 instruction set architecture called HPC-ACE2 (High Performance Computing-Arithmetic Computational Extensions 2) features two 256-bit wide SIMD units per core with advanced operation functions, and improves the computational throughput of the processor. HMC (Hybrid Memory Cube) allows a high memory bandwidth of 480GB/s per node and the one-processor-per-node architecture exploits the maximum memory performance. Tofu Interconnect 2 (Tofu2) is integrated into the SPARC64 XIfx processor and enhances node-to-node communication bandwidth to 12.5 GB/s per link with lower latency. Highly scalable Tofu2 enables a system configuration of over 100,000 nodes.

BF=0.48Byte/flops

http://img.jp.fujitsu.com/downloads/jp/jhpc/primehpc/primehpc-fx100-datasheet-en.pdf

FUJITSU Supercomputer PRIMEHPC FX100 Specifications

2,880nodes

CPU		Processor	SPARC64 XIfx
		Architecture	SPARC V9 + HPC-ACE2
		No. of cores	32 compute cores + 2 assistant cores
		Theoretical peak performance	Over 1 teraflops (double precision)
Node		Architecture	1 CPU per node
		Memory capacity	32 GB (HMC)
		Memory bandwidth	240 GB/s (read) + 240 GB/s (write)
		Interconnect	Tofu Interconnect 2
		Interconnect link bandwidth	12.5 GB/s x 2 (bi-directional) per link
Main unit		Form factor	19-inch 2U
		No. of nodes	12
System rack		No. of main units	Up to18 (216 nodes per rack)
Maximum configuration		No. of nodes	Over 100,000
		Theoretical peak performance	Over 100 petaflops (double precision)
		Total memory capacity	Over 3 PB

http://img.jp.fujitsu.com/downloads/jp/jhpc/primehpc/primehpc-fx100-datasheet-en.pdf

CX400/2550 systems

	Fujitsu Primergy CX400
OS	Red Hat Enterprise Linux
Processor	Intel Haswell
Number of Nodes	400
Total Peak Performance	400Tflops
Total memory capacity	51.2TiB
Peak Performance of a node	1Tflops
Memory capacity of a node	128GiB

CX400/270 systems

	Fujitsu Primergy CX400
OS	Red Hat Enterprise Linux 6.4
Processor	Intel IvyBridge (2.7GHz) 12 core
Number of Nodes	184(4,416cores)
Total Peak Performance	96Tflops
Total memory capacity	
Peak Performance of a node	518Gflops(24core)
Memory capacity of a node	128GiB

coprosessor	Xeon Phi3100 family(MIC)
Processor	Accelerator(Pcle G2) 1.1GHz 57cores
Number of Nodes(cores)	184(10,488cores)
Peak performance of a node	1,003Gflops(57cores)
Memory capacity of a node	6GiB

Visualization System for Pre/Post Process

- A Linux Server with Large Shared Memory
- WindowsPC's

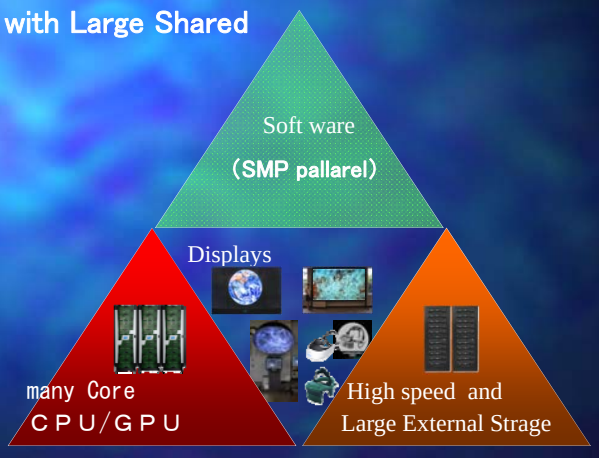

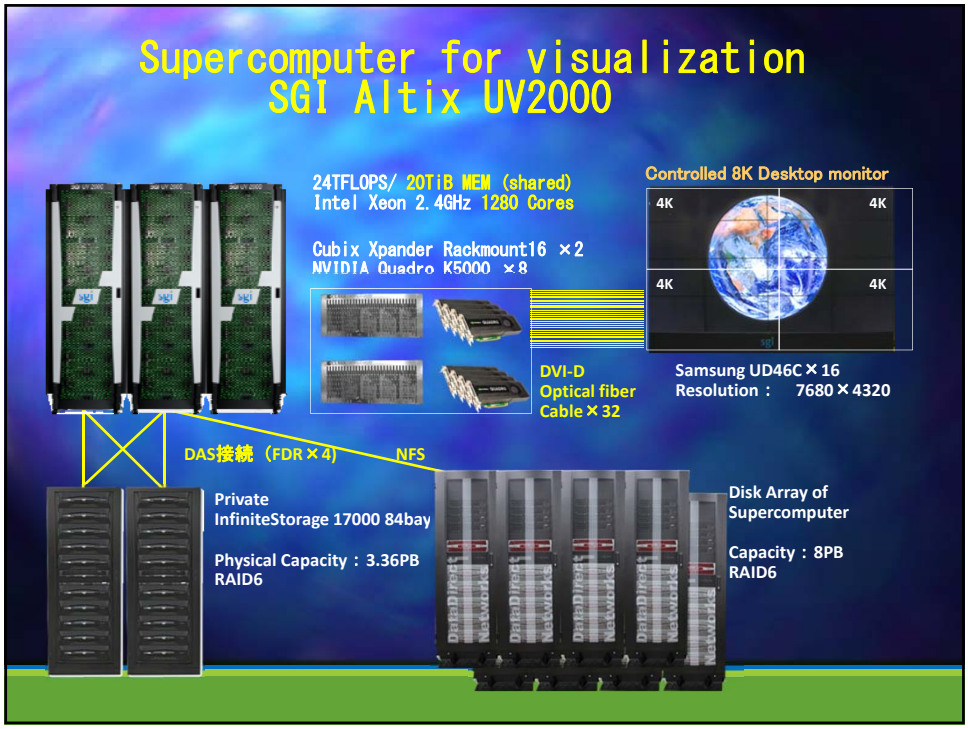
Supercomputer for visualization
SGI Altix UV2000
24TFLOPS/ 20TiB MEM (shared)
Intel Xeon 2.4GHz 1280 Cores
Cubix Xpander Rackmount16 ×2
NVIDIA Quadro K5000 ×8
DVI-D
Optical fiber
Cable×32
Controlled 8K Desktop monitor
4K
4K
4K
4K
Samsung UD46C×16
Resolution： 7680×4320
DAS接続 (FDR×4)
NFS
Private
InfiniteStorage 17000 84bay
Physical Capacity：3.36PB
RAID6
Disk Array of
Supercomputer
Capacity：8PB
RAID6


Visualizaion System
HPCI（Science Information Network:SINET）
Supercomputer systems of Nagoya Univ.（System LAN）
8K visualization System
185inch
8Ktile display
(46inch monitor 4×4)
Polarized 3D system
150inch
Poralized 3D FullHD DLP
Onsite Service
On site visualization system
Asterism IDX416VR
Connect my HardDick
Mixed Reality system
CANON MREAL Vicon
Video Camera+CG
Head Mounted Display
Dome Display Workstation
Fanoviewer
Wide view(about 180)
Remote visualization system
High resolution visualized system
Remote interactive operation

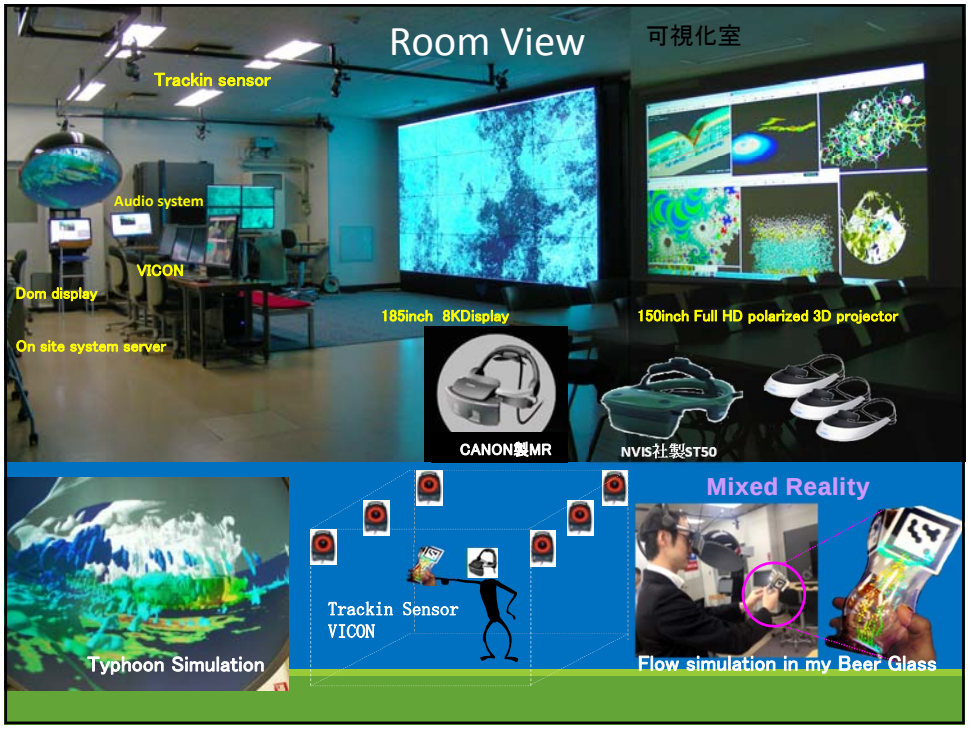
Room View
可視化室
Trackin sensor
Audio system
VICON
Dom display
On site system server
185inch 8KDisplay
150inch Full HD polarized 3D projector
CANON製MR
NVIS社製ST50
Mixed Reality
Typhoon Simulation
Trackin Sensor
VICON
Flow simulation in my Beer Glass

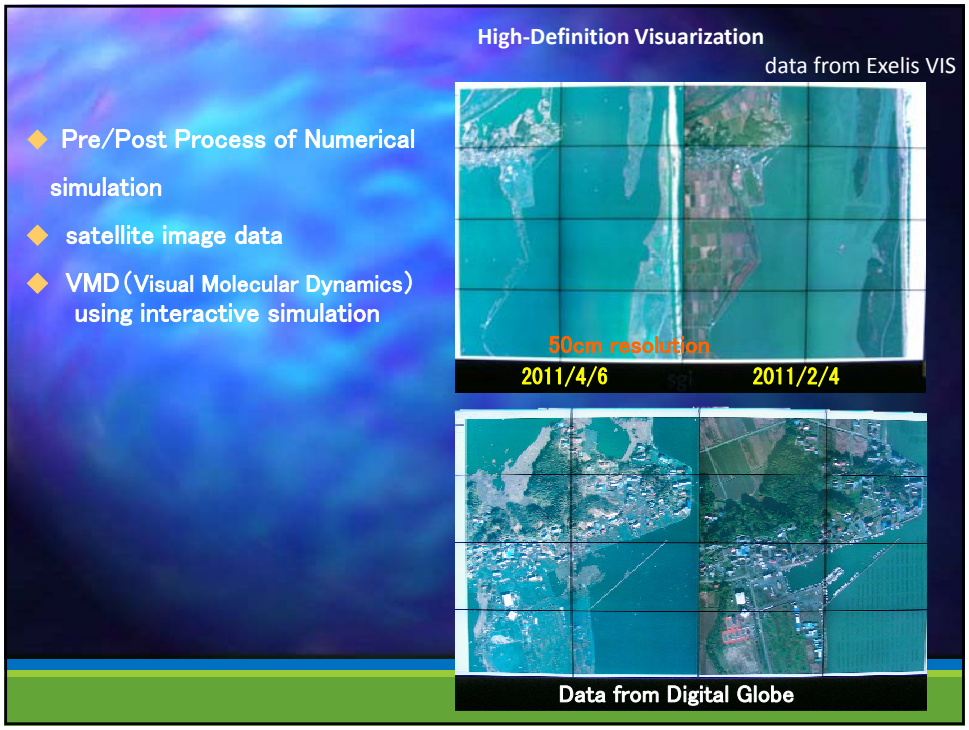
High-Definition Visuarization
data from Exelis VIS
Pre/Post Process of Numerical simulation
satellite image data
VMD (Visual Molecular Dynamics) using interactive simulation
50cm resolution
2011/4/6
2011/2/4
Data from Digital Globe

SGI UV2000　Software

システム名	SGI UV 2000
OS	SuSE Linux Enterprise Server 11
Development	Intel Composer XE 14.0.1
Math Library	Intel MKL 11.1
MPI	SGI MPT 2.08
Batch	PBS Professional 12.2
Softwares for Visualization	AVS Express Developer	PCE Enshight DR ParaView (parallel) IDL、ENVI
Software for remote 3D visualization	SGI NICE DCV 2013.0
Others	POV-Ray (SMP対応版) ffmpeg, ffplay, osgviewer, VMDなど
Fourier Transformation	FFTW 3.3
I/O library	HDF5 1.8.2, netcdf 4.2.1

Visualization　Examples

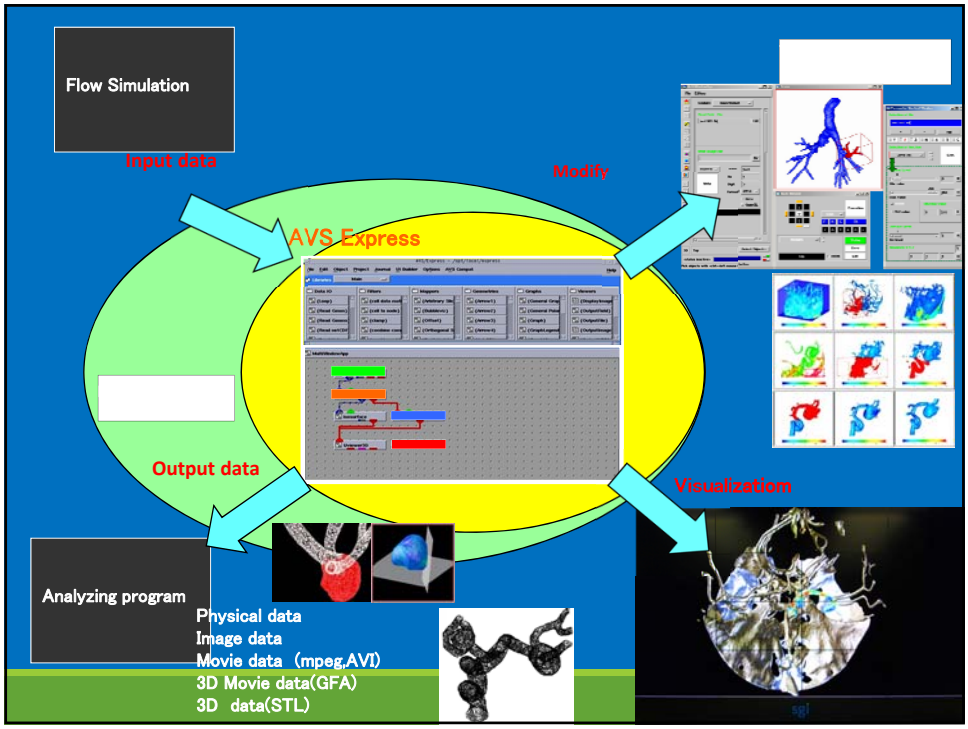
Flow Simulation
Input data
Modify
AVS Express
Output data
Visualization
Analyzing program
Physical data
Image data
Movie data (mpeg,AVI)
3D Movie data(GFA)
3D data(STL)

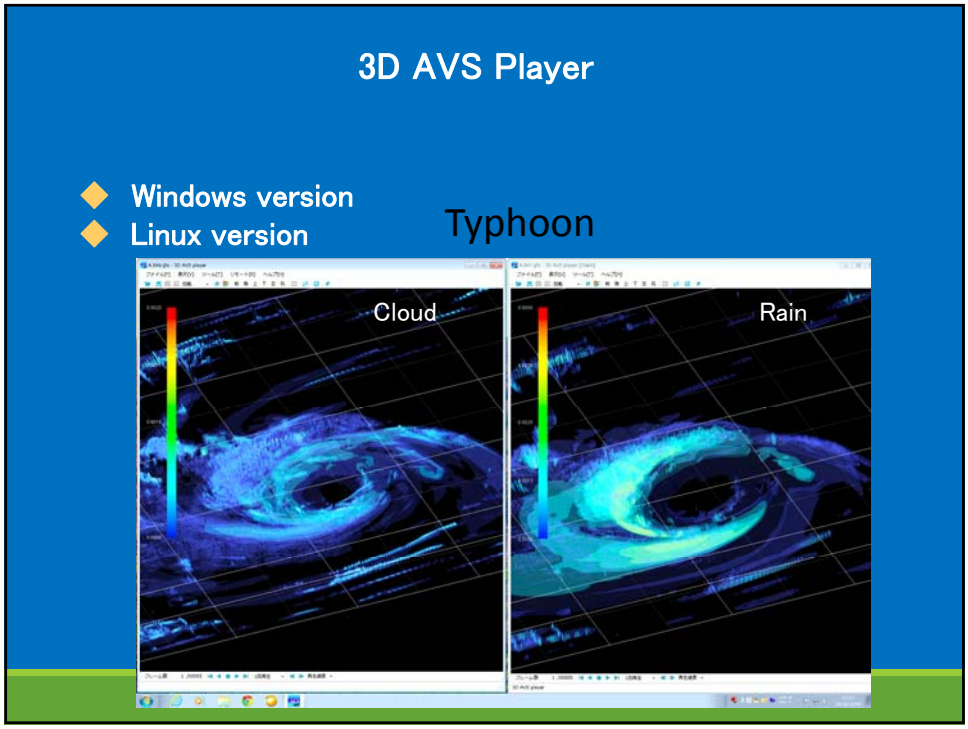
3D AVS Player
Windows version
Linux version
Typhoon
Cloud
Rain

Flow simulation in a Cerebral Aneurysm

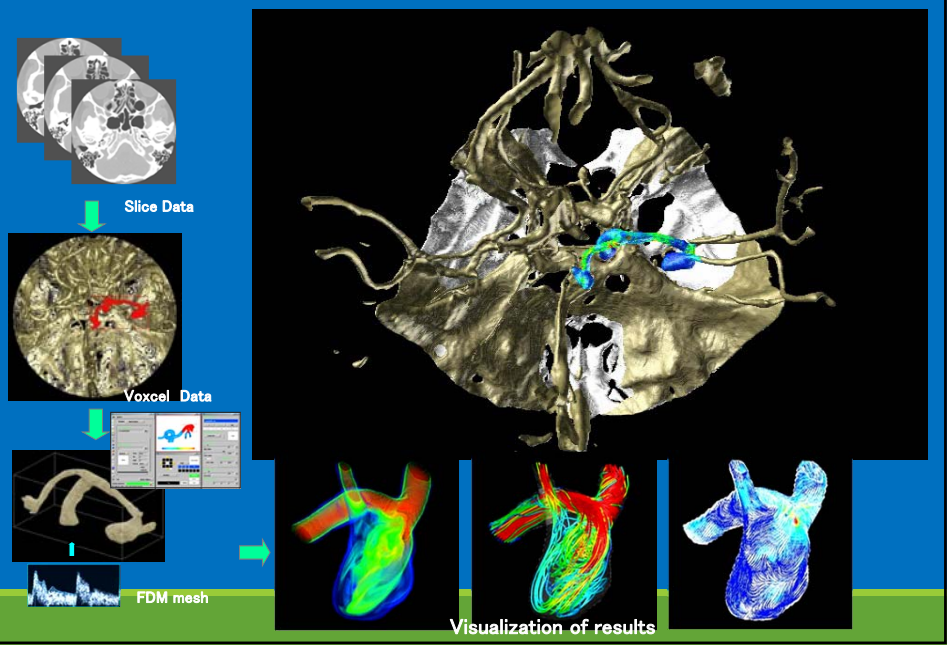

Acknowledgment

- FUJITSU
- SGI Japan
- NVIDIA
- CYBERNET SYSTEMS
- Exelis VIS