

System Packaging Solution for Future High Performance Computing

May 31, 2018

Shunichi Kikuchi

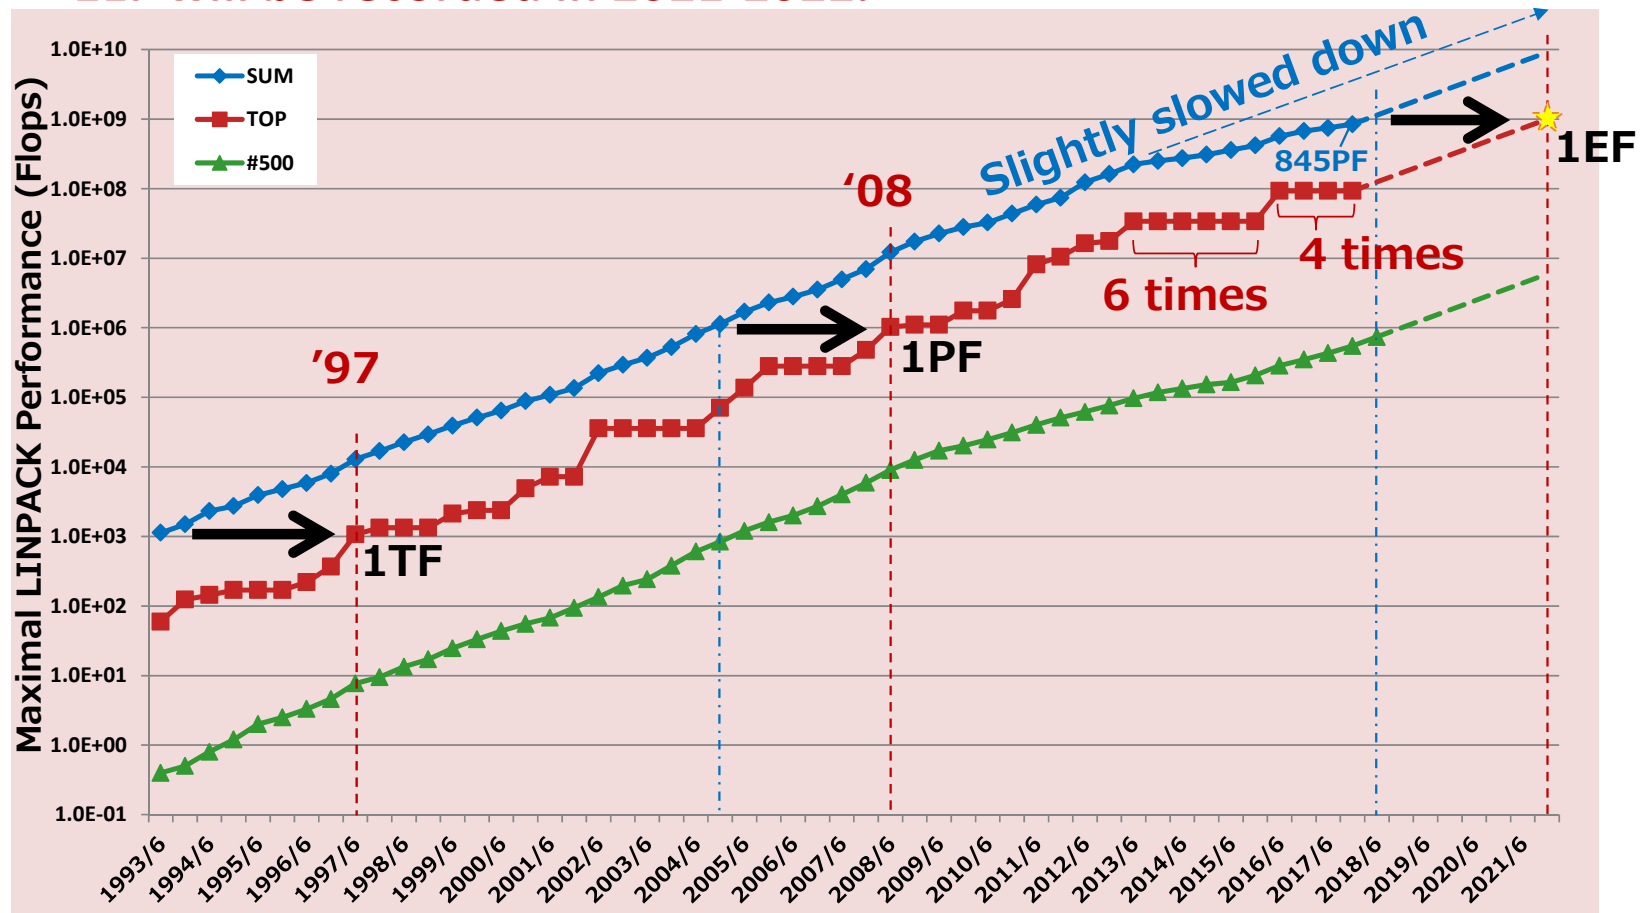
Fujitsu Limited

- Future HPC spec. from the transition of TOP500
- Possible System Packaging for the future

For the advent of Exascale Computing

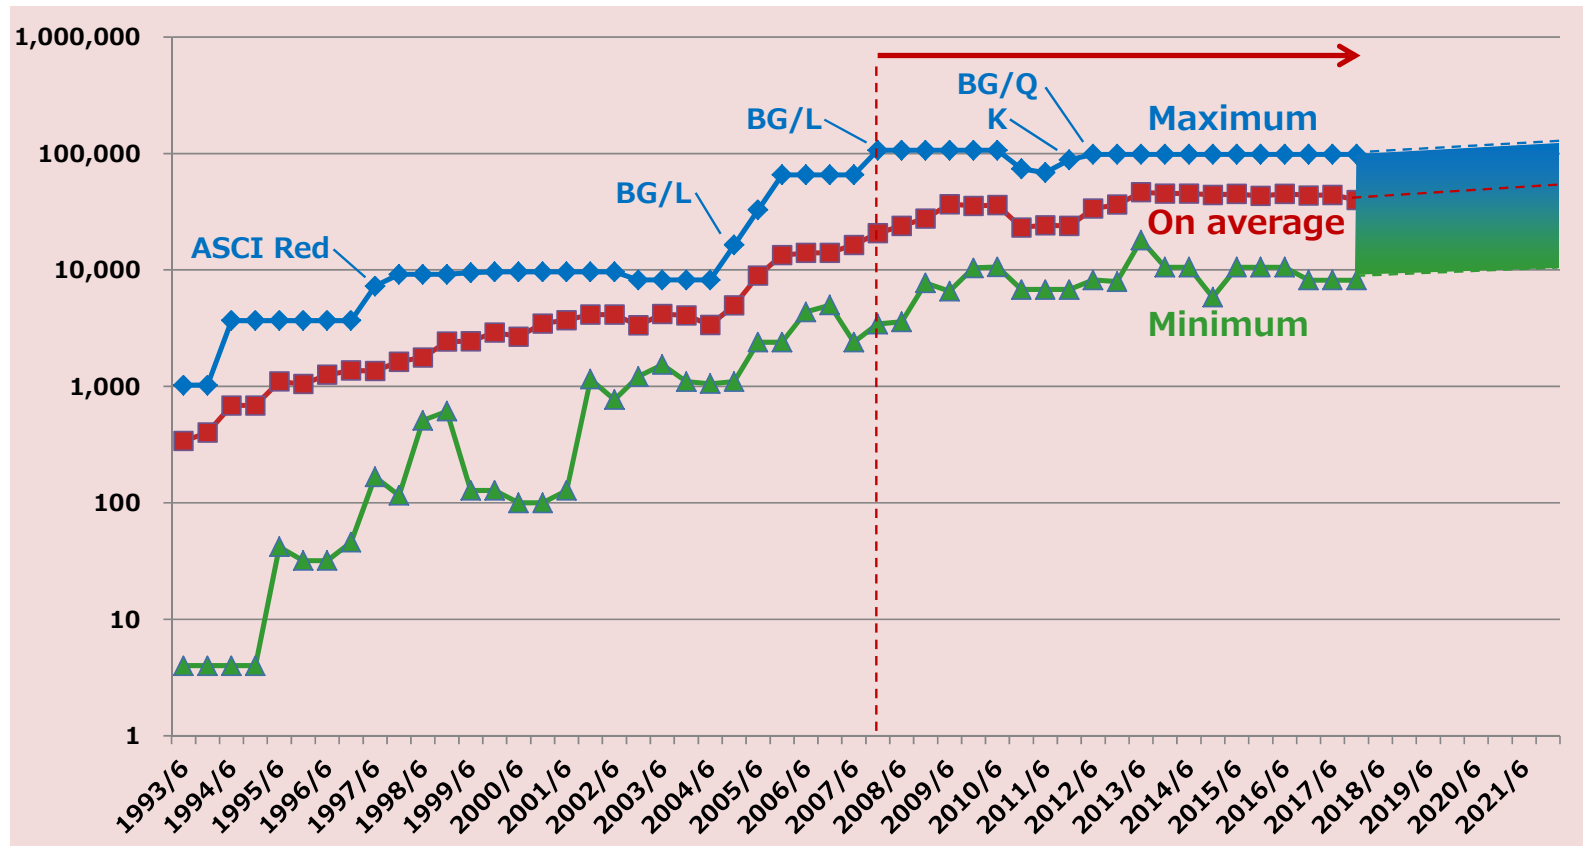
- A prediction from TOP500 history -

**“TOP”s have been recorded with 3.5-4 year delay from “SUM”s.
1EF will be recorded in 2021-2022.**




The number of CPUs and Accelerators (TOP10)

The number of them seems to have saturated since 2007.



Trend of the number of cores in system (TOP10)

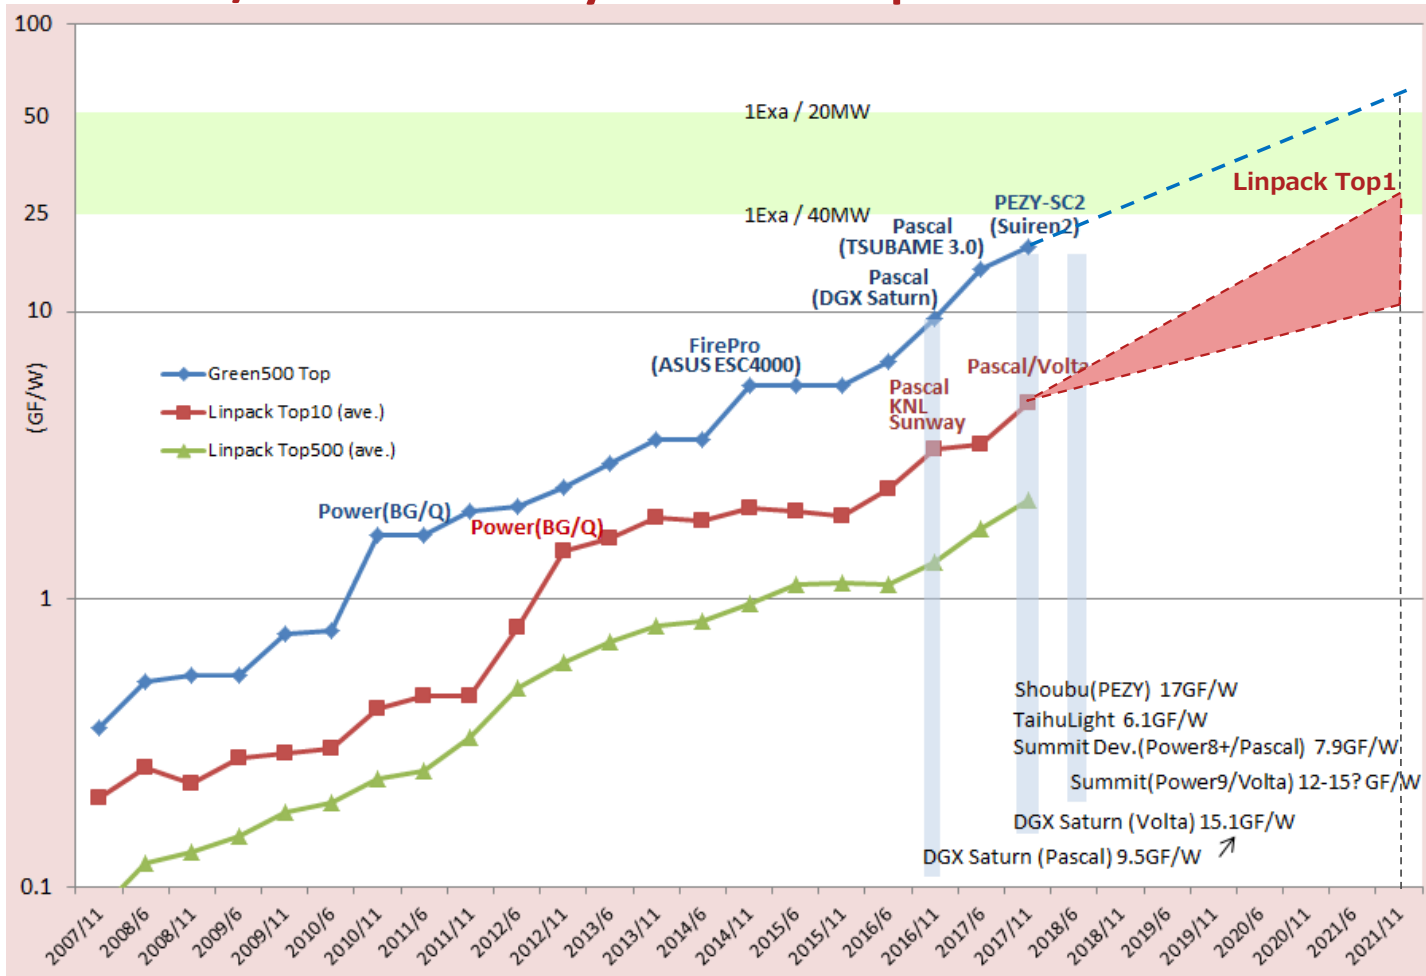


Performance improvement needs to increase the number of cores.

[Exascale era] { # of CPUs and Accelerators : 10,000 – 100,000+ pcs
Performance per each : 100,000 – 10,000 GFlops

Performance per power

25-50 GF/W will be a key factor to implement 1EF.



TOP10 feature @ Nov. 2017

There are three types of configurations.

CPU



6.1GF/W by 260core CPU



TaihuLight

No.1

NRPC



Sequoia

No.6

LLNL




K computer

No.10

©RIKEN


Accelerator(bootable)



Trinity

No.7

LANL



Cori

No.8

NERSC




Oak-forest PACS

No.9

JCAHPC

CPU+Accelerator



Tianhe-2

No.2

NSCC



Piz Daint

No.3


CSCS



Shoubu

No.4

JAMSTEC



Titan

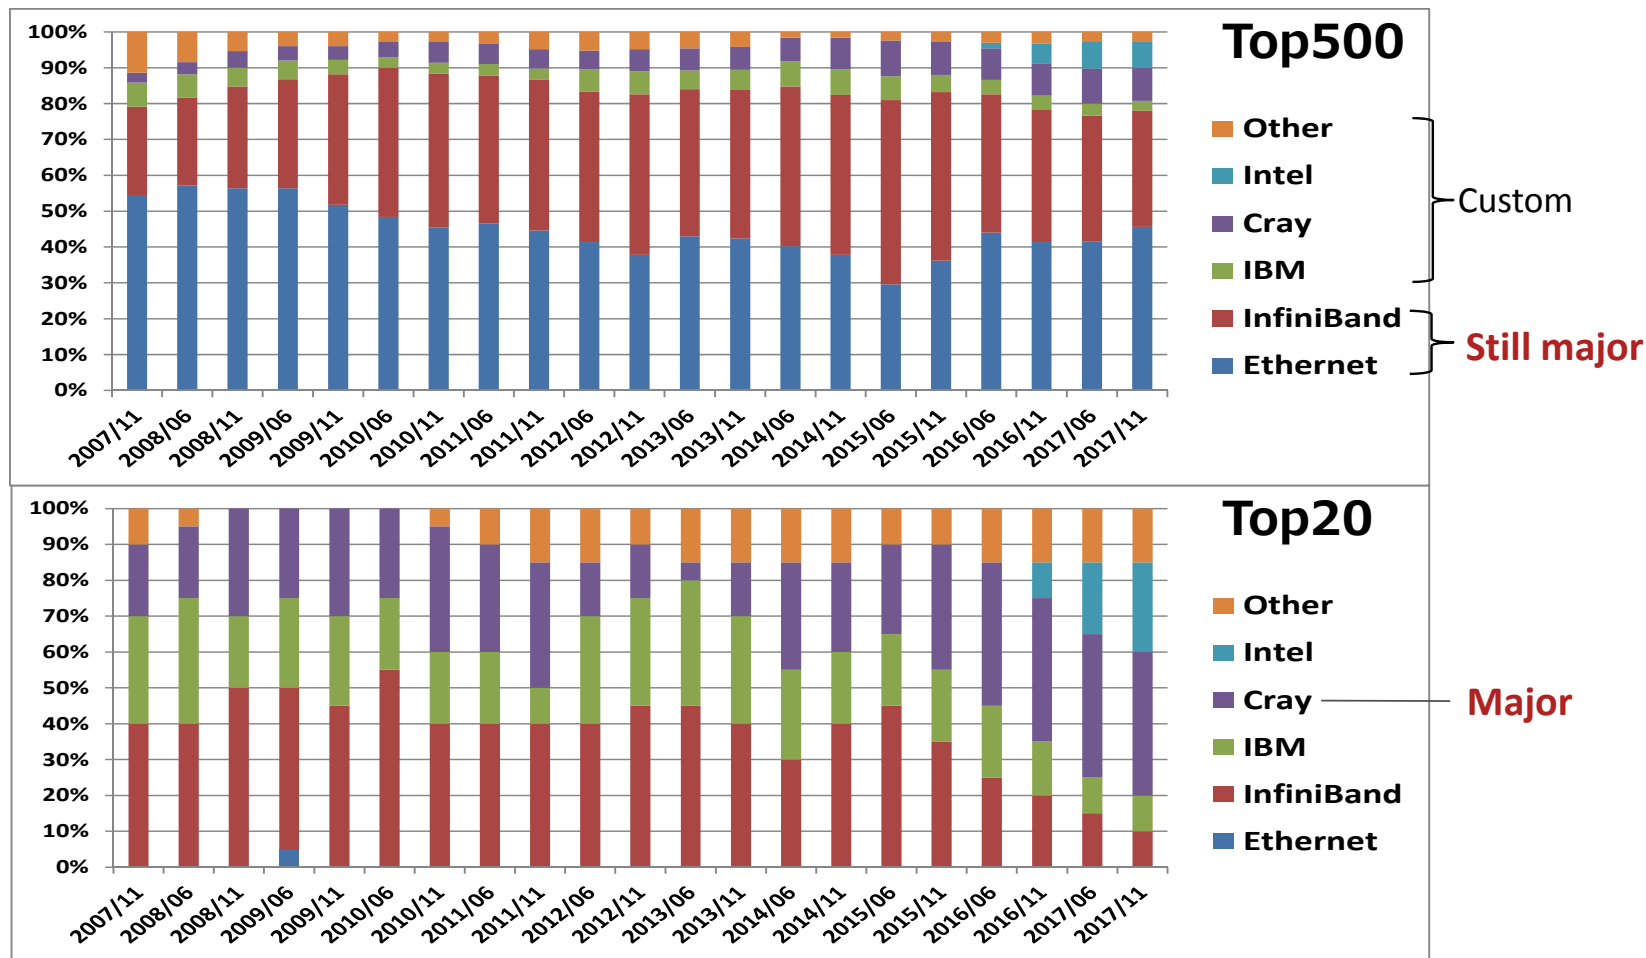
No.5

ORL

14.2GF/W by 1984coreAcc.



I/O interconnects in system

Both InfiniBand and Ethernet are popular as a whole.



Communication links for short-distance

- ACSI Purple -3K links @2.5G in 2005
- Roadrunner -40K links @5G in 2008
- Power 775 -60K fibers/rack @10G in 2011
- Sequoia -620K links @10G in 2011







Transition of interconnects in system

Optical interconnects have been applied between racks and between shelves. They will be near to CPUs and Accelerators, depending on system requirements, technology evolution (including costs).

Speed connected	-10Gbps	25Gbps	50Gbps-	Distance	Component
Example Systems	K computer FX10	FX100	Future system		
Between racks	Elec	Opto	Opto	- 100m	Cable
Between shelves	Elec	Opto	Opto	- 3m	Cable
Between boards in shelf	--	Opto	Opto	- 1m	Cable
	--	--	Opto		Cable
	Elec	--	Opto		Board, Back plane
On board	Elec	--	Elec	- 0.3m	Board
	--	--	Opto		Board(Fiber sheet, Waveguide)
	--	Elec/Opto	Elec/Opto		Board&cable
	--	--	Elec		Cable
	--	--	Opto		Cable


Interconnects on board




SPARC64™ Xlfx

Board-mounted optical assembly (12 lanes each)

Electrical connector (8 lanes)




Altera, Avago, 2012




PETRA/NEDO, 2016

Related technologies in the past ECTCs




336ch, 64Gb/s/mm²
(168*8Gb/s=1.344Tb/s)



348ch, 15Gb/s/mm²



96ch, 2.4Tb/s






Compass EOS@ECTC2014

IBM@ECTC2016

AIST/PETRA@ECTC2016

CPU/Accelerator chip trend (in general)

There seems no significant growth in clock frequency since 2007.
Die size has been reached to the manufacturing limit.
Heat dissipation may be controlled toward green computing.



Power delivery trend

At Package/Module Level


Mounted DDCs(ex. 48V→1V)



Exascaler Inc. @SC17



nVIDIA



At Datacenter Level

Creating electrical power from


Hydrogen
Natural gas etc.




Source: Daimler

Possible Processor Package and Module

High density package





Package substrate



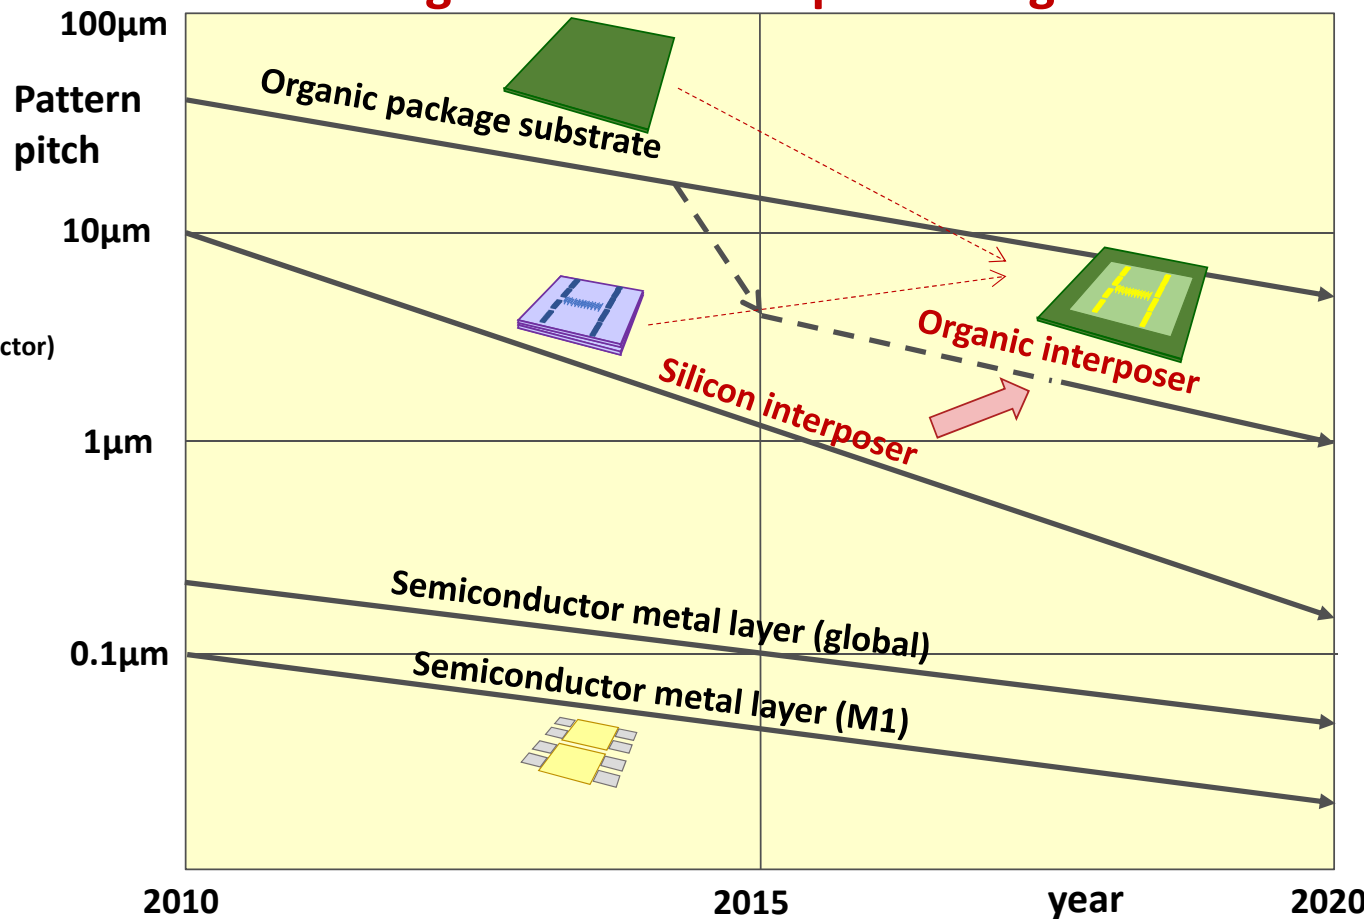


Silicon-Photronics

Standard I/F module



Transition of package substrate technology

Technology progress and design optimization may be able to mitigate the current patterning and cost issues.




Future prospect of system packaging

High density system packaging

Acceleration package

CPU/MEM package



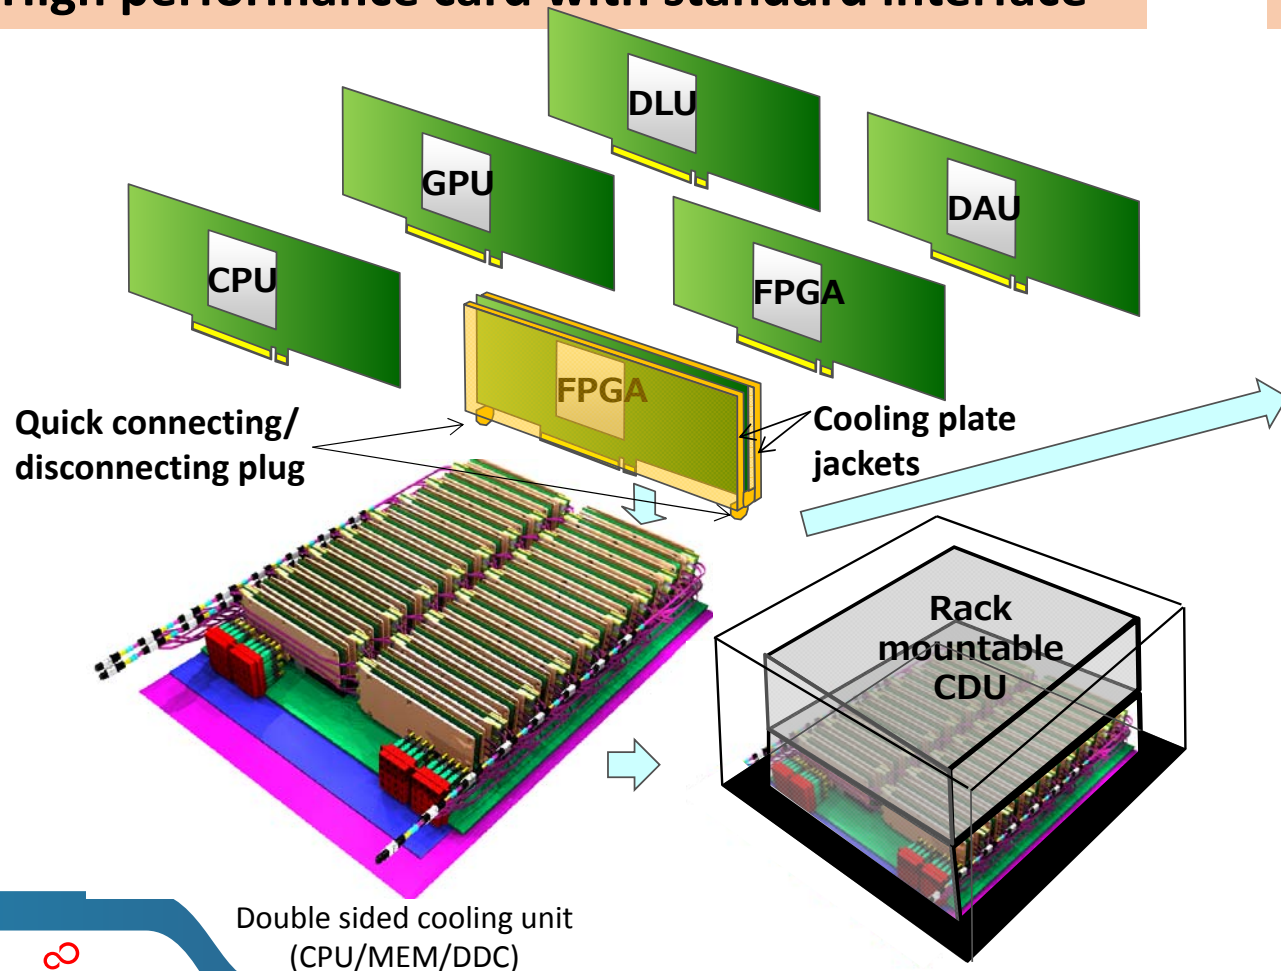
High density rack mounting



Future prospect of system packaging

High performance card with standard interface

Highly dense rack mounting




A wide range of choices should be prepared in order to match with customers' requirements :

Air-forced, liquid circulation, and liquid immersion

Coolant availability
Limitation of ambient temperature
Limitation of power consumption
Limitation of installation space


Air-forced cooling

- Heat sink
- Heat sink w/ heat pipe
- Vapor and liquid loop
- Vapor chamber
- Thermo-siphon



Liquid circulation cooling

- Hot water
- Chilled water
- Micro-channel
- Impinging



Liquid immersion cooling

- Submersion of all parts
- Natural convection
- Forced circulation
- Boiling




Cooling from the viewpoint of system packaging

High density main board




Liquid circulation



A lot of PCB units




Liquid immersion



Toward Big-data era and the beyond

If technology is continuously evolved, we may reach...

(*1) Source: 36.8PF, 3.2PBytes
Best of H+ Magazine 2008-2010 Volume 1



Human Brain
37 PF (*1)/20W=1.85x10⁶ GF/W



Source: FORTUNE: July 10, 2017


20MW /20W = 1Million Processors, 37PF*1M=37ZF@2039

200PF@2018, 10times/4yrs, 200*10^(21/4)=36ZF@2039

1 billion times as fast as Earth Simulator(36TF) in 2002-2004

Tera → Peta → Exa → Zeta

Toward Big-data era and the beyond



Which type of computing can be near to human brains?

Conventional HPC?

Quantum Computing?

Quantum Annealing?

Digital Annealing?

Or else?

Various device configurations for different issues

It is difficult to solve and find a better goal by using one type of HPC. A suitable system packaging solution can be chosen in the coming several years for each.



Health care



Traffic Relaxation



Sustainable Food and Agriculture



Reproducible Energy



Education for everyone



Sustainable Environment

Fujitsu's High-end HPC Development

- Fujitsu has provided HPC systems with original technologies, developed for over 40 years, to accelerate advanced research

K computer

The K computer continues to be competitive in various fields; from advanced research to manufacturing



© RIKEN



Gordon Bell Prize Finalist (2016)

HPCG

HPCG No.1(2017)



Graph500 No.1(2017)



PRIMEHPC FX10

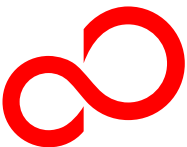


PRIMEHPC FX100

Post-K computer

RIKEN and Fujitsu are developing the Post-K to achieve superior application performance





FUJITSU

shaping tomorrow with you