The development of the application software for PEZY SC2 many-core processors

Ryutaro Himeno, **Toshikazu Ebisuzaki**, RIKEN
Junnichiro Makino, Kobe Univ.
Hide Sakaguchi, JAMSTEC
Tadashi Yamazaki, The University of Electro-Communications
Tadashi Ishikawa, KEK
Ken Kurokawa, National Institute of Genetics

RIKEN
(The Institute of Physical Chemical Research)
https://www.riken.jp/en/

- Only one institution, dedicated for Basic Science in Japan
- Founded in 1917 (102 years old in the next November)

E. Shibuzawa
Founder

U. Suzuki
Discoverer of
Vitamin B1

Y. Nishina
Kline-Nishina
Formula

H. Yukawa
Meson Theory
Novel Prize in
Physics 1949

S. Tomonaga
QED
Novel Prize in
Physics 1965

PEZY computing (Inc.)

- The **first venture company** for supercomputing in Japan
- Founded by Dr. Motoaki Saito, a medical doctor, who have worked for medical imaging.
- 2014 Spring: Start HPC business

PEZY stands for Peta, Exa, Zetta and Yotta

PEZY computing (Inc.)

- The **first venture company** for supercomputing in Japan
- Founded by Dr. Motoaki Saito, a medical doctor, who have worked for medical imaging.
- 2014 Spring: Start HPC business
- 2014 November: 2nd in Green500
- Green500: Eight times ranked in the past eight years

2014.11	#2	4.95Gflops/W	Suiren	KEK	
2015.07	#1	7.03GFlops/W	Shoubu	RIKEN	
	#2	6.84Gflops/W	Suiren Blue	KEK	
	#3	6.22Gflops/W	Suiren	KEK	
2015.10	#1	7.03Fflops/W	Shoubu	RIKEN	
	#2	6.22Gflops/W	Satsuki	RIKEN	
2016.06	#1	6.67Gflops/W	Shobu	RIKEN	
	#2	6.20Gflops/W	Satsuki	RIKEN	
2016.11	#1	6.67Gflops/W	Shobu	RIKEN	SC
2017.06	#1	14.05Gflops/W	Kukai	Yahoo! Japan	
2017.11	#1	17.01Gflops/W	Shobu B	RIKEN	SC2
	#2	16.76Gflops/W	Suiren	KEK	
	#2	16.7Gflops/W	Sakura	PEZY Comp.	
2018.06	#1	18.4Gflops/W	Shobu B	RIKEN	
	#2	16.8Gflops/W	Suiren 2	KEK	
	#3	16.7Gflops/W	Sakura	PEZY Comp.	
2018.11	#1	17.6Gflops/W	Shobu B	RIKEN	
2019.06	#1	17.6Gflops/W	Shobu B	RIKEN	

PEZY stands for Peta, Exa, Zetta and Yotta

ZettaScaler Supercomputers installed at PEZY, KEK, and Riken

Suiren（睡蓮） ZettaScaler-1.5
2014.10Install 2016.5 Upgrade
(32node to 48node)

Shoubu（菖蒲） ZettaScaler-1.6
2015.6Install 2016.5 Upgrade

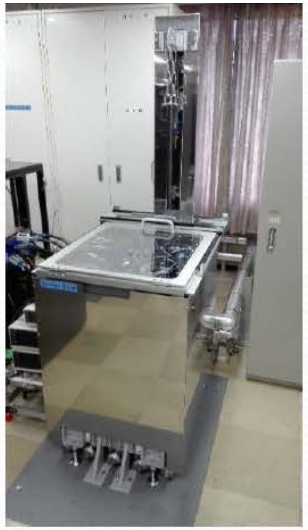

Suiren Blue（青睡蓮）
ZettaScaler-1.5
2015.5 Install
2016.5 upgrade

Ajisai（紫陽花）
ZettaScaler-1.6
2015.10Install 2016.5Upgrade

Satsuki（皐月） ZettaScaler-1.6
2016.5 Install

Sakura（さくら）
ZettaScaler-1.6
2016.5 Install

Green500: The latest Supercomputer ranking list as of June, 2019

Rank	TOP500 Rank	System	Rmax (TFlop/s)	Power (kW)	Power Efficieny (GFlops/watts)
1	472	**Shoubu system B** - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 , PEZY Computing / Exascaler Inc. Advanced Center for Computing and Communication, RIKEN Japan	1,063.3	60.4	17.604
2	470	**DGX SaturnV Volta** - NVIDIA DGX-1 Volta36, Xeon E5-2698v4 20C 2.2GHz, Infiniband EDR, NVIDIA Tesla V100 , Nvidia NVIDIA Corporation United States	1,070.0	97	15.113
3	1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	148,600.0	10,096	14.719
4	8	**AI Bridging Cloud Infrastructure (ABCI)** - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR , Fujitsu National Institute of Advanced Industrial Science and Technology (AIST) Japan	19,880.0	1,649.3	14.423
5	394	**MareNostrum P9 CTE** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Tesla V100 , IBM Barcelona Supercomputing Center Spain	1,145.0	81.0	14.131
6	25	**TSUBAME3.0** - SGI ICE XA, IP139-SXM2, Xeon E5-2680v4 14C 2.4GHz, Intel Omni-Path, NVIDIA Tesla P100 SXM2 , HPE GSIC Center, Tokyo Institute of Technology Japan	8,125.0	792.1	13.704
7	11	**PANGEA III** - IBM Power System AC922, IBM POWER9 18C 3.45GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Volta GV100 , IBM Total Exploration Production France	17,860.0	1,367	13.065
8	2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	94,640.0	7,438.3	12.723
9	43	**Advanced Computing System(PreE)** - Sugon TC8600, Hygon Dhyana 32C 2GHz, Deep Computing Processor, 200Gb 6D-Torus , Sugon Sugon China	4,325.0	380	11.382
10	23	**Taiwania 2** - QCT QuantaGrid D52G-4U/LC, Xeon Gold 6154 18C 3GHz, Mellanox InfiniBand EDR, NVIDIA Tesla V100 SXM2 , Quanta Computer / Taiwan Fixed Network / ASUS Cloud National Center for High Performance Computing Taiwan	9,000.0	797.5	11.285

#1 is PEZY-SC2 based system with 17.6 GFLOPS/W

#2 to #8 and #9 are dominated by NVIDIA V100/GV100/P100 15.1 GFLOPS/W is the highest

Note: Unfortunately, this ranking result was later delisted by Riken due to the expiration of the joint research agreement by the time of ranking submission in June

MIMD Processor, PEZY-SC and PEZY-SC2

Extremely bigger than typical Intel x86 CPU

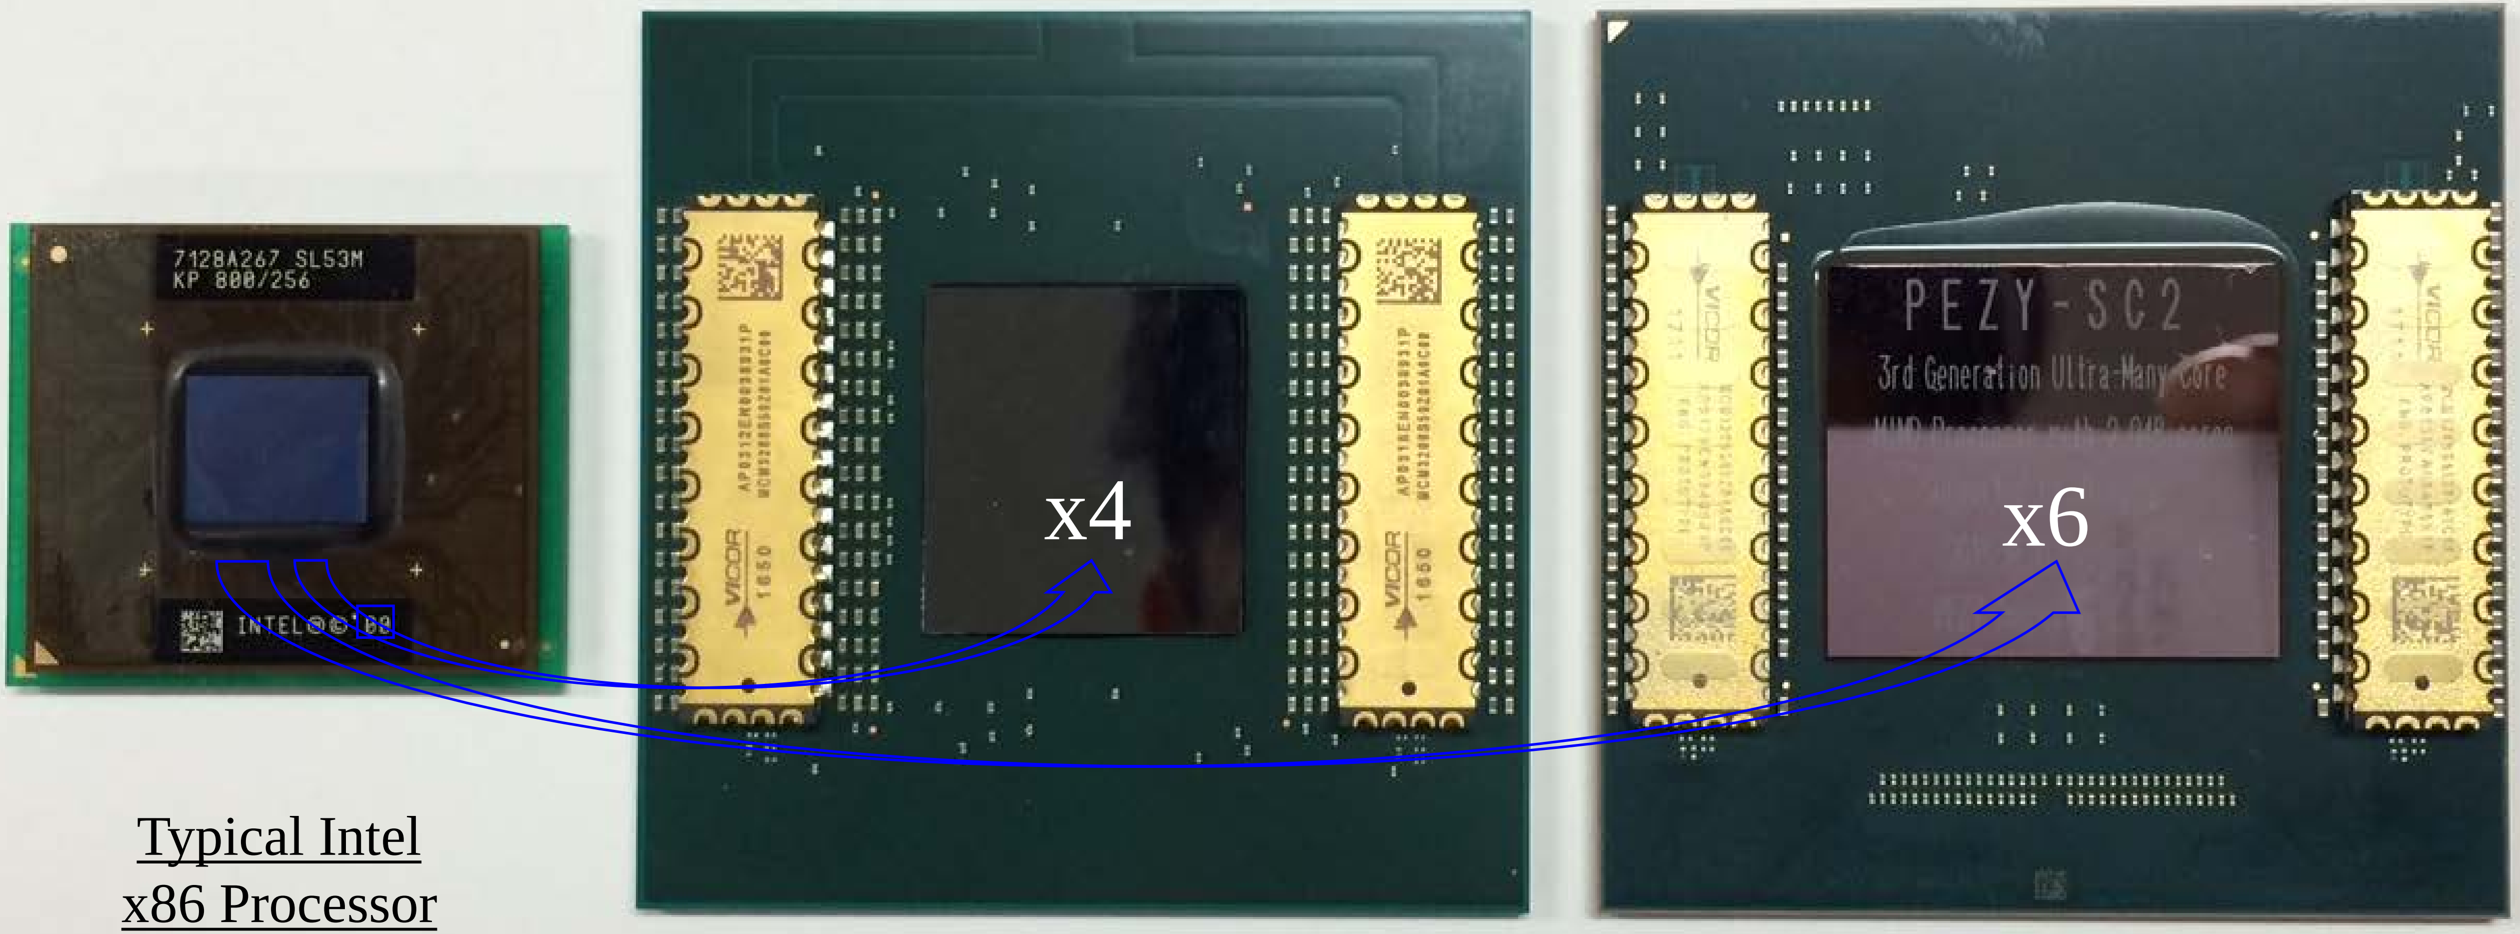

Typical Intel x86 Processor ($105mm^2$)

PEZY-SC ($410mm^2$)

PEZY-SC2 ($620mm^2$)

PEZY-SC2 Module Card

(4 of DDR4 DIMMs with 2 of PCIe Gen3 x16 port)

ZettaScaler-2.0 Brick

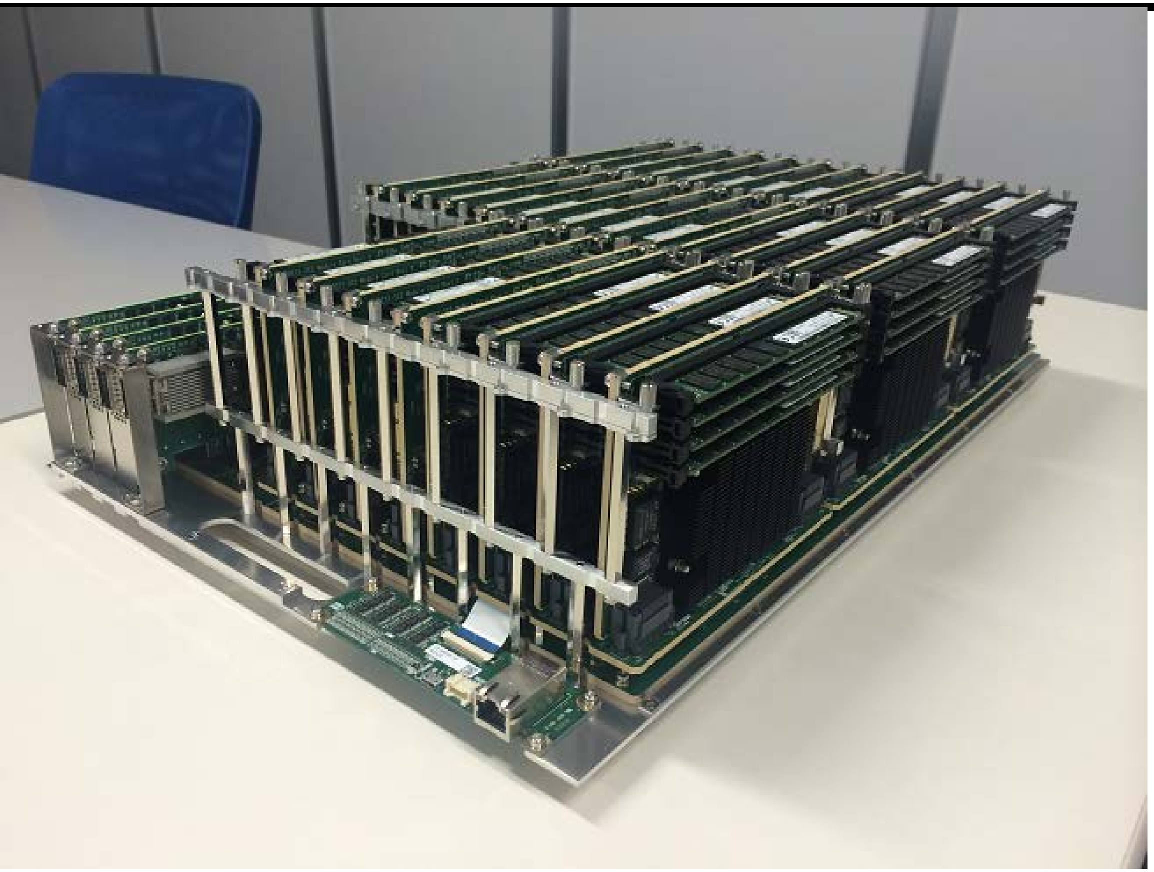

Rpeak) Brick
with 32 of PEZY-SC2 module
and 4 of EDR adaptor card

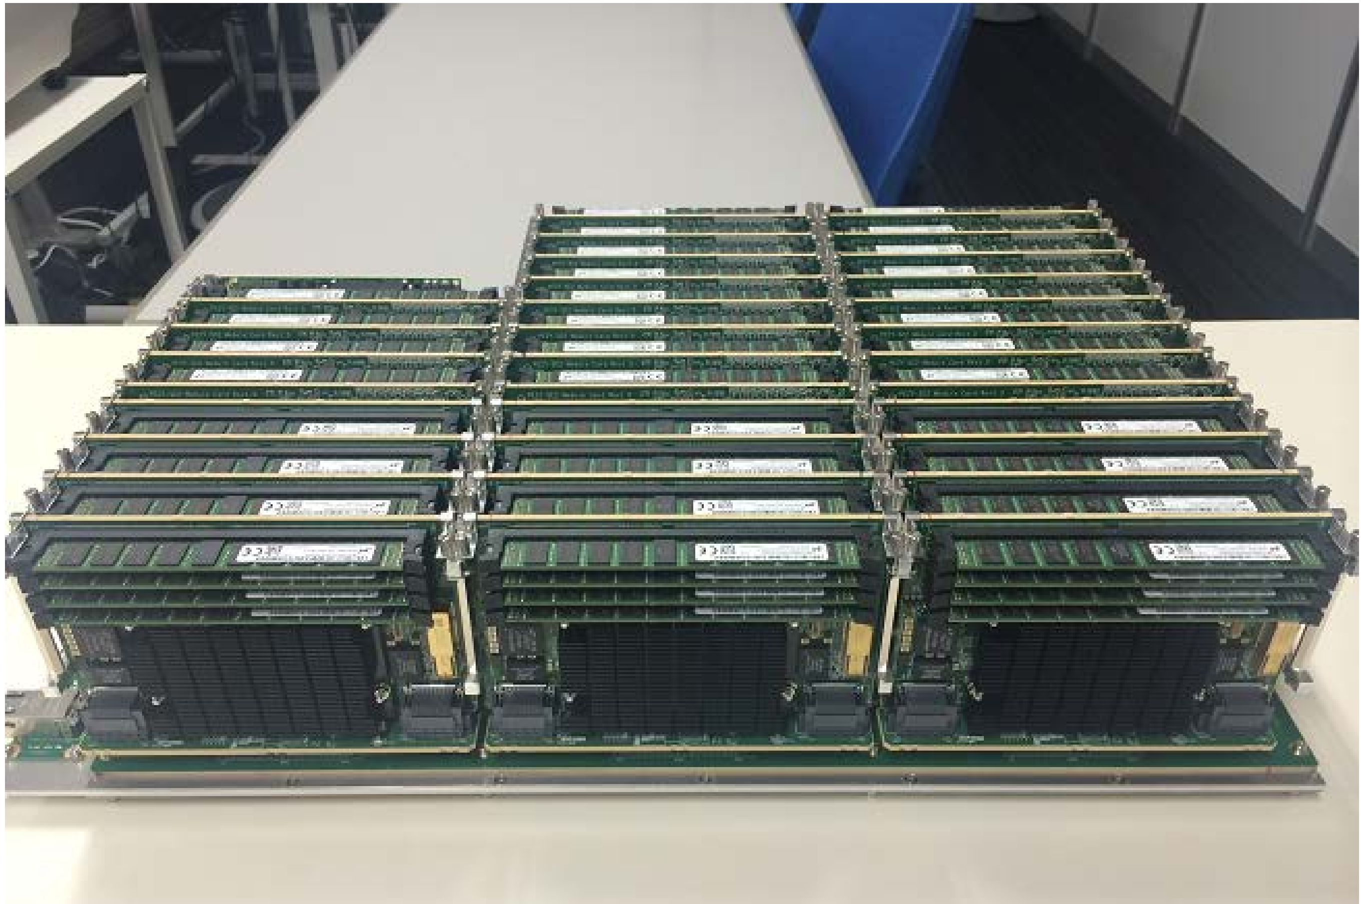

x 97 lanes) provide sufficient
bandwidth and flexibility of

PEZY-SCx Processor History and Plan

	2012	2014	2017		2019
Processor	PEZY-1	PEZY-SC	**PEZY-SC2**	scale	**PEZY-SC3**
Process Node	**40 nm**	**28 nm**	**16 nm**		**7 nm**
Die Size	118 mm^2	412 mm^2	620 mm^2	1.3	780 mm^2
Core Numbers	**512**	**1,024**	**2,048**	**2**	**4,096**
Core Voltage	1.2 V	0.9 V	0.8 V		0.7 V
Clock Frequency	433 MHz	533 MHz	700 MHz	1.7	1,200 MHz
DRAM-IO	DDR3	DDR4	DDR4		HBM2
Memory Bandwidth	48 GB/s	51 GB/s	77 GB/s	15.6	1,200 GB/s
Interface Bandwidth	12 GB/s	24 GB/s	32 GB/s	3	96 GB/s
DP Performance	**0.3 TFLOPS**	**0.9 TFLOPS**	**2.7 TFLOPS**	**7.3**	**19.7 TFLOPS**
Power Consumption	60 W	100 W	130 W	3.8	500 W
Power Efficiency	5.0 GFLOPS/w	9.0 GFLOPS/w	20.8 GFLOPS/w	1.9	38.4 GFLOPS/w
System Efficiency	**-**	**6.7 GFLOPS/w**	**17.6 GFLOPS/s**	**1.9**	**31.2 GFLOPS/s**

Development of PEZY-SC3 has been completed and is now ready for tape out

PEZY-SC3 will be taped out by the end of 2019 and will be in volume production from mid 2020

PEZY-SC3 will be the world biggest processor ever with highest performance and efficiency

Installed Supercomputers in the past

All Installed Supercomputers developed by PEZY Computing K.K. and ExaScaler K.K.

	System Name	Configuration	Installation	Rmax	Highest Top500	Gflops/W	Highest Green500	Comment	Site		Sales amount	Purpose
1	Suiren	ExaScaler-1.0	Oct., 2014	207	#365	6.22	#2	For Quantum Physics research	KEK	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
2	Shoubu	ExaScaler-1.4	June, 2015	354	#160	7.03	#1	3 times of #1 of Green500	RIKEN (HQ)	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
3	Suiren Blue	ExaScaler-1.4	June, 2015	194	#391	6.84	#2	For Quantum Physics research	KEK	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
4	Ajisai	ZettaScaler-1.6	June, 2015	187	-	6.80	-	Various HPC applications	RIKEN (HQ)	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
5	Satsuki	ZettaScaler-1.6	June, 2016	291	#486	6.20	#2	Various HPC applications	RIKEN (Kobe)	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
6	Fujitsu Test Sytem	ZettaScaler-1.8	Feb, 2017	-	-	-	-	-	Fujitsu (HQ)	IT Conglomerate	$8M	Not disclosed
7	Kukai	ZettaScaler-1.6	June, 2017	461	#466	14.05	#2	Commercially used sytem	Yahoo! Japan	Internet Service Provider	$4M	For daily mission critical application of Yahoo! Auction and many others
8	Gyoukou	ZettaScaler-2.2	June, 2017	20,200	#4	14.17	#5	First 20+ PFLOPS outside China	JAMSTEC	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
9	Shoubu System B	ZettaScaler-2.2	Oct., 2017	1,063	#259	17.60	#1	4 times of #1 of Green500	RIKEN (HQ)	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
10	Suiren2	ZettaScaler-2.2	Oct., 2017	788	#307	16.76	#2	For Quantum Physics research	KEK	National Research Laboratory	Funded by Government	Wide variety of High Performance Computing application use
11	Sakura	ZettaScaler-2.2	Oct., 2017	794	#276	16.66	#3	For own R&D use	PEZY Computing	Supercomputer/HPC R&D	-	PEZY-SC3 development
12	Kukai-2	ZettaScaler-2.4	May, 2018	730	-	15.80	-	Commercially used sytem	Yahoo! Japan	Internet Service Provider	$5M	For daily mission critical application of Yahoo! Auction and many others

ZettaScaler-1.8, the first Supercomputer with 1 PFLOPS+ per cubic meter in 2017

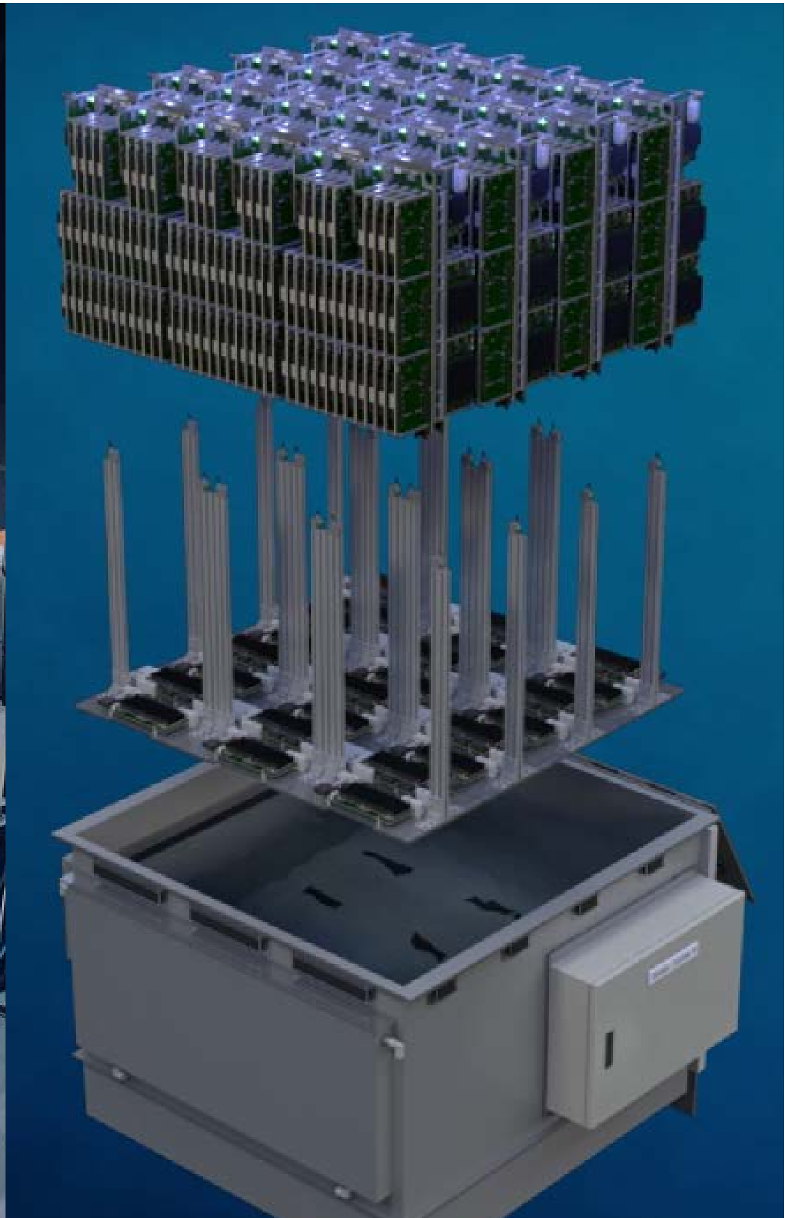

Performance increase in the same $1m^3$ tank

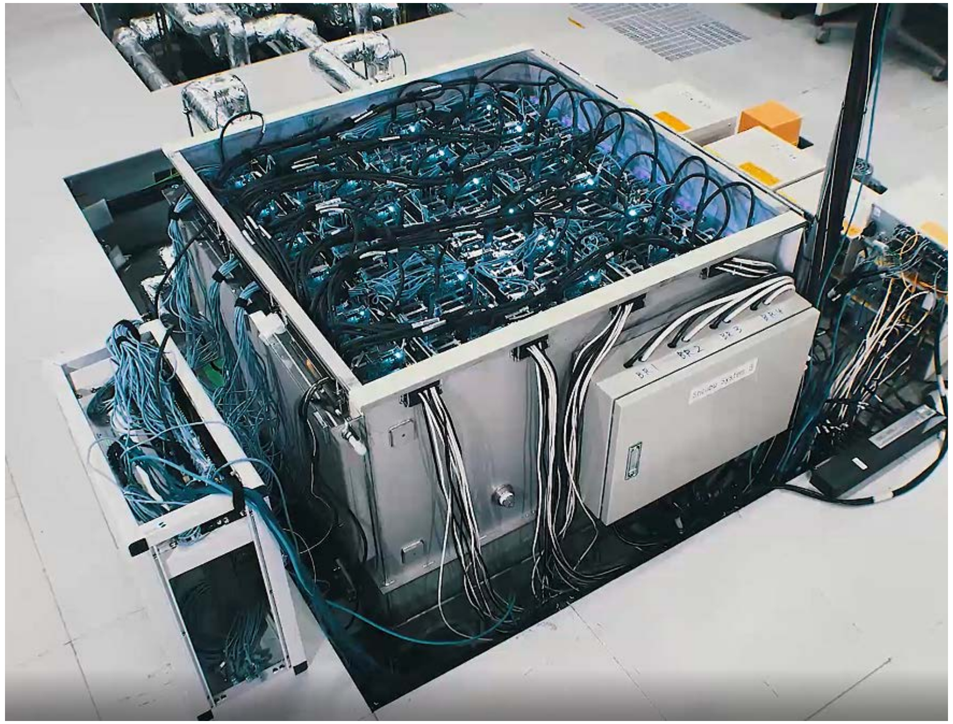

ZettaScaler-1.0: 0.15 PFLOPS/Tank (x0.4) in 2015

ZettaScaler-1.6: 0.25 PFLOPS/Tank (x1) in 2016

ZettaScaler-**2.0**: **1.5** PFLOPS/Tank (**x6**) in 2017

ZettaScaler-**3.0**: **3.0** PFLOPS/Tank (**x12**) in 2019

ZettaScaler-**4.0**: **6.0** PFLOPS/Tank (**x24**) in 2021


World 4th fastest Supercomputer (Nov, 2017)
World 7th fastest Supercomputer (June, 2019, if operated)
Gyoukou (暁光:The light of dawn)

The development applications on PEZY SC2 systems

1. Artificial cerebellum
2. Simulation of Particle system for Tsunami disaster
3. Middleware: FDPS and Formura
4. Genome analysis

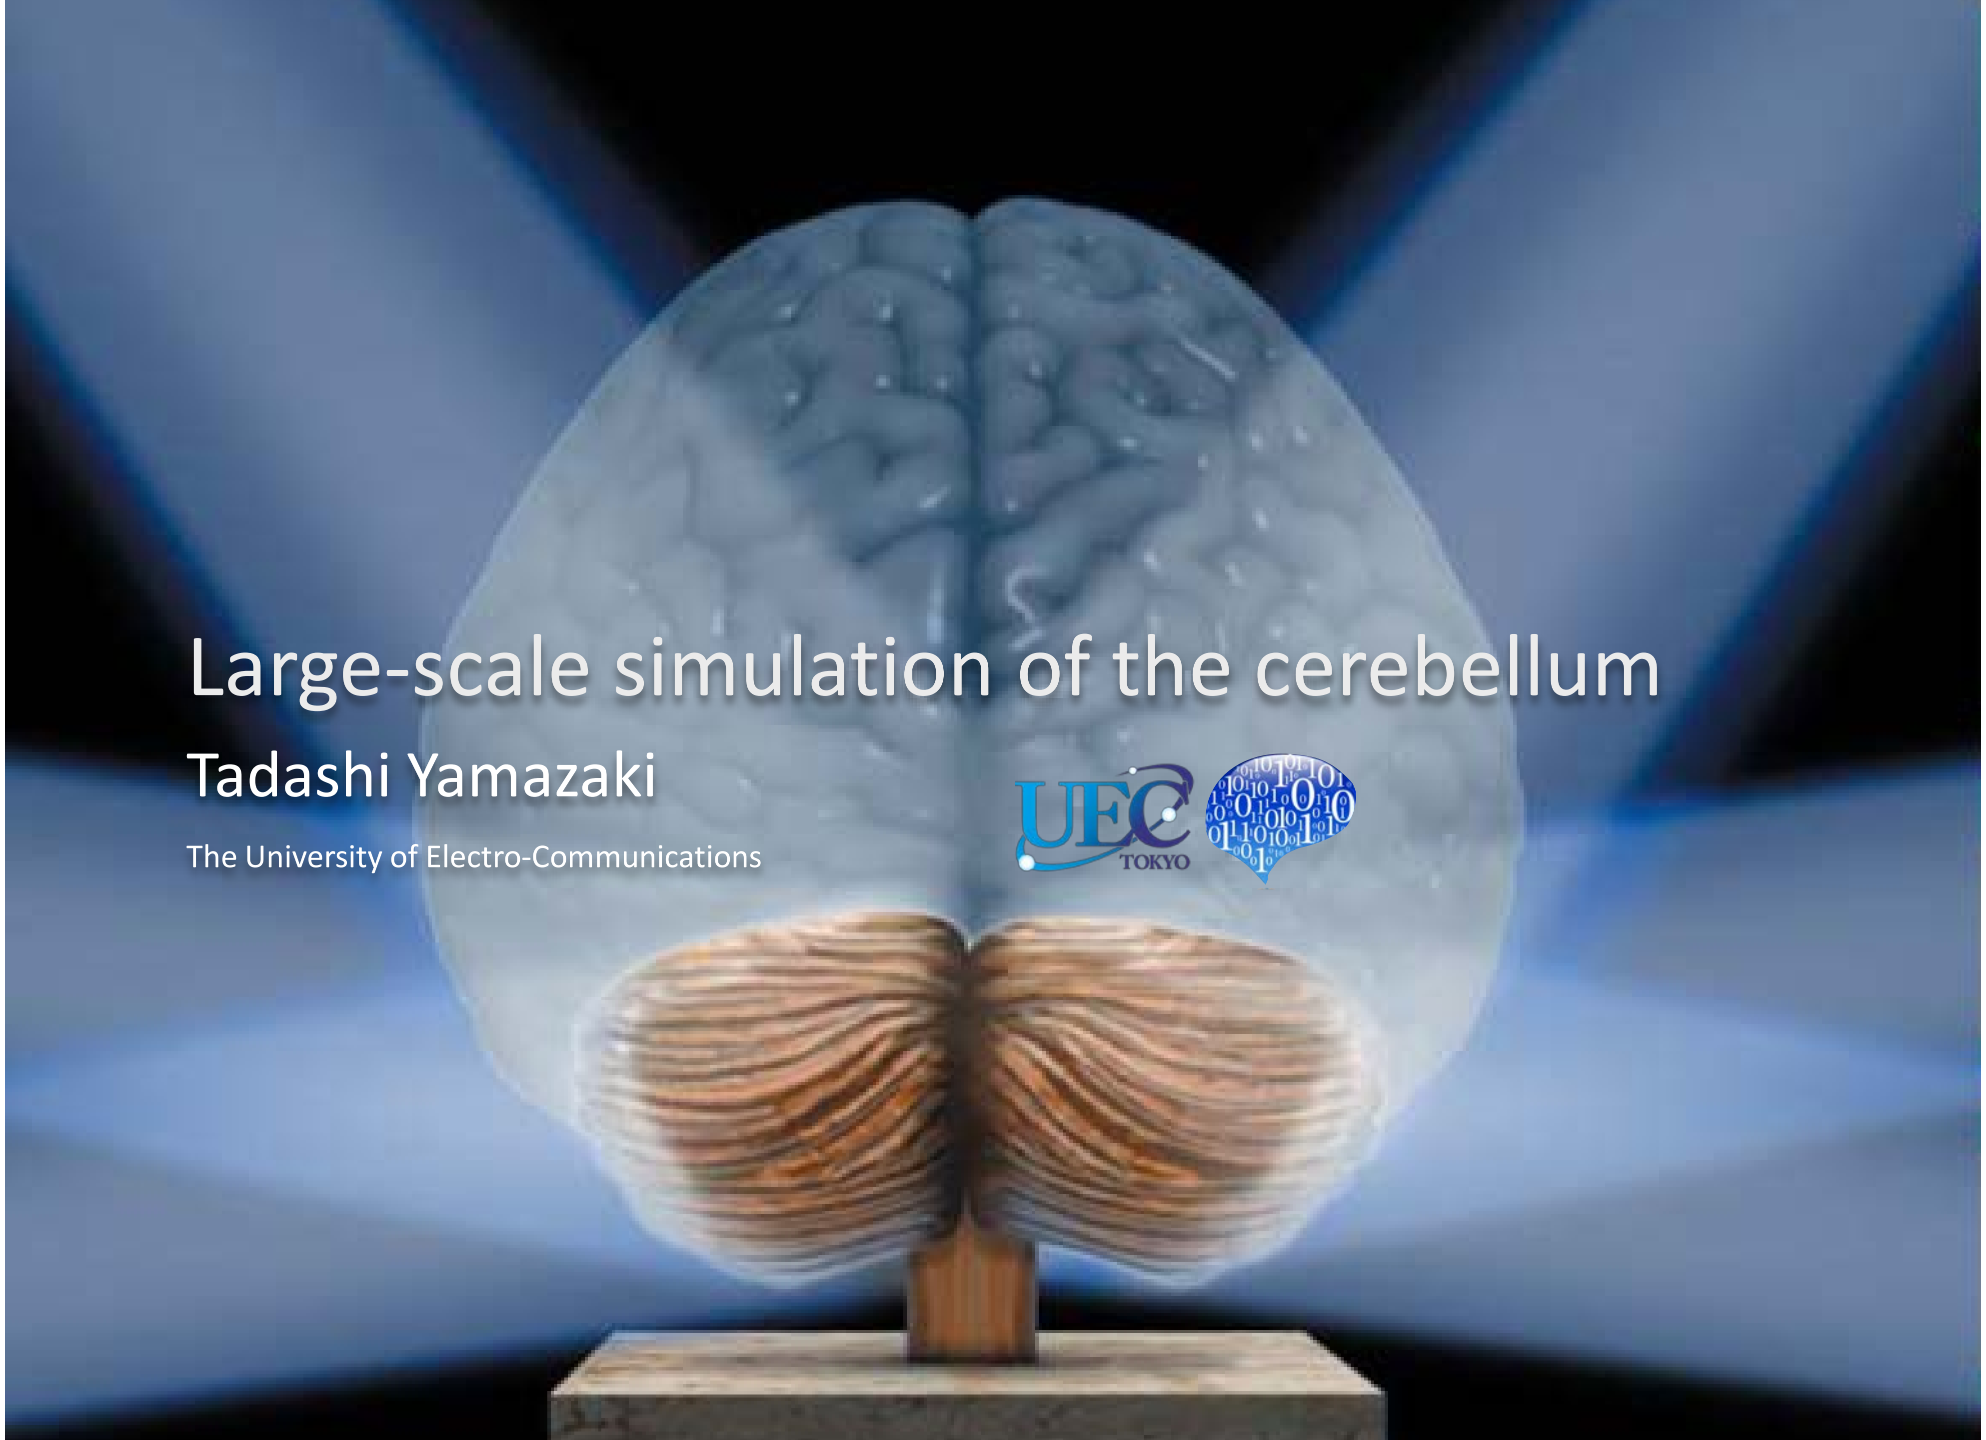

Large-scale simulation of the cerebellum

Tadashi Yamazaki

The University of Electro-Communications

Cerebrum vs Cerebellum

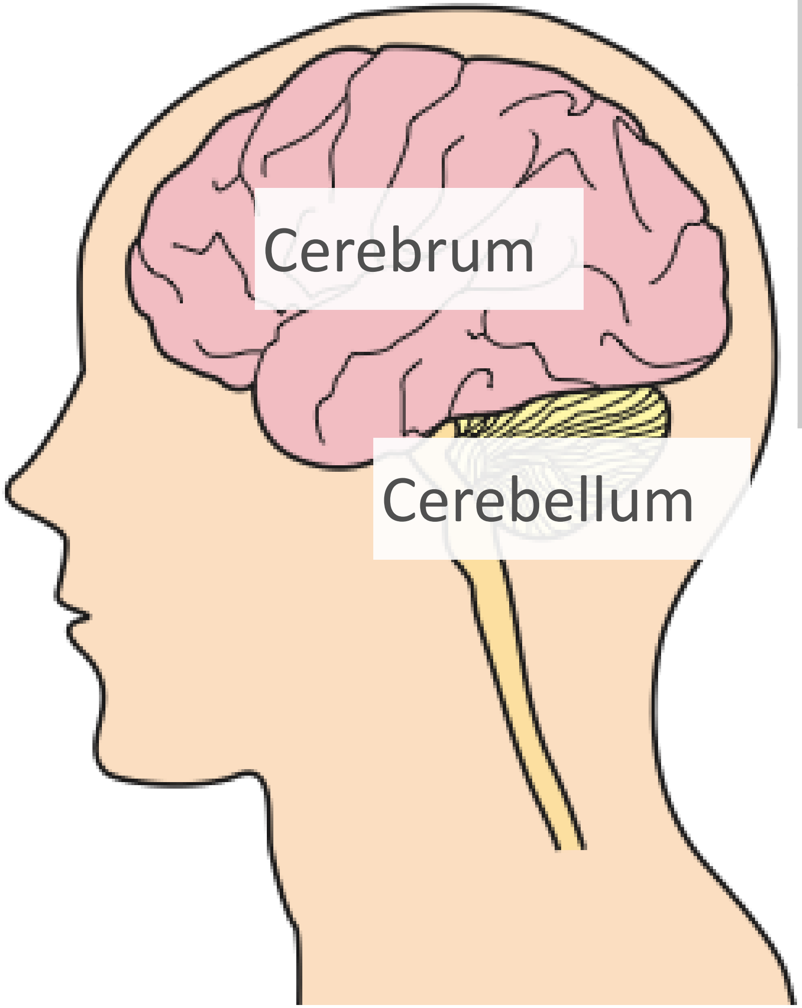

	Cerebrum	Cerebellum
Volume	80%	10%
Neuron number	1.6×10^{10} (19%)	6.9×10^{10} (80%)

Azevedo et al. J Comp Neurol (2009)

Cerebellar microcomplex model

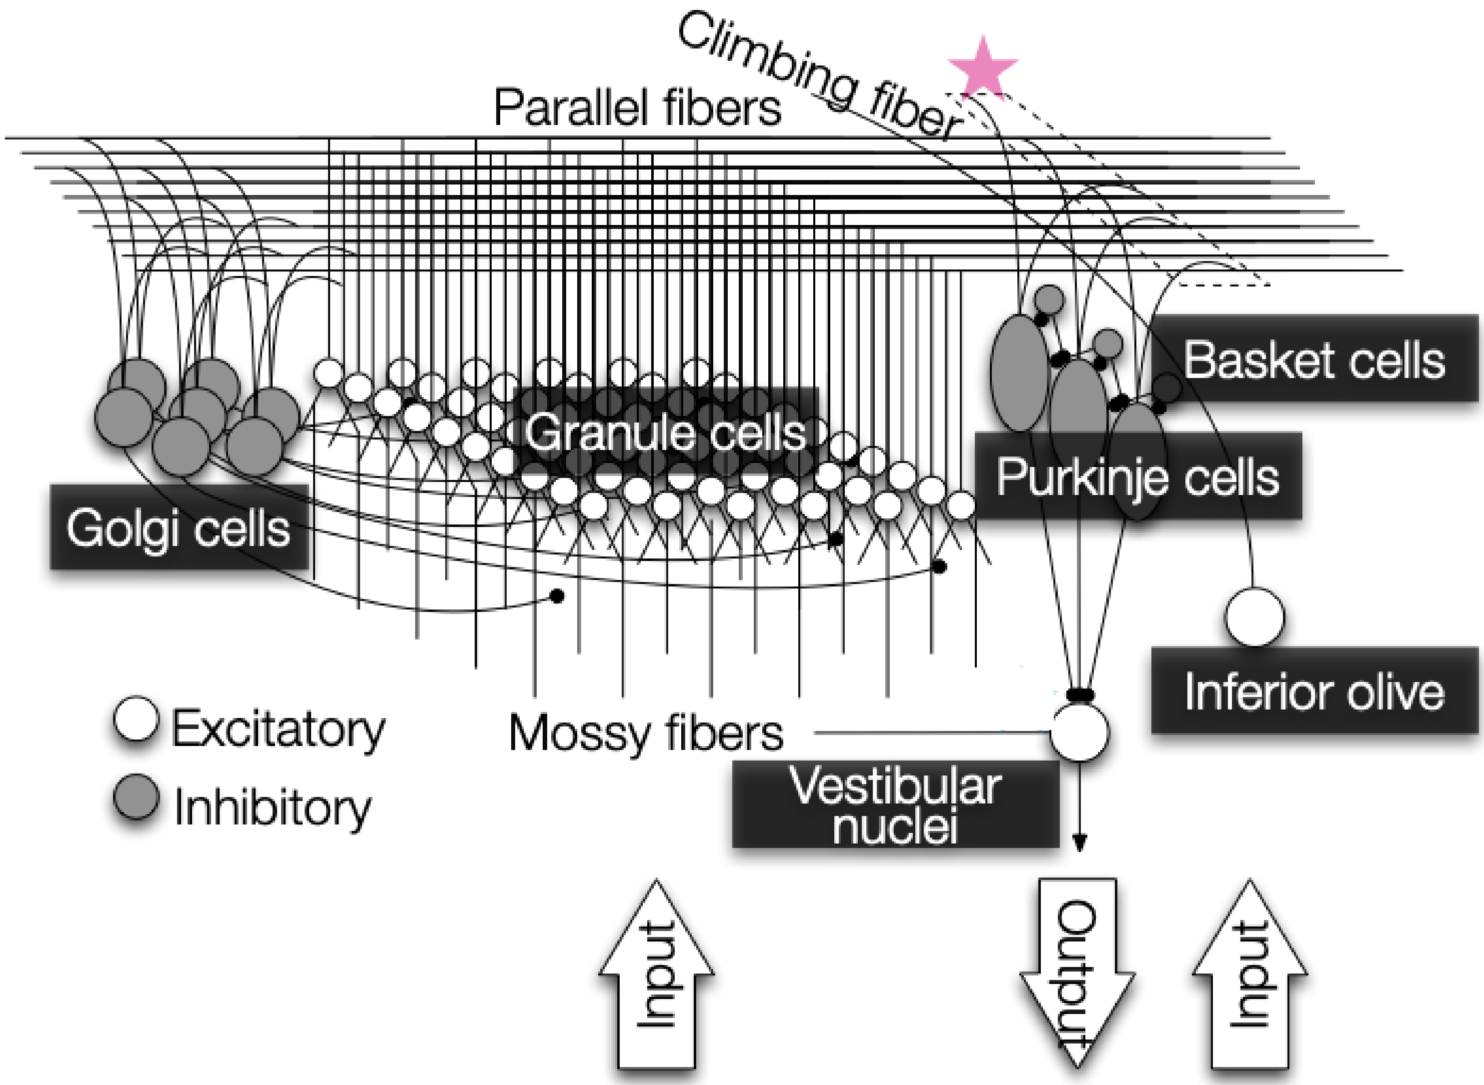

2 x 2 mm^2
1048576 granule cells
1024 GO
32 PKJ
32 BS
1 IO
1 VN
Spiking neuron models

Repeating copy-and-paste the circuit to build the entire cerebellum

Artificial cerebellum on Gyoukou

(Furusho, Yamazaki. In preparation)

Implementing 8 billion (= 8×10^9) spiking neurons

- Comparable with 2 monkeys' cerebella
- Used 7,921 out of 10,000 PEZY-SC2 processors

Realtime simulation

- Simulating cerebellar activity for 1 s within 1 s
- Δt = 1 ms

Online learning

- LTD/LTP at PF-PC synapses

General supervised learning machine

- Reservoir computing (Yamazaki, Tanaka. NN 2007)

Realtime simulation of motor learning

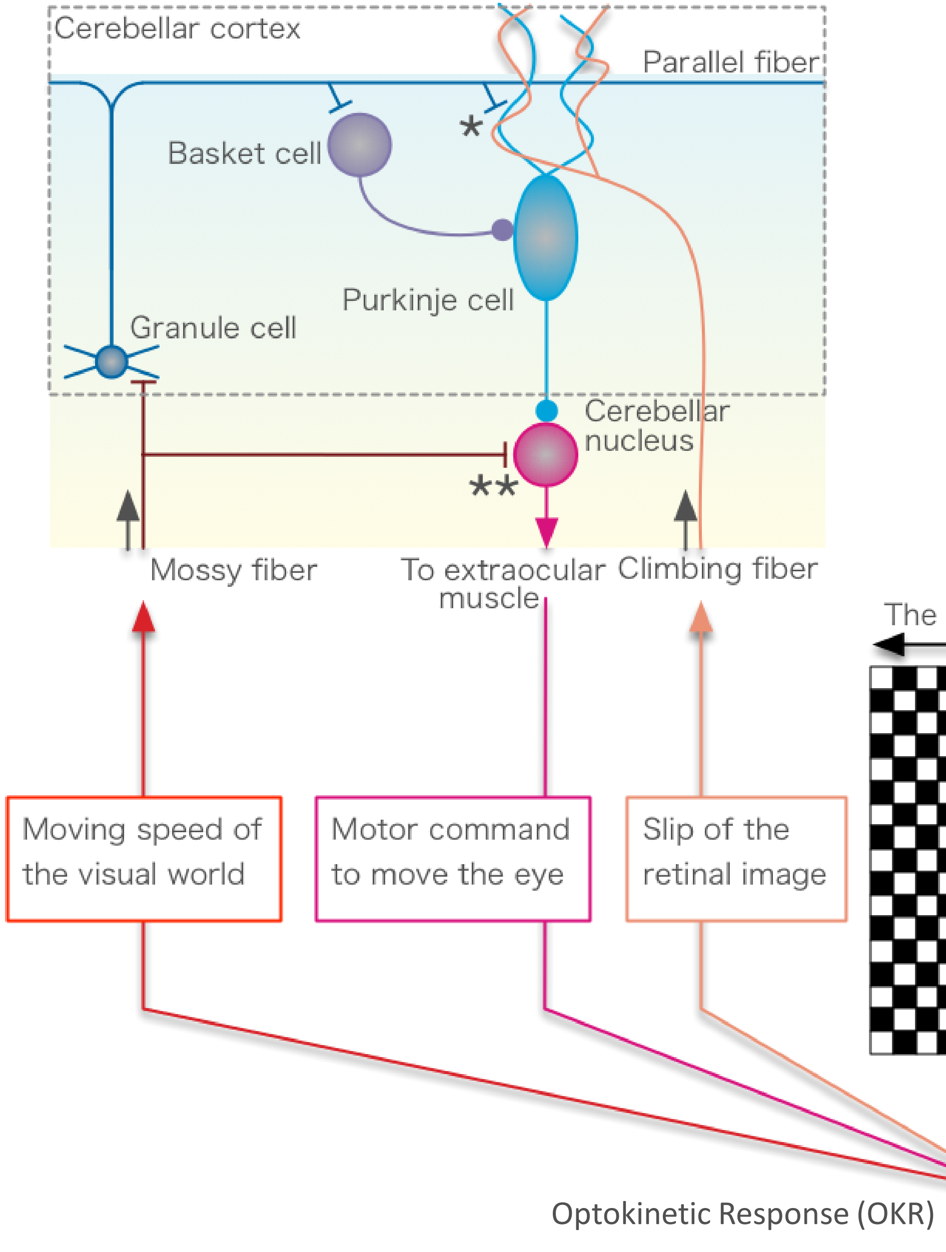

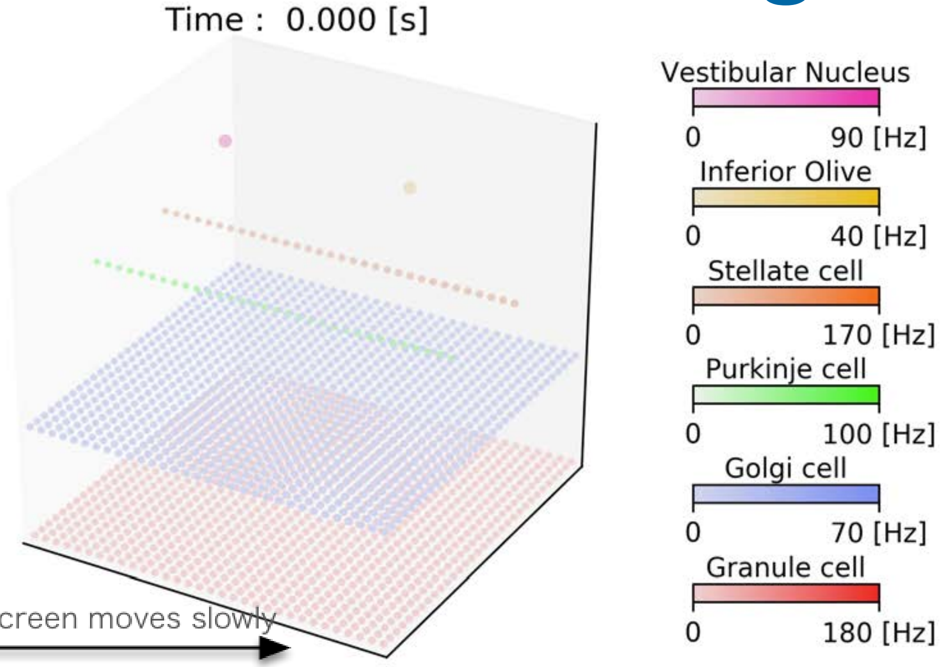

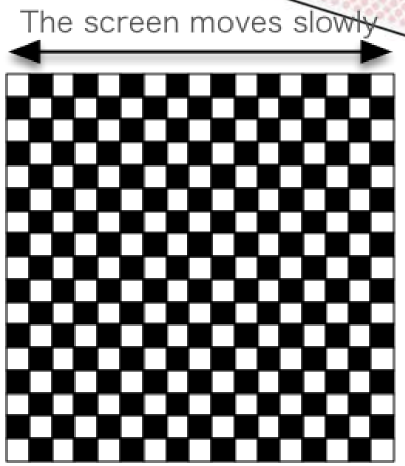

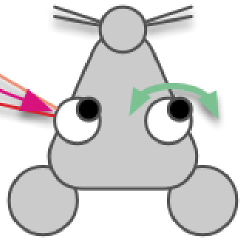

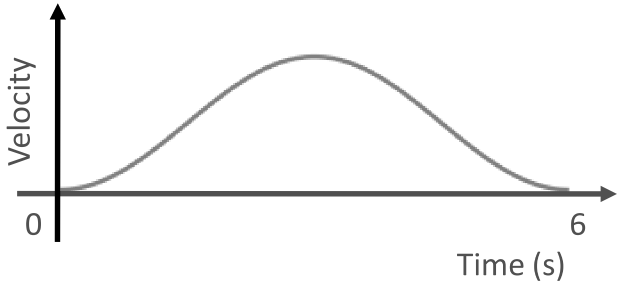

(Furusho, Yamazaki ICANN 2017)

Realtime gain adaptation

(Furusho, Yamazaki. In preparation)

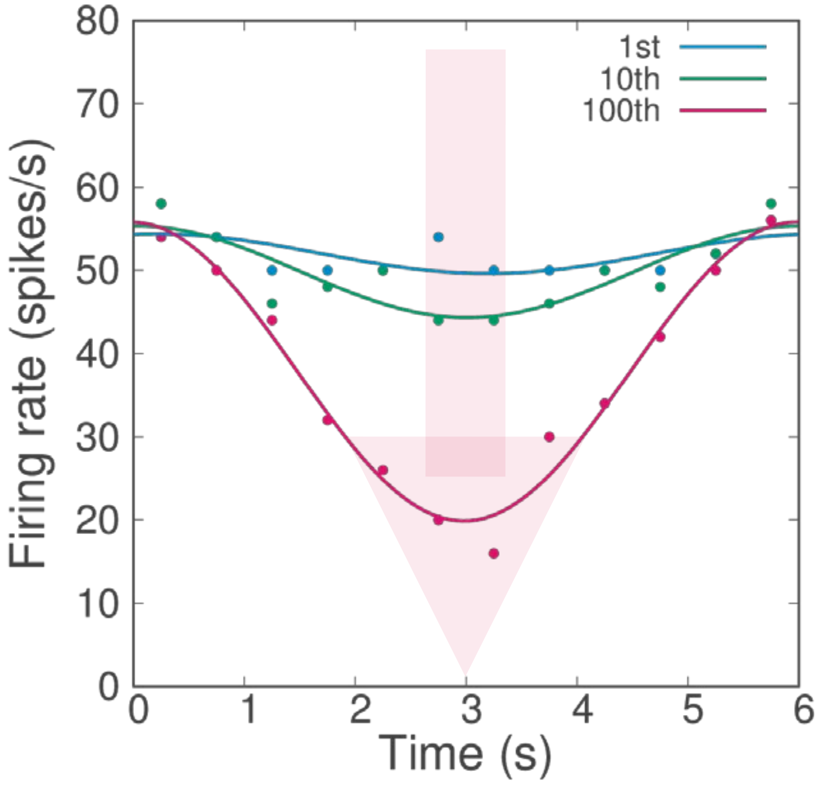

Purkinje cell

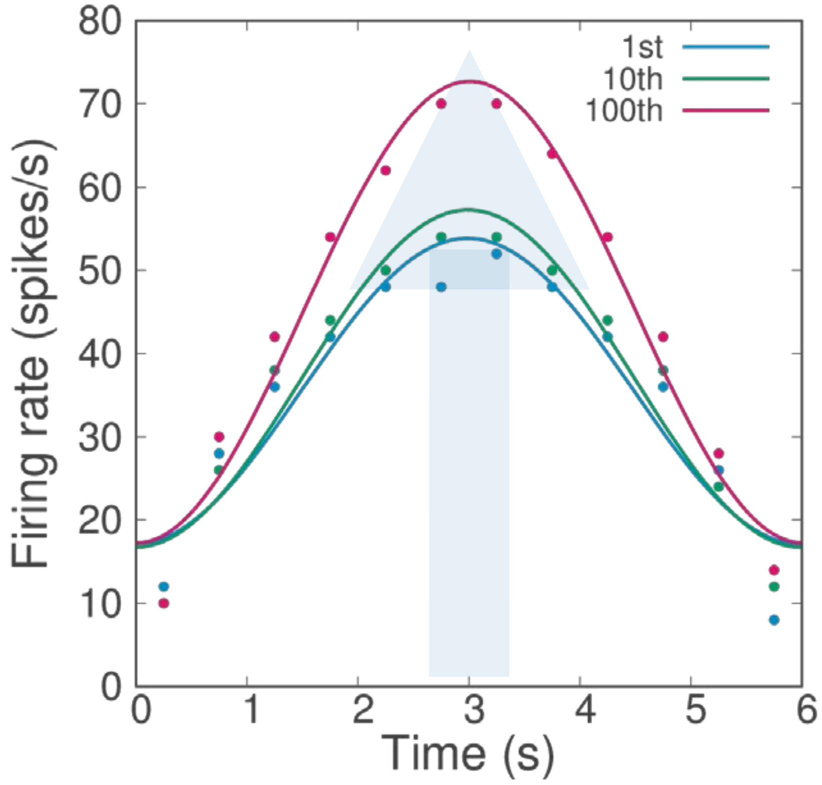

Vestibular nuclei

6 s simulation completes within 4.7 s !

Parallel supervised learning

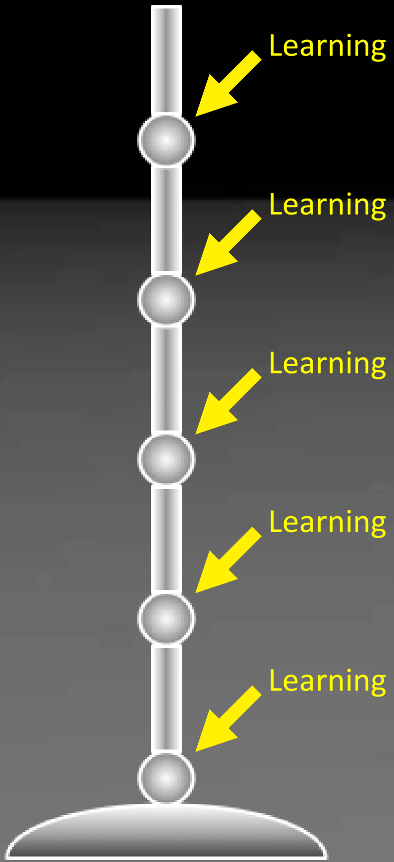

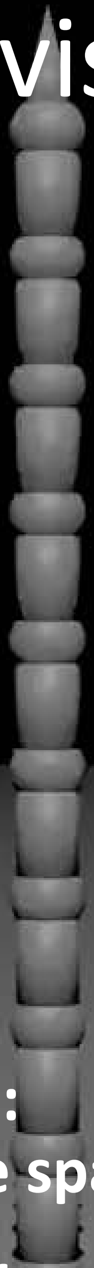

Standard:

State space size: $\vartheta_1 \times \theta_2 \times \theta_3 \times \theta_4 \times \ldots$

This study: (based on Otte et al. 2017)

State space size: $\vartheta_1 + \theta_2 + \theta_3 + \theta_4 + \ldots$

The development applications on PEZY SC2 systems

1. Artificial cerebellum
2. Simulation of Particle system for Tsunami disaster
3. Middleware: FDPS and Formura
4. Genome analysis

The Particle simulation method (SPH and DEM): JAMSTEC

- The **Smoothed Particle Hydrodynamics (SPH)** and **Discrete Element Method (DEM)** offer effective numerical applications of disaster, geodynamics and industrial processing simulations.
- Now, large simulation over 1billion particles is available with dynamic load balancing method.

[e.g. Furuichi and Nishiura, *Comput. Phys. Comm.*, 2017]

- The one of the remaining problems is the energy cost.

<u>One solution is to use the energy efficient supercomputer such as PEZY-SC.</u>

Off-load implementation of SPH and DEM with FDPS

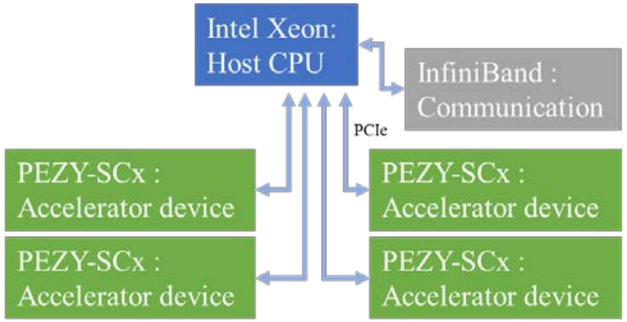

Design of PEZY-SC system

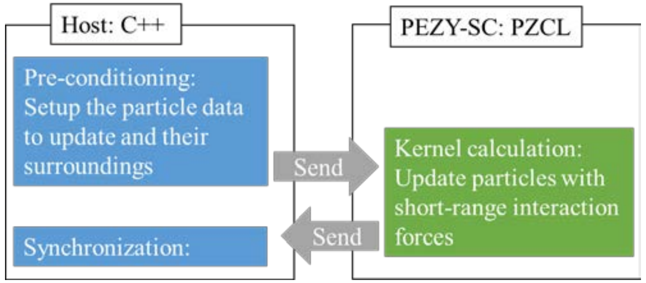

Schematic process of off-load implementation.

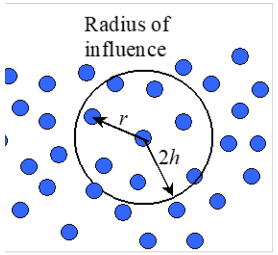

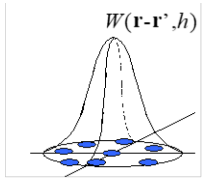

Compact support SPH kernel

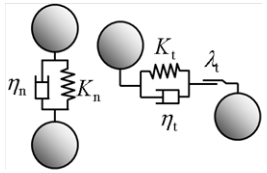

Contact forces of DEM model
(rolling friction is also considered)

※A single particle typically interacts with 64 ～512 and 12 ～256 particles for SPH and DEM, respectively.

・We implemented SPH and DEM with a help of **Framework for Developing Particle Simulator (FDPS)** [Iwasawa et al. 2016]. PZCL (open-CL like) programming is needed only for the kernel calculation and data transfer between the host and accelerators for off-load implementation.

・The **FDPS** supports SPH by default, but not DEM. We customized the FDPS for dealing with tangential forces of DEM.

SPH result by PEZY-SC for water dam break test

*Quintic Kernel, Viscos term of Crealy 1996, EOS of Tait's equation, explicit Euler time step, double precision, 0.4M particles, h = 2.1 l

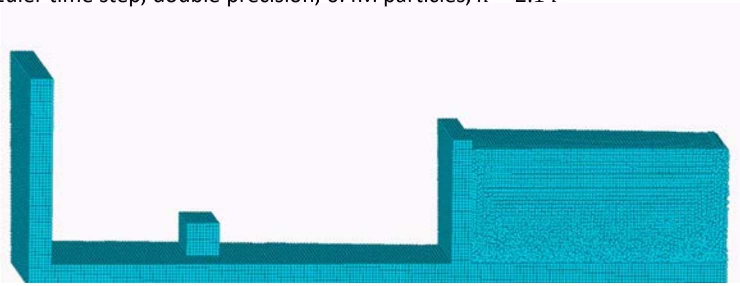

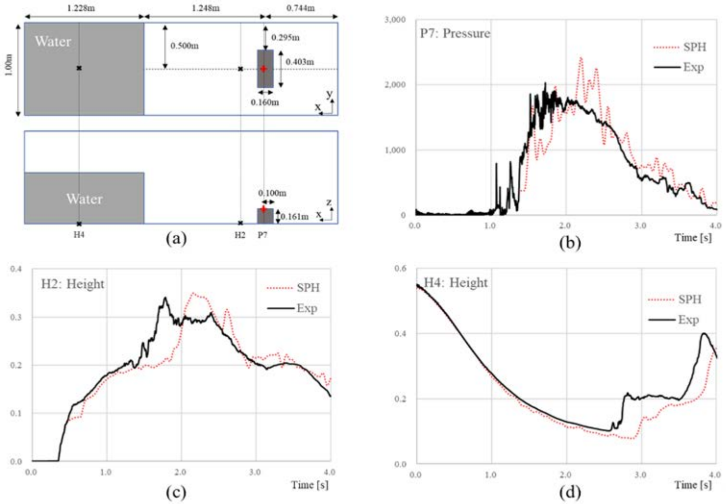

K.M.T. Kleefsman et. al., *J. Comp. Phys.* 206 (2005) 363-393

[Hosono and Furuichi, *ICCES*, 2019]

DEM result by PEZY-SC for power dam break test

* Contact force is Voigt model, which comprises the Hertz–Mindlin model with rolling friction model

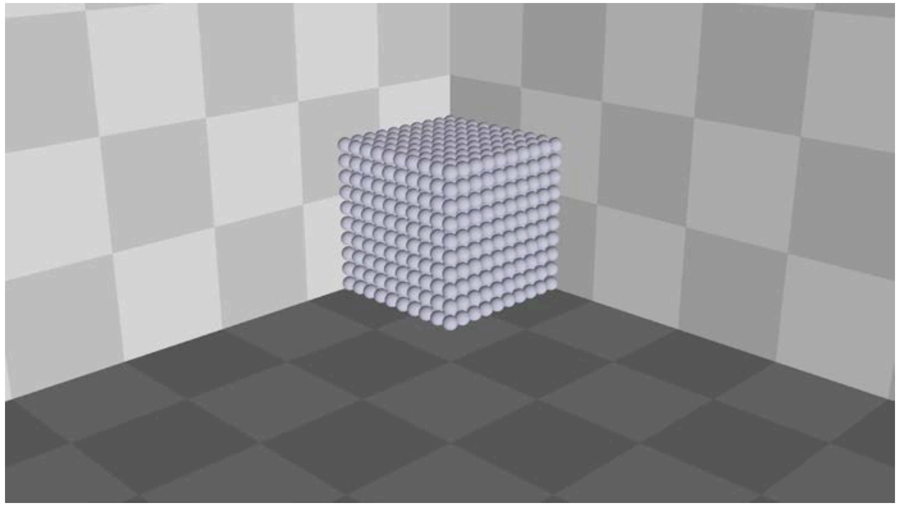

- General purpose SPH and DEM application are available on PEZY-SC!
- The FDPS based source code will be shared with GitHub.

The development applications on PEZY SC2 systems

1. Artificial cerebellum
2. Simulation of Particle system for Tsunami disaster
3. Middleware: FDPS and Formura

    Domain specific language for particle and lattice type codes.

4. Genome analysis

2. Middleware Development

- **Support for application developments for heterogenius many-core systems**
  - **Automatize tedious parallelization with frameworks**
    - MPI, OpenMP, SIMD
    - Effective use of Cache
    - Domain decomposition, load balance

Application developers use Middleware

- **Past/present achievements**
  - **FDPS** : Particle-based. High efficiency on both multi-core and accelerator systems: **Load balancing**
  - **Formura**: Grid. Automatically generates codes with temporal blocking: **Stencil calculation**

Kobe-U
Framework and DSL for large-scale simulation on PEZY-based systems

FDPS (**F**ramework for **D**eveloping **P**article **S**imulator)

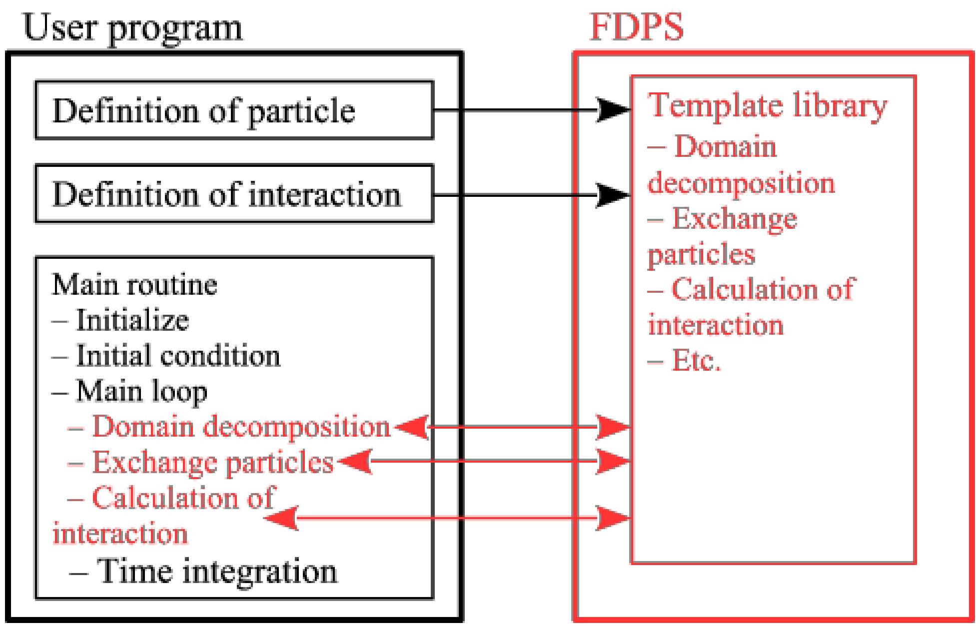

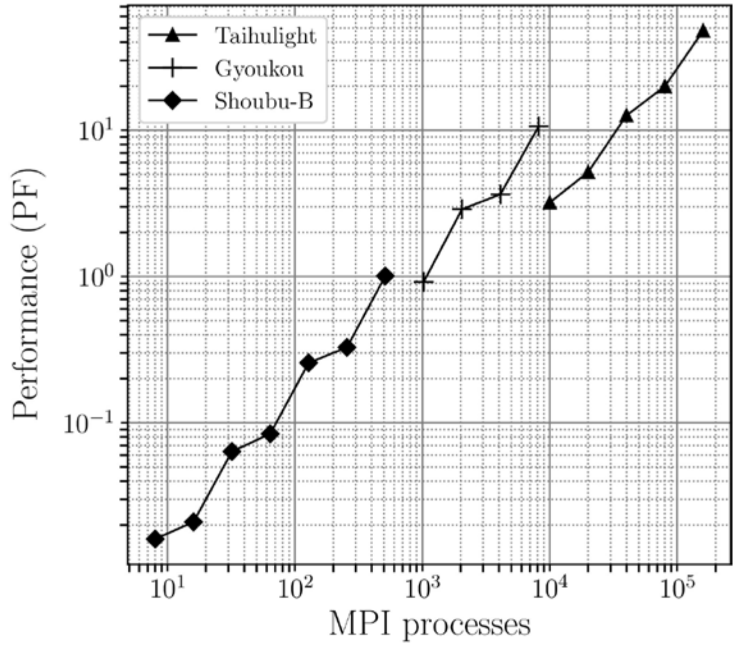

- Framework for particle-based simulations
- Generates highly scalable and efficient library functions from particle data strructure and interaction functions
- Generated code runs can use OpenMP, MPI, accelerators (GPUs and PEZY-SC)
- Simulation of planetary ring achieved the efficiency of 40% on Shoubu System B(512-node PEZY-SC2, efficiency on GYOUKOU is lower because it was sterned off just after our project started)

Kobe-U
Framework and DSL for large-scale simulation on PEZY-based systems

Formura (DSL for stencil computing)

- DSL for stencil computing
- Generates highly scalable and efficient code for explicit stencil calculation from high-level description of the numerical scheme
- Efficient practical implementation of temporal blocking
- Simulation of isotropic turbulence achieved the efficienct of 21% on Shoubu System B. 1.7x performance improvement by temporal blocking

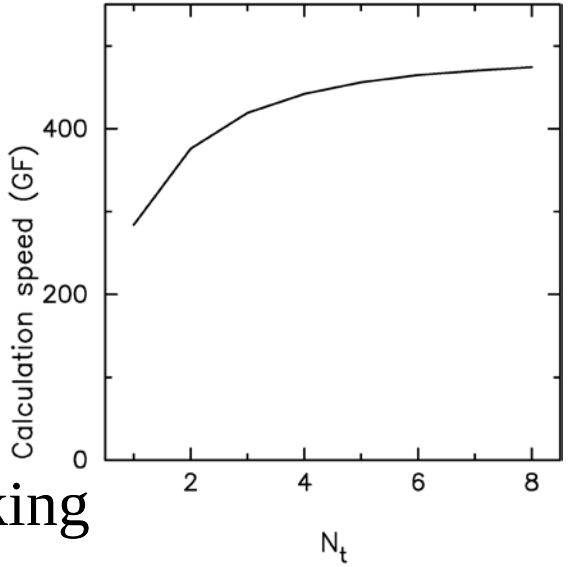

Input PDE to formura:

```
r[t,x,y,z]_t = -u[t,x,y,z]*r[t,x,y,z]_x - v[t,x,y,z]*r[t,x,y,z]_y
         - w[t,x,y,z]*r[t,x,y,z]_z
         - r[t,x,y,z]*(u[t,x,y,z]_x + v[t,x,y,z]_y + w[t,x,y,z]_z)
u[t,x,y,z]_t = -u[t,x,y,z]*u[t,x,y,z]_x - v[t,x,y,z]*u[t,x,y,z]_y
          - w[t,x,y,z]*u[t,x,y,z]_z
          - p[t,x,y,z]_x/r[t,x,y,z] + c*vis1[t,x,y,z]/r[t,x,y,z]
(v,w omitted)
p[t,x,y,z]_t = -u[t,x,y,z]*p[t,x,y,z]_x - v[t,x,y,z]*p[t,x,y,z]_y
               - w[t,x,y,z]*p[t,x,y,z]_z
          - gm*p[t,x,y,z]*(u[t,x,y,z]_x + v[t,x,y,z]_y + w[t,x,y,z]_z)
          - c2*(u[t,x,y,z]*vis1[t,x,y,z] + v[t,x,y,z]*vis2[t,x,y,z]
          + w[t,x,y,z]*vis3[t,x,y,z])
```

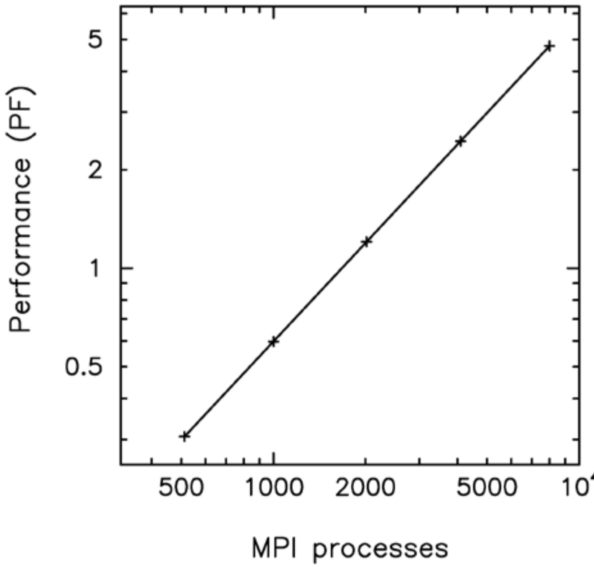

The development applications on PEZY SC2 systems

1. Artificial cerebellum
2. Simulation of Particle system for Tsunami disaster
3. Middleware: FDPS and Formura
4. Genome analysis

Full genome analysis for one human
Traditional System (e.g. Xeon)

Genome analysis workflow	Standard software	Process time (hour)
Raw reads from a sequencer		
Quality Control	Trimmomatic	
Alignment	BWA	2~4 hour
Mark duplicate	Picard MarkCuplicates	1~2 hour
Base quality score recalibrate	GATK: Base Recalibrator&ApplyBQSR	1~2 hour
Variant call	GATK: Haplotype Caller	2~6 hour

Total: ~14 hour

BWA-mem comparison

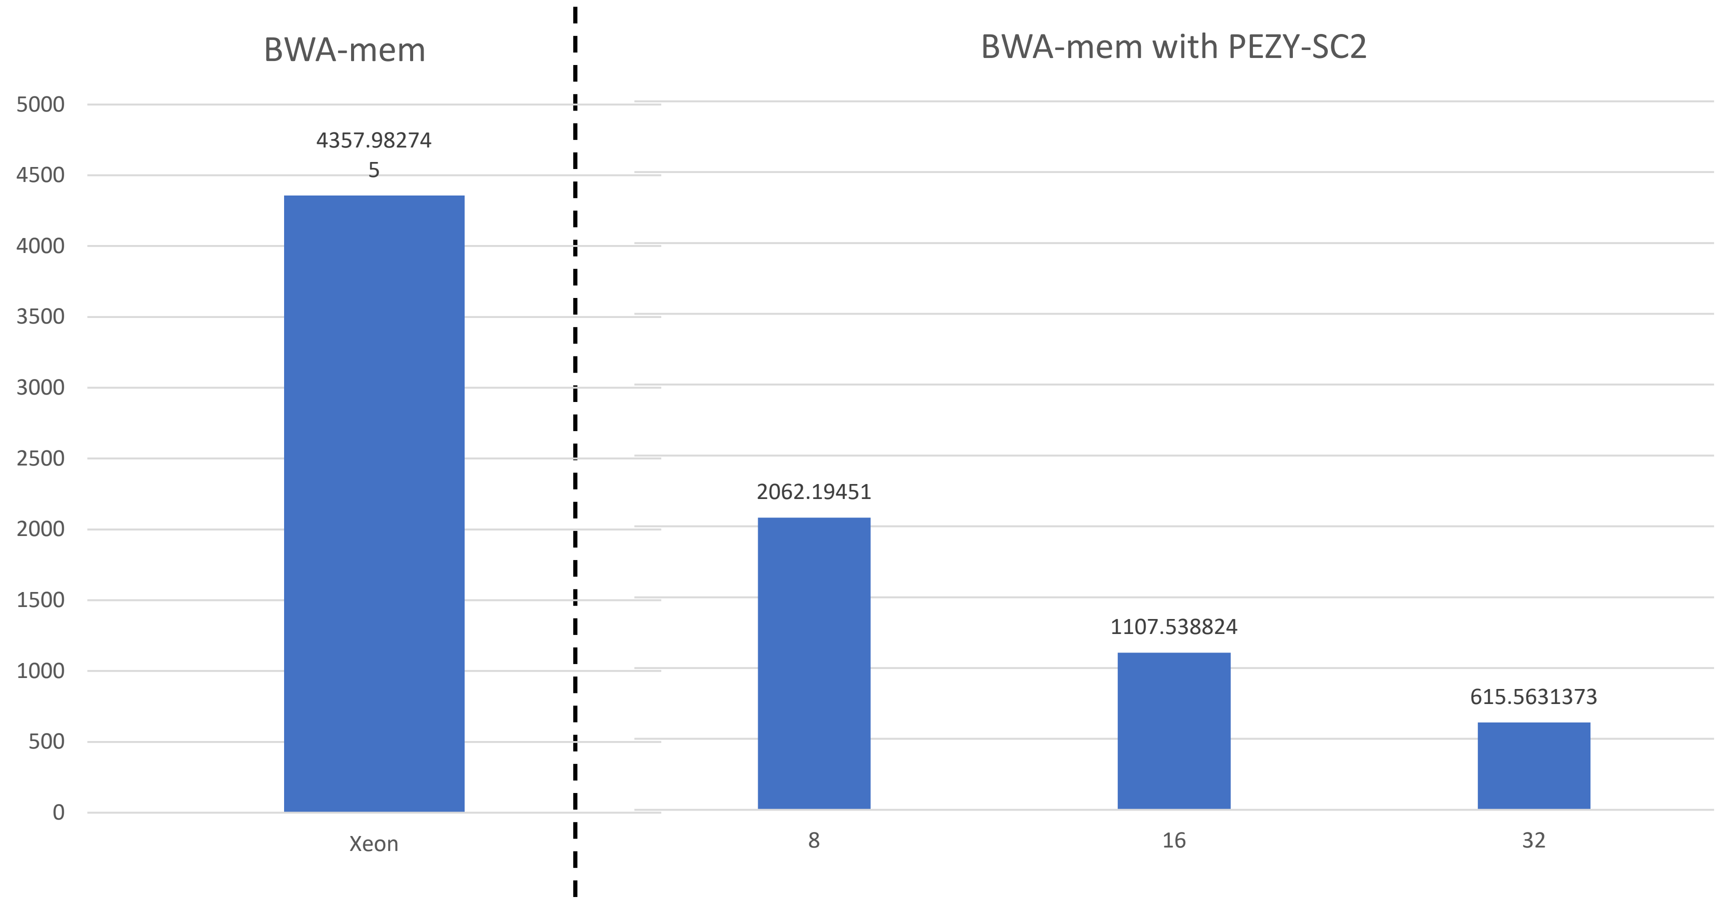

PEZY SC2 system

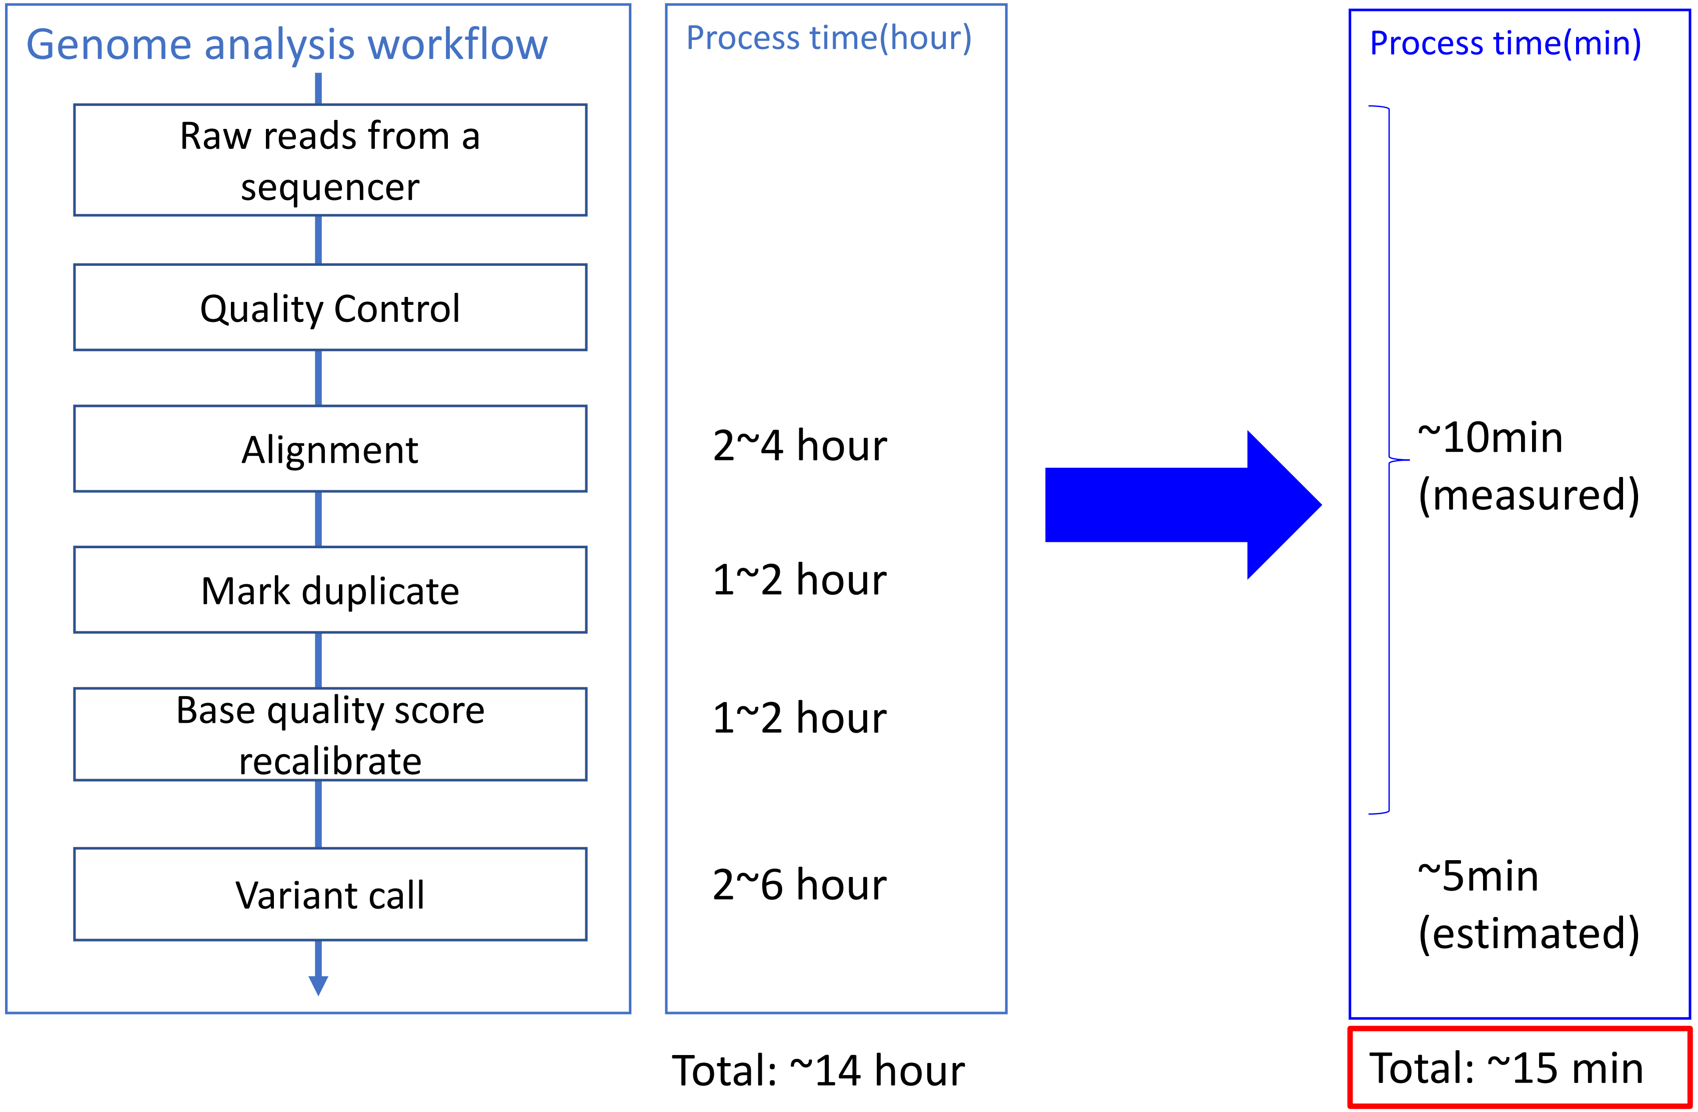

Human Genome Analysis Summary

- We have already ported alignment, mark duplicate and BQSR to PEZY-SC2.
- Aligner is BWA-mem.
  - BWA-mem performance is pretty good on our system.
- We are porting Variant Call now and will complete it within a month or so.
  - Based on GATK: HaplotypeCall.

- Quad PEZY-SC3 single board system will Complete whole genome analysis within 15 min.

Future (Next Year)

PEZY-SCx Processor History and Plan

	2012	2014	2017		2019
Processor	PEZY-1	PEZY-SC	**PEZY-SC2**	scale	**PEZY-SC3**
Process Node	**40 nm**	**28 nm**	**16 nm**		**7 nm**
Die Size	118 mm^2	412 mm^2	620 mm^2	1.3	780 mm^2
Core Numbers	**512**	**1,024**	**2,048**	**2**	**4,096**
Core Voltage	1.2 V	0.9 V	0.8 V		0.7 V
Clock Frequency	433 MHz	533 MHz	700 MHz	1.7	1,200 MHz
DRAM-IO	DDR3	DDR4	DDR4		HBM2
Memory Bandwidth	48 GB/s	51 GB/s	77 GB/s	15.6	1,200 GB/s
Interface Bandwidth	12 GB/s	24 GB/s	32 GB/s	3	96 GB/s
DP Performance	**0.3 TFLOPS**	**0.9 TFLOPS**	**2.7 TFLOPS**	**7.3**	**19.7 TFLOPS**
Power Consumption	60 W	100 W	130 W	3.8	500 W
Power Efficiency	5.0 GFLOPS/w	9.0 GFLOPS/w	20.8 GFLOPS/w	1.9	38.4 GFLOPS/w
System Efficiency	**-**	**6.7 GFLOPS/w**	**17.6 GFLOPS/s**	**1.9**	**31.2 GFLOPS/s**

Development of PEZY-SC3 has been completed and is now ready for tape out

PEZY-SC3 will be taped out by the end of 2019 and will be in volume production from mid 2020

PEZY-SC3 will be the world biggest processor ever with highest performance and efficiency

NVIDIA Voltas

	2017年		2019年		2021年
Processor	**PEZY-SC2**	scale	**PEZY-SC3**	scale	**PEZY-SC4**
Process Node	**16 nm**		**7 nm**		**5 nm**
Die Size	620 mm^2	1.3	780 mm^2	1.0	780 mm^2
Core Numbers	**2,048**	**2**	**4,096**	**1.5**	**6,144**
Core Voltage	0.8 V		0.7 V		0.6 V
Clock Frequency	700 MHz	1.7	1,200 MHz	1.2	1.400 MHz
DRAM-IO	DDR4		HBM2		HBM3
Memory Bandwidth	77 GB/s	15.6	1,200 GB/s	1.7	2,000 GB/s
Interface Bandwidth	32 GB/s	3	96 GB/s	3	192 GB/s
DP Performance	**2.7 TFLOPS**	**7.3**	**19.7 TFLOPS**	**1.9**	**36.9 TFLOPS**
Power Consumption	130 W	3.8	500 W	1.3	640 W
Power Efficiency	20.8 GFLOPS/w	1.9	38.4 GFLOPS/w	1.5	53.4 GFLOPS/w
System Efficiency	**17.6 GFLOPS/s**	**1.9**	**31.2 GFLOPS/s**	**1.6**	**48.9 GFLOPS/s**

2017年		2020年
Volta (V100)	scale	**Volta2 (V200)**
10 nm		**5 nm**
720 mm^2	1.0	720 mm^2
5,120	**1.6**	**8,192**
-		-
-		-
HBM2		HBM2
900 GB/s	1.3	1,200 GB/s
32 GB/s	2	64 GB/s
7.8 TFLOPS	**1.6**	**12.5 TFLOPS**
300 W	1.3	400 W
26.0 GFLOPS/w	1.2	31.3 GFLOPS/w
15.1 GFLOPS/s	**1.2**	**18.1 GFLOPS/s**

PEZY-SC4 will be produced in late 2021 with TSMC **5nm** process
PEZY-SC4 will be the world first processor to enable **ExaFLOPS** system

Volta2 (V200)
specs are estimated ones
and not confirmed yet

Volta2 (V200) will be produced in 2020, but will not be able to exceed even **PEZY-SC3** both in performance and power efficiency

PEZY-SC4 will have 4 times more performance and **2.7** times more system efficiency than **Volta2 (V200)** at half cost

PEZY-SC3

20.0 TFLOPS performance
: The first processor with over 10 TFLOPS (DP)

786mm^2 of monolithic silicon die size
: about 8 times bigger than typical Intel x86 CPUs

40.8 B transistors integrated
: about 2 times more transistor numbers than 2nd biggest processor (NVIDIA **V100**: 21.0B)

Comparison between other competitors

Fujitsu **A64FX**:	8.8B
Tesla **FSD**:	6.0B

In Japan, there are only 2 teams who can develop 7nm based large processor and less than10 teams all over the world

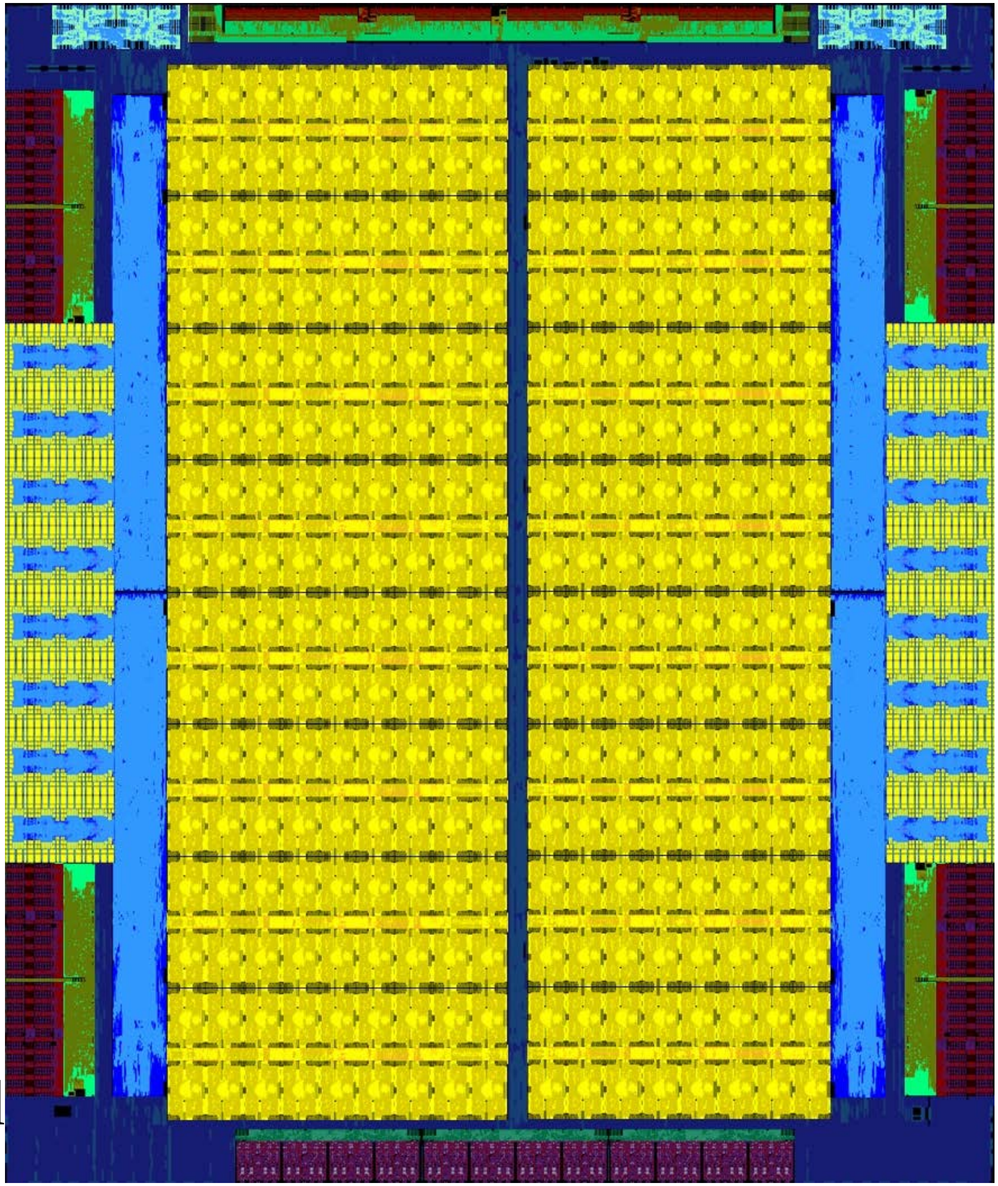

PEZY-SC3 20TF, 40.8B

Compare to FSD, the latest Tesla autopilot AI processor

PEZY-SC3
:786mm^2, 40.8B transistors

FSD die size is only one third of PEZY-SC3
Transistor number of FDS is one seventh

Tesla revealed the next generation
original autopilot AI processor, FSD

Silicon Die Size：3.0x
Transistor Number：
6.8x

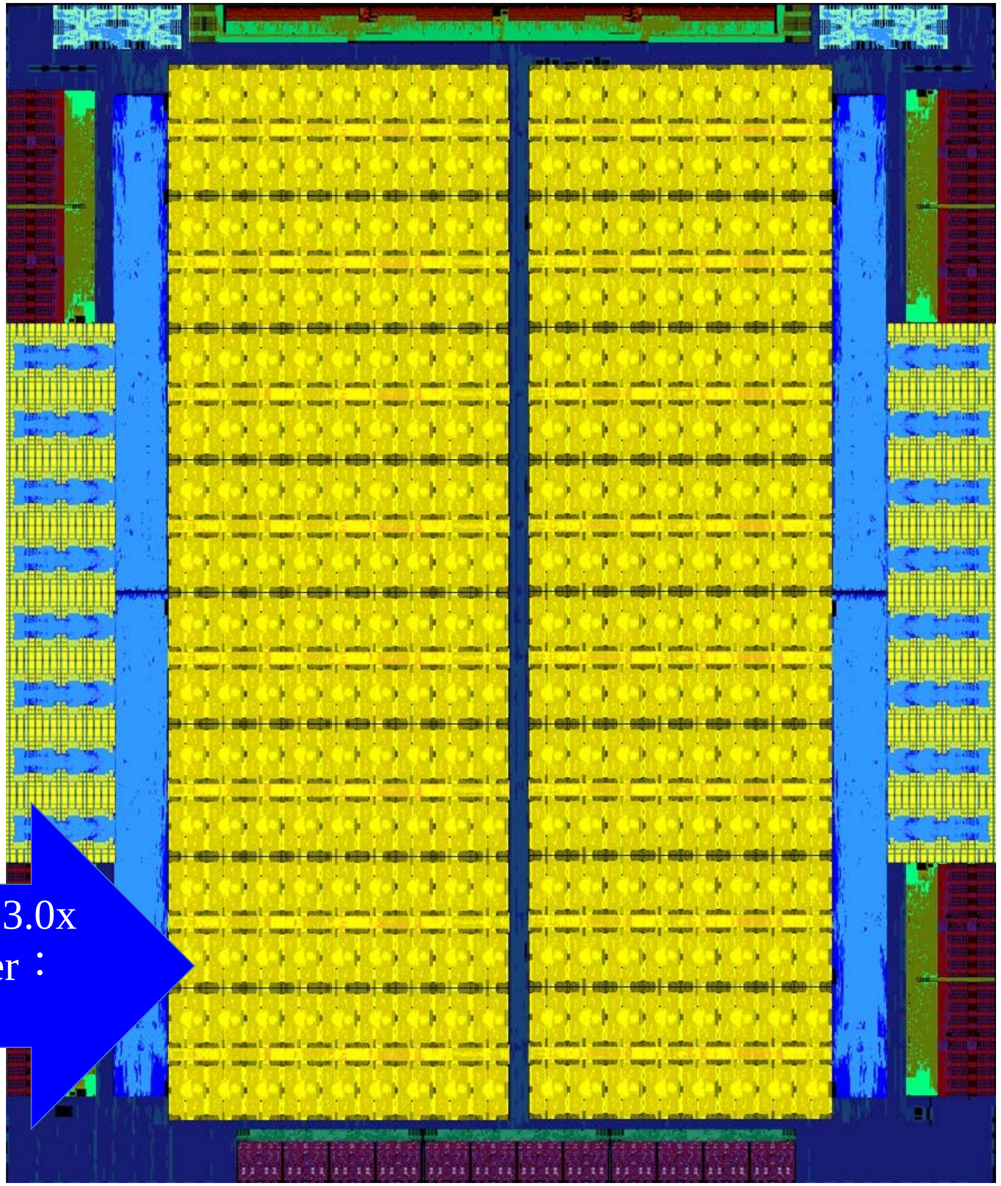

PEZY-SC3 20TF, 40.8B

Proprietary system board design
PEZY-SC3 x4 + AMD EPYC2 (Rome)

manufactured engineering sample board and modules for PEZY-SC3

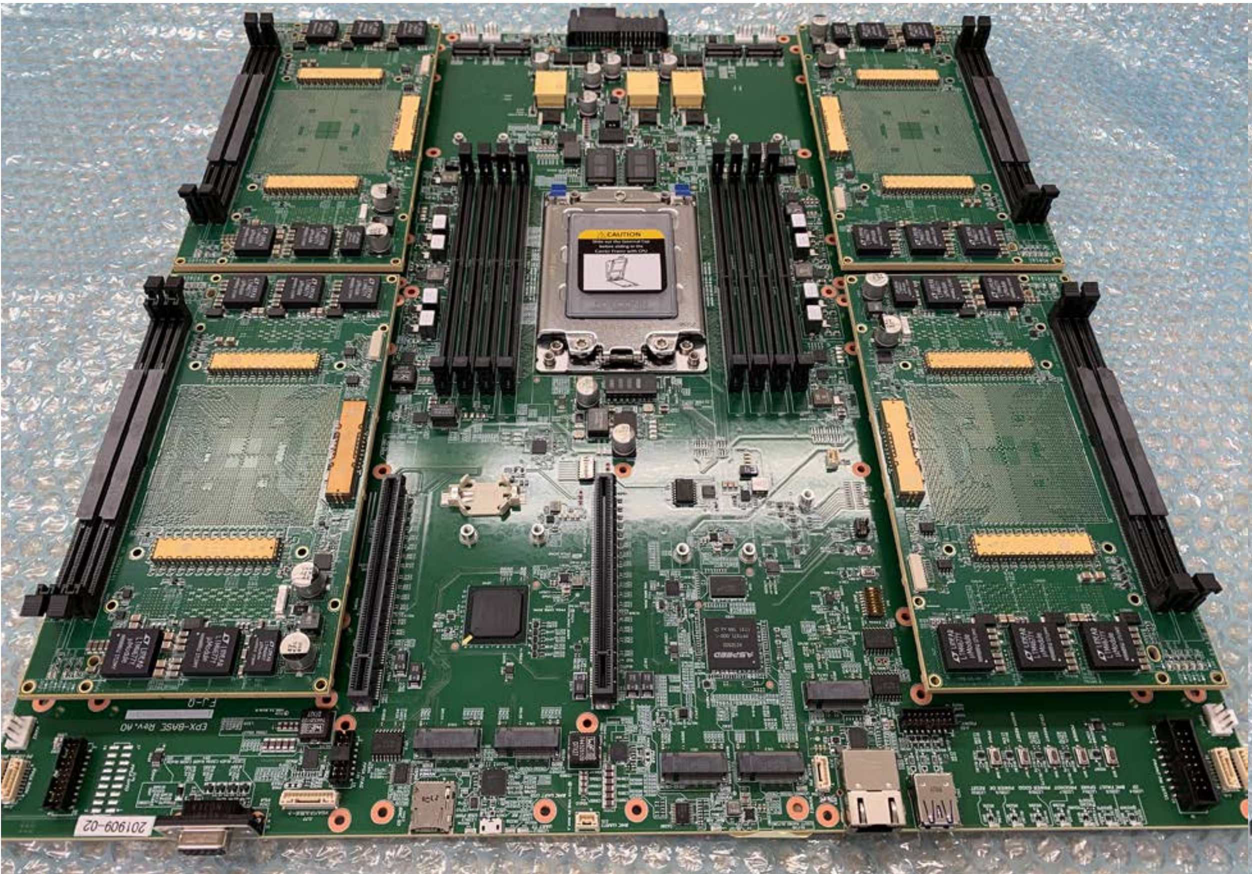
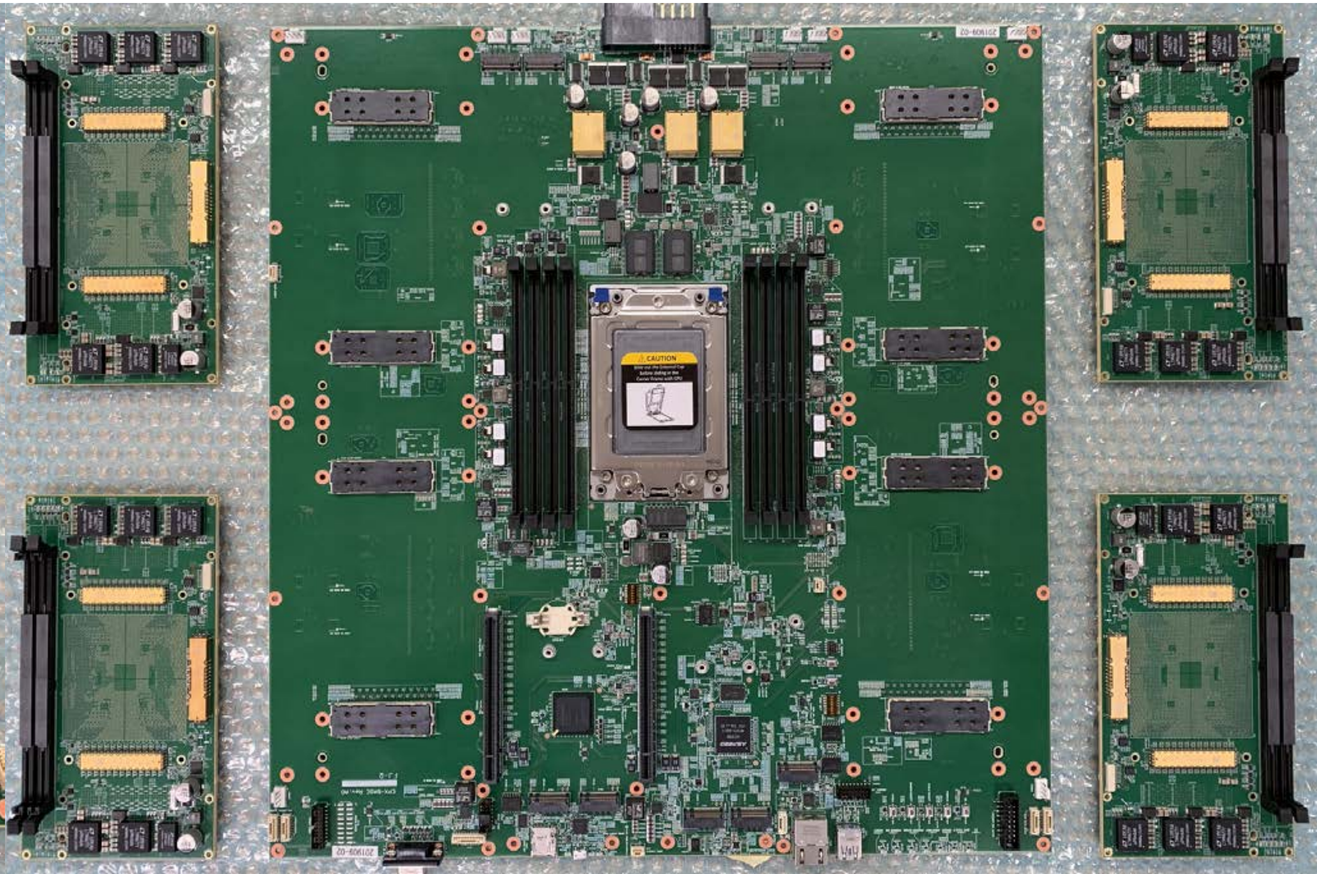

Estimated ZettaScaler-3.0 Specs

tanks configuration will provide about 100 PetaFLOPS (Rmax) and only consumes 4MW with the system cost of around $100M

ZettaScaler-3.0 single tank will have 40 nodes, 40 AMD EPYC2 (64 core), 160 of PEZY-SC3 with 48DC power

Single tank will have 3.2 PetaFLOPS (Rpeak) and 2.4 PetaFLOPS (Rmax) of DP performance

System power efficiency will be 30 GFLOPS/W or so and single tank requires 100kW range power

Liquid Immersion Cooling Tank (40 Node)

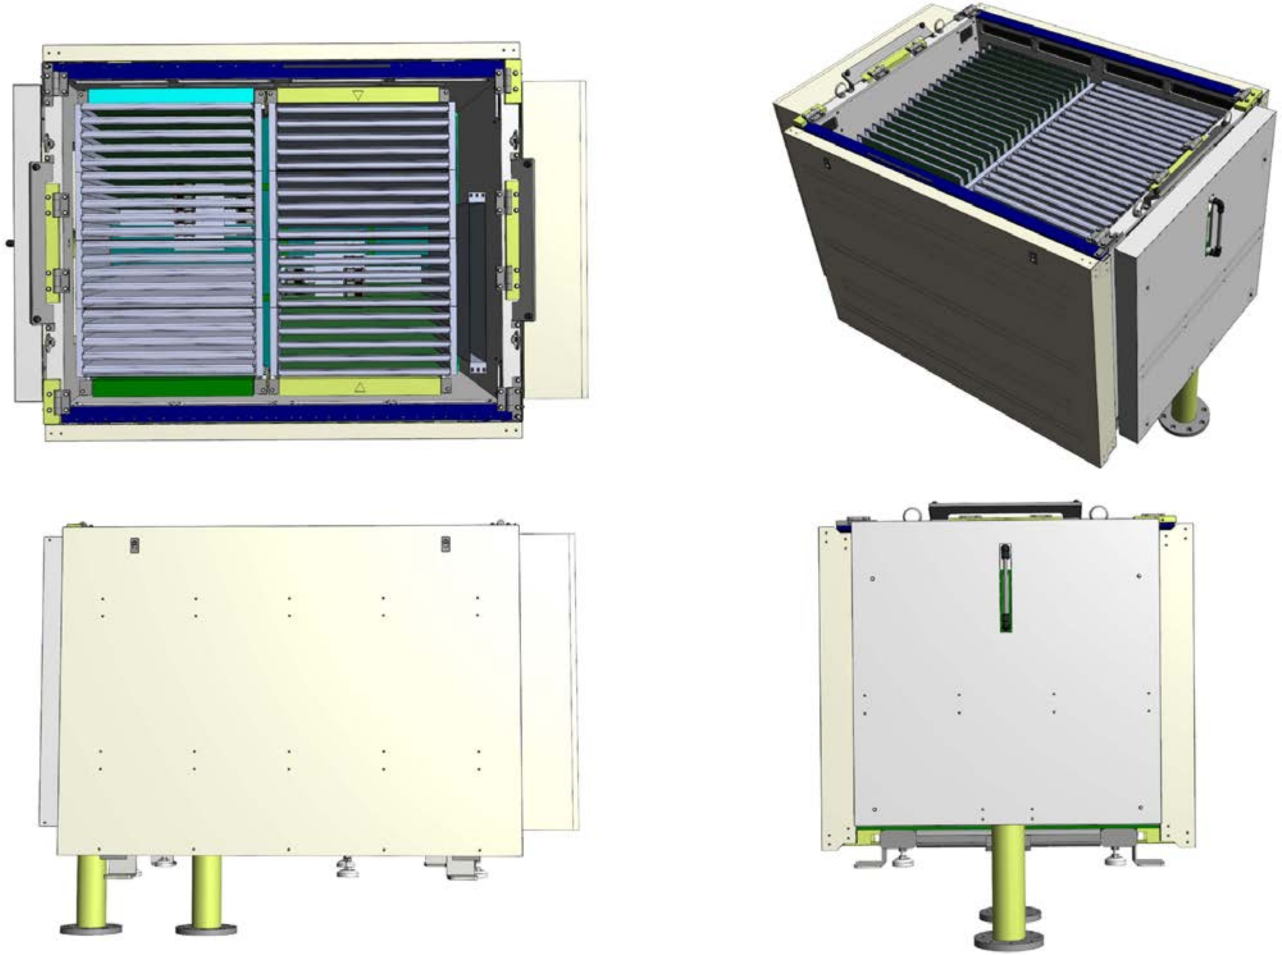

Summary

- PEZY many core processor **SC2**
  - **Deep many core 2048** cores per chip
  - **MIMD**
- **Application Development** for PEZY SC2 systems
  1. Artificial cerebellum
  2. Simulation of Particle system for Tsunami disaster
  3. Middleware: FDPS and Formura
  4. Genome analysis
- Brand new processor **PEZY SC3** and **Zetta Scaler-3.0**
  - 19.7 Tflops/chip 38.7 Gflops/w (SC3)
  - 3.2 PetaFLOPS per tank
  - 30-40 tanks for a 100 Pflops system (Zetta Scaler-3.0)
- International Conference: New Horizon of Supercomputing with many-core Processors
  - May 10-11, 2020, RIKEN Wako