

# **FPGAs for HPC: Open Challenges for Infrastructure and System Operation**

## **FPGA & GPU: compensate with each other toward perfect HPC hardware solution**

**Taisuke Boku**

**Director, Center for Computational Sciences**

**University of Tsukuba**

**taisuke@ccs.tsukuba.ac.jp**



# FPGA in HPC

## ■ Goodness of recent FPGA for HPC

- True **codesigning** with applications (essential)
- Programmability improvement – **HLS (High Level Synthesis): OpenCL, C, C++**
- High performance **interconnect**: ~100Gb (xn)
- **Precision control** is possible (ex. AI)
- Relatively low power

## ■ Problems

- Programmability: **OpenCL is not enough, not efficient**
- **Low standard FLOPS**: still cannot catch up to GPU  
-> “never try what GPU works well on”
- **Memory bandwidth**: 1-gen older than high end CPU/GPU  
-> be improved by HBM2 (Stratix10~)



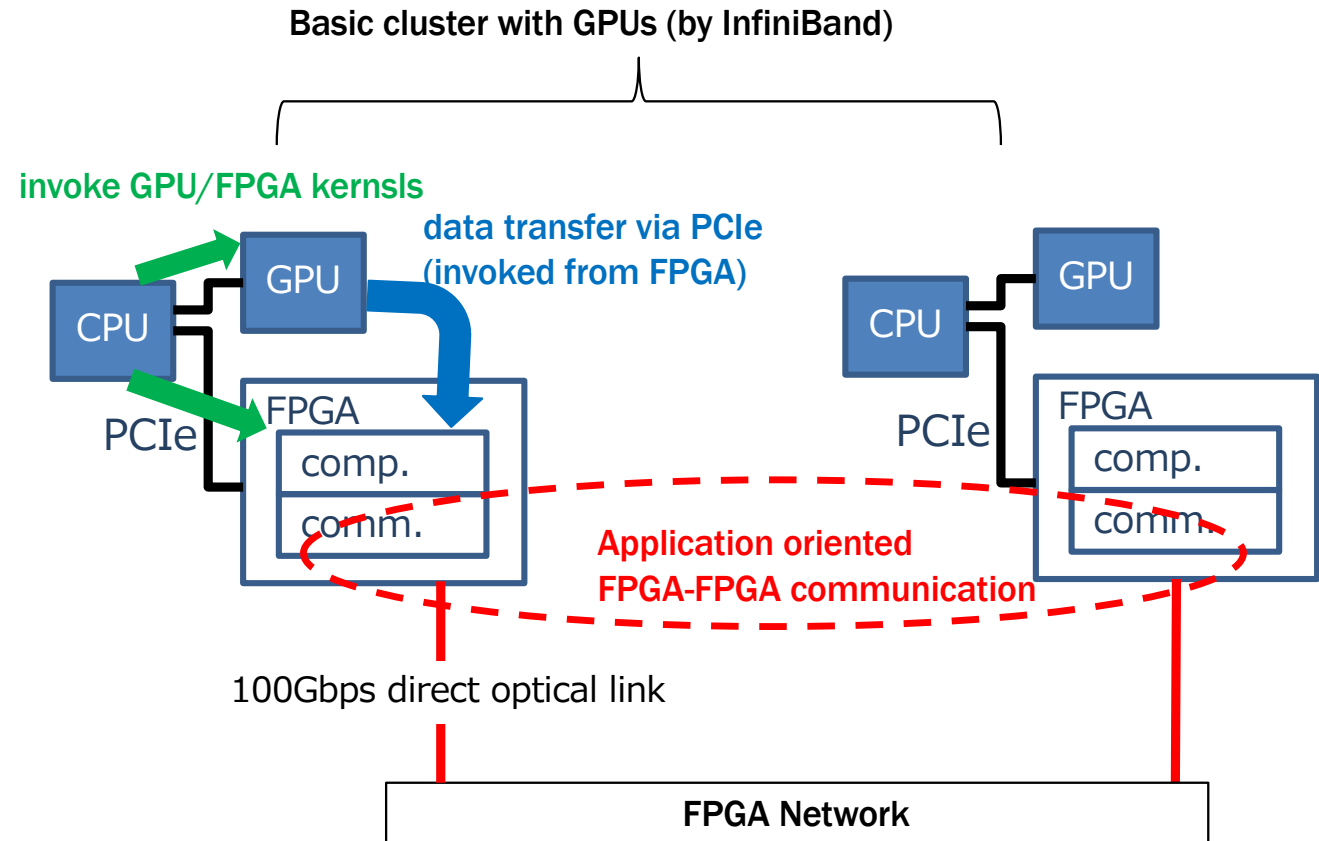


BittWare 520N with Intel Stratix10 FPGA equipped with 4x 100Gbps optical interconnection interfaces

# GPU vs FPGA as HPC solutions

device	GPU	FPGA
parallelization	SIMD (x multi-group)	pipeline (x multi-group)
standard FLOPS	😄😄 (1000x cores)	😄 (~100x pipeline)
conditional branch	😓 (warp divergence)	😄 (both direction)
memory	😄😄 (HBM2e)	😓 (DDR)→😄 (HBM2)
interconnect	😓 (via host facility)	😄😄 (own optical links)
programming	😄 (CUDA, OpenACC, OpenMP)	😓 (HDL)→😄 (HLS)
controllability	😓 (slave device of host CPU)	😄 (autonomic)
HPC applications	😄 (various fields)	😓 (not much)

# CHARM: Cooperative Heterogeneous Acceleration with Reconfigurable Multi-devices

multi-physics/multi-scale  
complicated problem



# Cygnus: world first multi-hybrid cluster with GPU+FPGA



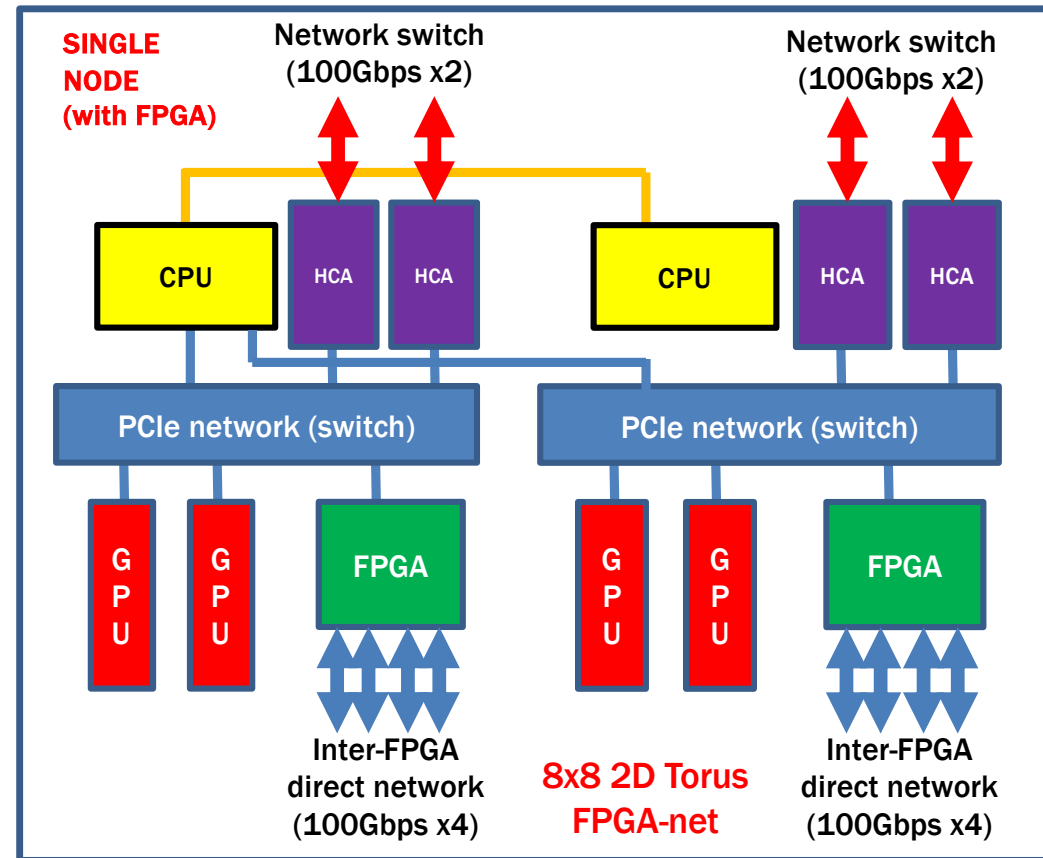
Cygnus supercomputer at Center for Computational Sciences, Univ. of Tsukuba (2019~)  
81 nodes in total including 32 “Albireo” nodes with GPU+FPGA (other “Deneb” nodes have GPU only)





# Single node configuration (Albireo)



- All nodes in Cygnus are equipped with both **IB EDR** and **FPGA-direct network**
- Some nodes (Albireo) are equipped with both FPGAs and GPUs, and other nodes (Deneb) are with GPUs only
- **GPU: NVIDIA V100 x4**  
**FPGA: Intel Stratix10 x2**







IB HDR100 x4  
⇒ HDR200 x2

100Gbps x4  
FPGA optical  
network

IB HDR200  
switch (for  
full-bisection  
Fat-Tree)

**Albireo node**




## Current research around Cygnus

- **FPGA-network:** CIRCUS (Communication Integrated Reconfigurable CompUting System)
  - direct interconnect facility among FPGA boards by multi-dimensional optical link (~100Gps) with router and OpenCL-ready API
  - pipelining all computation and communication seamlessly
- **GPU-FPGA DMA:** kicked by FPGA (without CPU)
  - PCIe-protocol base DMA engine to reduce multi-device high speed data transfer
- **Programming:** MHOAT (Multi-Hetero OpenACC Translator)
  - breaking single OpenACC code for multi-de vices (GPU + FPGA) for each compiler
  - FPGA: OpenACC -> OpenCL -> aocx (by OpenARC research compiler by ORNL)
  - GPU: OpenACC -> NVPTX (PGI compiler)
- **Application:** ARGOT, application on astrophysics for early-universe object generation
  - two main parts are executed by GPU and FPGA






# CIRCUS performance



Throughput (1hop~7hops)



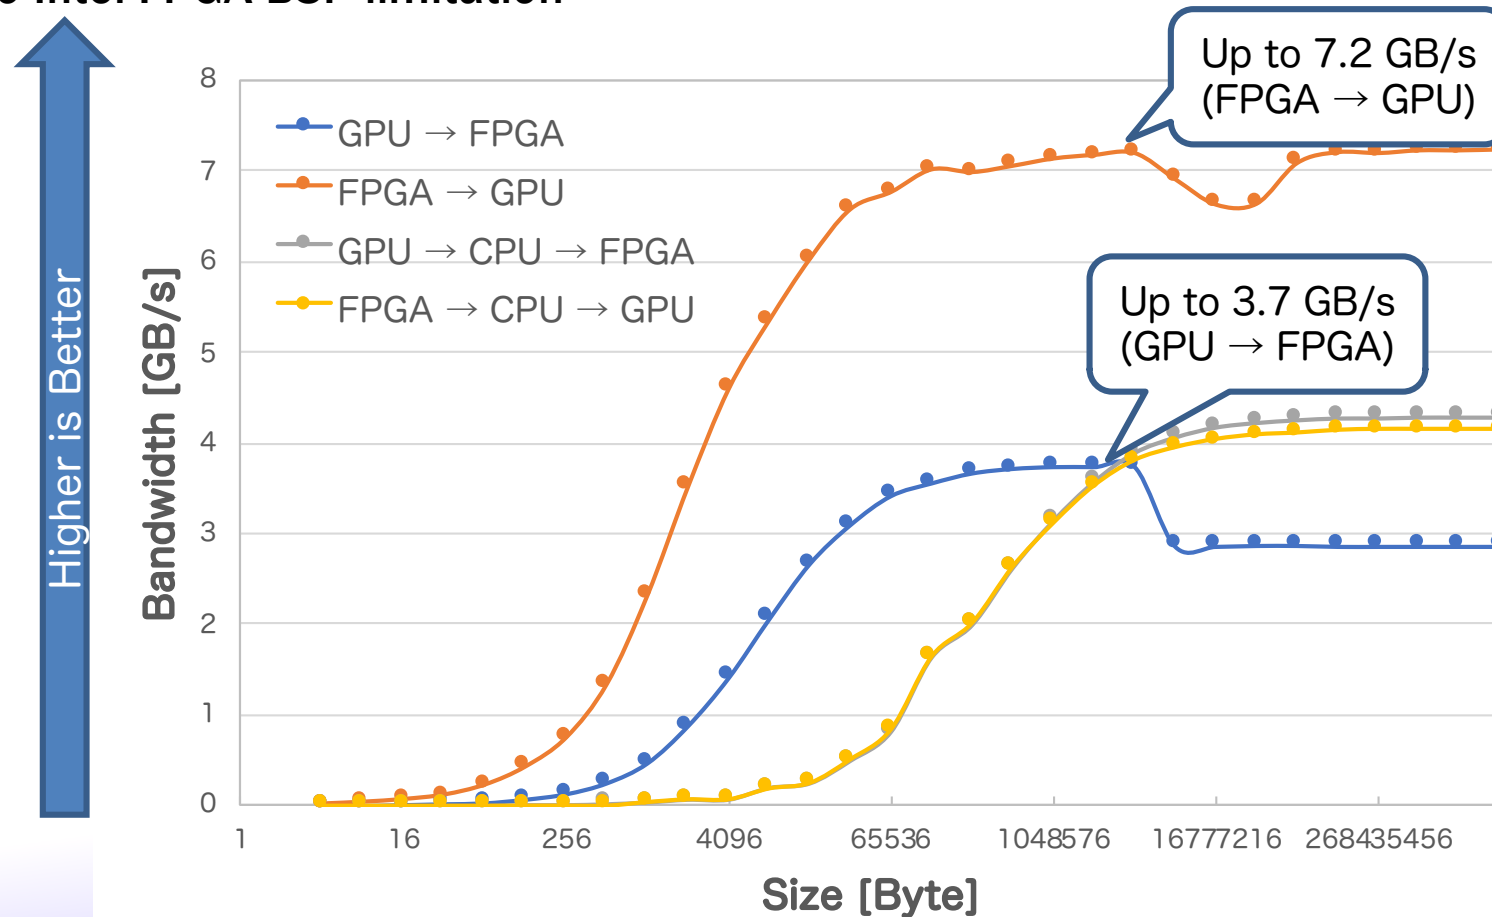
Latency (1hop~7hops)

Evaluated on up to 8 Bittware 520N FPGA boards in Cygnus supercomputer at CCS, University of Tsukuba [14]

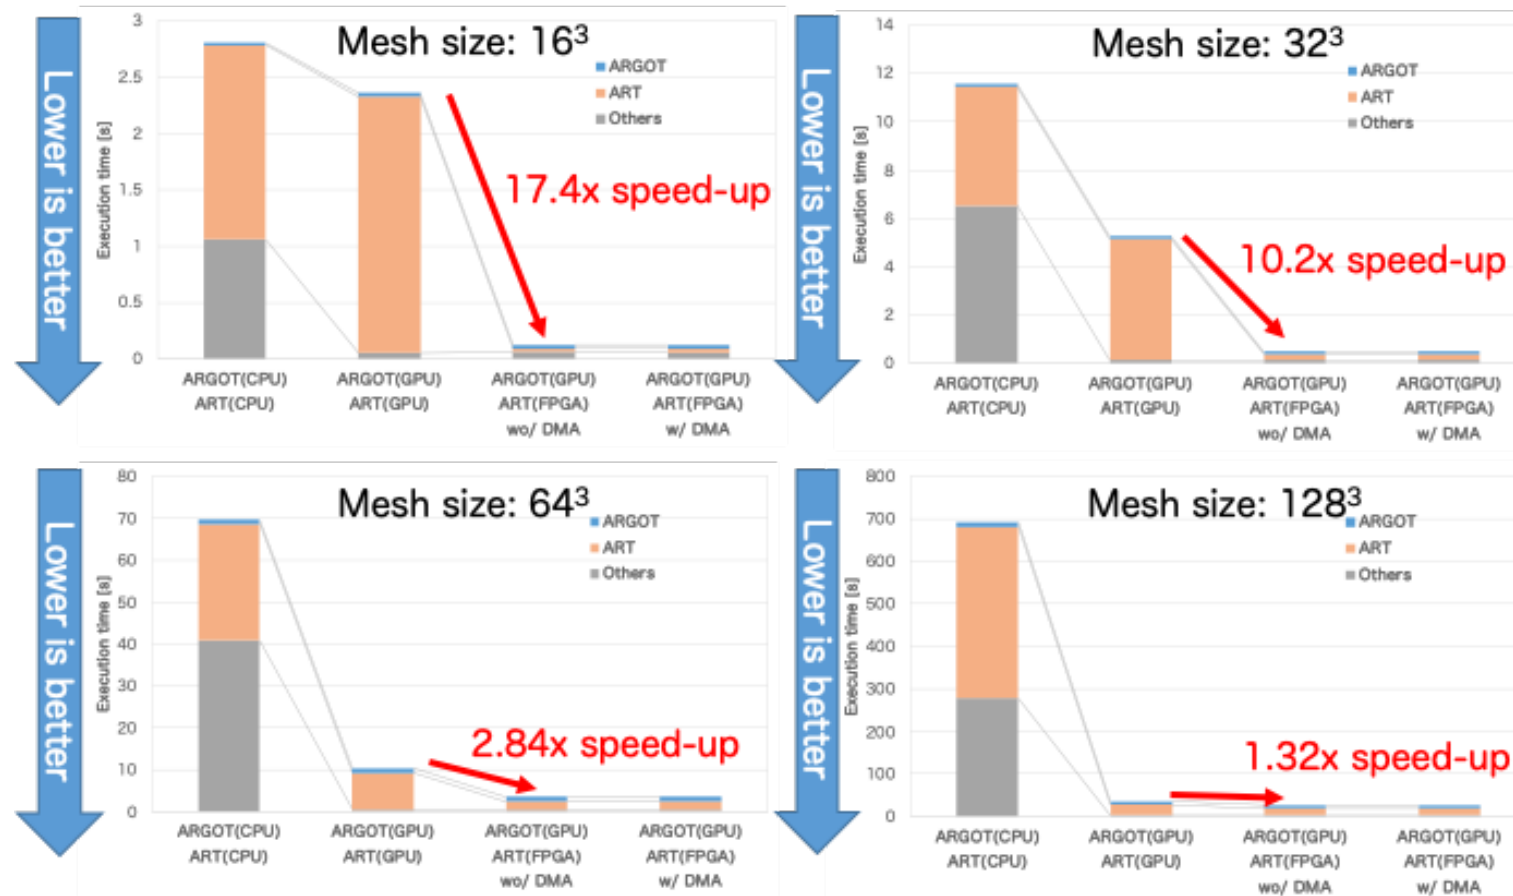
[14]: N. Fujita, et al., "Performance Evaluation of Pipelined Communication Combined with Computation in OpenCL Programming on FPGA", ASHES2020

# Bandwidth of GPU-FPGA DMA

- Data size: 4 ~ 2G Bytes (PCIe gen3x8\*: 7.8GB/s peak)  
\* due to Intel FPGA BSP limitation



# GPU-only vs GPU-FPGA coworking on astrophysics simulation ARGOT



Single node GPU (V100) + FPGA (Stratix10) evaluation on Cygnus

R. Kobayashi, et. al., "Accelerating Radiative Transfer Simulation with GPU-FPGA Cooperative Computation", ASAP2020, Jul. 2020



## Topics for Discussion

- What are the most promising feature or advantages of present or future HPC using FPGAs ?
  - ⇒ multi-hetero solution, **not FPGA stand-alone but coupling with other accelerators (GPUs) for perfect acceleration**
- What technology is missing in infrastructure and operation of HPC systems using FPGAs ?
  - ⇒ **programming framework** to release users from tough low-level programming, keeping a certain efficiency : ex. HLS or DSL ?
- How can we or should we tackle the above challenge in the community with academia and industries ?
  - ⇒ true **codesigning with application users, open source** facilities



# CFP for HPC FPGA Workshop 2021 (Cluster2021)

<https://sites.google.com/view/hpcfpga2021/>

- **HPC FPGA Workshop 2021 in IEEE Cluster 2021 (Sep. 7th, 2021, virtual)**
  - topic: FPGA for HPC (exactly with same as this session!)
  - any hardware, software, middleware, application on FPGA utilization toward HPC
  - proceedings will be published as IEEE Cluster2021 workshop volume
- **Important dates**
  - paper submission due: **June 25th ⇒ to be extended ???**
  - acceptance notification: July 19th
  - camera ready due: July 30th
  - workshop: September 7th
- **Organization**
  - Taisuke Boku (chair), Martin Herbordt (PC chair), Franck Cappello, Kentaro Sano