様　式　Ｃ－１９、Ｆ－１９－１、Ｚ－１９　（共通）

科学研究費助成事業　　研究成果報告書

令和　６ 年　 ６ 月 １１ 日現在

機関番号：８２４０１
研究種目：基盤研究(B)（一般）
研究期間：2019～2023
課題番号：１９Ｈ０４１１９

研究課題名（和文）ExaPath: Hierarchical Routing for Next-Gen Supercomputers and Beyond

研究課題名（英文）ExaPath: Hierarchical Routing for Next-Gen Supercomputers and Beyond

研究代表者
ドンケ　イェンス（Domke, Jens）

国立研究開発法人理化学研究所・計算科学研究センター・チームリーダー

研究者番号：７０８１５４８０

交付決定額（研究期間全体）：（直接経費）　　13,200,000円

研究成果の概要（和文）：現代社会は、人工知能などのサイエンス分野に対応するため、ますます大きな計算能力を求めている。最近では、スーパーコンピューターがより大規模なシステムへとスケールアウトし始めている。これらのシステムのバックボーンは、適切なルーティングを必要とする相互接続ネットワークである。このプロセスは、グーグルマップがどの道を走ればいいかを教えてくれるのに似ている。スーパーコンピューターでは、ルーティングがメッセージに進むべき道を指示する。我々のプロジェクトは、何千台、何百万台ものコンピュータを接続する非常に複雑なネットワークのための新しいルーティング・アプローチを設計することを目的としている。

研究成果の学術的意義や社会的意義
Our developed routing algorithms, and methods to make supercomputer interconnects faster, will help other scientists to accelerate their workflows. Meaning, with optimal routing, the supercomputers can finish more scientific simulations, and hence the scientists can get more results in shorter time.

研究成果の概要（英文）：Modern society is seeking for ever-increasing compute power to serve science fields such as artificial intelligence (e.g. ChatGPT), weather forecast, airplane and supernova simulations, and medicine discovery. Certain physical limitations prevent us from making computer chips much faster, and hence supercomputers and other tightly coupled compute systems started to scale-out to larger and larger systems. The backbone of all these architectures is the interconnection network, which must be "routed" correctly to increase the effectiveness of the entire system. This process is similar to GoogleMaps telling us which roads to drive. In the supercomputers, the routing tells the messages which path to take. However, the state-of-the-art routing algorithms cannot keep pace with the future scale-out and hardware trends, and hence new algorithms have to be invented. Our project aims to design novel routing approaches for highly complex networks which connect thousands or millions of computers.

研究分野：Supercomputers

キーワード：HPC interconnects　routing algorithms　network design　artificial intelligance　message passing

様　式　Ｃ－１９、Ｆ－１９－１（共通）

１．研究開始当初の背景

With Moore's law slowly coming to an end, high-performance computers (HPC) and data center (DC) installations started to scale-out to larger systems. It can be expected that the number of network endpoints will grow significantly which emphasizes the role of the interconnection network as one of the most critical components in a supercomputer. The network is largely controlled by routing algorithms which determine how to forward packets. These routing algorithms have to balance multiple, partially conflicting, requirements. For example, they shall provide the best forwarding strategy (a NP-hard problem) while minimizing the runtime of the routing in order to quickly react to failures of network components.

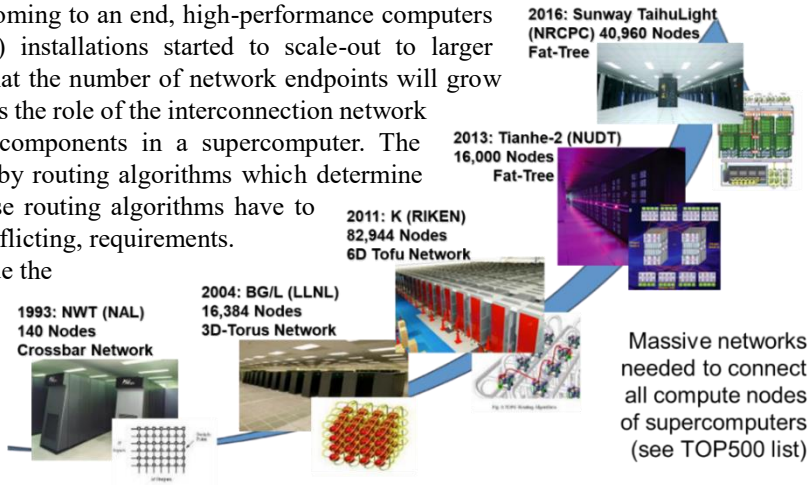

２．研究の目的

Interconnection networks, message passing libraries, and routing algorithms are the backbone of any connected but distributed computing infrastructure – be it traditional supercomputers, data centers running artificial intelligence (AI) workloads, internet & IoT, or the approach of edge computing – which are all essential for a modern society. The 1st research objective is the invention and development of a novel type of routing algorithms, so called hierarchical routing algorithms, for the upcoming scalability challenges in these connected infrastructures, with special focus on supercomputing. Furthermore, since innovation thrives from collaboration and open access, the 2nd objective of this research is the development of a community-driven, open-source routing library, whose framework shall allow for easy research and development and rigor testing, and ease of use by integration into other software frameworks. Lastly, due to the emergence of artificial intelligence for image recognition (ResNet) and large-language models (GPT), the 3rd purpose of this study is the research, porting, and tuning (with special focus on topologies, routing, and message passing optimizations) of AI frameworks for Fugaku and other Japanese HPC systems.

３．研究の方法

The research plan is subdivided into three major research and development and engineering tasks, so that the PI, with the support of research assistants (RAs), interns and collaborators, is able to tackle the following R&D challenges in parallel. Furthermore, the rapid growth of the importance of AI for the Japanese society and research community led to the addition of two more tasks during the project phase:

(1) Survey (hierarchical) static/adaptive routing algorithm for state-of-the-art technologies; and
(2) Development of open-source routings with focus on interoperability and novel topologies; and
(3) Routing validation and verification infrastructure to aid an effective research progress; and
(4) Porting of AI frameworks to Fugaku and analyze their routing and message passing needs; and
(5) Optimizing MPI algorithms for AI via tuned routing, task placement, and message scheduling.

４．研究成果

(1) Achievements (researchers and students funded by this KAKENHI grant highlighted in bold font)

Peer-reviewed Publications and Posters List

① **J. Domke**, S. Matsuoka, **I.R. Ivanov**, Y. Tsushima, T. Yuki, A. Nomura, S. Miura, N. McDonald, D.L. Floyd, N. Dube, "The First Supercomputer with HyperX Topology: A Viable Alternative to Fat-Trees?," peer-reviewed short paper presented at the 2019 IEEE 26th Symposium on High-Performance Interconnects (HOTI 26), Aug. 2019.

② **J. Domke**, S. Matsuoka, **I.R. Ivanov**, Y. Tsushima, T. Yuki, A. Nomura, S. Miura, N. McDonald, D.L. Floyd, N. Dube, "HyperX Topology: First at-scale Implementation and Comparison to the Fat-Tree," in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '19, (Piscataway, NJ, USA), IEEE Press, Nov. 2019.

③ **I.R. Ivanov**, **J. Domke**, A. Nomura, **T. Endo**, "Improved failover for HPC interconnects through localised routing restoration," Poster presented at The 3rd R-CCS International Symposium (RCCS-IS3), Kobe, Japan, Feb. 2021.

④ M. Besta, **J. Domke**, M. Schneider, M. Konieczny, S.D. Girolamo, T. Schneider, A. Singla, T. Hoefler, "High-Performance Routing with Multipathing and Path Diversity in Supercomputers and Data Centers," IEEE Transactions on Parallel and Distributed Systems, vol. 32, no. 4, pp. 943-959, 2021.
⑤ **J. Domke**, "A64FX – Your Compiler You Must Decide!," in Proceedings of the 2021 IEEE International Conference on Cluster Computing (CLUSTER), EAHPC Workshop, (Portland, Oregon, USA), IEEE Computer Society, Sept. 2021.
⑥ **I.R. Ivanov**, **J. Domke**, **T. Endo**, "Automatic translation of CUDA code into high performance CPU code using LLVM IR transformations," Poster presented at The 4rd R-CCS International Symposium (RCCS-IS4), Kobe, Japan, Feb. 2022.
⑦ W.S. Moses, **I.R. Ivanov**, **J. Domke**, **T. Endo**, J. Doerfert, O. Zinenko, "High-Performance GPU-to-CPU Transpilation and Optimization via High-Level Parallel Constructs," Poster presented at 2022 LLVM Developers' Meeting, San Jose, USA, Nov. 2022.
⑧ **I.R. Ivanov**, W.S. Moses, **J. Domke**, **T. Endo**, "Parallel Optimizations and Transformations of GPU Kernels Using a High-Level representation in MLIR/Polygeist," Poster presented at IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2023, Feb. 2023.
⑨ W.S. Moses, **I.R. Ivanov**, **J. Domke**, **T. Endo**, J. Doerfert, O. Zinenko, "High-Performance GPU-to-CPU Transpilation and Optimization via High-Level Parallel Constructs," in Proceedings of the 28th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPoPP '23, (New York, NY, United States), Association for Computing Machinery, Feb. 2023.
⑩ **I. R. Ivanov**, O. Zinenko, **J. Domke**, **T. Endo**, W. S. Moses, "Retargeting and Respecializing GPU Workloads for Performance Portability," in Proceedings of the IEEE/ACM International Symposium on Code Generation and Optimization (CGO'24), Feb. 2024.
⑪ **N. Blach**, M. Besta, D.D. Sensi, **J. Domke**, H. Harake, S. Li, P. Iff, M. Konieczny, K. Lakhotia, A. Kubicek, M. Ferrari, F. Petrini, T. Hoefler, "A High-Performance Design, Implementation, Deployment, and Evaluation of The Slim Fly Network," in Proceedings of the 21st USENIX Symposium on Networked Systems Design and Implementation (NSDI'24), April 2024.

Currently being Reviewed and Planned Publication

① **K. Fan**, **J. Domke**, S. Ba, S. Kumar, "TuNA: Parameterized Hierarchical Algorithms for Non-uniform All-to-all" under review at 53rd International Conference on Parallel Processing (ICPP '24)
② M. Besta, **N. Blach**, P. Iff, M. Schneider, A. Maissen, S.D. Girolamo, **J. Domke**, J. Krattenmacher, A. Singla, K. Lakhotia, F. Petrini, T. Hoefler, "EvalNet: A Practical Toolchain for Generation, Analysis, and Simulation of Large-Scale Networks" preparing for submission at 39th IEEE International Parallel & Distributed Processing Symposium (IPDPS '25)

Invited Talks (excluding presentations in conjunction with conference publications)

① **J. Domke** "First At-Scale HyperX Implementation: A Compelling Alternative to Fat-Trees?" in High Performance Consortium for Advanced Scientific and Technical Computing (HP-CAST 32), June 2019.
② **J. Domke** "The Bright Future for HPC Interconnects -- Opportunities, Challenges, and Misconceptions in Deployment and Management of Large-Scale Networks" in Focus Session: Leveraging Silicon Photonics in HPC to Meet Future Exascale Needs in 36th ISC High Performance (ISC '21), June 2021.
③ **J. Domke** "MocCUDA: Running Cuda Codes on Fugaku" in 13th JLESC Workshop, Sept. 2021.
④ **J. Domke** "MocCUDA: Running Cuda Codes on Fugaku" in SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP '22), Seattle, Washington, USA, Feb. 2022.
⑤ **J. Domke** "Working with Proxy-Applications: Interesting Findings, Lessons Learned, and Future Directions" in Benchmarking in the Data Center: Expanding to the Cloud (workshop) held in conjunction with PPoPP 2022: Principles and Practice of Parallel Programming 2022, April 2022.
⑥ **J. Domke** "Octopodes A candidate to replace Mini Apps and Motifs?" in 14th JLESC WS, Sept. 2022.
⑦ **I. R. Ivanov** "Optimization of CUDA GPU Kernels and Translation to AMDGPU in Polygeist/MLIR" in 2023 LLVM Developers' Meeting. Student Talk
⑧ **I. R. Ivanov** "GPU Kernel Compilation in Polygeist/MLIR" in 2023 LLVM Developers' Meeting GPU Offloading Workshop. Lightning Talk

Collaborations

① We established a long-term collaboration with Prof. T. Endo at Tokyo Institute of Technology to collaborate with him and his students on hierarchical routings and other network topics
② We continued collaboration with Prof. T. Hoefler at ETH Zurich to work on Slimfly testbeds
③ We established various long-term collaboration with multiple researchers from Massachusetts Institute of Technology, Google LLC, Argonne and Lawrence Livermore National Laboratories, and University of Illinois at Chicago to work on multi-node AI training/ inference, GPU/CPU portability, and network optimizations (esp. routing and MPI allreduce and alltoall for AI)

Training and Researcher Mentoring

① Co-supervise **I.R. Ivanov** (Tokyo Tech) from 2019 to now; Bachelor's Thesis: "Improved failover for HPC interconnects through localised routing restoration" and Master's Thesis: "Optimizations and Transformations of Parallel Code via High Level Intermediate Representation"; **I.R. Ivanov was selected for the Junior Research Associate Program at RIKEN Center for Computational Science**

② Co-supervised J. Bokstaller (ETH Zürich) in 2019; Bachelor's Thesis: "Design and Implementation of Multipath Switching in InfiniBand Slimfly Networks"
③ Co-supervised **N. Blach** (ETH Zürich) in 2021; Bachelor's Thesis: "Multipath Routing for Low-Diameter Network Topologies on InfiniBand Architecture"
④ Hosted multiple RIKEN Internship students: **K. Fan** (University of Illinois at Chicago), **N. Blach** (ETH Zürich) and **F. Hoerold** (ETH Zürich) in 2023; **K. Fan and F. Hoerold became both long-term remote trainees at RIKEN Center for Computational Science**

Hereafter, the content of a few selected peer-reviewed scientific papers and academic publications will be provided. Due to the large number of publications related to this project, the following list is abbreviated and is focusing on the works with the highest research and societal impact:

(2) Content of Scientific Publication "High-Performance Routing with Multipathing and Path Diversity in Supercomputers and Data Centers" related to project task 1:

To facilitate high-performance routing in modern networks, we analyze existing routing protocols and architectures, focusing on how well they exploit the diversity of minimal and non-minimal paths. We first develop a taxonomy of different forms of support for multipathing and overall path diversity. Then, we analyze how existing routing schemes support this path diversity, see the Figure on the right. Among others, we consider multipathing with both shortest and non-shortest paths, support for disjoint paths, or enabling adaptivity. To address the ongoing convergence of HPC and "Big Data" domains, we consider routing protocols developed for both HPC systems and for data centers as well as general clusters. Thus, we cover architectures and protocols based on Ethernet, InfiniBand, and other HPC networks such as Myrinet. Our review will foster developing future high-performance multipathing routing protocols in supercomputers and data centers.

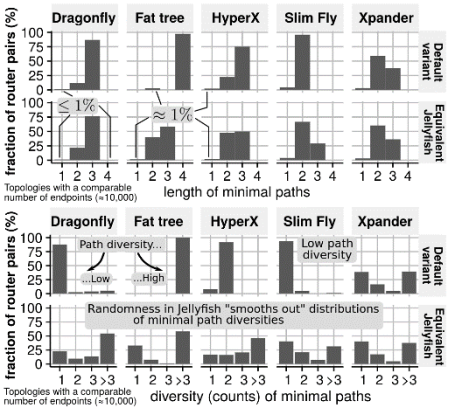

(3) Content of Scientific Publication "HyperX Topology: First at-scale Implementation and Comparison to the Fat-Tree" related to project task 2:

The de-facto standard topology for modern HPC systems and data-centers are Folded Clos networks, commonly known as Fat-Trees. The number of network endpoints in these systems is steadily increasing. The switch radix increase is not keeping up, forcing an increased path length in these multi-level trees that will limit gains for latency-sensitive applications. Additionally, today's Fat-Trees force the extensive use of active optical cables which carries a prohibitive cost-structure at scale. We built the world's first 3 Pflop/s supercomputer with two separate networks, a 3-level Fat-Tree and a 12x8 HyperX. This dual-plane system allows us to perform a side-by-side comparison using a broad set of benchmarks. We show that the HyperX, together with our novel communication pattern-aware routing, can challenge the performance of, or even outperform, traditional Fat-Trees.

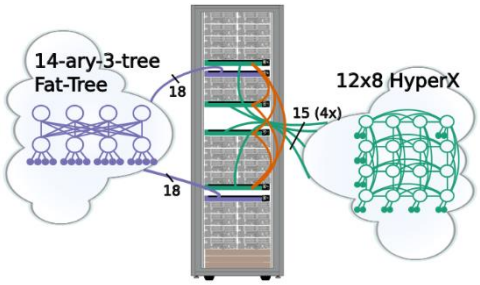

(4) Content of Scientific Publication "EvalNet: A Practical Toolchain for Generation, Analysis, and Simulation of Large-Scale Networks" related to project task 3:

Our EvalNet seamlessly integrates different tools for network simulation. It adapts a flow-level evaluation tool, called TopoBench, that enables approximating throughput with arbitrary traffic patterns and with specially designed worst-case scenario that puts particularly large stress on the interconnect. It also incorporates two complementary packet-level simulators, htsim and Booksim2. Compared to competing solutions which require clusters themselves to simulate scales of 100k or more, such as ROSS/CODES, EvalNet enhances the design of htsim to enable much more scalability and we show how to simulate networks with up to 1M servers on a PC laptop. Figure 20.5 demonstrates just a few of the capabilities of EvalNet. The left side shows the histograms for lengths of the shortest path within state-of-the-art topologies, which are analytically derived by the framework, while the right side of the Figure shows the achievable flow-completion times (FCT) for various realistic network topologies and supercomputer sizes collected by the network simulation environment.

EvalNet, highlighted in the Figure below, is straightforwardly extensible with new topologies and network evaluation measures. The framework (publicly available after acceptance of publication) will facilitate research into today's and future large-scale network designs, and can be used to rapidly deliver a large number of results which will help system designers in their co-design efforts and push the boundaries

in large-scale network deployments.

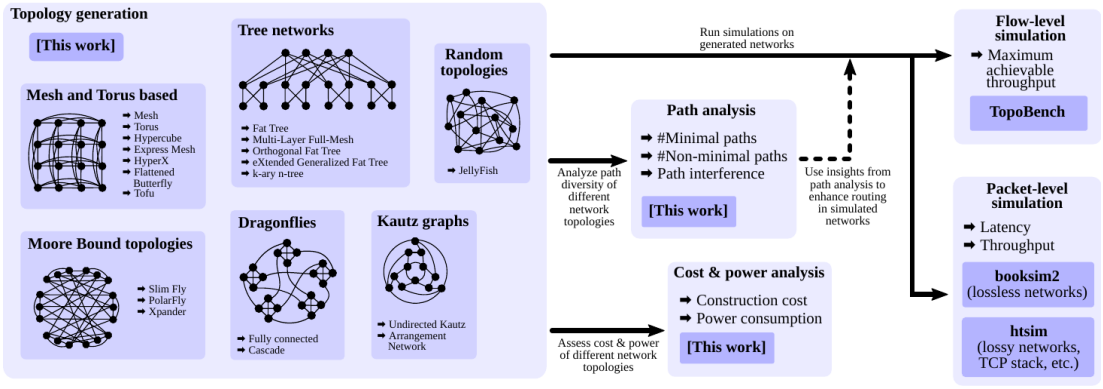

(5) Content of Scientific Publication "High-Performance GPU-to-CPU Transpilation and Optimization via High-Level Parallel Constructs" related to project task 4:

A “shortcoming” of Fugaku when it comes to AI frameworks is the lack of GPUs. Unlike HPC codes, these frameworks have been designed with (Nvidia) GPUs and InfiniBand interconnect as first class citizens. Hence, we proposed an alternative approach that automatically translates programs written in one programming model (CUDA), into another (CPU threads) based on Polygeist/MLIR. We evaluated our framework by transpiling and optimizing the CUDA Rodinia benchmark suite for a multi-core CPU and TofuD interconnection network and achieve a 58% speedup (geomean) over handwritten OpenMP code, as can be seen in the thread scaling study on the left side of the following Figure.

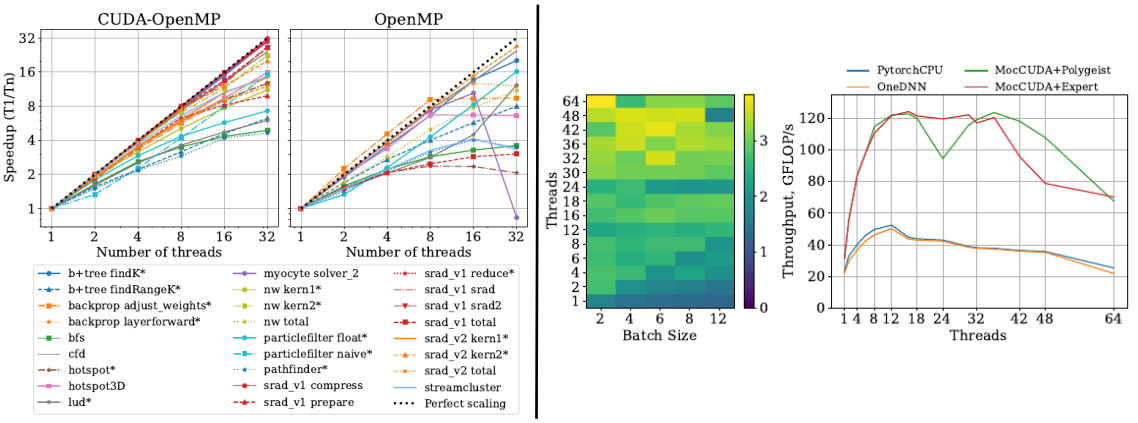

Further, we demonstrate with the right Figure how CUDA kernels from PyTorch can efficiently run and scale on the CPU-only Supercomputer Fugaku without user intervention. Our PyTorch compatibility layer, called MocCUDA, which emulates the existence (and some functionality) of Nvidia GPUs is making use of transpiled CUDA PyTorch kernels to outperform PyTorch’s CPU native backend by 2.7x.

(6) Content of Scientific Publication "Parameterized hierarchical algorithms for non-uniform all-to-all" related to project task 5:

Scaling some HPC workloads, such as computing Fast Fourier Transforms, graph neural networks (GNN) training, and simulating quantum computers, can be impeded by inefficient MPI_Alltoallv communications. Current algorithms overlook the intricacies of modern HPC system architectures, such as the substantial performance difference between inter-node and intra-node communication.

Our work introduces a highly parameterizable algorithm, called TuNA, that considers key factors such as the hierarchical architecture of HPC systems, the number of data exchange rounds, and the communication burst size. Our algorithm efficiently addresses the trade-off between bandwidth and latency that existing implementations fail to optimize. We show, in the Figure below, a performance improvement over the state-of-the-art implementations by factors of 42x and 138x on Polaris and Fugaku, respectively.

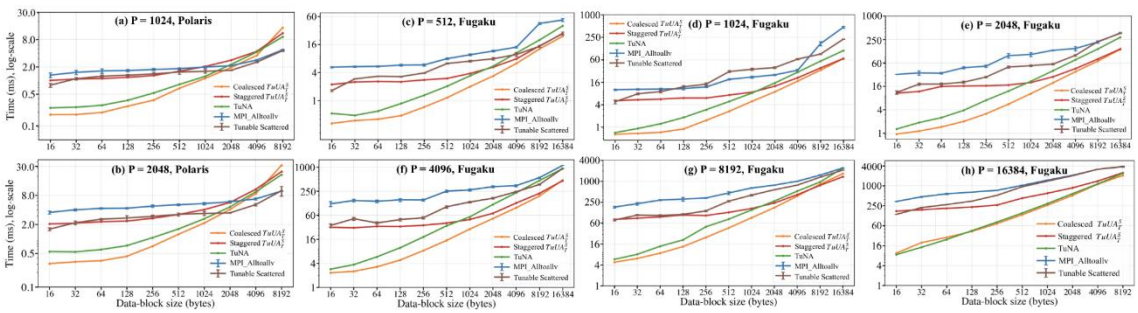

５．主な発表論文等

〔雑誌論文〕　計9件（うち査読付論文　9件／うち国際共著　8件／うちオープンアクセス　0件）

１．著者名 Moses William S.、Ivanov Ivan R.、Domke Jens、Endo Toshio、Doerfert Johannes、Zinenko Oleksandr	４．巻 0
２．論文標題 High-Performance GPU-to-CPU Transpilation and Optimization via High-Level Parallel Constructs	５．発行年 2023年
３．雑誌名 28th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPoPP '23	６．最初と最後の頁 119-134
掲載論文のDOI（デジタルオブジェクト識別子） 10.1145/3572848.3577475	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

１．著者名 I.R. Ivanov, W.S. Moses, J. Domke, T. Endo	４．巻 0
２．論文標題 Parallel Optimizations and Transformations of GPU Kernels Using a High-Level representation in MLIR/Polygeist	５．発行年 2023年
３．雑誌名 IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2023	６．最初と最後の頁 1
掲載論文のDOI（デジタルオブジェクト識別子） なし	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

１．著者名 W.S. Moses, I.R. Ivanov, J. Domke, T. Endo, J. Doerfert, O. Zinenko	４．巻 0
２．論文標題 High-Performance GPU-to-CPU Transpilation and Optimization via High-Level Parallel Constructs	５．発行年 2022年
３．雑誌名 2022 LLVM Developers' Meeting	６．最初と最後の頁 1
掲載論文のDOI（デジタルオブジェクト識別子） なし	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

１．著者名 J. Domke	４．巻 0
２．論文標題 A64FX - Your Compiler You Must Decide!	５．発行年 2021年
３．雑誌名 2021 IEEE International Conference on Cluster Computing (CLUSTER), EAHPC Workshop	６．最初と最後の頁 1-5
掲載論文のDOI（デジタルオブジェクト識別子） なし	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

1．著者名 I.R. Ivanov, J. Domke, T. Endo	4．巻 0
2．論文標題 Automatic translation of CUDA code into high performance CPU code using LLVM IR transformations	5．発行年 2022年
3．雑誌名 The 4rd R-CCS International Symposium (RCCS-IS4)	6．最初と最後の頁 1
掲載論文のDOI（デジタルオブジェクト識別子） なし	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 －

1．著者名 Maciej Besta, Jens Domke, Marcel Schneider, Marek Konieczny, Salvatore Di Girolamo, Timo Schneider, Ankit Singla, Torsten Hoefler	4．巻 32
2．論文標題 High-Performance Routing With Multipathing and Path Diversity in Ethernet and HPC Networks	5．発行年 2021年
3．雑誌名 IEEE Transactions on Parallel and Distributed Systems	6．最初と最後の頁 1-14
掲載論文のDOI（デジタルオブジェクト識別子） 10.1109/TPDS.2020.3035761	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

1．著者名 Ivan R. Ivanov, Jens Domke, Akihiro Nomura, Toshio Endo	4．巻 0
2．論文標題 Improved failover for HPC interconnects through localised routing restoration	5．発行年 2021年
3．雑誌名 The 3rd R-CCS International Symposium (RCCS-IS3)	6．最初と最後の頁 -
掲載論文のDOI（デジタルオブジェクト識別子） なし	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

1．著者名 Domke Jens、Matsuoka Satoshi、Ivanov Ivan R.、Tsushima Yuki、Yuki Tomoya、Nomura Akihiro、Miura Shin'ichi、McDonald Nie、Floyd Dennis L.、Dube Nicolas	4．巻 SC'19
2．論文標題 HyperX Topology: First at-scale Implementation and Comparison to the Fat-Tree	5．発行年 2019年
3．雑誌名 Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis	6．最初と最後の頁 40:1-40:23
掲載論文のDOI（デジタルオブジェクト識別子） 10.1145/3295500.3356140	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

１．著者名 Domke Jens、Matsuoka Satoshi、Radanov Ivan、Tsushima Yuki、Yuki Tomoya、Nomura Akihiro、Miura Shin'ichi、McDonald Nic、Floyd Dennis Lee、Dube Nicolas	４．巻 HOTI'26
２．論文標題 The First Supercomputer with HyperX Topology: A Viable Alternative to Fat-Trees?	５．発行年 2019年
３．雑誌名 2019 IEEE Symposium on High-Performance Interconnects (HOTI)	６．最初と最後の頁 4
掲載論文のDOI（デジタルオブジェクト識別子） 10.1109/HOTI.2019.00013	査読の有無 有
オープンアクセス オープンアクセスではない、又はオープンアクセスが困難	国際共著 該当する

〔学会発表〕　計10件（うち招待講演　2件／うち国際学会　7件）

１．発表者名 J. Domke
２．発表標題 Working with Proxy-Applications: Interesting Findings, Lessons Learned, and Future Directions
３．学会等名 Benchmarking in the Data Center: Expanding to the Cloud (workshop) held in conjunction with PPoPP 2022: Principles and Practice of Parallel Programming 2022（国際学会）
４．発表年 2022年

１．発表者名 J. Domke
２．発表標題 Octopodes A candidate to replace Mini Apps and Motifs?
３．学会等名 14th JLESC Workshop
４．発表年 2022年

１．発表者名 J. Domke
２．発表標題 MocCUDA: Running Cuda Codes on Fugaku
３．学会等名 SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP ’22)（国際学会）
４．発表年 2022年

１．発表者名
Jens Domke

２．発表標題
MocCUDA: Running CUDA codes on Fugaku

３．学会等名
12th JLESC Workshop（国際学会）

４．発表年
2021年

１．発表者名
Jens Domke

２．発表標題
The Bright Future for HPC Interconnects -- Opportunities, Challenges, and Misconceptions in Deployment and Management of Large-Scale Networks

３．学会等名
Focus Session: Leveraging Silicon Photonics in HPC to Meet Future Exascale Needs in 36th ISC High Performance (ISC '21)（国際学会）

４．発表年
2020年

１．発表者名
Ivan R. Ivanov

２．発表標題
Improved failover for HPC interconnects through localised routing restoration

３．学会等名
The 3rd R-CCS International Symposium (RCCS-IS3)（国際学会）

４．発表年
2021年

１．発表者名
Domke Jens

２．発表標題
HyperX Topology: First at-scale Implementation and Comparison to the Fat-Tree

３．学会等名
International Conference for High Performance Computing, Networking, Storage and Analysis (SC'19)（国際学会）

４．発表年
2019年

１．発表者名
Domke Jens

２．発表標題
The First Supercomputer with HyperX Topology: A Viable Alternative to Fat-Trees?

３．学会等名
2019 IEEE Symposium on High-Performance Interconnects（国際学会）

４．発表年
2019年

１．発表者名
Domke Jens

２．発表標題
The First Supercomputer with HyperX Topology: A Viable Alternative to Fat-Trees?

３．学会等名
The 179th R-CCS Cafe（招待講演）

４．発表年
2019年

１．発表者名
Domke Jens

２．発表標題
First At-Scale HyperX Implementation: A Compelling Alternative to Fat-Trees?

３．学会等名
High Performance Consortium for Advanced Scientific and Technical Computing (HP-CAST 32)（招待講演）

４．発表年
2019年

〔図書〕　計0件

〔産業財産権〕

〔その他〕
MocCUDA
https://gitlab.com/domke/MocCUDA
Repo for thesis of localised routing restoration:
https://gitlab.com/ivanradanov/localisedrerouting
TSUBAME2 HyperX experiment
https://gitlab.com/domke/t2hx

6．研究組織

	氏名 （ローマ字氏名） （研究者番号）	所属研究機関・部局・職 （機関番号）	備考
研究分担者	遠藤　敏夫 (Endo Toshio) (80396788)	東京工業大学・学術国際情報センター・教授 (12608)	

	氏名 （ローマ字氏名） （研究者番号）	所属研究機関・部局・職 （機関番号）	備考
研究協力者	イヴァノヴ　イヴァン (Ivanov Ivan Radanov)	東京工業大学・学術国際情報センター・学生 (12608)	
研究協力者	プラグ　ニルス (Blach Nils)	スイス連邦工科大学チューリッヒ校・Computer Science Department・学生	
研究協力者	ハロルヅ　ファヴェオ (Hoerold Faveo)	スイス連邦工科大学チューリッヒ校・Computer Science Department・学生	
研究協力者	ボクスタラー　ジョナス (Bokstaller Jonas)	スイス連邦工科大学チューリッヒ校・Computer Science Department・学生	
研究協力者	ファン　ケ (Fan Ke)	イリノイ大学シカゴ校・Department of Computer Science・学生	
研究協力者	ベスタ　マチェイ (Besta Maciej)	スイス連邦工科大学チューリッヒ校・Computer Science Department・研究員	

６．研究組織（つづき）

	氏名 （ローマ字氏名） （研究者番号）	所属研究機関・部局・職 （機関番号）	備考
研究協力者	モーセ　ウィリアム (Moses William S.)	イリノイ大学アーバナ・シャンペーン校・Computer Science Department・助教授	

７．科研費を使用して開催した国際研究集会

〔国際研究集会〕　計0件

８．本研究に関連して実施した国際共同研究の実施状況

共同研究相手国	相手方研究機関