

FPGA-Accelerated Analytics: From Single Nodes to Clusters

Other titles in Foundations and Trends® in Databases

Algorithmic Aspects of Parallel Data Processing

Paraschos Koutris, Semih Salihoglu and Dan Suciu

ISBN: 978-1-68083-406-2

Data Infrastructure for Medical Research

Thomas Heinis and Anastasia Ailamaki

ISBN: 978-1-68083-348-5

Main Memory Database Systems

Franz Faerber, Alfons Kemper, Per-Ake Larson, Justin Levandoski,

Thomas Neumann and Andrew Pavlo

ISBN: 978-1-68083-324-9

Query Processing on Probabilistic Data: A Survey

Guy Van den Broeck and Dan Suciu

ISBN: 978-1-68083-314-0

Big Graph Analytics Platforms

Da Yan, Yingyi Bu, Yuanyuan Tian and Amol Deshpande

978-1-68083-242-6

FPGA-Accelerated Analytics: From Single Nodes to Clusters

Zsolt István

IMDEA Software Institute

Spain

zsolt.istvan@imdea.org

Kaan Kara

Oracle Labs

Switzerland

kaan.kara@oracle.com

David Sidler

Microsoft Corporation

USA

david.sidler@microsoft.com

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Databases

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

Z. István, K. Kara and D. Sidler. *FPGA-Accelerated Analytics: From Single Nodes to Clusters*. Foundations and Trends[®] in Databases, vol. 9, no. 2, pp. 101–208, 2020.

ISBN: 978-1-68083-735-3

© 2020 Z. István, K. Kara and D. Sidler

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Databases
Volume 9, Issue 2, 2020
Editorial Board

Editor-in-Chief

Joseph M. Hellerstein
University of California at Berkeley
United States

Editors

Anastasia Ailamaki
EPFL

Peter Bailis
Stanford University

Mike Cafarella
University of Michigan

Michael Carey
University of California, Irvine

Surajit Chaudhuri
Microsoft Research

Minos Garofalakis
Technical University Crete

Ihab Ilyas
University of Waterloo

Christopher Olston
Google

Jignesh Patel
University of Wisconsin

Chris Re
Stanford University

Gerhard Weikum
Max Planck Institute Saarbrücken

Editorial Scope

Topics

Foundations and Trends® in Databases publishes survey and tutorial articles in the following topics:

- Data Models and Query Languages
- Query Processing and Optimization
- Storage, Access Methods, and Indexing
- Transaction Management, Concurrency Control and Recovery
- Deductive Databases
- Parallel and Distributed Database Systems
- Database Design and Tuning
- Metadata Management
- Object Management
- Trigger Processing and Active Databases
- Data Mining and OLAP
- Approximate and Interactive Query Processing
- Data Warehousing
- Adaptive Query Processing
- Data Stream Management
- Search and Query Integration
- XML and Semi-Structured Data
- Web Services and Middleware
- Data Integration and Exchange
- Private and Secure Data Management
- Peer-to-Peer, Sensornet and Mobile Data Management
- Scientific and Spatial Data Management
- Data Brokering and Publish/Subscribe
- Data Cleaning and Information Extraction
- Probabilistic Data Management

Information for Librarians

Foundations and Trends® in Databases, 2020, Volume 9, 4 issues. ISSN paper version 1931-7883. ISSN online version 1931-7891. Also available as a combined paper and online subscription.

Contents

1	Introduction	2
2	Background	6
2.1	What Are FPGAs?	6
2.2	FPGAs vs. ASICs	9
2.3	Where to Deploy Acceleration?	11
3	FPGAs in Datacenters and Clouds	14
3.1	The Path to Adoption	14
3.2	Interfacing with FPGA Accelerators	16
3.3	Database Integration Case Studies	23
4	Designing Accelerators with FPGAs	32
4.1	Design Guidelines	33
4.2	Core SQL: Like and Regular Expressions	35
4.3	Core SQL: Group-by Aggregation	42
4.4	Core SQL: Where Predicates and Filtering	52
4.5	Machine Learning: K-Means Clustering	56
4.6	Machine Learning: Stochastic Gradient Descent	65
4.7	Distributed Joins: Data Partitioning	79

5	Future Challenges and Architectures	86
5.1	Datacenter and Cloud Architecture Trends	86
5.2	Device Trends	87
5.3	Sharing FPGAs and Accelerators	89
5.4	Programming FPGAs	91
6	Closing Remarks	93
	References	94

FPGA-Accelerated Analytics: From Single Nodes to Clusters

Zsolt István¹, Kaan Kara² and David Sidler³

¹*IMDEA Software Institute, Spain; zsolt.istvan@imdea.org*

²*Oracle Labs, Switzerland; kaan.kara@oracle.com*

³*Microsoft Corporation, USA; david.sidler@microsoft.com*

ABSTRACT

In this monograph, we survey recent research on using reconfigurable hardware accelerators, namely, Field Programmable Gate Arrays (FPGAs), to accelerate analytical processing. Such accelerators are being adopted as a way of overcoming the recent stagnation in CPU performance because they can implement algorithms differently from traditional CPUs, breaking traditional trade-offs. As such, it is timely to discuss their benefits in the context of analytical processing, both as an accelerator within a single node database and as part of distributed data analytics pipelines. We present guidelines for accelerator design in both scenarios, as well as, examples of integration within full-fledged Relational Databases. We do so through the prism of recent research projects that explore how emerging compute-intensive operations in databases can benefit from FPGAs. Finally, we highlight future research challenges in programmability and integration, and cover architectural trends that are propelling the rapid adoption of accelerators in datacenters and the cloud.

1

Introduction

Big Data has been instrumental to our lives in the last decade and has led to scientific insights, the boom of Machine Learning and the emergence of novel on-line services. Datacenters that are hosting such data-intensive applications are facing however an important challenge: the amount of data that needs to be stored and processed is increasing at an exponential rate whereas traditional processor performance has been stagnating for years. As Moore's law and Dennard scaling taper off, CPU transistor counts are growing less year-on-year and, due to heat dissipation issues, it is becoming challenging to power on all parts of CPUs at high clock rates at the same time, resulting in "dark silicon". The stagnation of CPU performance, however, presents opportunities as well: CPUs can now incorporate more heterogeneous components without having to keep them powered on at all times. In addition, for many workloads, replacing CPUs partially or entirely with specialized hardware has become the more economic choice. These devices can utilize transistors more efficiently for the tasks at hand, delivering more performance for the same energy budget.

Driven by the above described trends, and in order to keep up with growing data sizes, data processing and management applications have

been becoming increasingly distributed. Today, applications built with platforms such as Apache Spark, routinely span hundreds of server nodes. While helpful in reducing compute bottlenecks, this distribution brings new data movement bottlenecks at various levels of the software and hardware architecture. In this monograph we focus on how specialized hardware accelerators can provide an answer to the compute stagnation problem and showing how they can be also helpful in reducing data movement bottlenecks by placing them in the right location within the computer architecture.

There are many different technologies one could use for building accelerators but one particular technology stands out as a middle ground between energy efficiency and versatility: Field Programmable Gate Arrays (FPGAs). FPGAs make it possible to express algorithms in ways that are fundamentally different from CPUs or GPUs: FPGAs have no instruction sets and functionality is laid out directly as circuitry. As opposed to traditional CPUs that concentrate on-chip memory into layers of caches, FPGAs have small SRAM memories distributed throughout the device that can be configured flexibly and can be co-located with computation. These differences to CPUs result in a higher performance in the same or lower energy footprint, making FPGAs attractive both as accelerators and as energy-efficient replacements of software solutions. FPGAs are *reconfigurable* meaning that they can implement different functionality over time and can be reprogrammed from software. Once programmed, they act as integrated circuits (ASICs), bringing significant improvements in energy efficiency when compared to CPUs.

Today, FPGAs are available in most clouds (e.g., Amazon F1 Instances, [n.d.](#); Firestone, [2016](#); Putnam, [2014](#); Weerasinghe *et al.*, [2016](#)) as accelerators. As a prominent example, the Microsoft Catapult project uses FPGAs to create programmable network-interface cards (NICs) to offload tenant network functions from the CPU. Meanwhile, programmable co-processors based on FPGAs are also becoming more common, for instance, with projects such as Intel Xeon+FPGA (Gupta, [2015](#)). Storage devices/nodes are also increasingly more programmable,

with examples such as Samsung's SmartSSD¹ and Amazon Redshift AQUA.

Several early projects of FPGA-based database acceleration proposed PCIe-attached accelerator cards a decade ago, demonstrating that FPGAs are able to speed up projection and selection (Dennl *et al.*, 2012; Salami *et al.*, 2017; Sukhwani *et al.*, 2013; Wang *et al.*, 2016a; Woods *et al.*, 2013), aggregation (Dennl *et al.*, 2013; Salami *et al.*, 2017), joins and even sorting (Casper and Olukotun, 2014; Sukhwani *et al.*, 2013; Zhang *et al.*, 2016), by an order of magnitude when compared to commonly used row-stores, such as MySQL and PostgreSQL. However, in many previous systems, once all integration costs were factored in, in particular the cost of data transfers over PCIe, the benefits were significantly reduced. In the meantime, however, the bandwidth of interconnects (PCIe, NVMe, etc.) and networks has been increasing at a steady pace, making FPGA acceleration once again an attractive option. Furthermore, due to increased data sizes, there are more and more distributed databases that suffer from various data movement bottlenecks, for instance, when retrieving data from distributed storage nodes or when shuffling tuples between compute nodes for a join operation. Specialized hardware can be used to move computation closer to the source, and by reducing data sizes through filtering or transformations close to storage, memory, etc., reduce data movement bottlenecks.

In this monograph we explore what is required for integrating these accelerators with software systems (focusing on database management systems) and how one should design the acceleration functionality to ensure that the overall speedups in the system are worth the additional complexity. The monograph provides a detailed look into several representative examples of FPGA-accelerated databases and the internals of accelerated SQL operators, Machine Learning operators and data shuffling operators.

Monograph Structure. In this section (Section 1), we outlined the reasons why FPGAs and similar specialized hardware have become not only economically feasible for database acceleration but, in many

¹<https://samsungsemiconductor-us.com/smartssd/index.html>.

cases, even necessary to keep up with data growth. We also provided an intuition on how, thanks to a fundamentally different execution model from CPUs and GPUs, these devices offer benefits in efficiency. The rest of the monograph is structured as follows:

- In Section 2 we present a background on FPGAs and highlight their differences with ASICs. Readers already familiar with FPGAs and their development work flow might consider skipping this section.
- Section 3 covers salient aspects of integrating FPGAs in data processing systems, in various locations of the software and hardware architecture. We describe several representative examples of how FPGAs are integrated with databases and what features the frameworks that enable this need to provide.
- In Section 4 we turn our attention to individual operators from the domain of core SQL acceleration, Machine Learning and Distributed Joins. We present both high level guidelines to help future FPGA programmers design circuits adequate to the task and deep dive into the design and performance characteristics of representative operators, implemented by the authors.
- Section 5 describes the remaining challenges related to programming and sharing FPGA accelerators more easily in datacenters and the cloud. This section also provides a peek into the future of re-programmable hardware and the opportunities this will bring for database management systems.

References

- Achronix Vectorpath with Speedster7t FPGAs (n.d.). <https://www.achronix.com/vectorpath/>.
- Agron, J. and D. Andrews (2009). “Building heterogeneous reconfigurable systems with a hardware microkernel”. In: *Proceedings of the 7th IEEE/ACM International Conference on Hardware/Software Codesign and System Synthesis*. ACM. 393–402.
- Aho, A. V. and J. D. Ullman (1992). *Foundations of Computer Science*. Computer Science Press.
- Alonso, G., Z. Istvan, K. Kara, M. Owaida, and D. Sidler (2019). “doppioDB 1.0: Machine learning inside a relational engine”. *IEEE Data Engineering Bulletin*. 42(2): 19–31.
- Amazon F1 Instances (n.d.). aws.amazon.com/ec2/instance-types/f1/.
- Asiatici, M., N. George, K. Vipin, S. A. Fahmy, and P. Ienne (2017). “Virtualized execution runtime for FPGA accelerators in the cloud”. *IEEE Access*. 5: 1900–1910.
- Balkesen, C., J. Teubner, G. Alonso, and M. T. Özsu (2013). “Main-memory hash joins on multi-core CPUs: Tuning to the underlying hardware”. In: *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE. 362–373.
- Barber, R., G. Lohman, I. Pandis, V. Raman, R. Sidle, G. Attaluri, N. Chainani, S. Lightstone, and D. Sharpe (2014). “Memory-efficient hash joins”. *Proceedings of the VLDB Endowment*. 8(4): 353–364.

- Barthels, C., S. Loesing, G. Alonso, and D. Kossmann (2015). “Rack-scale in-memory join processing using RDMA”. In: *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM. 1463–1475.
- Barthels, C., I. Müller, T. Schneider, G. Alonso, and T. Hoeffler (2017). “Distributed join algorithms on thousands of cores”. *Proceedings of the VLDB Endowment*. 10(5): 517–528.
- Bispo, J., I. Sourdis, J. M. Cardoso, and S. Vassiliadis (2006). “Regular expression matching for reconfigurable packet inspection”. In: *2006 IEEE International Conference on Field Programmable Technology*. IEEE. 119–126.
- Blott, M., K. Karras, L. Liu, K. Vissers, J. Bär, and Z. István (2013). “Achieving 10 Gbps line-rate key-value stores with FPGAs”. In: *5th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 13)*.
- Böhm, C., M. Perdacher, and C. Plant (2017). “Multi-core K-means”. In: *Proceedings of the 2017 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics. 273–281.
- Boppana, V., S. Ahmad, I. Ganusov, V. Kathail, V. Rajagopalan, and R. Wittig (2015). “UltraScale+ MPSoC and FPGA families”. In: *2015 IEEE Hot Chips 27 Symposium (HCS)*. IEEE. 1–37.
- Bosshart, P., D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker (2014). “P4: Programming protocol-independent packet processors”. *ACM SIGCOMM Computer Communication Review*. 44(3): 87–95.
- Bottou, L. (2010). “Large-scale machine learning with stochastic gradient descent”. In: *Proceedings of COMPSTAT’2010*. Springer. 177–186.
- Byma, S., J. G. Steffan, H. Bannazadeh, A. L. Garcia, and P. Chow (2014). “FPGAs in the cloud: Booting virtualized hardware accelerators with Openstack”. In: *2014 IEEE 22nd Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE. 109–116.
- Casper, J. and K. Olukotun (2014). “Hardware acceleration of database operations”. In: *Proceedings of the 2014 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM. 151–160.

- Caulfield, A. M., E. S. Chung, A. Putnam, H. Angepat, I. Fowers, M. Haselman, S. Heil, M. Humphret, P. Kaur, J.-Y. Kim, D. Lo, T. Massengill, K. Ovtcharov, M. Papamichael, L. Woods, S. Lanka, D. Chiou, and D. Burger (2016). “A cloud-scale acceleration architecture”. In: *MICRO-49: The 49th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM. 1–13.
- Chalamalasetti, S. R., K. Lim, M. Wright, A. AuYoung, P. Ranganathan, and M. Margala (2013). “An FPGA memcached appliance”. In: *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. ACM. 245–254.
- Chen, F., Y. Shan, Y. Zhang, Y. Wang, H. Franke, X. Chang, and K. Wang (2014). “Enabling FPGAs in the cloud”. In: *Proceedings of the 11th ACM Conference on Computing Frontiers*. ACM. 3.
- Chung, E., J. Fowers, K. Ovtcharov, M. Papamichael, A. Caulfield, T. Massengill, M. Liu, D. Lo, S. Alkalay, M. Haselman, and M. Abeydeera (2018). “Serving DNNs in real time at datacenter scale with project brainwave”. *IEEE Micro*. 38(2): 8–20.
- Da Silva, H. C., F. Pisani, and E. Borin (2016). “A comparative study of SYCL, OpenCL, and OpenMP”. In: *2016 International Symposium on Computer Architecture and High Performance Computing Workshops (SBAC-PADW)*. IEEE. 61–66.
- Dennl, C., D. Ziener, and J. Teich (2012). “On-the-fly composition of FPGA-based SQL query accelerators using a partially reconfigurable module library”. In: *2012 IEEE 20th Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE. 45–52.
- Dennl, C., D. Ziener, and J. Teich (2013). “Acceleration of SQL restrictions and aggregations through FPGA-based dynamic partial reconfiguration”. In: *2013 IEEE 21st Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE. 25–28.
- Dhanachandra, N., K. Mangle, and Y. J. Chanu (2015). “Image segmentation using K-means clustering algorithm and subtractive clustering algorithm”. *Procedia Computer Science*. 54: 764–771.

- Dhawan, U. and A. DeHon (2015). “Area-efficient near-associative memories on FPGAs”. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*. 7(4): 30.
- Do, J., Y.-S. Kee, J. M. Patel, C. Park, K. Park, and D. J. DeWitt (2013). “Query processing on smart SSDs: Opportunities and challenges”. In: *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. ACM. 1221–1230.
- Dragojević, A., D. Narayanan, O. Hodson, and M. Castro (2014). “Farm: Fast remote memory”. In: *11th USENIX Symposium on Networked Systems Design and Implementation*. USENIX Association. 401–414.
- Dragojević, A., D. Narayanan, E. B. Nightingale, M. Renzelmann, A. Shamis, A. Badam, and M. Castro (2015). “No compromises: Distributed transactions with consistency, availability, and performance”. In: *Proceedings of the 25th Symposium on Operating Systems Principles*. 54–70.
- Eran, H., L. Zeno, M. Tork, G. Malka, and M. Silberstein (2019). “NICA: An infrastructure for inline acceleration of network applications”. In: *2019 USENIX Annual Technical Conference*. 345–362.
- Esmailzadeh, H., E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger (2011). “Dark silicon and the end of multicore scaling”. In: *2011 38th Annual International Symposium on Computer Architecture (ISCA)*. IEEE. 365–376.
- Estlick, M., M. Leeser, J. Theiler, and J. J. Szymanski (2001). “Algorithmic transformations in the implementation of K-means clustering on reconfigurable hardware”. In: *Proceedings of the 2001 ACM/SIGDA Ninth International Symposium on Field Programmable Gate Arrays*. ACM. 103–110.
- Firestone, D. (2016). “SmartNIC: Accelerating azure’s network with FPGAs on OCS servers”. In: *Open Compute Summit 2016*.
- Francisco, P. (2011). “The netezza data appliance architecture: A platform for high performance data warehousing and analytics.” *IBM Redbooks*.
- Fukuda, E. S., H. Inoue, T. Takenaka, D. Kim, T. Sadahisa, T. Asai, and M. Motomura (2014). “Caching memcached at reconfigurable network interface”. In: *2014 24th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE. 1–6.

- Gaide, B., D. Gaitonde, C. Ravishankar, and T. Bauer (2019). “Xilinx adaptive compute acceleration platform: Versal™ architecture”. In: *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. 84–93.
- Gao, M. and C. Kozyrakis (2016). “HRL: Efficient and flexible reconfigurable logic for near-data processing”. In: *IEEE International Symposium on High-Performance Computer Architecture*. IEEE. 126–137.
- Gokhale, M., J. Frigo, K. McCabe, J. Theiler, C. Wolinski, and D. Lavenier (2003). “Experience with a hybrid processor: K-means clustering”. *The Journal of Supercomputing*. 26(2): 131–148.
- Gupta, P. K. (2015). “Xeon+FPGA platform for the data center”. In: *Fourth Workshop on the Intersections of Computer Architecture and Reconfigurable Logic*. Vol. 119.
- Gupta, S., A. Agrawal, K. Gopalakrishnan, and P. Narayanan (2015). “Deep learning with limited numerical precision”. In: *International Conference on Machine Learning*. 1737–1746.
- Hadian, A. and S. Shahrivari (2014). “High performance parallel K-means clustering for disk-resident datasets on multi-core CPUs”. *The Journal of Supercomputing*. 69(2): 845–863.
- Happe, M., A. Traber, and A. Keller (2015). “Preemptive hardware multitasking in ReconOS”. In: *International Symposium on Applied Reconfigurable Computing*. Springer. 79–90.
- He, Z., D. Sidler, Z. István, and G. Alonso (2018). “A flexible K-means operator for hybrid databases”. In: *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE. 368–3683.
- IBM/Netezza (2011). “The netezza data appliance architecture: A platform for high performance data warehousing and analytics”. <http://www.redbooks.ibm.com/abstracts/redp4725.html>.
- Ismail, A. and L. Shannon (2011). “FUSE: Front-end user framework for O/S abstraction of hardware accelerators”. In: *2011 IEEE 19th Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE. 170–177.

- Istvan, Z., L. Woods, and G. Alonso (2014). “Histograms as a side effect of data movement for big data”. In: *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. ACM. 1567–1578.
- István, Z., G. Alonso, M. Blott, and K. Vissers (2015). “A hash table for line-rate data processing”. *ACM TRETS*. 8(2): 13.1–13.15.
- István, Z., D. Sidler, G. Alonso, and M. Vukolic (2016). “Consensus in a box: Inexpensive coordination in hardware”. In: *13th USENIX Symposium on Networked Systems Design and Implementation*. USENIX Association. 425–438.
- István, Z., D. Sidler, and G. Alonso (2017). “Caribou: Intelligent distributed storage”. *Proceedings of the VLDB Endowment*. 10(11): 1202–1213.
- István, Z., G. Alonso, and A. Singla (2018). “Providing multi-tenant services with FPGAs: Case study on a key-value store”. In: *28th International Conference on Field Programmable Logic and Applications*. IEEE. 119–124.
- Iturbe, X., K. Benkrid, C. Hong, A. Ebrahim, R. Torrego, I. Martinez, T. Arslan, and J. Perez (2013). “R3TOS: A novel reliable reconfigurable real-time operating system for highly adaptive, efficient, and dependable computing on FPGAs”. *IEEE Transactions on Computers*. 62(8): 1542–1556.
- Jamesmanoharan, J., S. H. Ganesh, M. L. P. Felciah, and A. K. Shafreenbanu (2014). “Discovering students’ academic performance based on GPA using K-means clustering algorithm”. In: *World Congress on Computing and Communication Technologies*. IEEE. 200–202.
- Jo, I., D.-H. Bae, A. S. Yoon, J. U. Kang, S. Cho, D. D. Lee, and J. Jeong (2016). “YourSQL: A high-performance database system leveraging in-storage computing”. *Proceedings of the VLDB Endowment*. 9(12): 924–935.
- Jun, S.-W., M. Liu, and K. E. Fleming (2014). “Scalable multi-access flash store for big data analytics”. In: *Proceedings of the 2014 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. 55–64.

- Junqueira, F. P., B. C. Reed, and M. Serafini (2011). “Zab: High-performance broadcast for primary-backup systems”. In: *2011 IEEE/IFIP 41st International Conference on Dependable Systems & Networks (DSN)*. IEEE. 245–256.
- Kalia, A., M. Kaminsky, and D. G. Andersen (2014). “Using RDMA efficiently for key-value services”. In: *Proceedings of the 2014 ACM Conference on SIGCOMM*. 295–306.
- Kara, K., D. Alistarh, G. Alonso, O. Mutlu, and C. Zhang (2017). “FPGA-accelerated dense linear machine learning: A precision-convergence trade-off”. In: *2017 IEEE 25th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE. 160–167.
- Kara, K., Z. Wang, C. Zhang, and G. Alonso (2020). “doppioDB 2.0: Hardware techniques for improved integration of machine learning into databases”. *Proceedings of the VLDB Endowment*. 12(12): 1818–1821.
- Kim, J. K., Z. Zhang, and J. A. Fessler (2011). “Hardware acceleration of iterative image reconstruction for X-ray computed tomography”. In: *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 1697–1700.
- Klimovic, A., C. Kozyrakis, E. Thereksa, B. John, and S. Kumar (2016). “Flash storage disaggregation”. In: *Proceedings of the Eleventh European Conference on Computer Systems*. ACM. 1–15.
- Knodel, O., P. R. Genssler, and R. G. Spallek (2017). “Migration of long-running tasks between reconfigurable resources using virtualization”. *ACM SIGARCH Computer Architecture News*. 44(4): 56–61.
- Koeplinger, D., M. Feldman, R. Prabhakar, Y. Zhang, S. Hadjis, R. Fiszel, T. Zhao, L. Nardi, A. Pedram, C. Kozyrakis, and K. Olukotun (2018). “Spatial: A language and compiler for application accelerators”. In: *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 296–311.

- Koo, G., K. K. Matam, I. Te, H. V. K. G. Narra, J. Li, H.-W. Tseng, S. Swanson, and M. Annavaram (2017). “Summarizer: Trading communication with computing near storage”. In: *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture. MICRO-50 '17*. Cambridge, Massachusetts. 219–231.
- Kuhring, L., E. Garcia, and Z. István (2019). “Specialize in moderation—building application-aware storage services using FPGAs in the datacenter”. In: *11th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 19)*. USENIX Association.
- Kumm, M. and P. Zipf (2014). “Pipelined compressor tree optimization using integer linear programming”. In: *2014 24th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE. 1–8.
- Lang, H., V. Leis, M.-C. Albutiu, T. Neumann, and A. Kemper (2013). “Massively parallel NUMA-aware hash joins”. In: *In Memory Data Management and Analysis*. Springer. 3–14.
- Lavasani, M., H. Angepat, and D. Chiou (2014). “An FPGA-based in-line accelerator for memcached”. *IEEE Computer Architecture Letters*. 13(2): 57–60.
- Lesser, B., M. Mücke, and W. N. Gansterer (2011). “Effects of reduced precision on floating-point SVM classification accuracy”. *Procedia Computer Science*. 4: 508–517.
- Li, Y., K. Zhao, X. Chu, and J. Liu (2013). “Speeding up K-means algorithm by GPUs”. *Journal of Computer and System Sciences*. 79(2): 216–229.
- Li, B., Z. Ruan, W. Xiao, Y. Lu, Y. Xiong, A. Putnam, E. Chen, and L. Zhang (2017). “KV-direct: High-performance in-memory key-value store with programmable NIC”. In: *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM. 137–152.
- Liu, W.-C., J.-L. Huang, and M. S. Chen (2005). “Kacu: K-means with hardware centroid-updating”. In: *Conference, Emerging Information Technology 2005*. IEEE.
- Loesing, S., M. Pilman, T. Etter, and D. Kossmann (2015). “On the design and scalability of distributed shared-data databases”. In: *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM. 663–676.

- Lübbers, E. and M. Platzner (2009). “ReconOS: Multithreaded programming for reconfigurable computers”. *ACM Transactions on Embedded Computing Systems (TECS)*. 9(1): 8.
- Mai, L., L. Rupprecht, A. Alim, P. Costa, M. Migliavacca, P. Pietzuch, and A. L. Wolf (2014). “NetAgg: Using middleboxes for application-specific on-path aggregation in data centres”. In: *Proceedings of the 10th ACM International Conference on Emerging Networking Experiments and Technologies*. 249–262.
- Manegold, S., P. Boncz, and M. Kersten (2002). “Optimizing main-memory join on modern hardware”. *IEEE TKDE*. 14(4): 709–730.
- Mitchell, C., Y. Geng, and J. Li (2013). “Using one-sided RDMA reads to build a fast, CPU-efficient key-value store”. In: *USENIX Annual Technical Conference*. USENIX Association. 103–114.
- Morales-Villanueva, A., R. Kumar, and A. Gordon-Ross (2016). “Configuration prefetching and reuse for preemptive hardware multitasking on partially reconfigurable FPGAs”. In: *2016 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. 1505–1508.
- Mücke, M., B. Lesser, and W. N. Gansterer (2010). “Peak performance model for a custom precision floating-point dot product on FPGAs”. In: *European Conference on Parallel Processing*. Springer. 399–406.
- Nakahara, H., T. Sasao, and M. Matsuura (2011). “A regular expression matching circuit based on a decomposed automaton”. In: *Reconfigurable Computing: Architectures, Tools and Applications*. Springer.
- Nurvitadhi, E., D. Kwon, A. Jafari, A. Boutros, J. Sim, P. Tomson, H. Sumbul, G. Chen, P. Knag, R. Kumar, and R. Krishnamurthy (2019). “Evaluating and enhancing intel® stratix® 10 FPGAs for persistent real-time AI”. In: *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. 119–119.
- Oliver, N., R. R. Sharma, S. Chang, B. Chitlur, E. Garcia, J. Grecco, A. Grier, N. Ijih, Y. Liu, P. Marolia, and H. Mitchel (2011). “A reconfigurable computing system based on a cache-coherent fabric”. In: *2011 International Conference on Reconfigurable Computing and FPGAs (ReConFig)*. IEEE. 80–85.
- OneAPI (2020). “The OneAPI specification”. URL: www.oneapi.com/spec/ (accessed 06/2020).

- OpenCAPI OC-Accel (n.d.). <https://github.com/OpenCAPI/oc-accel>.
- Oracle (2015). “Software in silicon: What it does and why”. <https://community.oracle.com/docs/DOC-932216>.
- Oskin, M., F. T. Chong, and T. Sherwood (1998). *Active Pages: A Computation Model for Intelligent Memory*. Vol. 26. IEEE Computer Society.
- Ousterhout, J., P. Agrawal, D. Erickson, C. Kozyrakis, J. Leverich, D. Mazières, S. Mitra, A. Narayanan, G. Parulkar, M. Rosenblum, and S. M. Rumble (2010). “The case for RAMClouds: Scalable high-performance storage entirely in DRAM”. *ACM SIGOPS Operating Systems Review*. 43(4): 92–105.
- Owaida, M., D. Sidler, K. Kara, and G. Alonso (2017). “Centaur: A framework for hybrid CPU-FPGA databases”. In: *2017 IEEE 25th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE. 211–218.
- Perez-Garcia, A. N., G. M. Tornez-Xavier, L. M. Flores-Nava, F. Gómez-Castañeda, and J. A. Moreno-Cadenas (2014). “Multilayer perceptron network with integrated training algorithm in FPGA”. In: *2014 11th International Conference on Electrical Engineering, Computing Science and Automatic Control (CCE)*. IEEE. 1–6.
- Pirk, H., J. Giceva, and P. Pietzuch (2019). “Thriving in the no man’s land between compilers and databases”.
- Pirk, H., O. Moll, M. Zaharia, and S. Madden (2016). “Voodoo—a vector algebra for portable database performance on modern hardware”. *Proceedings of the VLDB Endowment*. 9(14): 1707–1718.
- Plattner, C. and G. Alonso (2004). “Ganymed: Scalable replication for transactional web applications”. In: *ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*. Berlin, Heidelberg: Springer. 155–174.
- Polychroniou, O. and K. A. Ross (2014). “A comprehensive study of main-memory partitioning and its application to large-scale comparison and radix-sort”. In: *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. ACM. 755–766.
- Putnam, A. (2014). “Large-scale reconfigurable computing in a microsoft datacenter”. In: *Hot Chips 26 Symposium (HCS), 2014 IEEE*. IEEE. 1–38.

- Putnam, A., A. M. Caulfield, E. S. Chung, D. Chiou, K. Constantinides, J. Demme, H. Esmaeilzadeh, J. Fowers, G. P. Gopal, J. Gray, and M. Haselman (2014). “A reconfigurable fabric for accelerating large-scale datacenter services”. In: *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*. IEEE. 13–24.
- Rabieah, M. B. and C.-S. Bouganis (2016). “FPGASVM: A framework for accelerating kernelized support vector machine”. In: *Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications*. 68–84.
- Recht, B., C. Re, S. Wright, and F. Niu (2011). “Hogwild: A lock-free approach to parallelizing stochastic gradient descent”. In: *Advances in Neural Information Processing Systems*. 693–701.
- Root, C. and T. Mostak (2016). “MapD: A GPU-powered big data analytics and visualization platform”. In: *ACM SIGGRAPH 2016 Talks*. 1–2.
- Saegusa, T. and T. Maruyama (2006). “An FPGA implementation of K-means clustering for color images based on Kd-tree”. In: *2006 International Conference on Field Programmable Logic and Applications*. IEEE. 1–6.
- Salami, B., G. A. Malazgirt, O. Arcas-Abella, A. Yurdakul, and N. Sonmez (2017). “AxleDB: A novel programmable query processing platform on FPGA”. *Microprocessors and Microsystems*. 51: 142–164.
- Schuh, S., X. Chen, and J. Dittrich (2016). “An experimental comparison of thirteen relational equi-joins in main memory”. In: *Proceedings of the 2016 International Conference on Management of Data*. ACM. 1961–1976.
- Schuhknecht, F. M., P. Khanchandani, and J. Dittrich (2015). “On the surprising difficulty of simple things: The case of radix partitioning”. *Proceedings of the VLDB Endowment*. 8(9): 934–937.
- Sidler, D., G. Alonso, M. Blott, K. Karras, K. Vissers, and R. Carley (2015). “Scalable 10 Gbps TCP/IP stack architecture for reconfigurable hardware”. In: *2015 IEEE 23rd Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE. 36–43.

- Sidler, D., Z. István, and G. Alonso (2016). “Low-latency TCP/IP stack for data center applications”. In: *26th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE. 1–4.
- Sidler, D., Z. István, M. Owaida, and G. Alonso (2017). “Accelerating pattern matching queries in hybrid CPU-FPGA architectures”. In: *Proceedings of the 2017 ACM International Conference on Management of Data*. ACM. 403–415.
- Sidler, D., Z. Wang, M. Chiosa, A. Kulkarni, and G. Alonso (2020). “StRoM: Smart remote memory”. In: *Proceedings of the Fifteenth European Conference on Computer Systems*. ACM.
- Sitaridi, E. and K. Ross (2016). “GPU-accelerated string matching for database applications”. *The VLDB Journal*. 25: 719–740.
- Sukhwani, B., H. Min, M. Thoennes, P. Dube, B. Iyer, B. Brezzo, D. Dillenberger, and S. Asaad (2012). “Database analytics acceleration using FPGAs”. In: *Proceedings of the 21st International Conference on Parallel Architectures and Compilation Techniques*. ACM. 411–420.
- Sukhwani, B., H. Min, M. Thoennes, P. Dube, B. Brezzo, S. Asaad, and D. E. Dillenberger (2013). “Database analytics: A reconfigurable-computing approach”. *IEEE Micro*. 34(1): 19–29.
- Tang, Q. Y. and M. A. Khalid (2016). “Acceleration of K-means algorithm using Altera SDK for OpenCL”. *ACM TRETTS*. 10(1): 6.
- Tang, Z., K. Liu, J. Xiao, L. Yang, and Z. Xiao (2017). “A parallel k-means clustering algorithm based on redundance elimination and extreme points optimization employing MapReduce”. *Concurrency and Computation: Practice and Experience*. 29(20): e4109.
- Taylor, M. B. (2012). “Is dark silicon useful? Harnessing the four horsemen of the coming dark silicon apocalypse”. In: *DAC Design Automation Conference 2012*. IEEE. 1131–1136.
- Teubner, J. and L. Woods (2013). “Data processing on FPGAs”. *Morgan & Claypool Synthesis Lectures on Data Management*. 5(2): 1–118.
- Teubner, J., L. Woods, and C. Nie (2012). “Skeleton automata for FPGAs: Reconfiguring without reconstructing”. In: *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. ACM. 229–240.

- Thin, T. N., S. Kittitornkun, and S. Tomiyama (2007). “Applying cuckoo hashing for FPGA-based pattern matching in NIDS/NIPS”. In: *2007 International Conference on Field-Programmable Technology*. IEEE. 121–128.
- Vaishnav, A., K. D. Pham, and D. Koch (2018). “A survey on FPGA virtualization”. In: *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE. 131–1317.
- Vaishnav, A., K. D. Pham, and D. Koch (2019). “Heterogeneous resource-elastic scheduling for CPU+FPGA architectures”. In: *Proceedings of the 10th International Symposium on Highly-Efficient Accelerators and Reconfigurable Technologies*. 1–6.
- Wang, X. and M. Leeser (2007). “K-Means clustering for multispectral images using floating-point divide”. In: *15th Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM 2007)*. IEEE. 151–162.
- Wang, Z., B. He, and W. Zhang (2015). “A study of data partitioning on OpenCL-based FPGAs”. In: *2015 25th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE. 1–8.
- Wang, Z., J. Paul, H. Y. Cheah, B. He, and W. Zhang (2016a). “Relational query processing on OpenCL-based FPGAs”. In: *2016 26th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE.
- Wang, Z., S. Zhang, B. He, and W. Zhang (2016b). “Melia: A MapReduce framework on OpenCL-based FPGAs”. *IEEE Transactions on Parallel and Distributed Systems*. 27(12): 3547–3560.
- Weerasinghe, J., F. Abel, C. Hagleitner, and A. Herkersdorf (2015). “Enabling FPGAs in hyperscale data centers”. In: *2015 IEEE 12th Intl. Conf. on Ubiquitous Intelligence and Computing and 2015 IEEE 12th Intl. Conf. on Autonomic and Trusted Computing and 2015 IEEE 15th Intl. Conf. on Scalable Computing and Communications and Its Associated Workshops (UIC-ATC-ScalCom)*. IEEE. 1078–1086.
- Weerasinghe, J., R. Polig, F. Abel, and C. Hagleitner (2016). “Network-attached FPGAs for data center applications”. In: *2016 International Conference on Field-Programmable Technology (FPT)*. IEEE. 36–43.

- Weiss, R. (2012). “A technical overview of the oracle exadata database machine and exadata storage server”. *Oracle White Paper*. Oracle Corporation, Redwood Shores.
- Woods, L., J. Teubner, and G. Alonso (2010). “Complex event detection at wire speed with FPGAs”. *Proceedings of the VLDB Endowment*. 3(1–2): 660–669.
- Woods, L., G. Alonso, and J. Teubner (2013). “Parallel computation of skyline queries”. In: *IEEE 21st Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE. 1–8.
- Woods, L., Z. István, and G. Alonso (2014). “Ibex – An intelligent storage engine with support for advanced SQL off-loading”. *PVLDB*. 7(11): 963–974.
- Woods, L., G. Alonso, and J. Teubner (2015). “Parallelizing data processing on FPGAs with shifter lists”. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*. 8(2): Article 7.
- Xu, S., S. Lee, S.-W. Jun, M. Liu, and J. Hicks (2016). “BlueCache: A scalable distributed flash-based key-value store”. *Proceedings of the VLDB Endowment*. 10(4): 301–312.
- Yang, Y.-H. E., W. Jiang, and V. K. Prasanna (2008). “Compact architecture for high-throughput regular expression matching on FPGA”. In: *Proceedings of the 4th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*. ACM. 30–39.
- You, Y., X. Lian, J. Liu, H.-F. Yu, I. S. Dhillon, J. Demmel, and C.-J. Hsieh (2016). “Asynchronous parallel greedy coordinate descent”. In: *Advances in Neural Information Processing Systems*. 4682–4690.
- Zamanian, E., C. Binnig, T. Harris, and T. Kraska (2017). “The end of a myth: Distributed transactions can scale”. *Proceedings of the VLDB Endowment*. 10(6): 685–696.
- Zhang, C., R. Chen, and V. Prasanna (2016). “High throughput large scale sorting on a CPU-FPGA heterogeneous platform”. In: *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE. 148–155.
- Zhang, H., J. Li, K. Kara, D. Alistarh, J. Liu, and C. Zhang (2017a). “Zipml: Training linear models with end-to-end low precision, and a little bit of deep learning”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org. 4035–4043.

- Zhang, J., Y. Xiong, N. Xu, R. Shu, B. Li, P. Cheng, G. Chen, and T. Moscibroda (2017b). “The feniks FPGA operating system for cloud computing”. In: *Proceedings of the 8th Asia-Pacific Workshop on Systems*. ACM. 22.
- Zhu, Z., A. X. Liu, F. Zhang, and F. Chen (2018). “FPGA resource pooling in cloud computing”. *IEEE Transactions on Cloud Computing*. DOI: [10.1109/TCC.2018.2874011](https://doi.org/10.1109/TCC.2018.2874011).