Example Quote from Citations

Jason Cong
UCLA
ACM/IEEE Fellow

of different optimization strategies. [20] and [25] provide valuable guidances for hardware designers to make good use of various optimization strategies. However, since they do not come up with an automation solution, accelerator developers still have to manually conduct design space exploration. Other studies like [11] provide a

Cong, Jason, Peng Wei, Cody Hao Yu, and Peng Zhang. “Automated accelerator generation and optimization with composable, parallel and pipeline architecture." **DAC, 2018.**

Gustavo Alonso
ETHZ
ACM/IEEE Fellow

to be 312 Million tuples/s for 511 partitions. The fastest to date FPGA data partitioning implementation is presented by Wang et al. [37] with 256 Million tuples/s for 8192 partitions. They improved an existing OpenCL implementation of a partitioner and deployed it on an FPGA. The data to

Kara, Kaan, Jana Giceva, and Gustavo Alonso. "Fpga-based data partitioning." **SIGMOD 2017.**

Example Quote from Citations

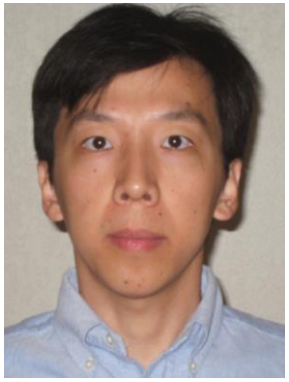

Yun Liang

PKU

a general performance model for accelerator design. Wang et al. [16] present an optimization framework for OpenCL programs for Altera FPGAs by employing a coarse-grained performance model. The optimization capability is funda-

Wang, Shuo, Yun Liang, and Wei Zhang. “Flexcl: An analytical performance model for opencl workloads on flexible fpgas.” **DAC, 2017.**