GraFF: A Multi-FPGA System with Memory Semantic Fabric for Scalable Graph Processing

Xu Zhang, Yisong Chang, Tianyue Lu,

Ke Liu, Ke Zhang, Mingyu Chen

Institute of Computing Technology, Chinese Academy of Science

Background and Motivation

1、Huge number of edges across FPGA boundaries leads to Inter-FPGA synchronization.

2、Long inter-FPGA fabric latency introduces severe overhead.

Q: Is it possible to fully overlap processing and synchronization?

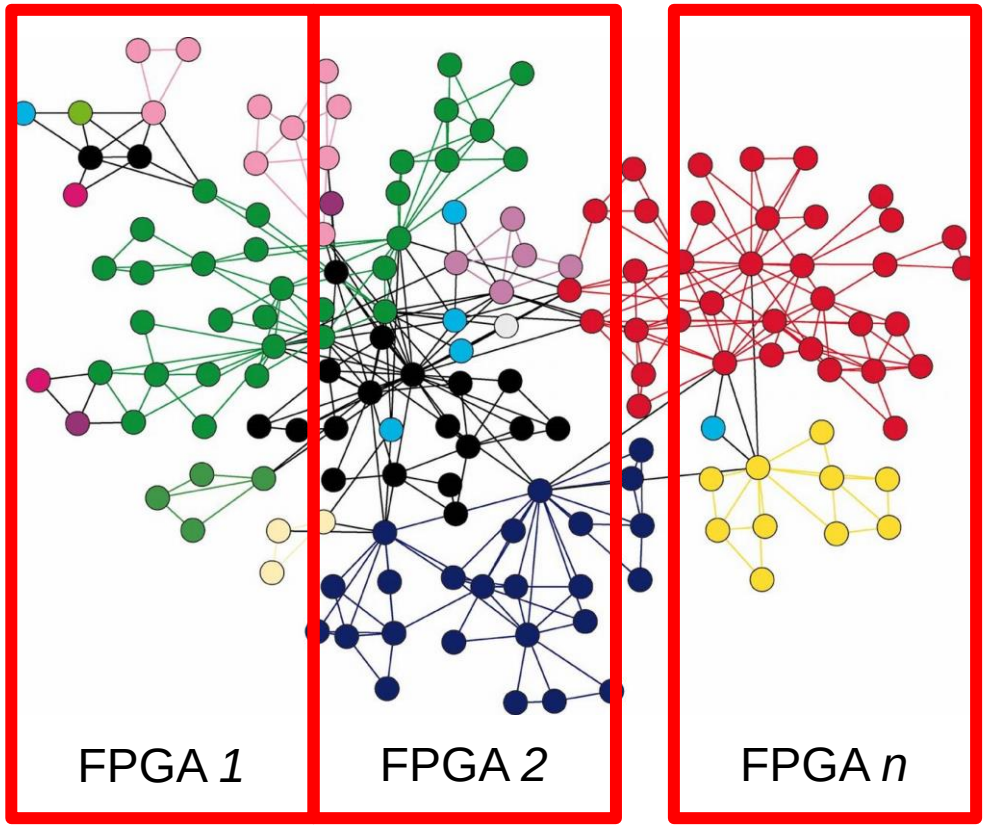

中科院计算所
Institute of Computing Technology, Chinese Academy of Sciences

Solution #1

- Splitting synchronization data into fine-grained high-concurrent memory semantic transactions.
  - atomic update
- Remote nodes receive transactions and update memory simultaneously.

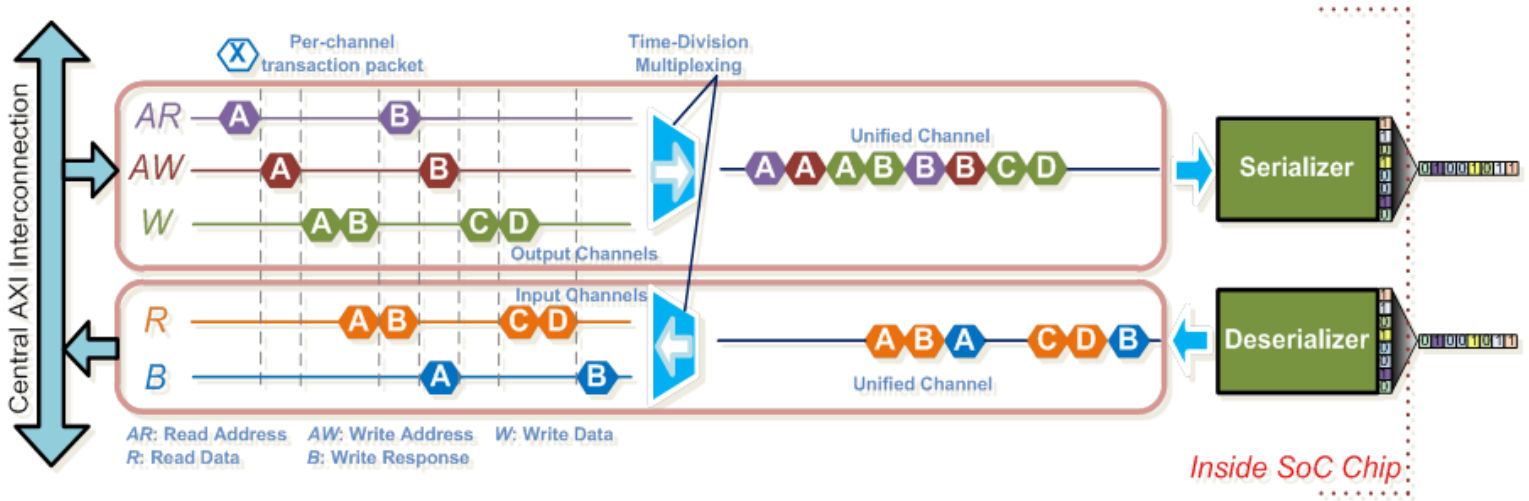

Institute of Computing Technology, Chinese Academy of Sciences

Solution #2

- GAS module splits superstep into Gather, Apply, Scatter phases.

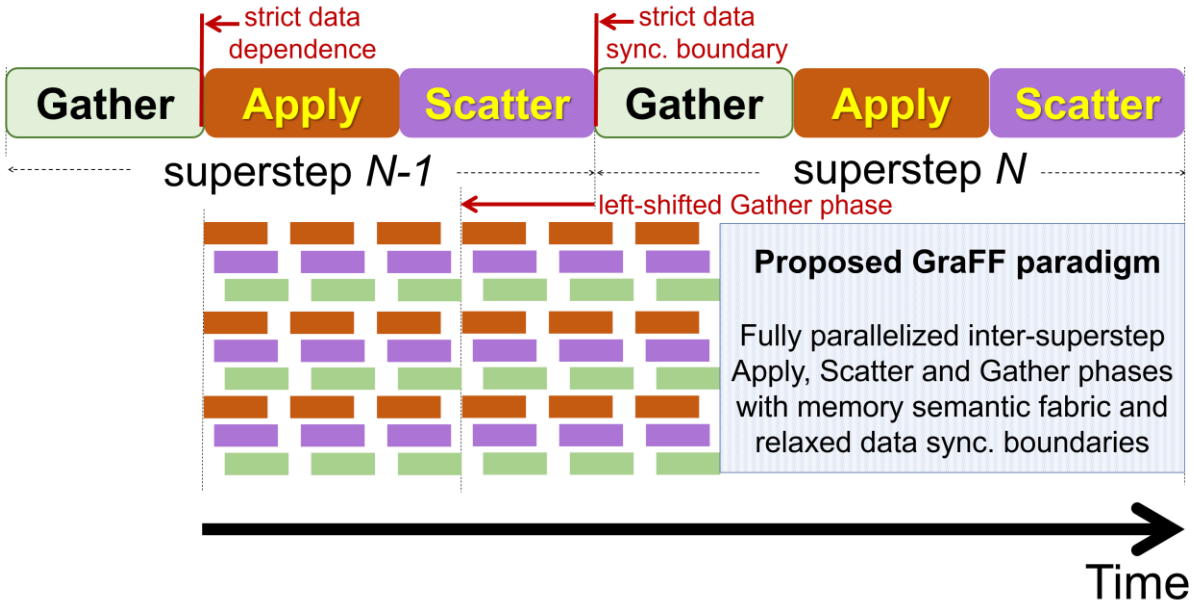

- Overlap synchronization (Apply, Scatter) in the superstep N-1 and processing (Gather) in the superstep N!

中科院计算所
Institute of Computing Technology, Chinese Academy of Sciences

GraFF Design Overview

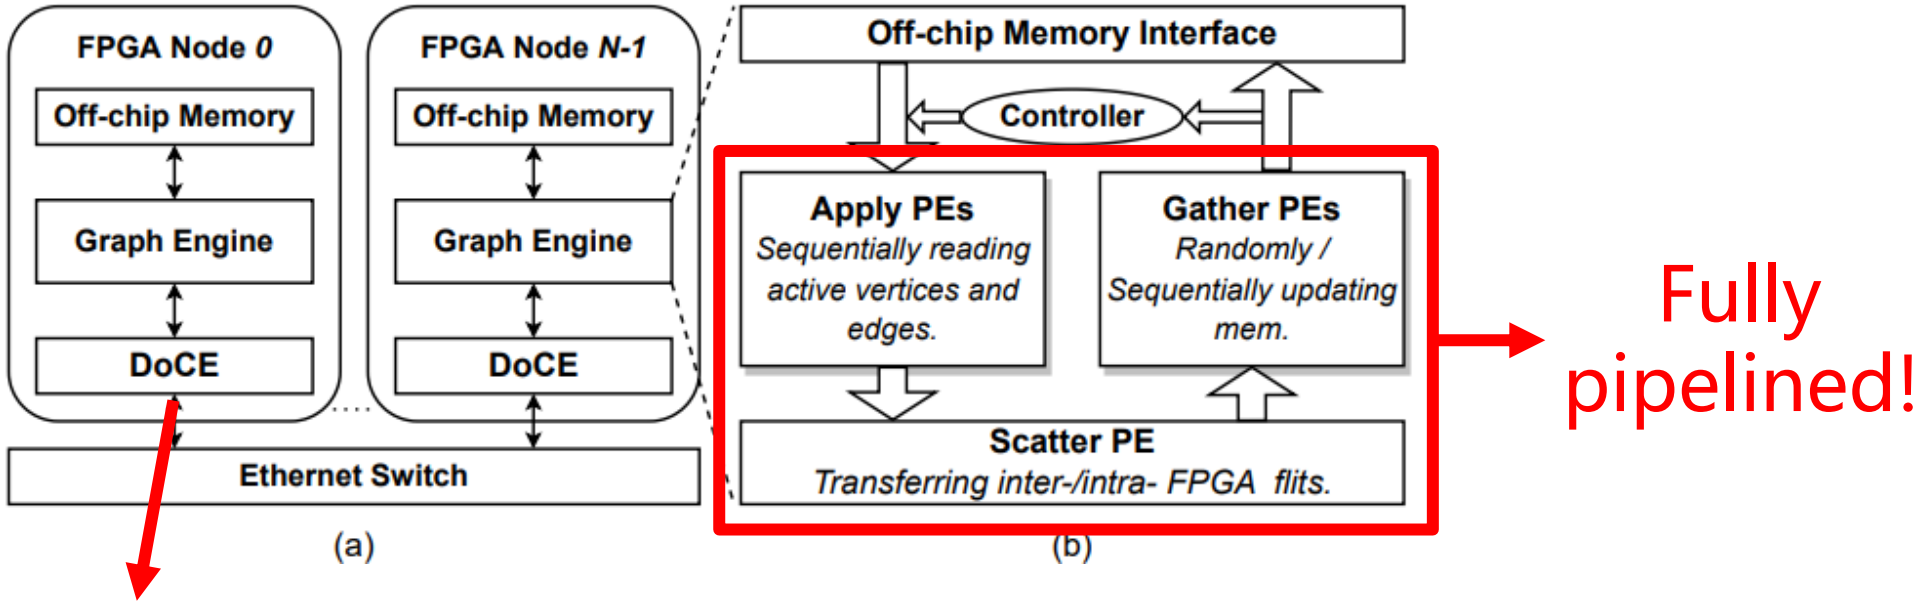

Fully pipelined!

DoCE converts
ARM AMBA AXI protocol to
memory semantic protocol

Institute of Computing Technology, Chinese Academy of Sciences

Preliminary Results

Custom FPGA Board with a Xilinx Zynq UltraScale+ MPSoC (ZynqMP) Chip

- **GraFF runs in 200MHz**

System	FPGA chip	# Nodes	Throughput (GTEPS)	
			BFS	PageRank
GraFF	XCZU19EG	4	6.23	8.40
GraVF-M *	XCKU060	4	5.49	4.62

Workload	BFS GTEPS (speedup)		
	1x FPGA	2x FPGA	3x FPGA
Soc-LiveJournal	0.65 (1.00)	1.30 (2.00)	2.65 (4.00)
RMAT-24	1.60 (1.00)	3.11 (1.94)	6.23 (3.89)

- **GraFF indicates 1.13x-4.52x performance improvement**

*Nina Engelhardt and Hayden K.-H. So. 2019. GraVF-M: Graph Processing System Generation for Multi-FPGA Platforms. ACM Trans. Reconfigurable Technol. Syst. 12, 4, Article 21 (December 2019), 28 pages. https://doi.org/10.1145/3357596

中科院计算所
Institute of Computing Technology, Chinese Academy of Sciences

GraFF: A Multi-FPGA System with Memory Semantic Fabric for Scalable Graph Processing

Thanks for listening!

Xu Zhang, **Yisong Chang**, **Tianyue Lu**,
Ke Liu, **Ke Zhang**, **Mingyu Chen**

Institute of Computing Technology, Chinese Academy of Science

中国科学院计算技术研究所
Institute of Computing Technology,Chinese Academy of Sciences

中科院计算所
Institute of Computing Technology, Chinese Academy of Sciences

System	FPGA chip	# Nodes	Throughput (GTEPS)	
			BFS	PageRank
GraFF	XCZU19EG	4	6.23	8.40
GraVF-M	XCKU060	4	5.49	4.62
ForeGraph	XCVU190	4	1.46	1.86
FDGLib	XCU250	32	2.50	2.36