MaSSIF: Massively Scalable Secure Computation Infrastructure Using FPGAs

Northeastern University

Miriam Leeser, Mehmet Gungor, Kai Huang, Stratis Ioannidis

MOTIVATION

- Most computation is done on cloud with user private data
- Secure Function Evaluation (SFE) is needed to protect privacy of user data
- Accelerating compute times helps to make SFE practical
- Cloud services provide FPGA infrastructure so that SFE can be done on the cloud while protecting user's privacy
- We accelerate garbled circuits in the cloud

CHALLENGE

- Garbled circuits are complex especially for large problems
- Configuring an FPGA for each circuit or problem is impractical and time consuming

CONTRIBUTIONS

- We design a general hardware implementation for FPGAs that uses an "overlay" instead of designing specific hardware for each garbled circuit
- We propose a method to adopt and configure the mapping of a new user problem to FPGA hardware by remapping address registers; there is no need to generate a new hardware design.
- Our approach is flexible and can be applied to large problems while achieving hardware speed up compared to a software implementation
- Our hybrid design makes efficient use of both on-chip and off chip memory supporting large examples

EXPERIMENT RESULTS

- Our design gain speed up against software implementation is up to **28x** for million gate examples

GARBLER TIMING DDR VS HYBRID MEMORY DESIGN. ALL UNITS ARE MS.

	total gates	4 AND 4 XOR DDR	4 AND 4 XOR hybrid	speedup	8 AND 8 XOR	speedup (over 4 and 4)
4bit 5×5 MM	15500	45.48	29.47	1.54	26.42	1.12
8bit 5×5 MM	63000	184.23	111.74	1.65	96.61	1.16
4bit 10×10 MM	126000	368.22	283.86	1.30	242.55	1.17
8bit 10×10 MM	508000	1487.21	1180.49	1.26	1067.35	1.11
12bit 10×10 MM	1146000	3234.93	2570.84	1.26	2356.41	1.09
16bit 10×10 MM	2040000	5636.27	4606.83	1.22	4185.36	1.10
4bit 20×20 MM	1016000	3153.26	2571.50	1.23	2346.86	1.10
8bit 20×20 MM	4080000	12638.08	10226.60	1.24	9378.26	1.09

SOFTWARE VS BEST FPGA IMPLEMENTATION. ALL UNITS ARE MS

	software	8 AND 8 XOR	speedup
4bit 5×5 MM	659.08	26.42	24.95
8bit 5×5 MM	2684.03	96.61	27.78
4bit 10×10 MM	5391.43	242.55	22.23
8bit 10×10 MM	22031.15	1067.35	20.7
12bit 10×10 MM	49906.86	2356.41	21.18
16bit 10×10 MM	89392.44	4185.36	21.35
4bit 20×20 MM	44466.74	2346.86	18.95

Timing for total system with python (unit:ms)

applications	Total	garbler	evaluator	gt transfer
16Bit Adder	4.933	3.14	1.793	4.8×10^{-4}
30Bit Ham	18.032	11.686	6.341	4.8×10^{-3}
50Bit Ham	27.811	19.370	8.433	8×10^{-3}
8Bit a*b	30.361	20.473	9.792	9.6×10^{-3}
16Bit a*b	126.366	85.389	40.937	0.0397
32Bit a*b	515.867	349.713	165.993	0.161
64Bit a*b	2066.394	1393.873	671.871	0.650
4Bit Sort10 Number	287.957	182.825	105.065	0.067
4Bit 5x5 Mat Mult	978.663	659.079	319.272	0.312
8Bit 5x5 Mat Mult	4003.290	2684.033	1317.993	1.264
4Bit 10x10 Mat Mult	7984.151	5391.426	2592.123	0.602
8Bit 10x10 Mat Mult	32587.864	22031.146	10546.542	10.176
4Bit 20x20 Mat Mult	65173.249	43868.833	21284.064	20.352

Timing for total system with software garbler and FPGA garbler in ms

applications	Total(garbler sw)	Total(garbler FPGA)	Speed Up
16Bit Adder	4.933	2.406	2.41
30Bit Ham	18.032	7.290	2.47
50Bit Ham	27.811	9.991	2.78
8Bit a*b	30.361	11.33	2.68
16Bit a*b	126.366	46.817	2.70
32Bit a*b	515.867	189.910	2.72
64Bit a*b	2066.394	768.183	2.69
4Bit Sort10 Number	287.957	120.599	2.39
4Bit 5x5 Mat Mult	978.663	365.063	2.68
8Bit 5x5 Mat Mult	4003.290	1503.485	2.66
4Bit 10x10 Mat Mult	7984.151	2960.941	2.70
8Bit 10x10 Mat Mult	32587.864	12043.928	2.71
4Bit 20x20 Mat Mult	65173.249	24271.066	2.69

MM: Matrix multiplication
total system: time : Time to garble, transfer and evaluate

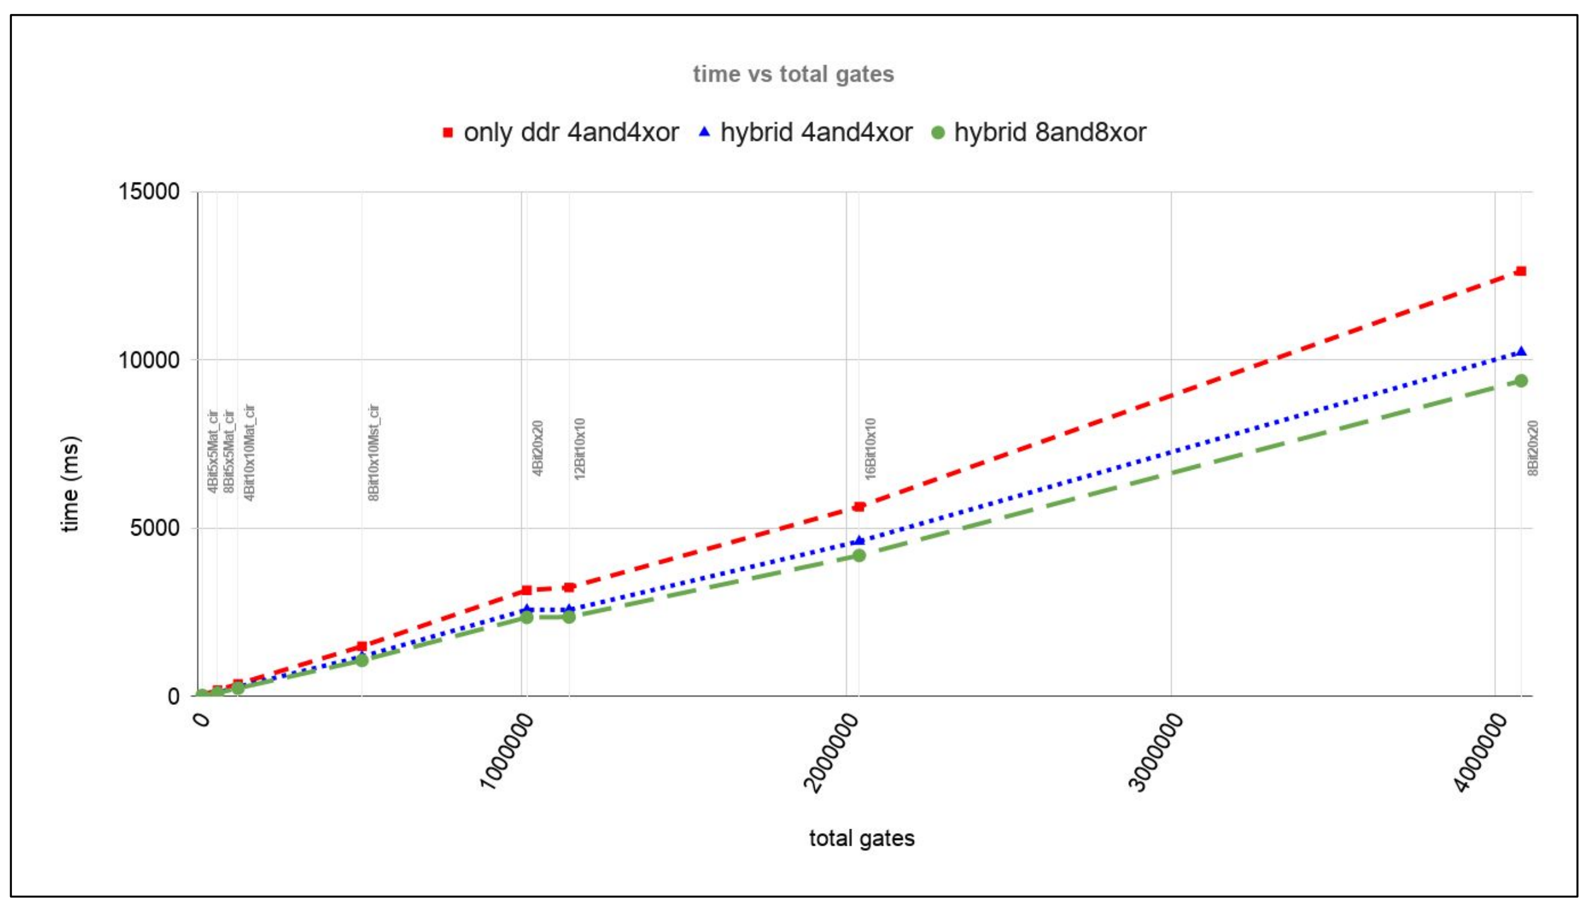

MaSSIF repo link : https://github.com/RCL-lab/NU_MASSIF

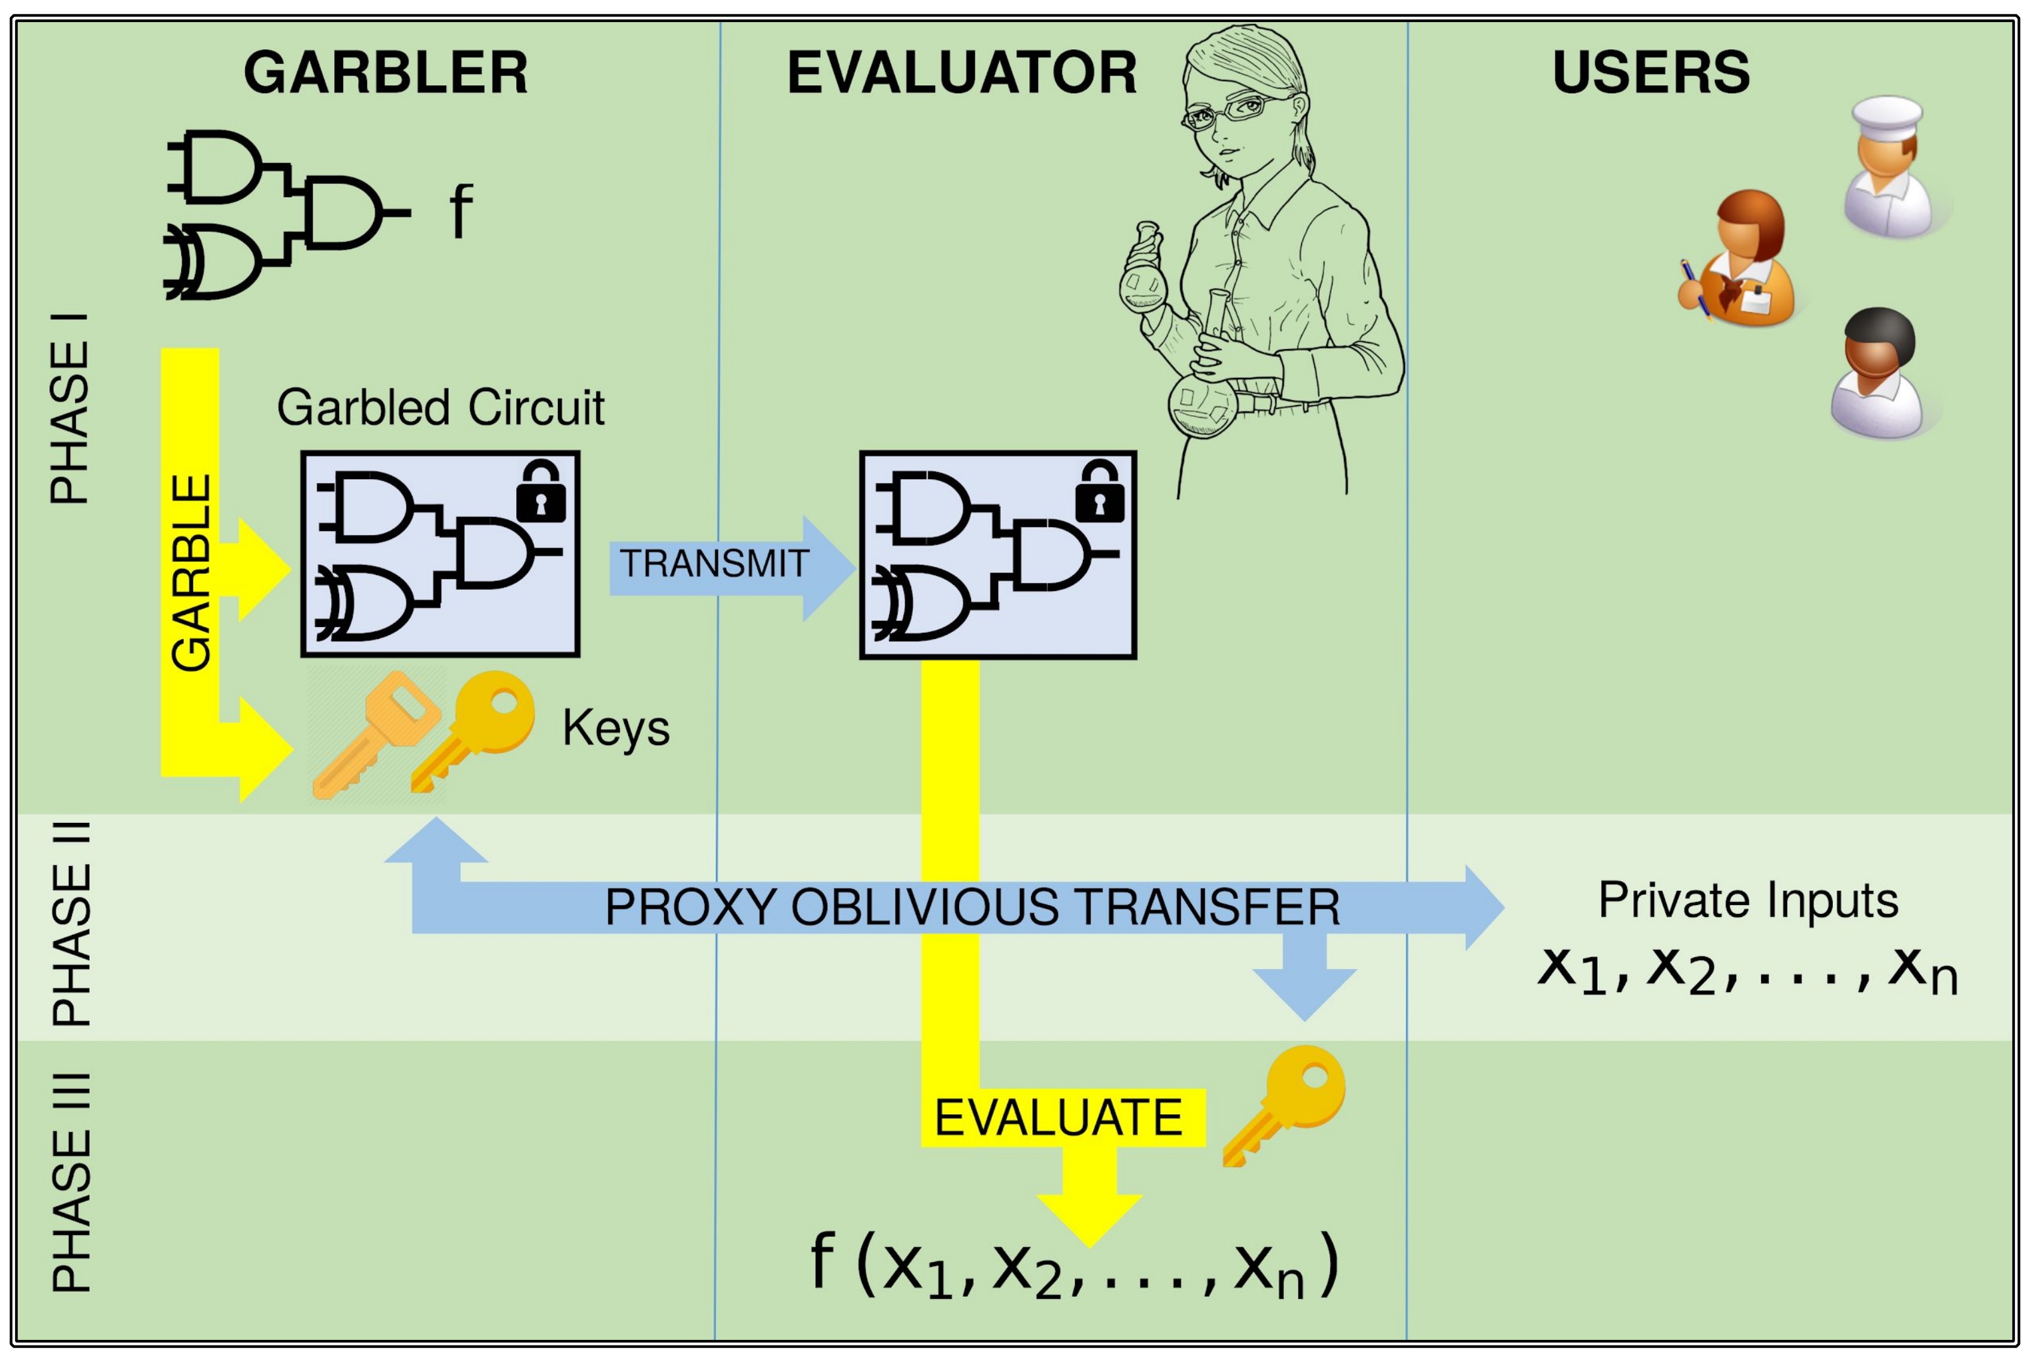

WORKFLOW AND SYSTEM ARCHITECTURE

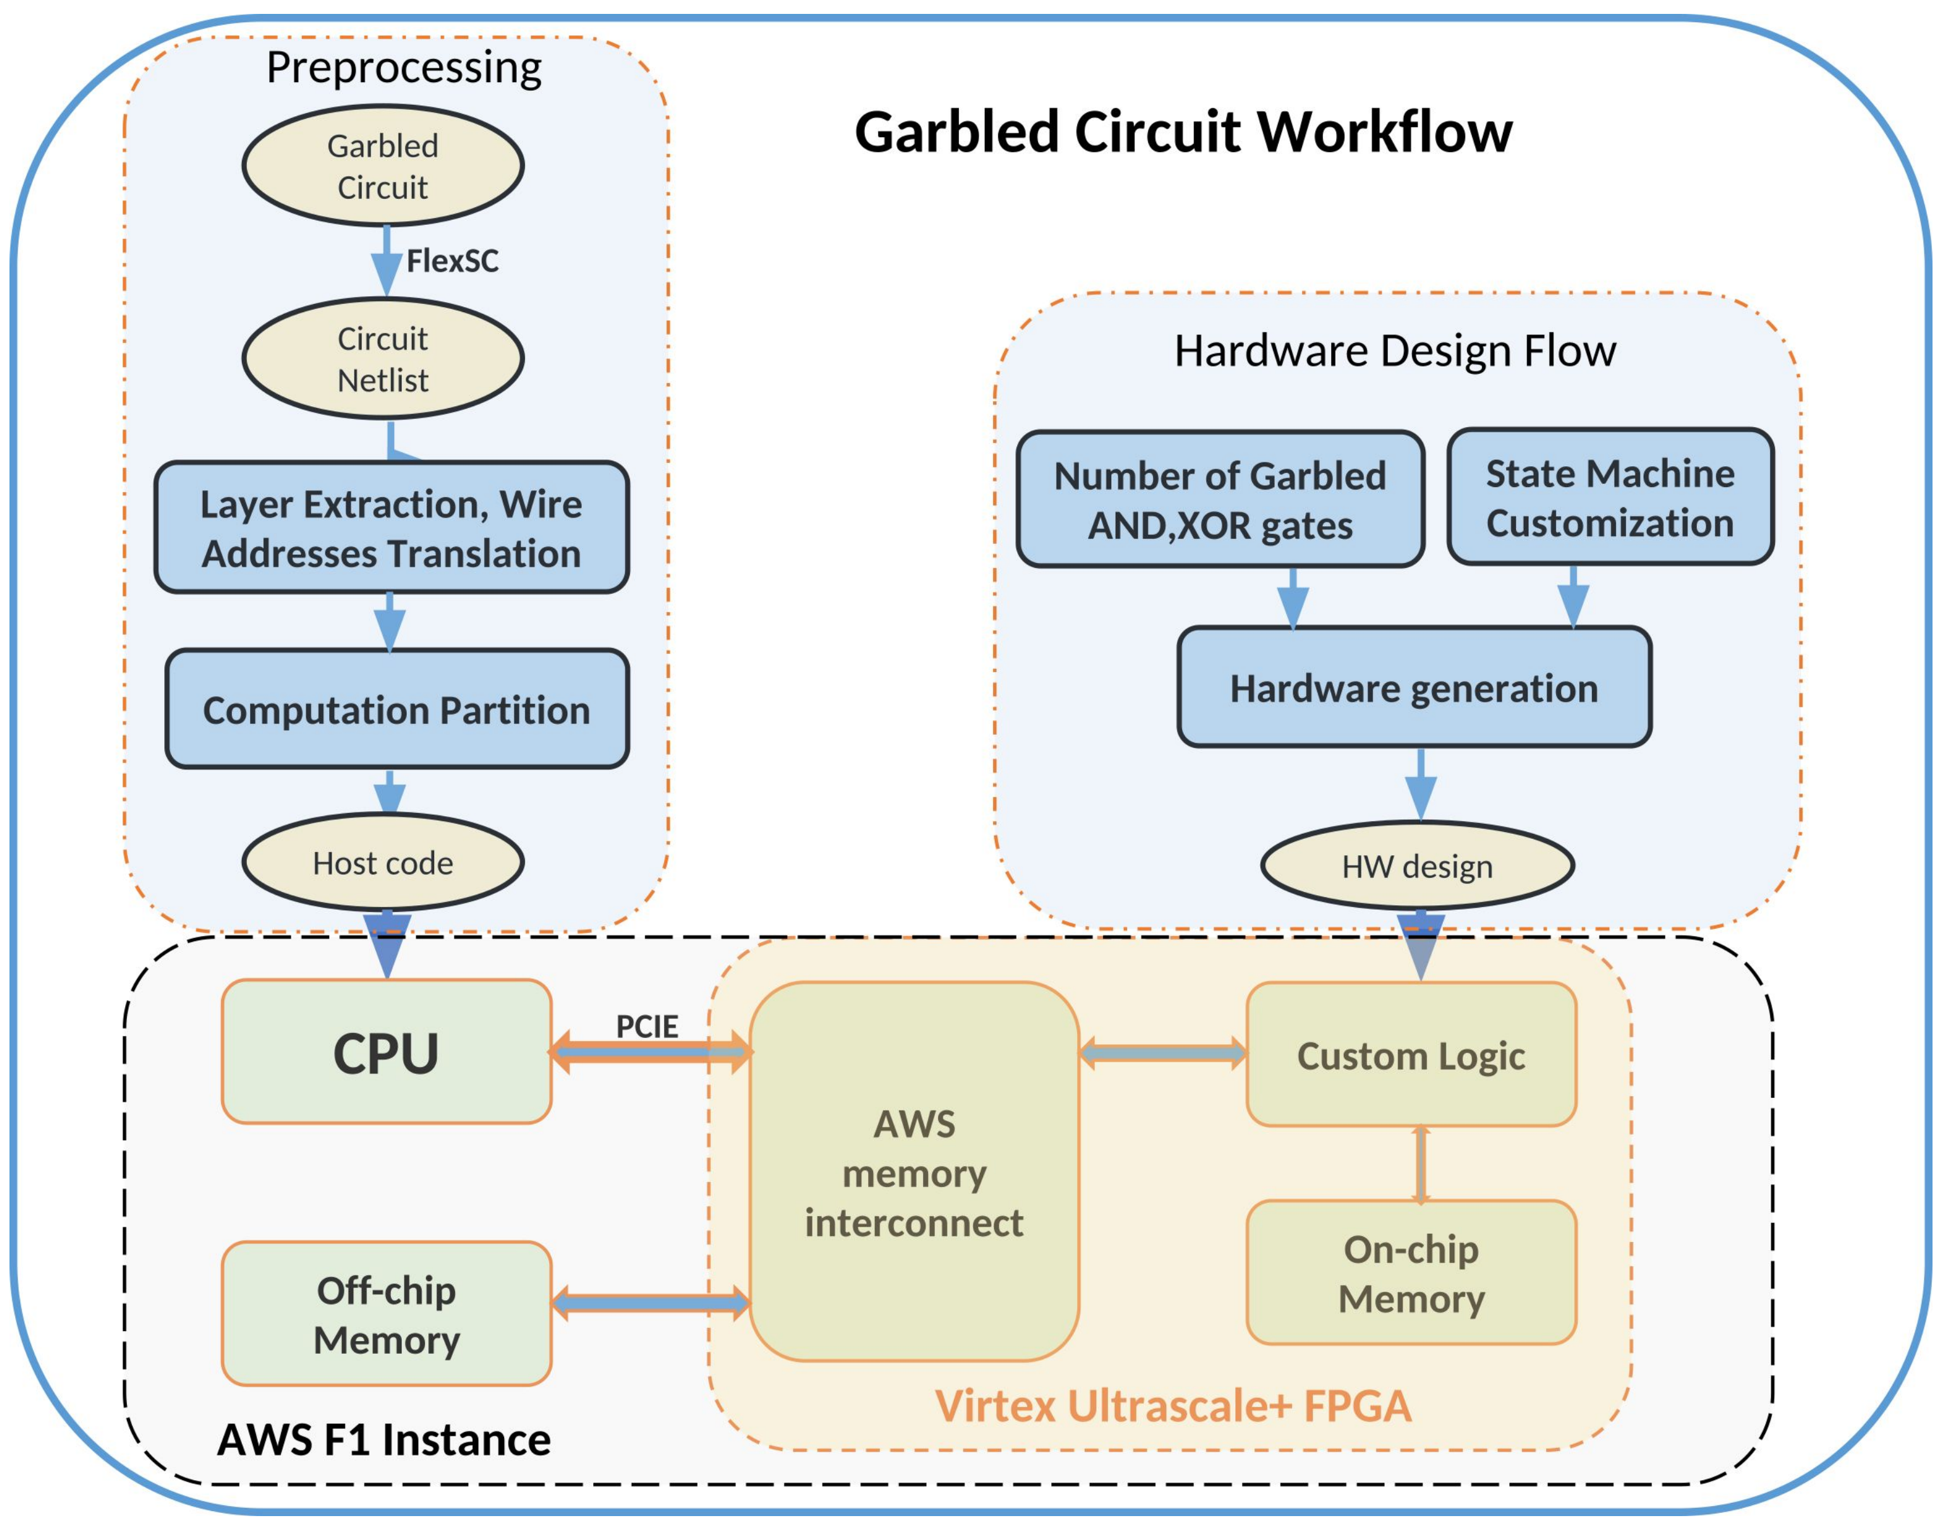

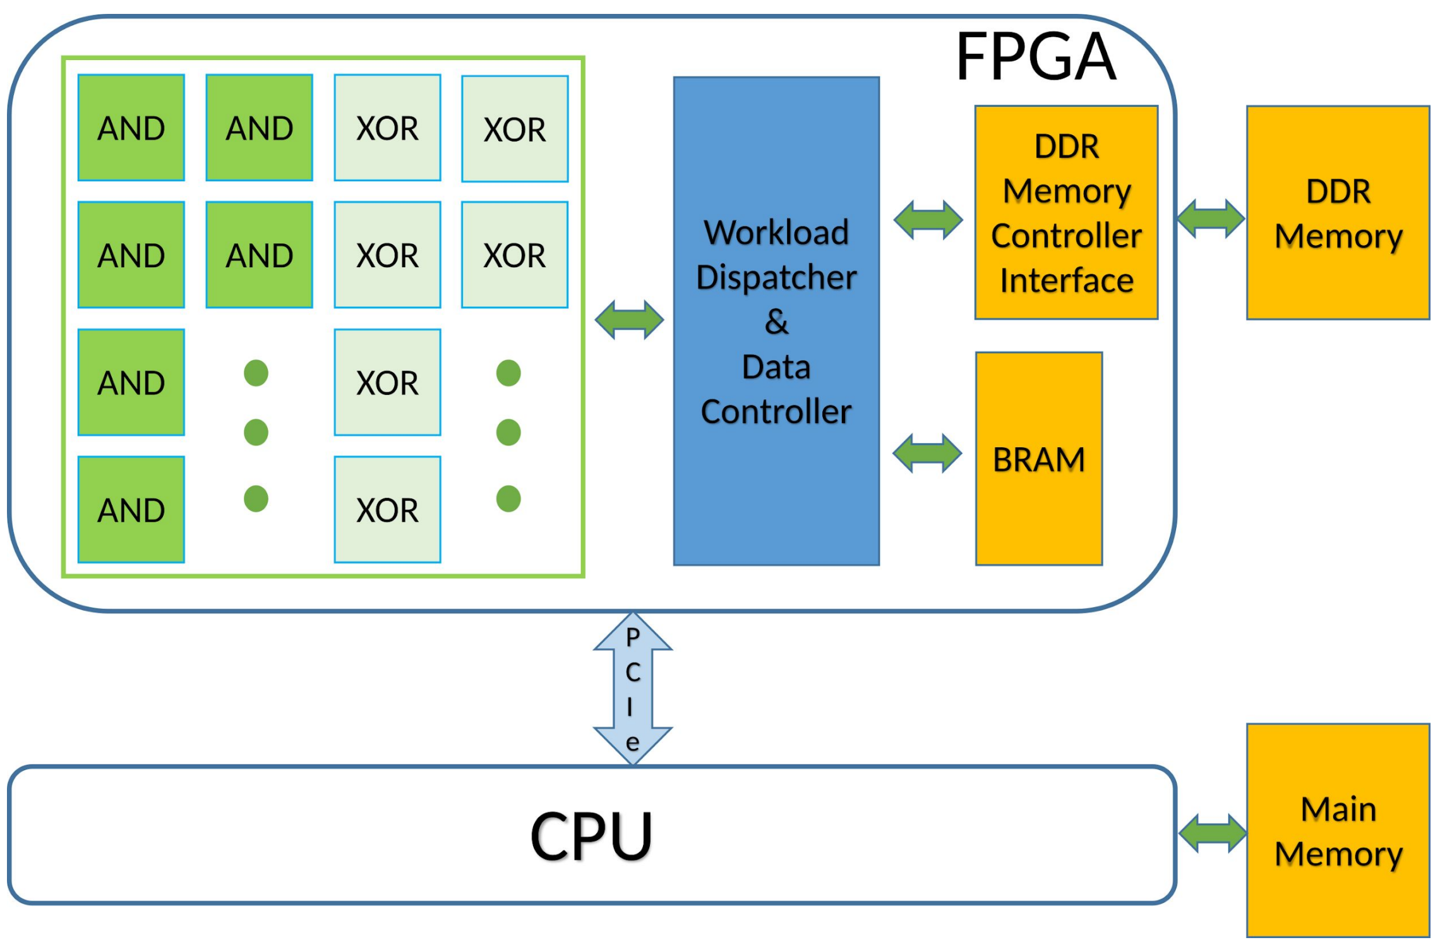

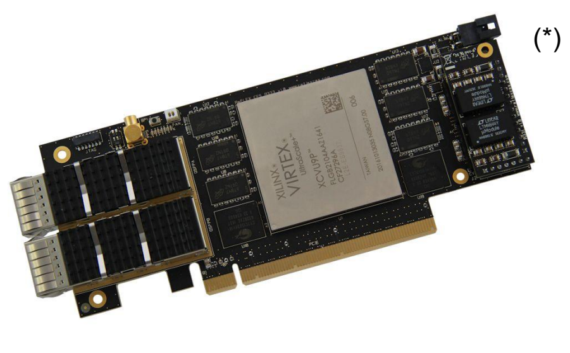

(**) (***)

PUBLICATIONS

- M. Gungor, K. Huang, X. Fang, S. Ioannidis, and M. Leeser, "Garbled circuits in the cloud using FPGA enabled nodes," in 2019 IEEE High Performance Extreme Computing Conference (HPEC) .IEEE, 2019,pp. 1–6.
- M. Leeser, M. Gungor, K. Huang, X. Fang and S. Ioannidis, "Accelerating Large Garbled Circuits on an FPGA-enabled Cloud", under submission.

Acknowledgment:
This research was supported by NSF grant CNS-1717213, and in part by a grant from Amazon Web Services, which supported our experimentation

(*) https://twinind.com/files/8315/0647/0723/XpressVUP_image1.jpg
(**) https://upload.wikimedia.org/wikipedia/commons/9/93/Amazon_Web_Services_Logo.svg
(***)https://d1.awsstatic.com/f1.ce891e39434406bc54217c6cd1385201d915b40a.JPG