

Nikolaos Alachiotis, Dimitris Theodoropoulos, and Dionisios Pnevmatikatos

Computer Architecture and VLSI Systems Laboratory

Institute of Computer Science

Foundation for Research and Technology - Hellas (FORTH)

{nalachio, dtheodor, pnevmati}@ics.forth.gr

Introduction

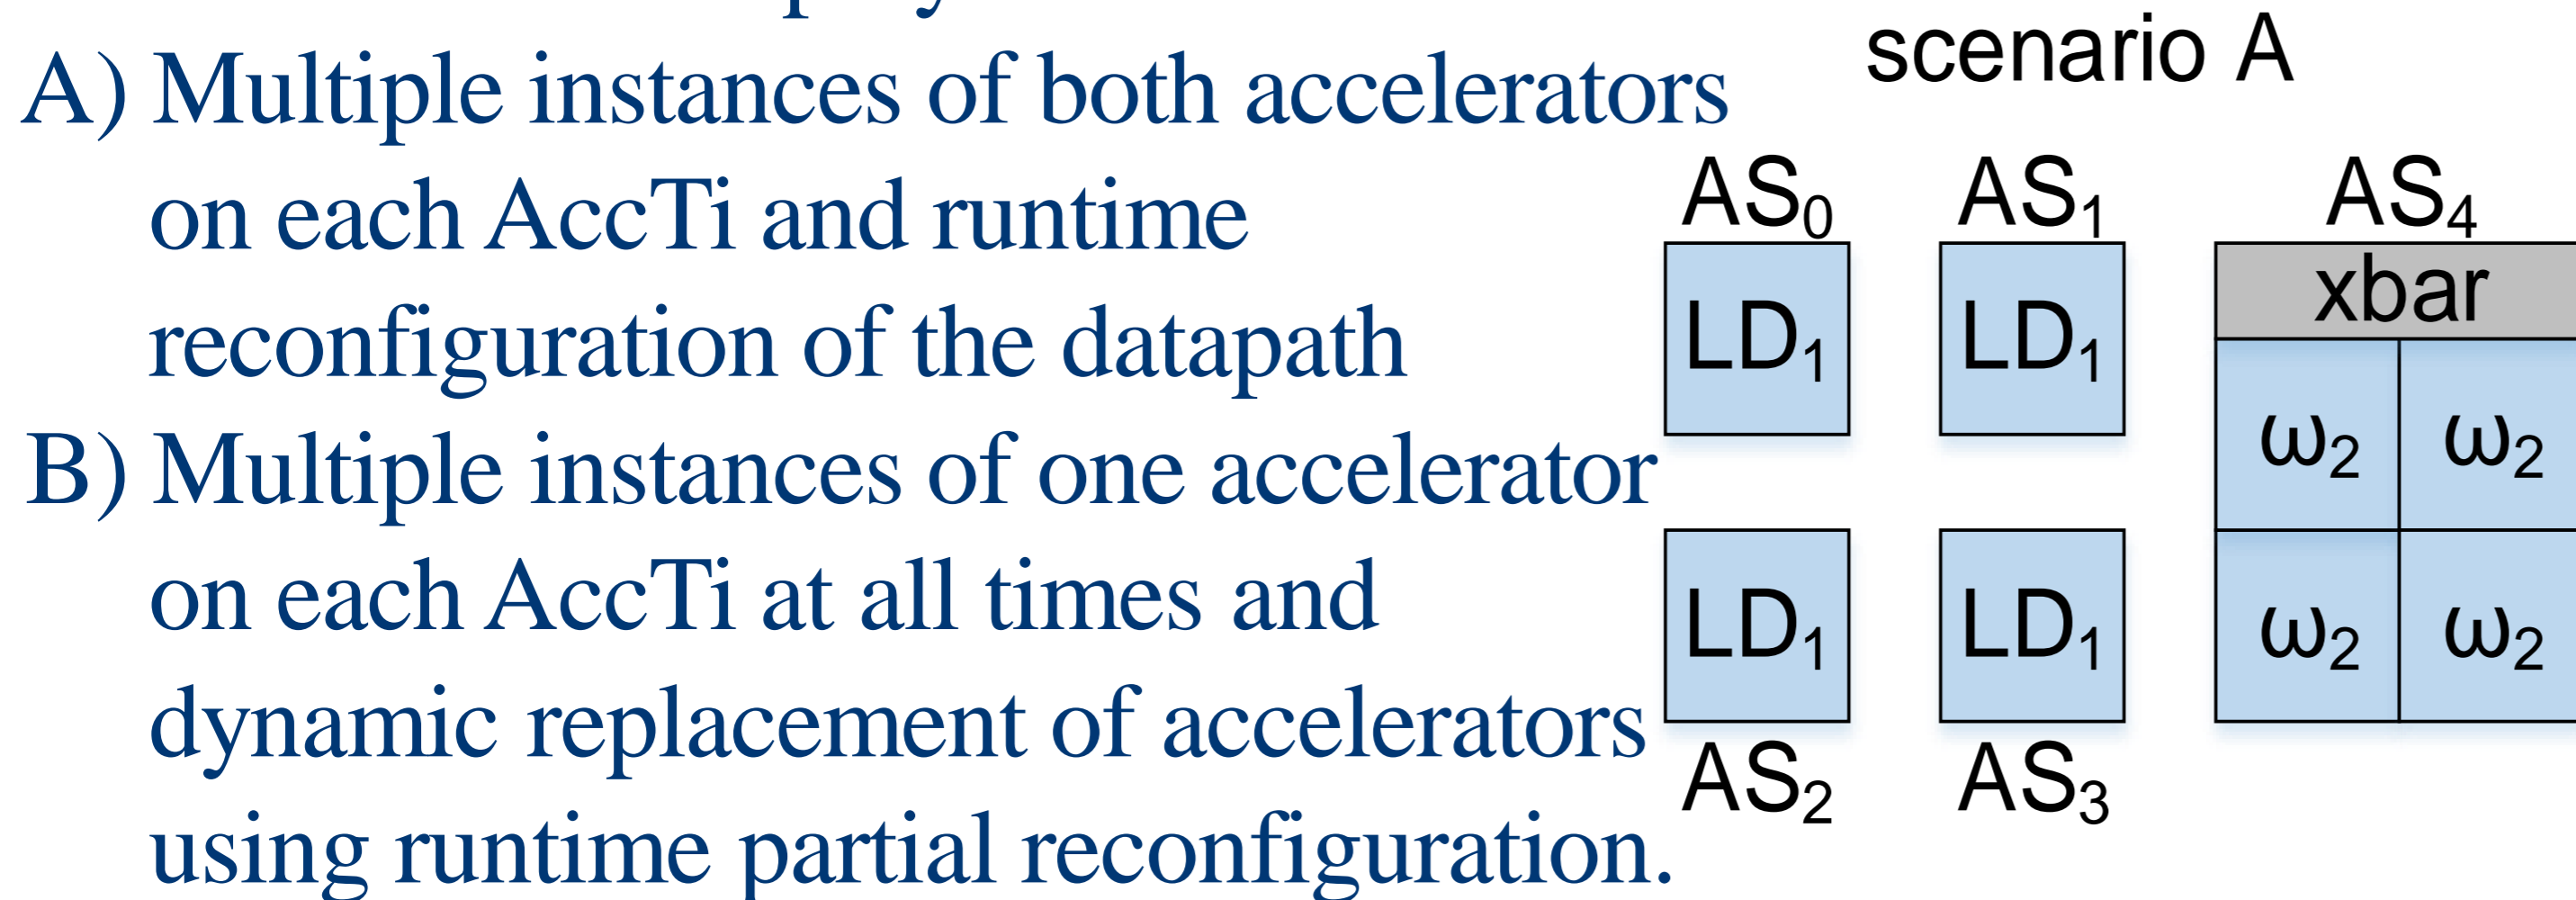
- Datacenters typically exhibit predefined resource proportionality, comprising fixed amounts of compute, memory, and peripherals.
- A broad range of applications, with varying compute and memory requirements, need to be efficiently executed on them.
- An increasingly explored solution relies on **datacenter disaggregation**, with resources organized as large homogeneous pools, allocated on a per-VM basis.

Contribution

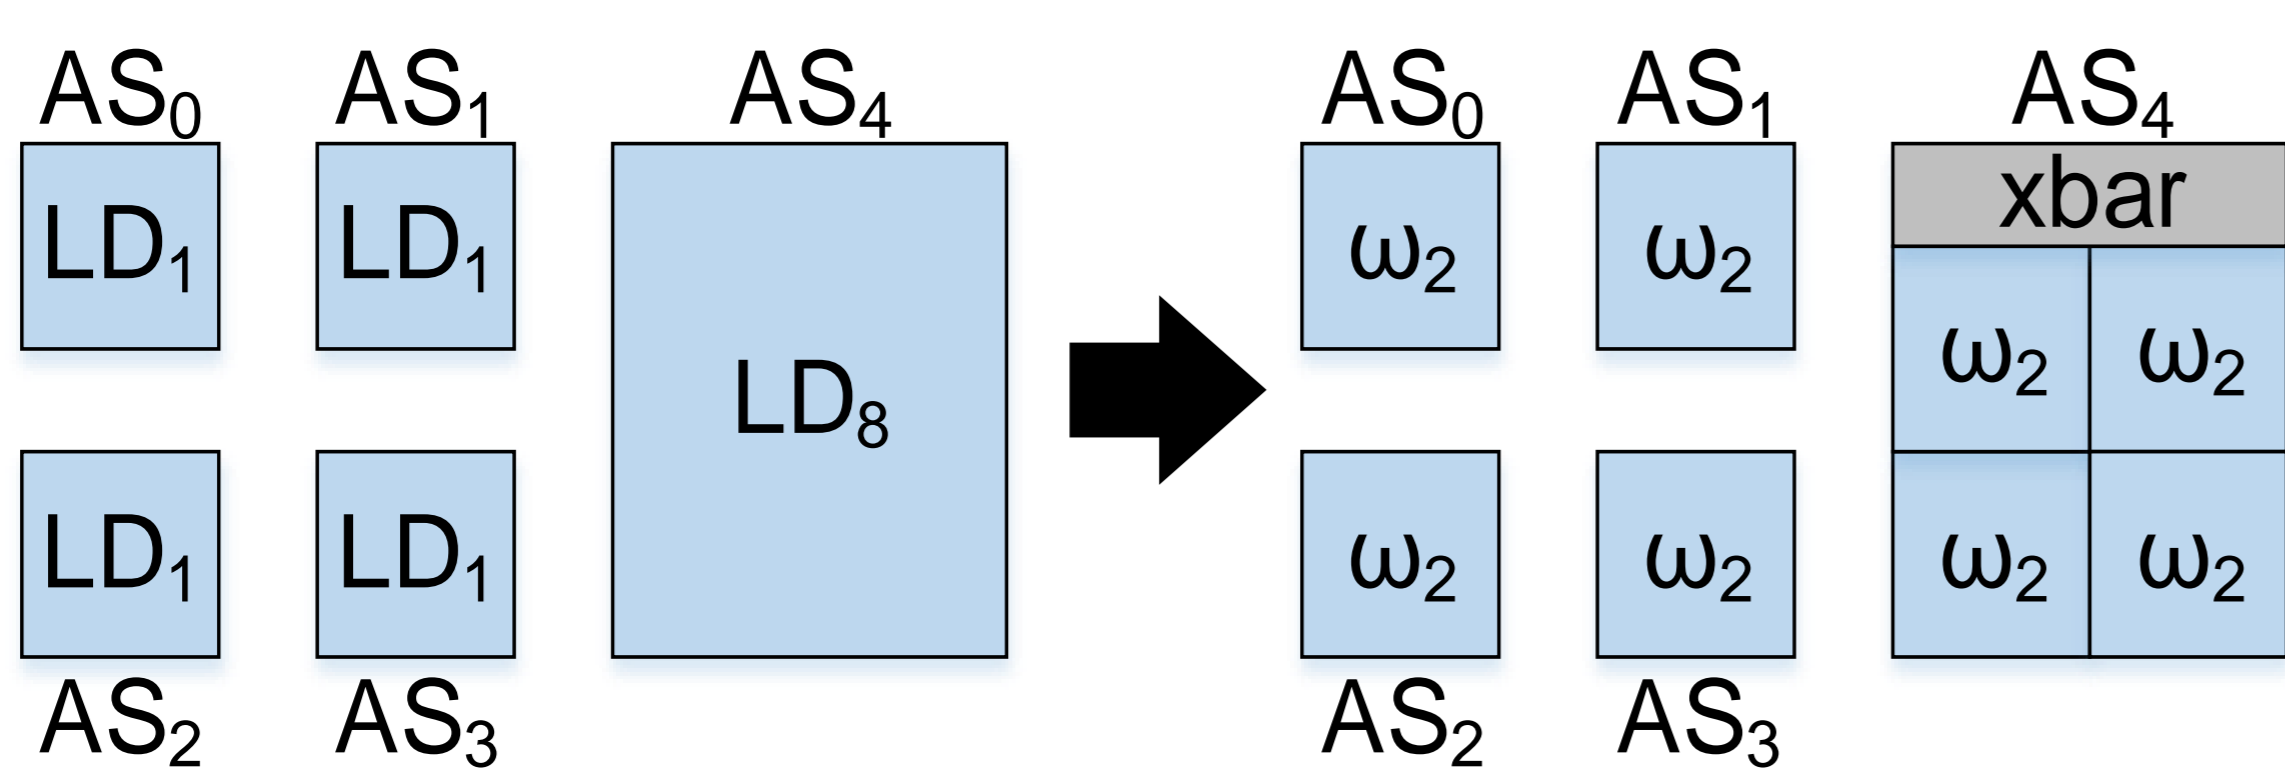
- A reconfigurable acceleration architecture based on accelerator tiles, with two levels of configuration capability: a) partially reconfigurable regions, and b) runtime-configurable intra-datapath accelerator interconnect.
- Generic accelerator deployment solution and application on a large-scale population genomics algorithm for detecting traces of positive selection.

Application Mappings and Evaluation

- Two custom accelerators: LD[1] and HLS-generated ω [2].
- Two accelerator deployment scenarios:



scenario B



2AccTis vs 12 Xeon CPU cores @ 2.6GHz Dataset # = #sequences:#SNPs

	D1=10 ³ :10 ⁴	D2=10 ⁴ :10 ⁴	D3=10 ⁴ :10 ⁵
OmegaPlus: LD	2.6s (14.54mLD/s)	19.1s (2.63MLD/s)	308.2s (2.89mLD/s)
OmegaPlus: ω	32.9s (85.63mLD/s)	21.5s (74.22mLD/s)	7.1s (75.03mLD/s)
OmegaPlus: Total	35.9s	43.7s	358.7s
Accel. Scen. A: LD	0.78s (48.39mLD/s)	9.89s (5.09mLD/s)	175.1s (5.08mLD/s)
Accel. Scen. A: ω	4.41s (639.29mLD/s)	2.66s (600.09mLD/s)	0.81s (660.17mLD/s)
Accel. Scen. A: Total	5.59s / 6.42x	15.65s / 2.79x	218.8s / 1.64x
Accel. Scen. B: LD	0.57s (65.76mLD/s)	6.67s (7.52mLD/s)	126.8s (7.03mLD/s)
Accel. Scen. B: ω	5.84s (481.99mLD/s)	3.89s (409.47mLD/s)	1.37s (389.65mLD/s)
Accel. Scen. B: Total	6.81s / 5.27x	13.66s / 3.19x	171.57s / 2.09x

Bibliography

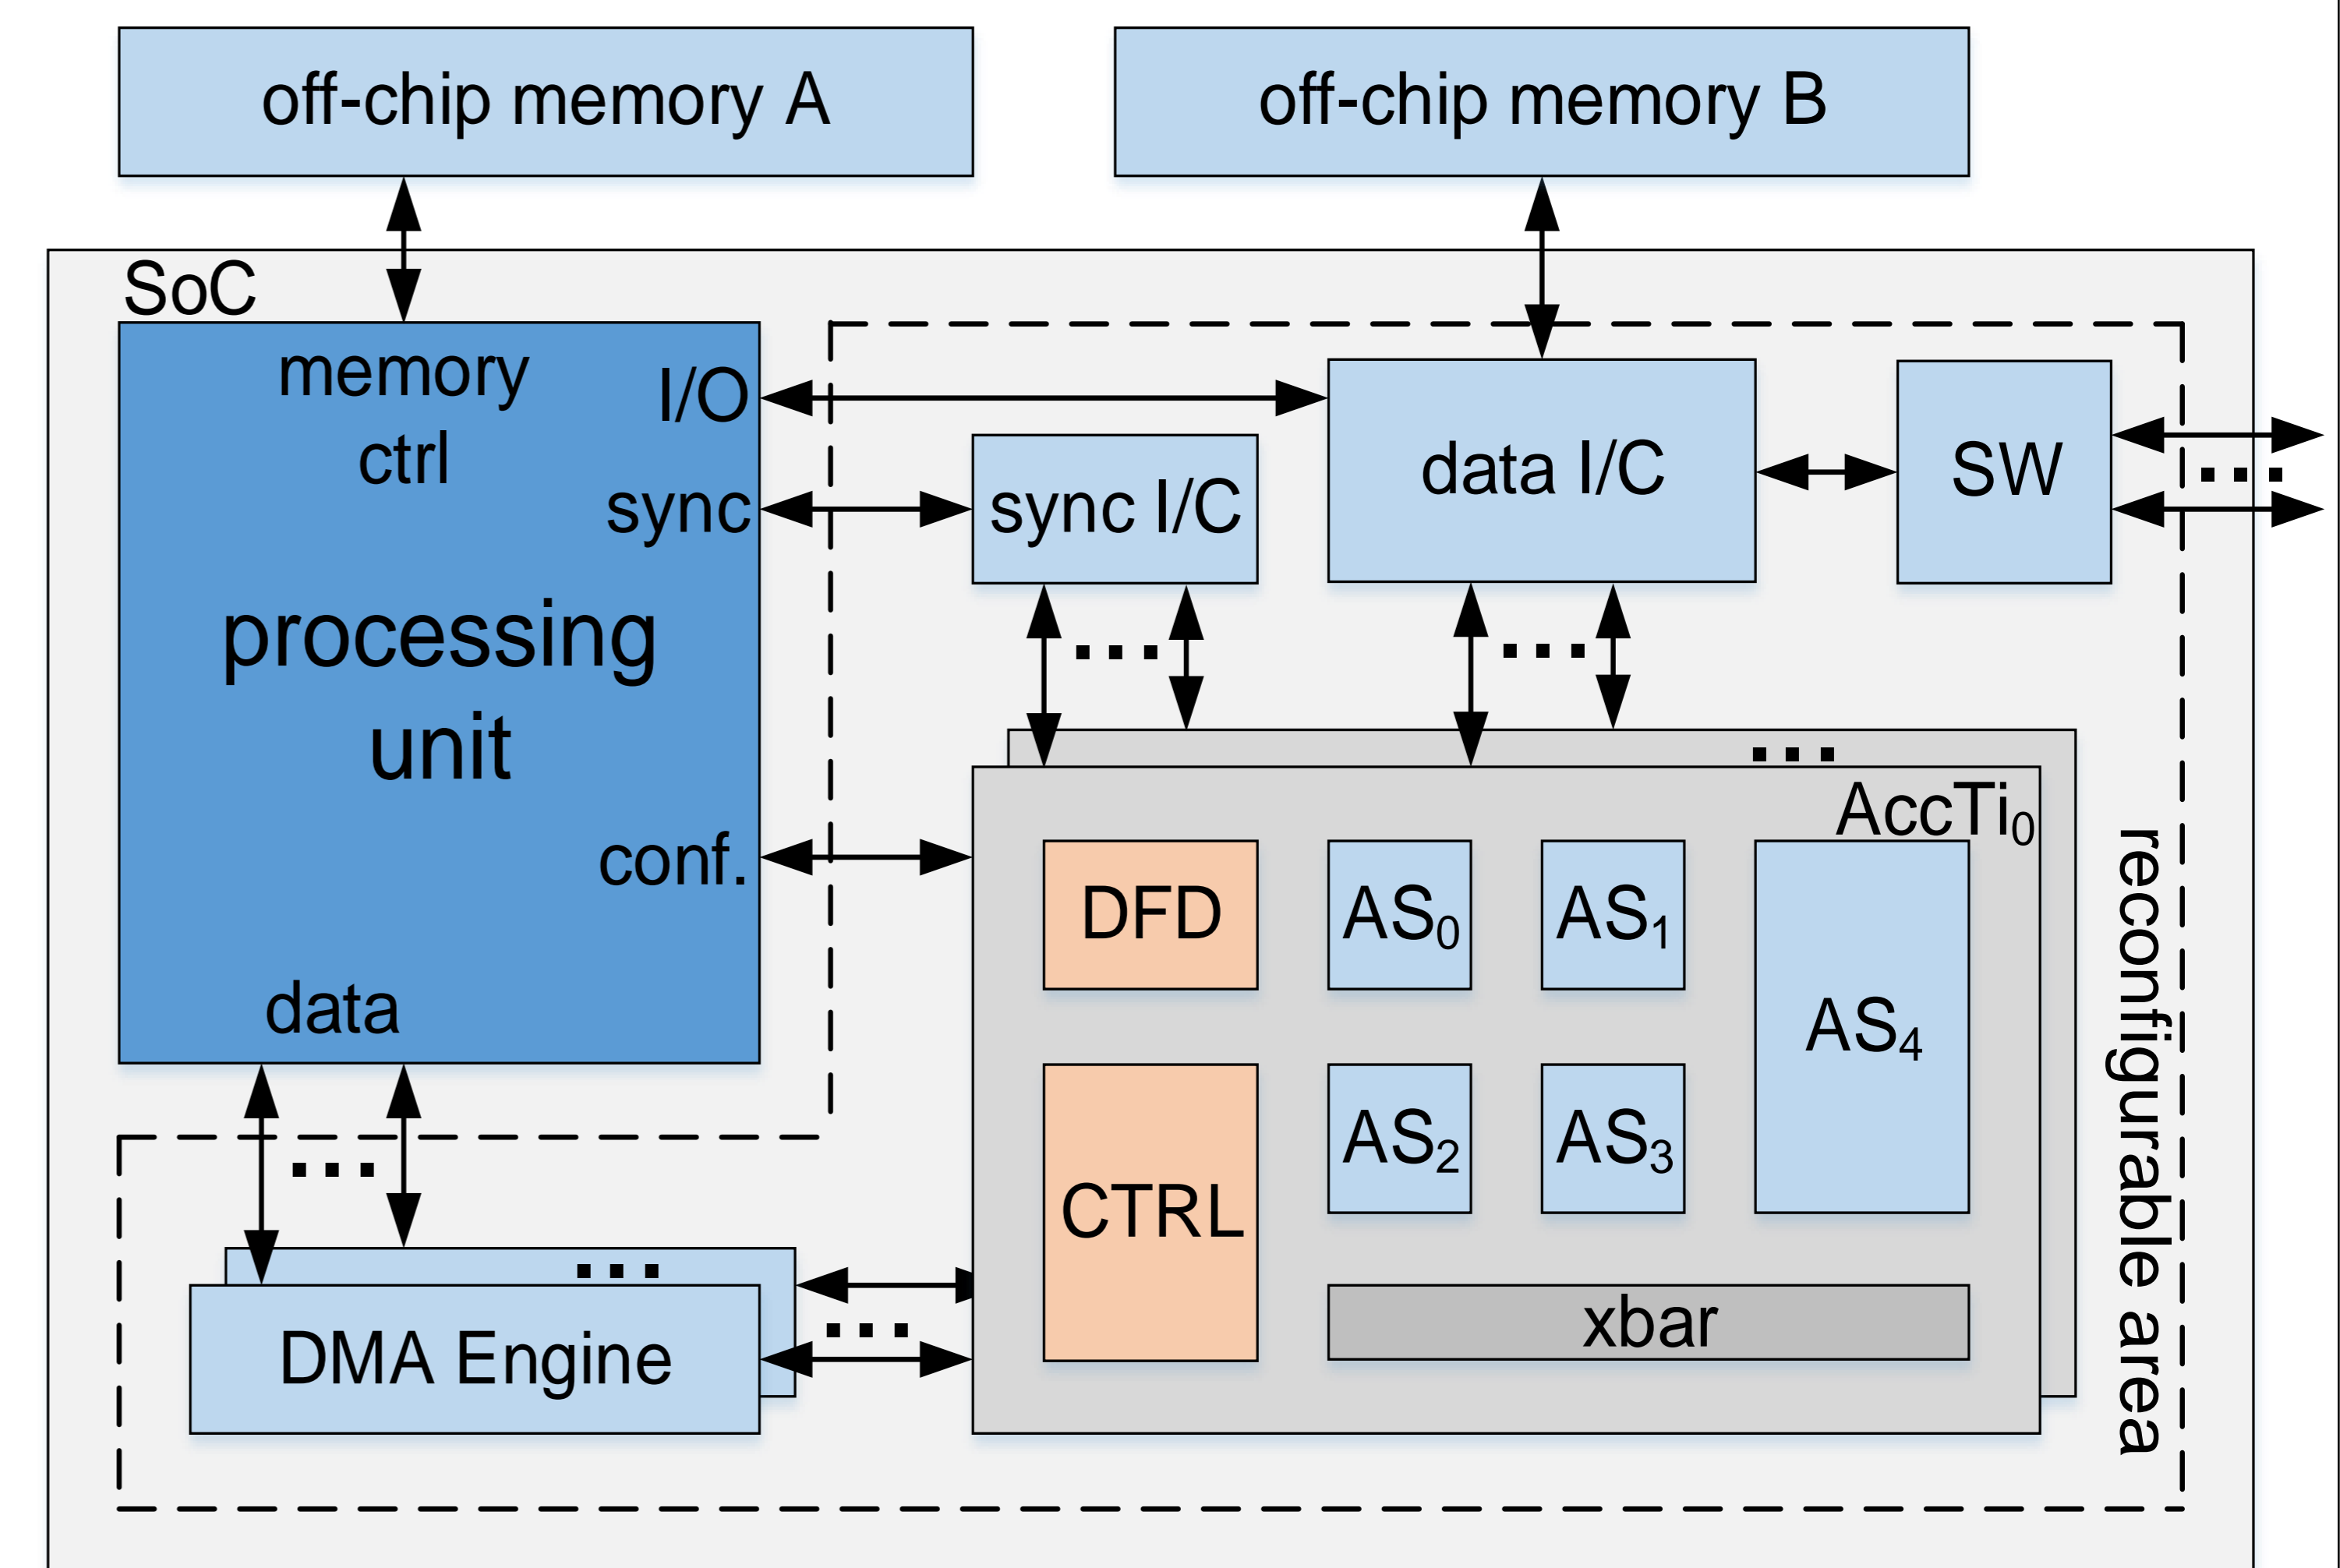
- [1] Alachiotis N. and Weisz G. , "High Performance Linkage Disequilibrium: FPGAs Hold the Key" FPGA 2016.
- [2] Kim Y. and Nielsen R. , "Linkage Disequilibrium as a signatures of selective sweeps", Genetics, 2004.

Bioinformatics Case Study

- A selective sweep is the reduction of genetic variation in a genomic region due to positive selection.
- Next-generation sequencing is generating a plethora of whole genomes, pushing compute and memory demands of **selective sweep detection** to unprecedentedly high levels.
- Deploying custom hardware accelerators is an effective solution that can help bridge the gap between the increasing data accumulation pace and the computational capacity of existing processing platforms.


Accelerator-Tile Architecture

Custom accelerators can be deployed in each Accelerator Slot per Acceleration Tile and form reprogrammable datapaths.



Multi-FPGA evaluation platform

Three ZC706 boards, connected with SMA and SFP cables



Acknowledgments

This work has been supported in part by EU H2020 ICT project dRedBox, contract #687632