

The Case for Custom Parallel Memories: an Application-Centric Analysis

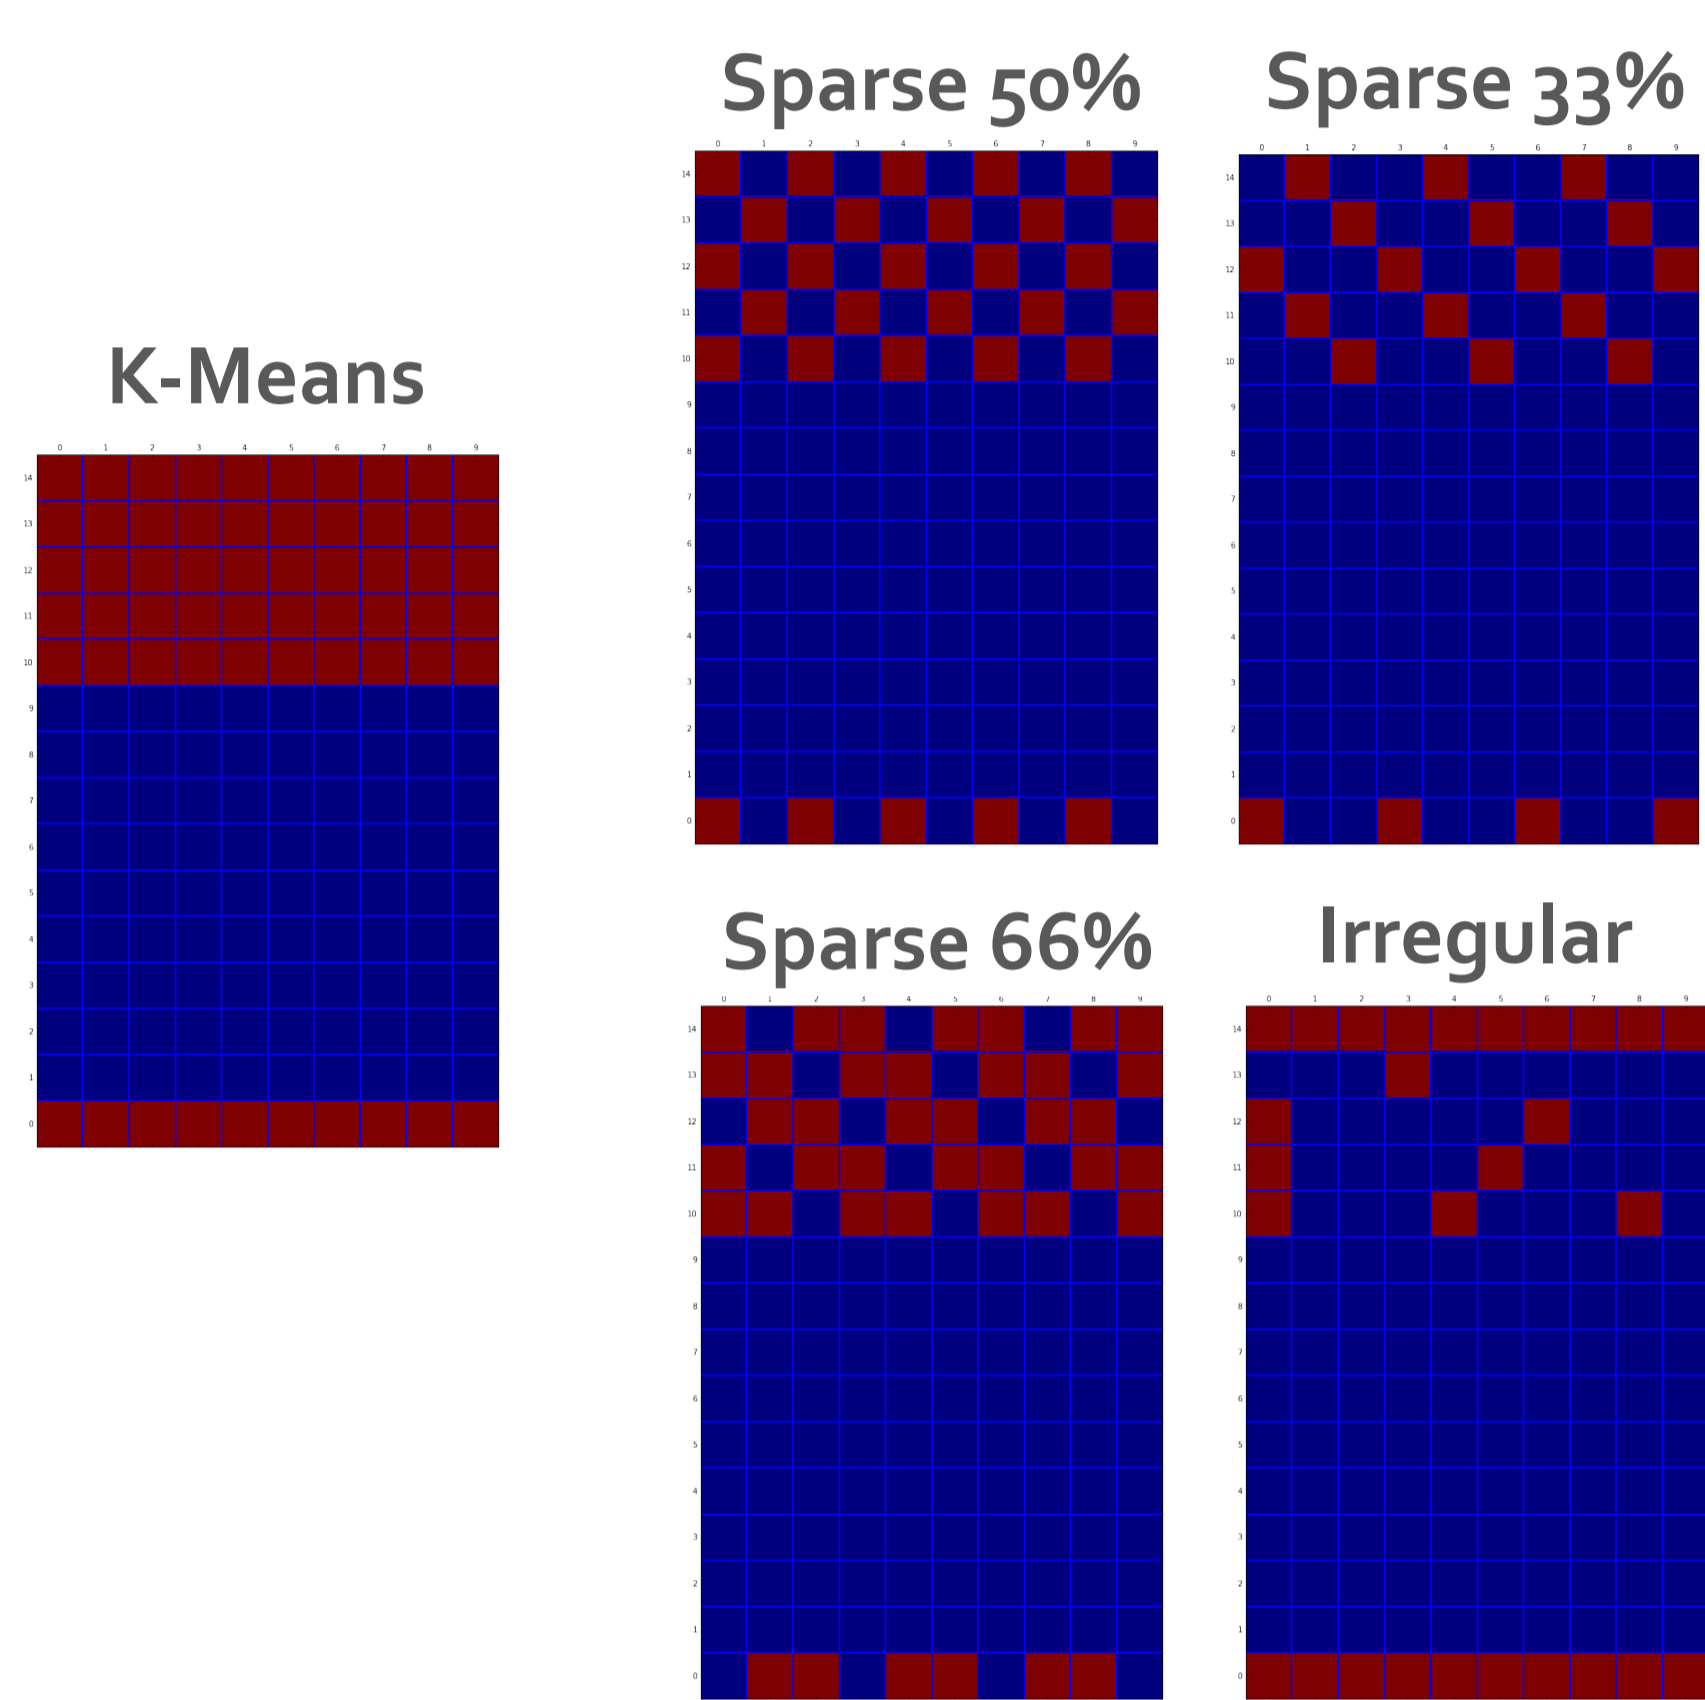
Giulio Stramondo, Cătălin Bogdan Ciobanu, Ana Lucia Varbanescu, Cees de Laat



The main idea

For many applications, the performance bottleneck is memory throughput. In this work, we propose to use FPGAs for building parallel custom memory systems to accelerate such applications.

Application Access Patterns

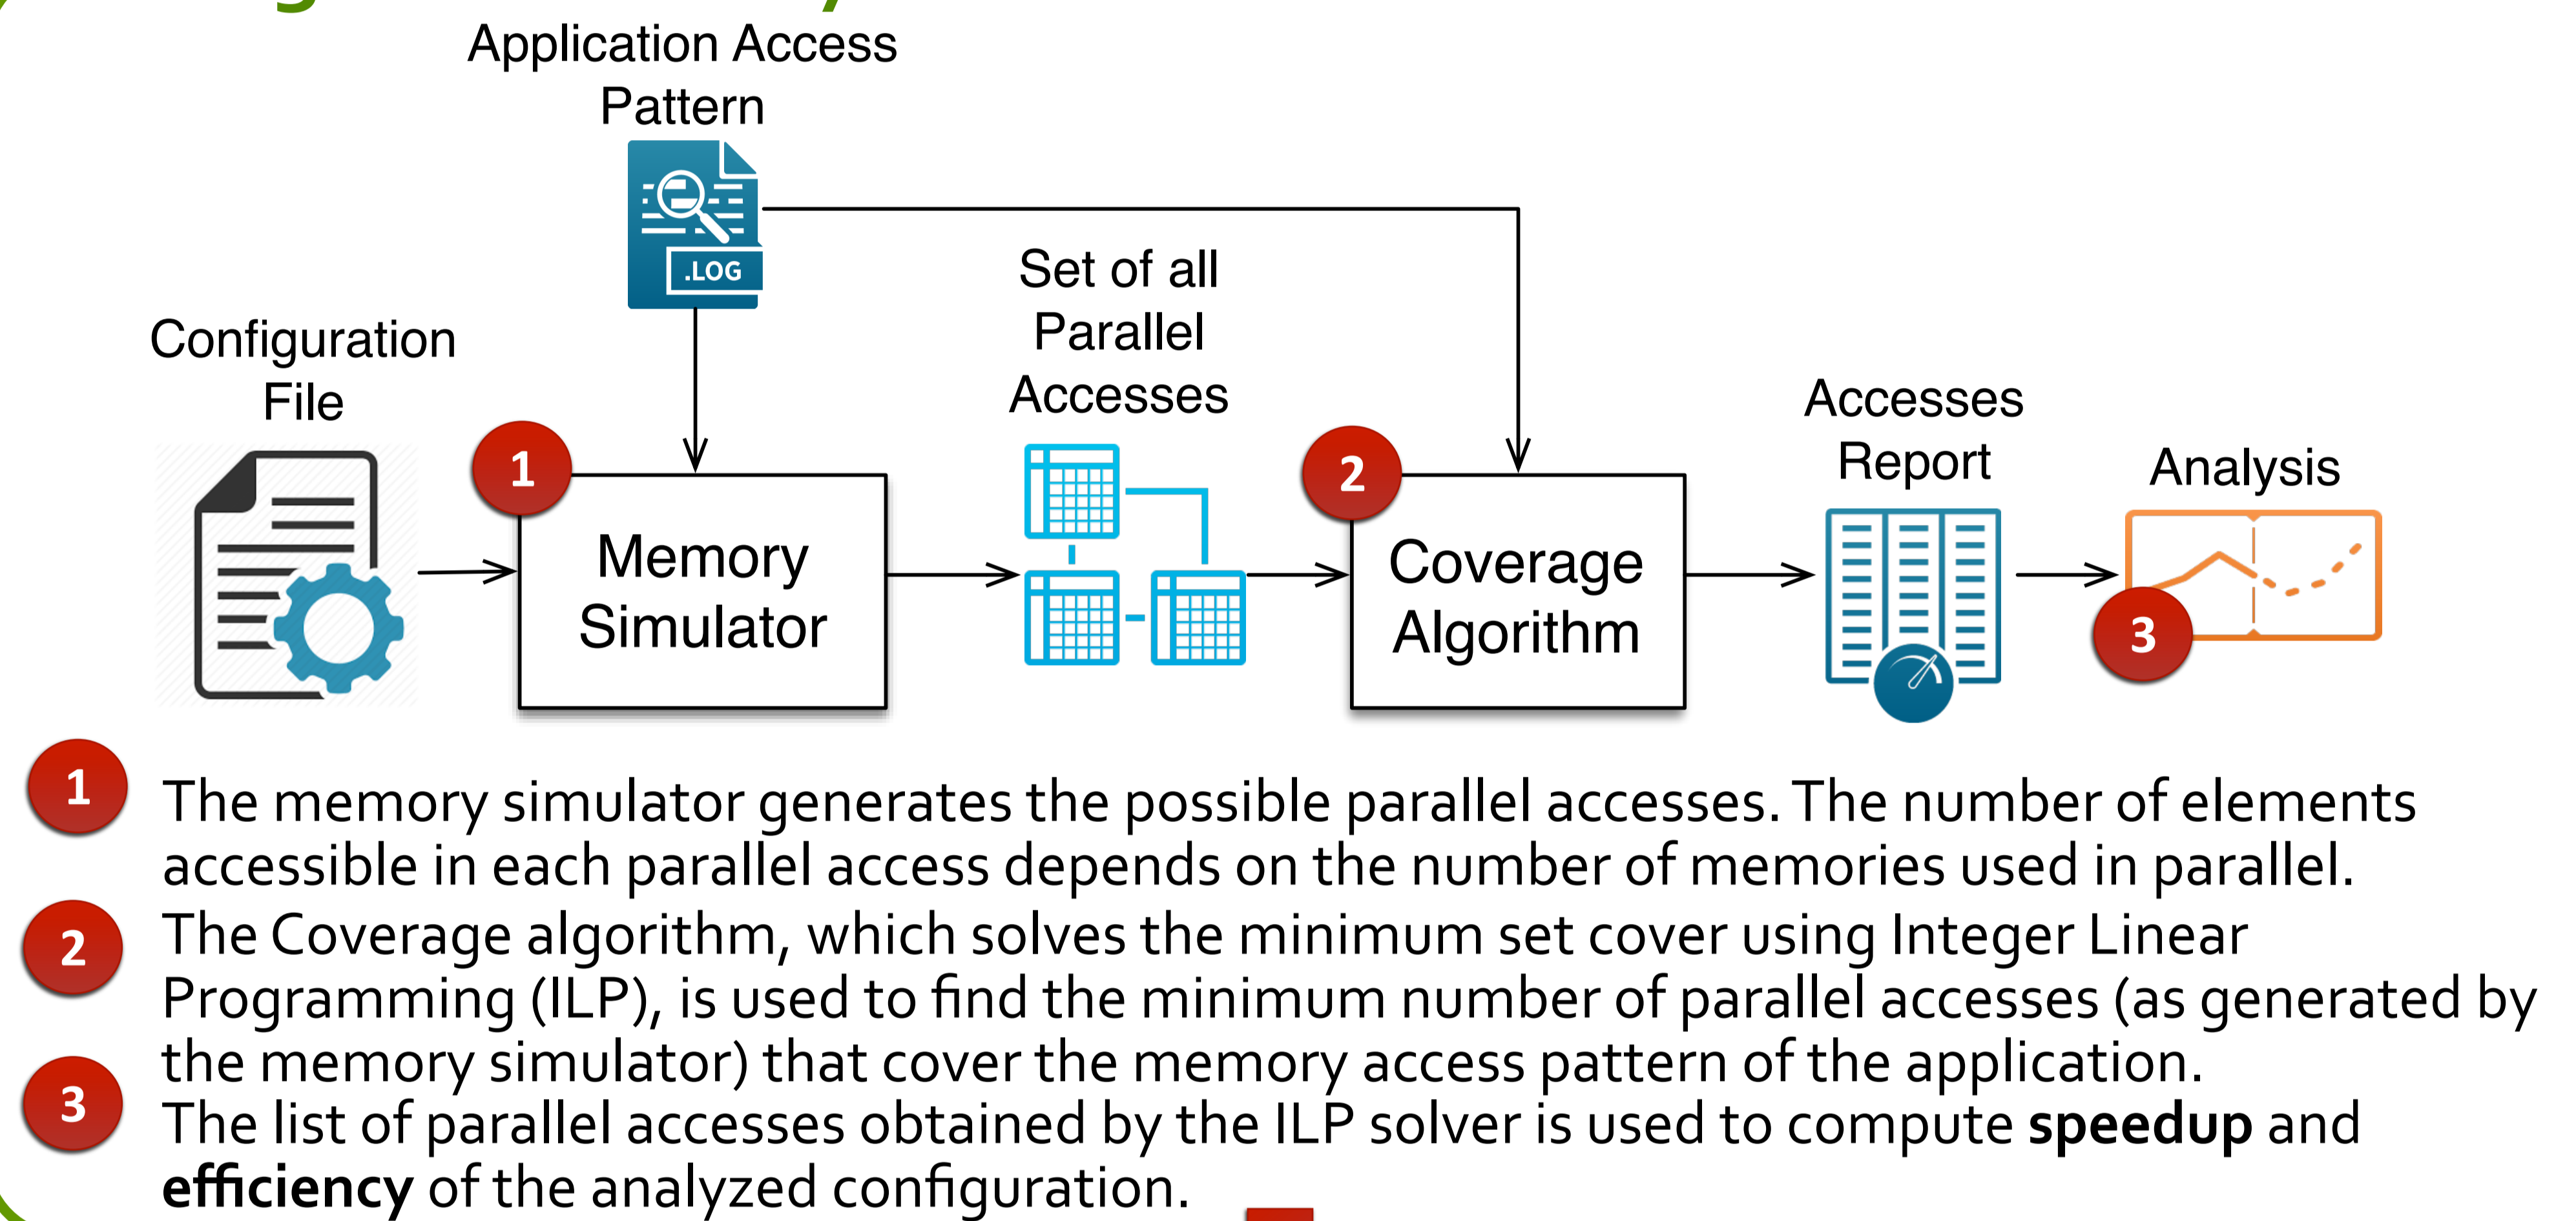


Metrics

$$\text{Speedup} = \frac{\# \text{Sequential Accesses}}{\# \text{Parallel Accesses}}$$

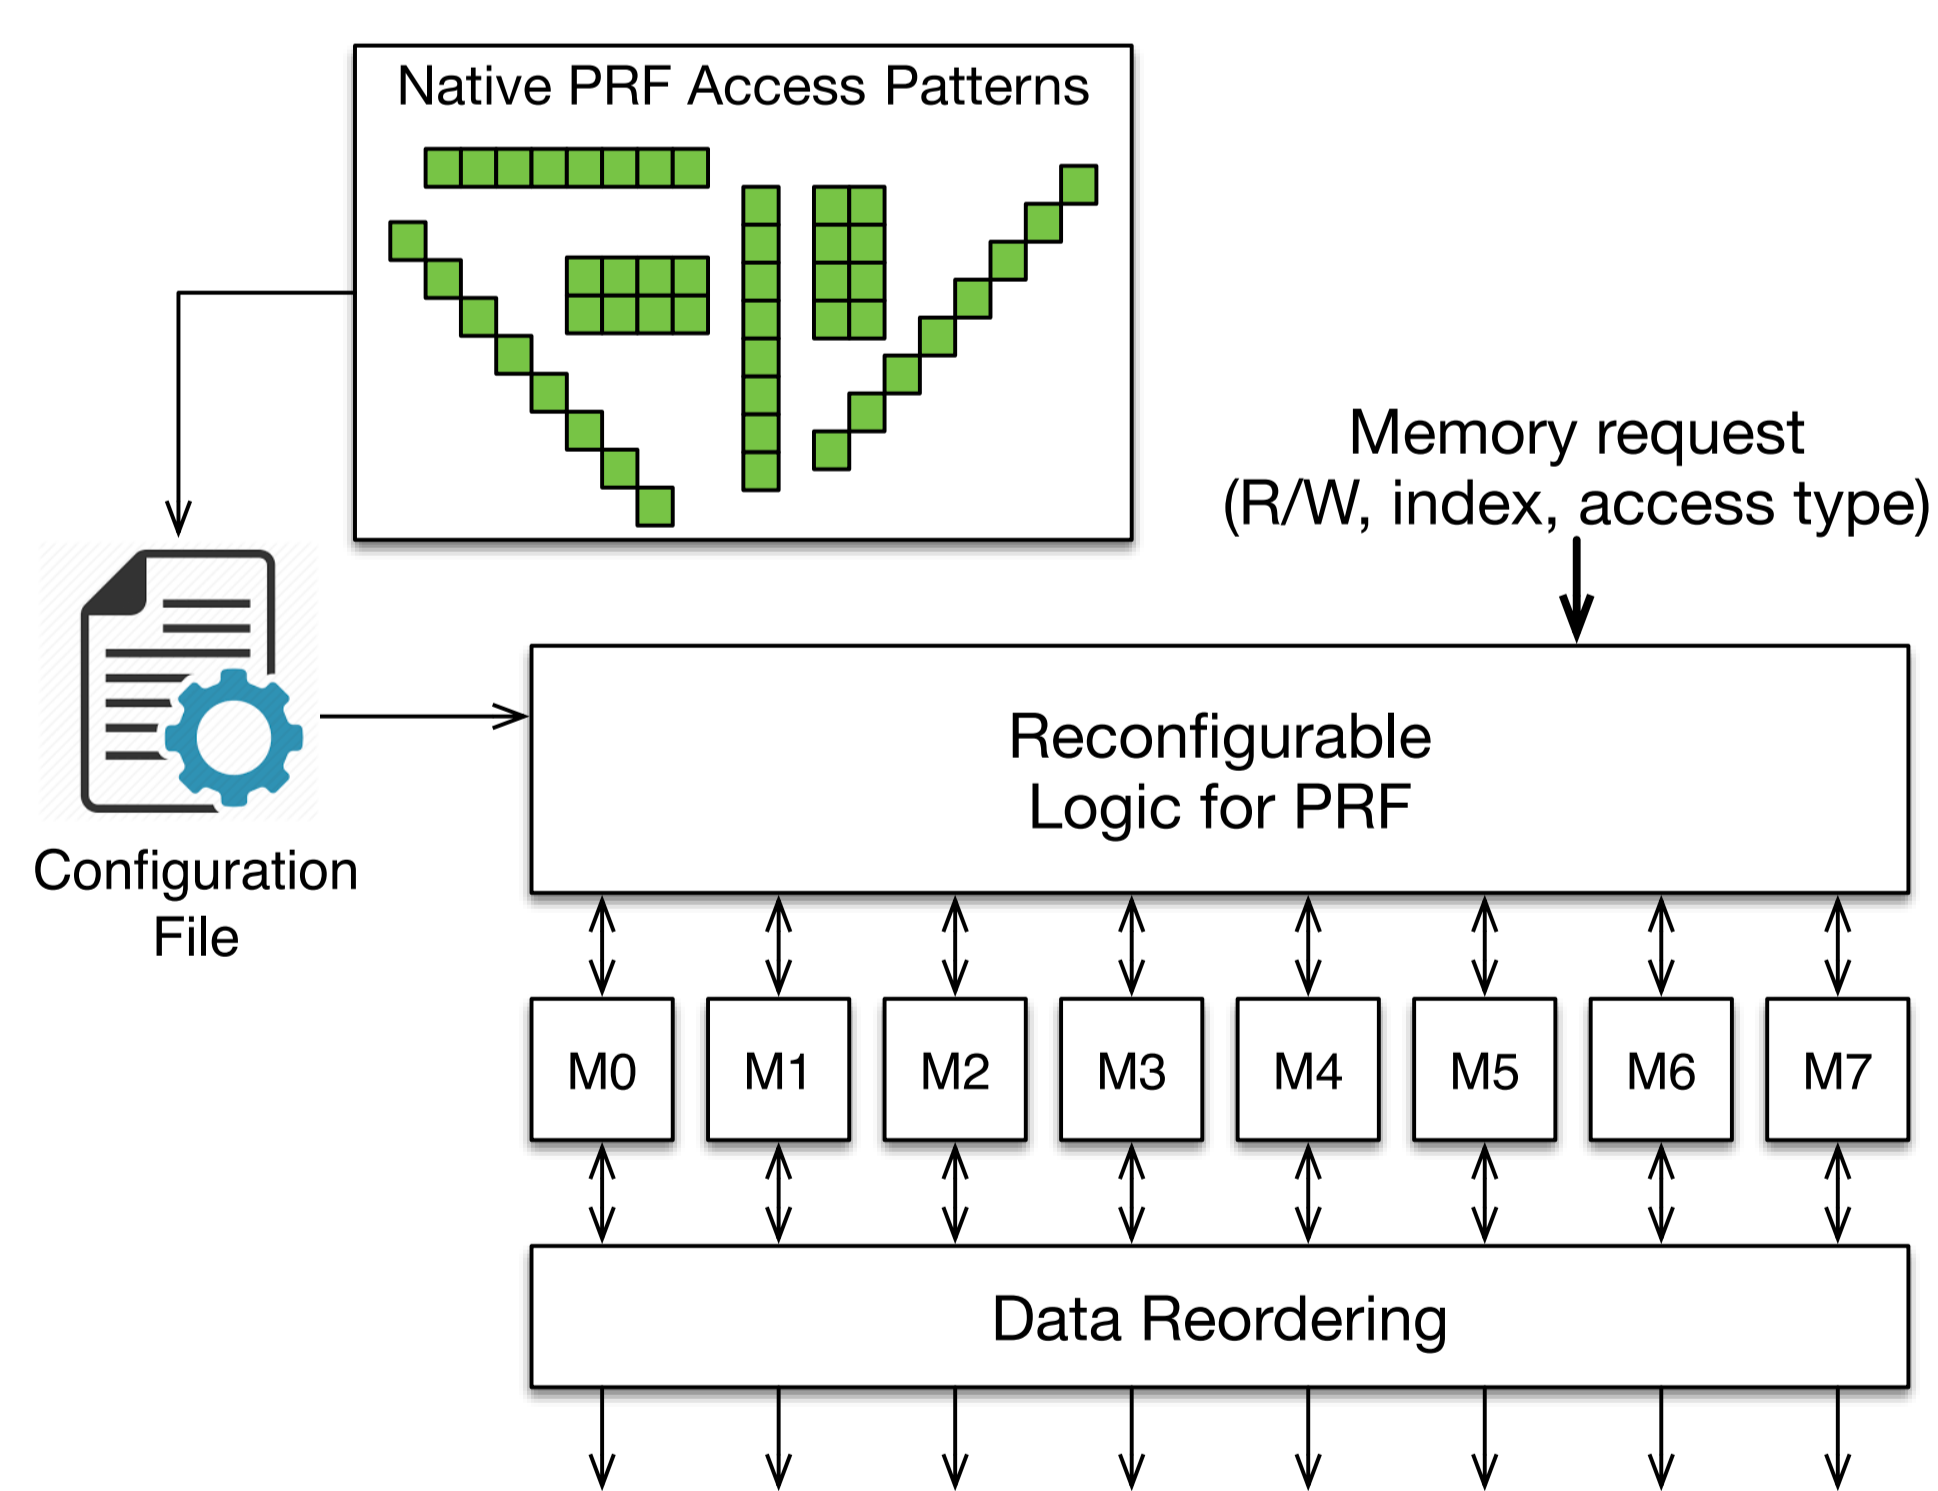
$$\text{Efficiency} = \frac{\text{ElementsReadByApplication}}{\text{ElementsReadByParallelMemory}}$$

Configuration Analysis




Parallel Reconfigurable Memory

We use the PRF [1] as a 2D reconfigurable scratchpad parallel memory.




A configuration for the PRF specifies the number of memory used in parallel and (the subset of) the patterns to be supported.

Speed-up



- (1) Parallel memory systems provide speed-up in all cases.
- (2) Combining multiple, different PRF access patterns is beneficial for speed-up.
- (3) The "staircase" effect indicates that there are multiple configurations possible to achieve the same speed-up.

Efficiency



- (1) Diversifying the supported PRF access patterns increases efficiency.
- (2) The design of a parallel memory system should maximize efficiency to avoid useless memory ops; thus, for the same speed-up, the design with the highest efficiency should be selected.

Take Home Message

Using parallel memories is not trivial. Our approach offers an integrated semi-automatic approach to customize the memory system for your application, and evaluate its potential performance.

More info ?



Get in touch!
g.stramondo@uva.nl
c.b.ciobanu@uva.nl
a.l.varbanescu@uva.nl



References

- [1] C. Ciobanu. **Customizable Register Files for Multidimensional SIMD Architectures**. PhD thesis, Delft University of Technology, Delft, Netherlands, March 2013.
- [2] Krste Asanovic, Ras Bodik, Bryan Christopher Catanzaro, Joseph James Gebis, Parry Husbands, Kurt Keutzer, David A Patterson, William Lester Plishker, John Shalf, Samuel Webb Williams, et al. **The landscape of parallel computing research: A view from Berkeley**. Technical report, Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley, 2006.
- [3] Shuai Che, Michael Boyer, Jiayuan Meng, David Tarjan, Jeremy W Sheaffer, Sang-Ha Lee, and Kevin Skadron. **Rodinia: A benchmark suite for heterogeneous computing**. In Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on, pages 44–54. IEEE, 2009.
- [4] Richard M Karp. **Reducibility among combinatorial problems**. In Complexity of computer computations, pages 85–103. Springer, 1972.