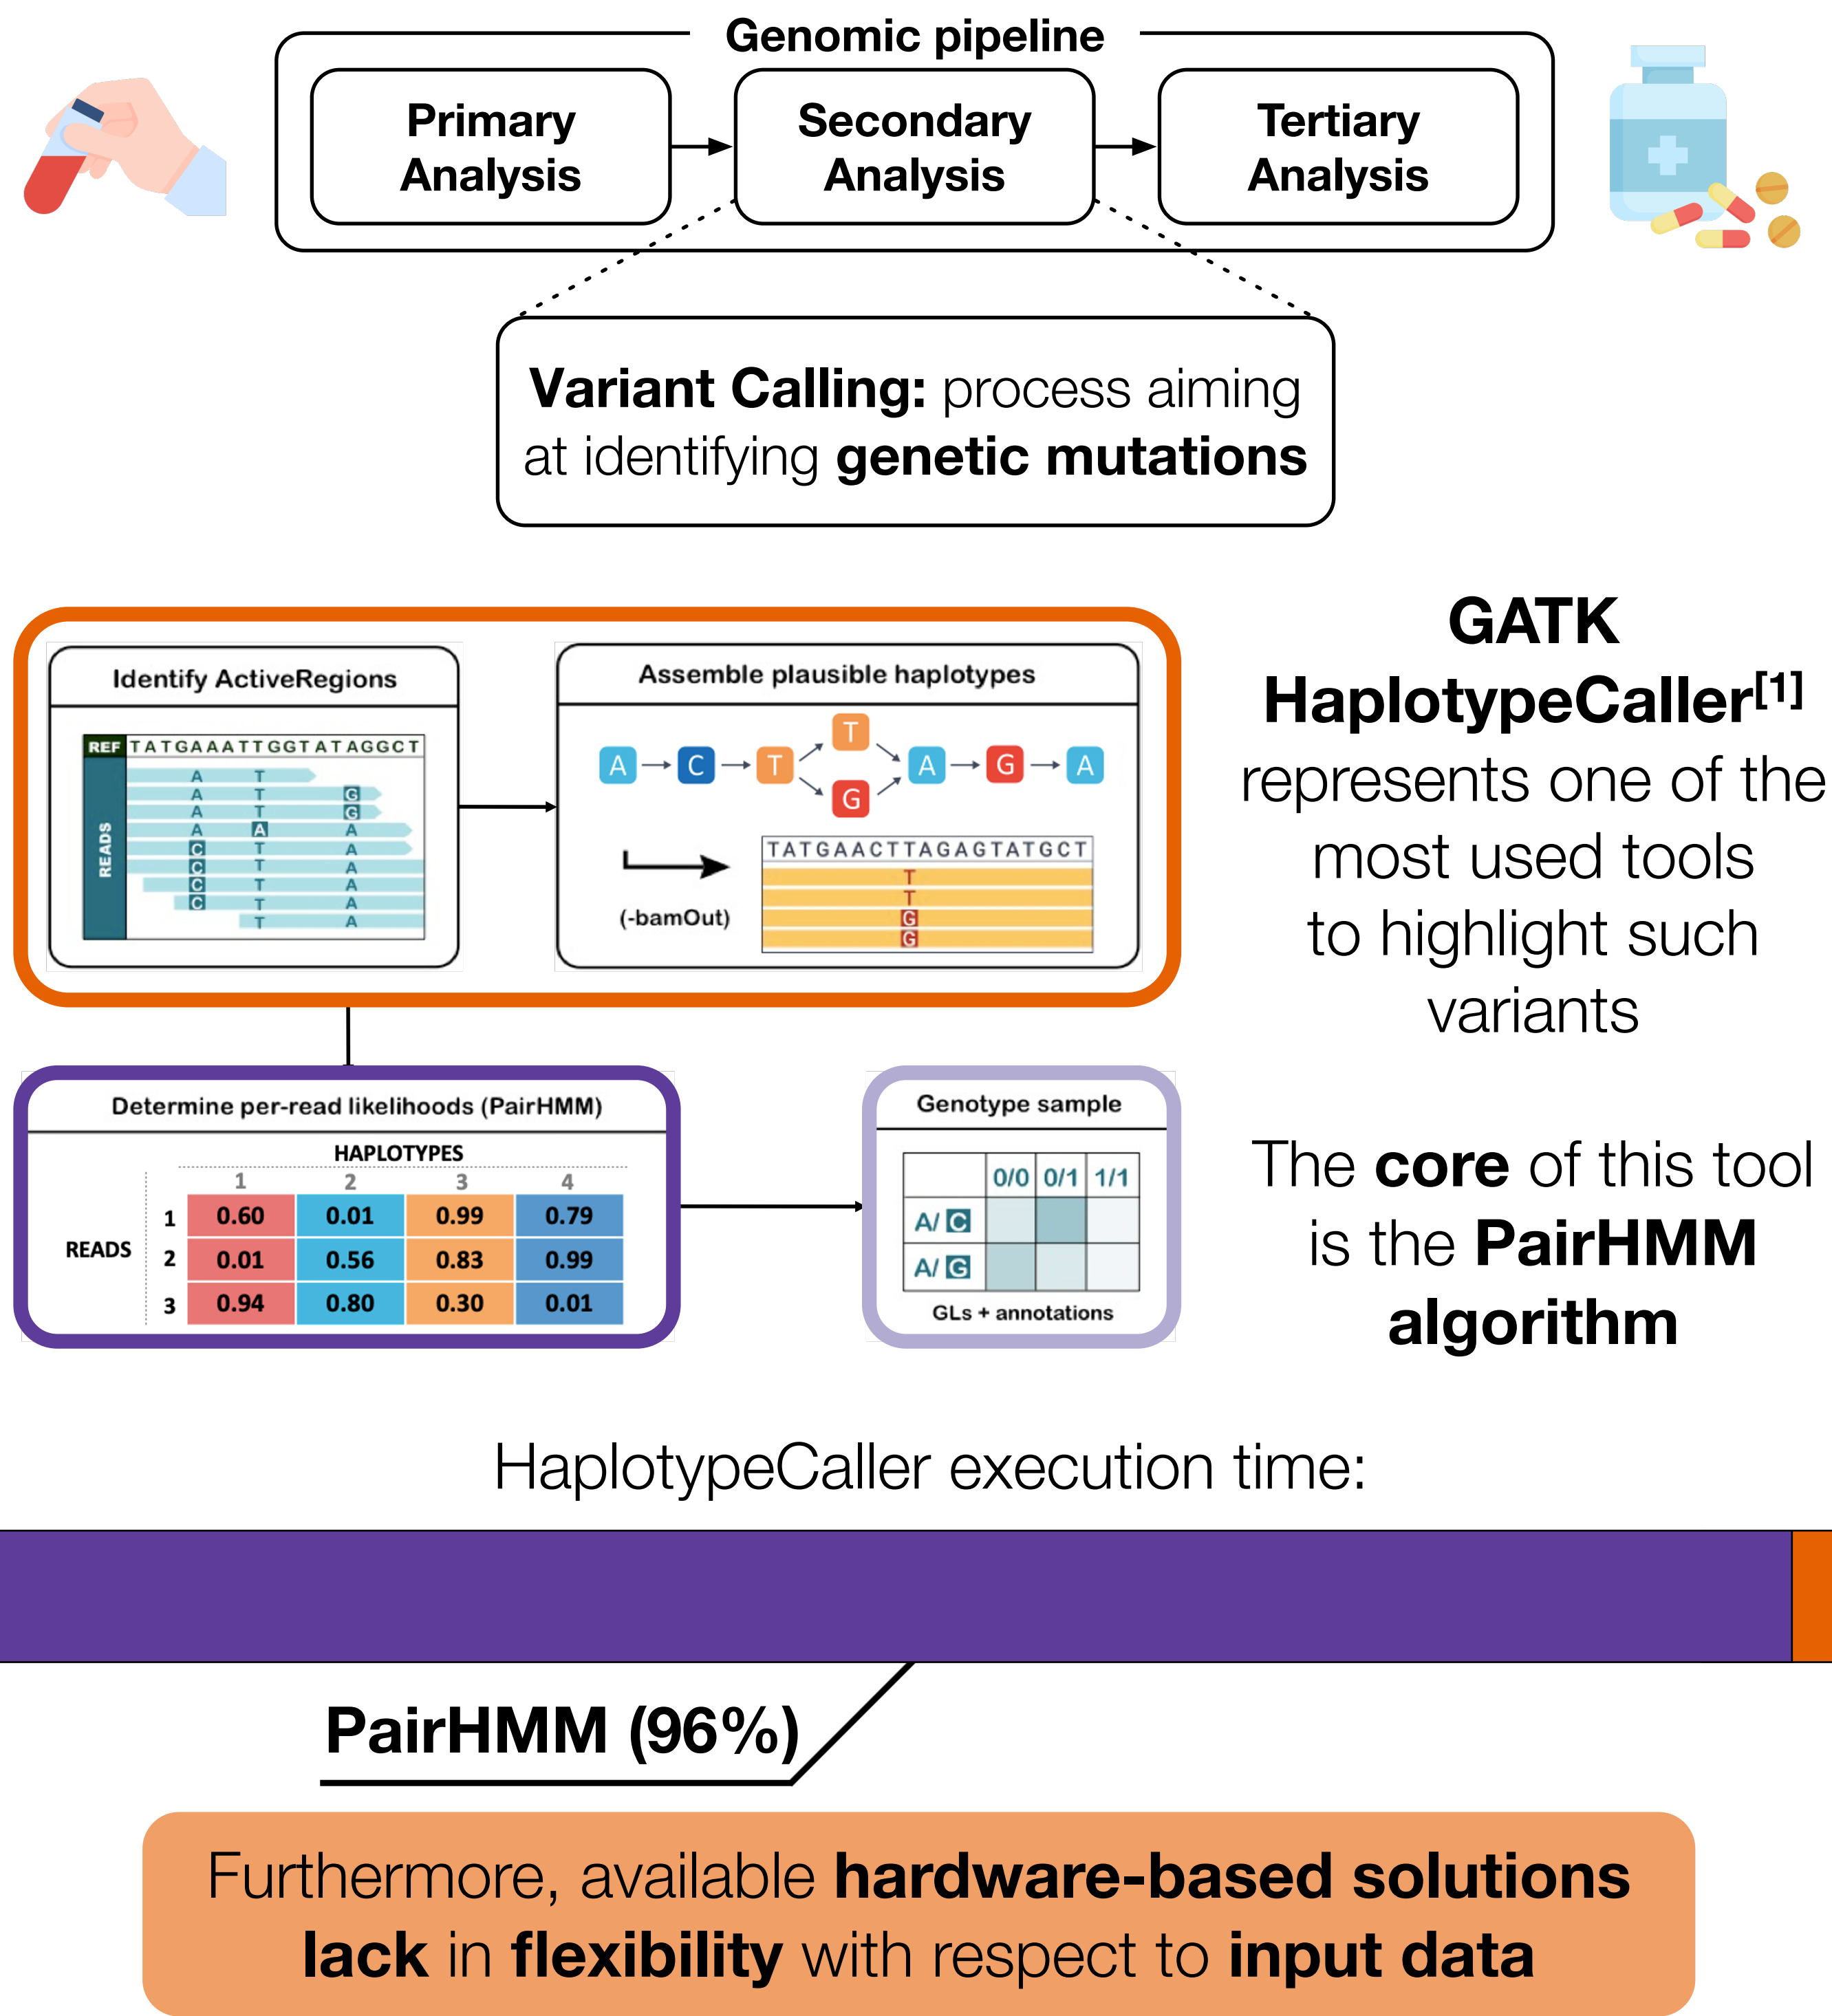


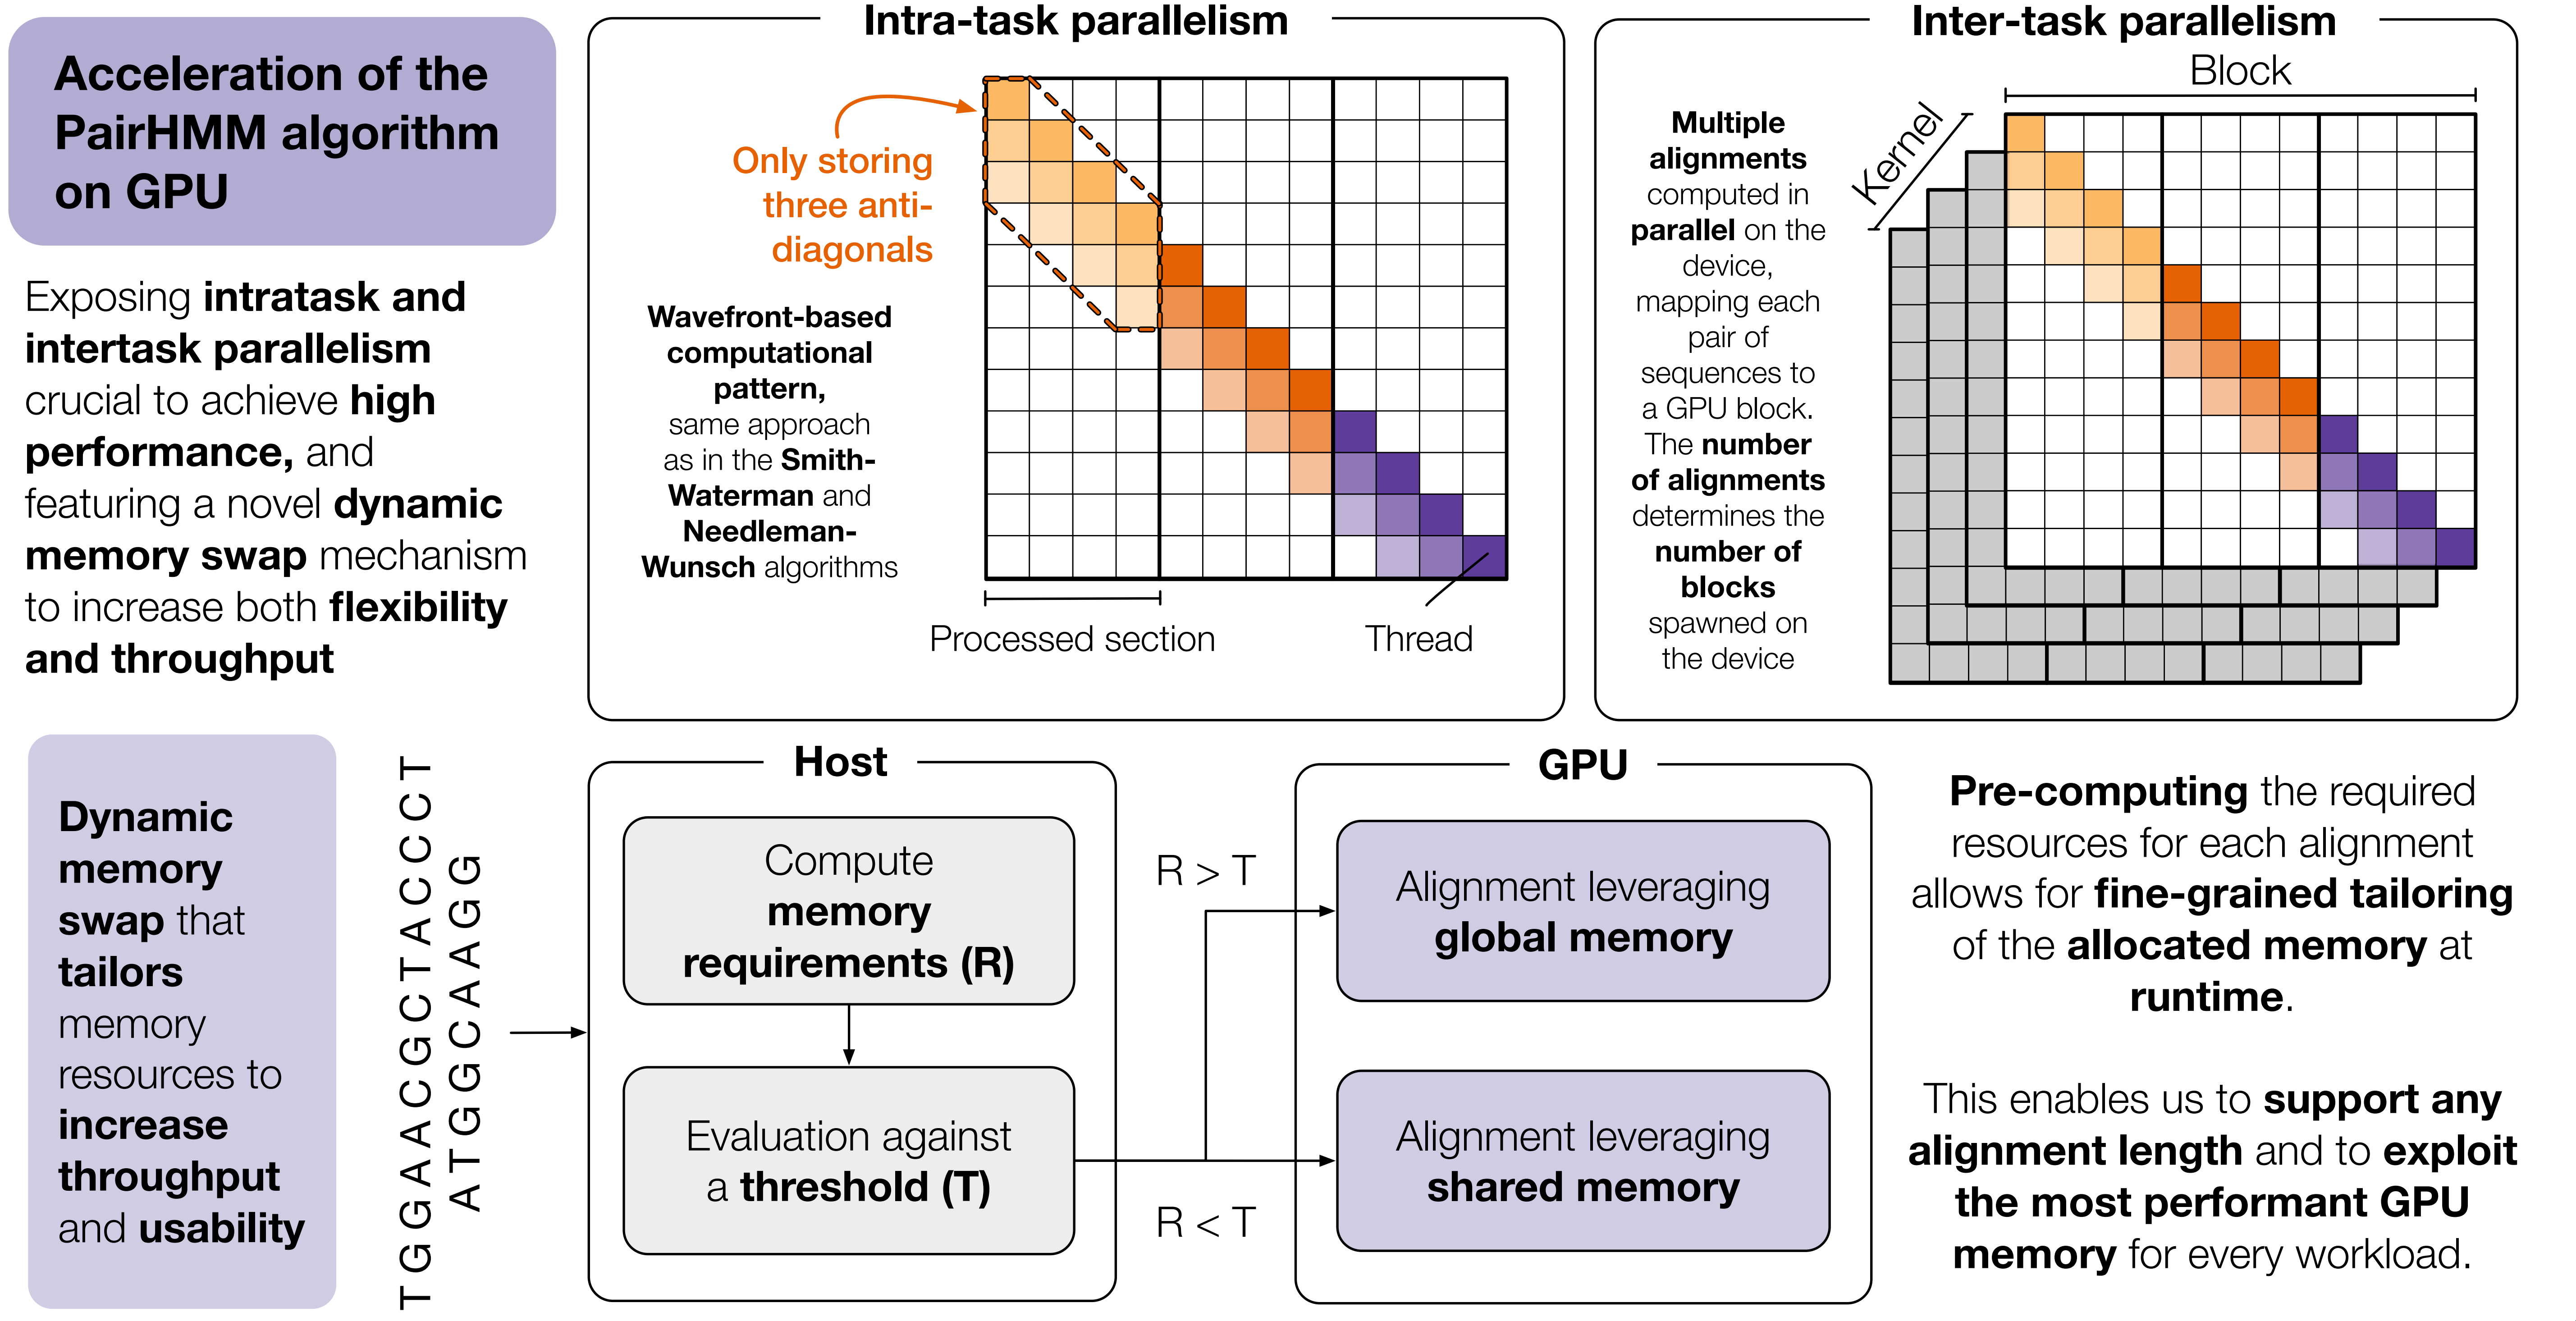
A Methodology for Accelerating Variant Calling on GPU

Beatrice Branchini, Alberto Zeni, Marco D. Santambrogio
 beatrice.branchini@pnnl.gov, {alberto.zeni, marco.santambrogio}@polimi.it

Context Definition

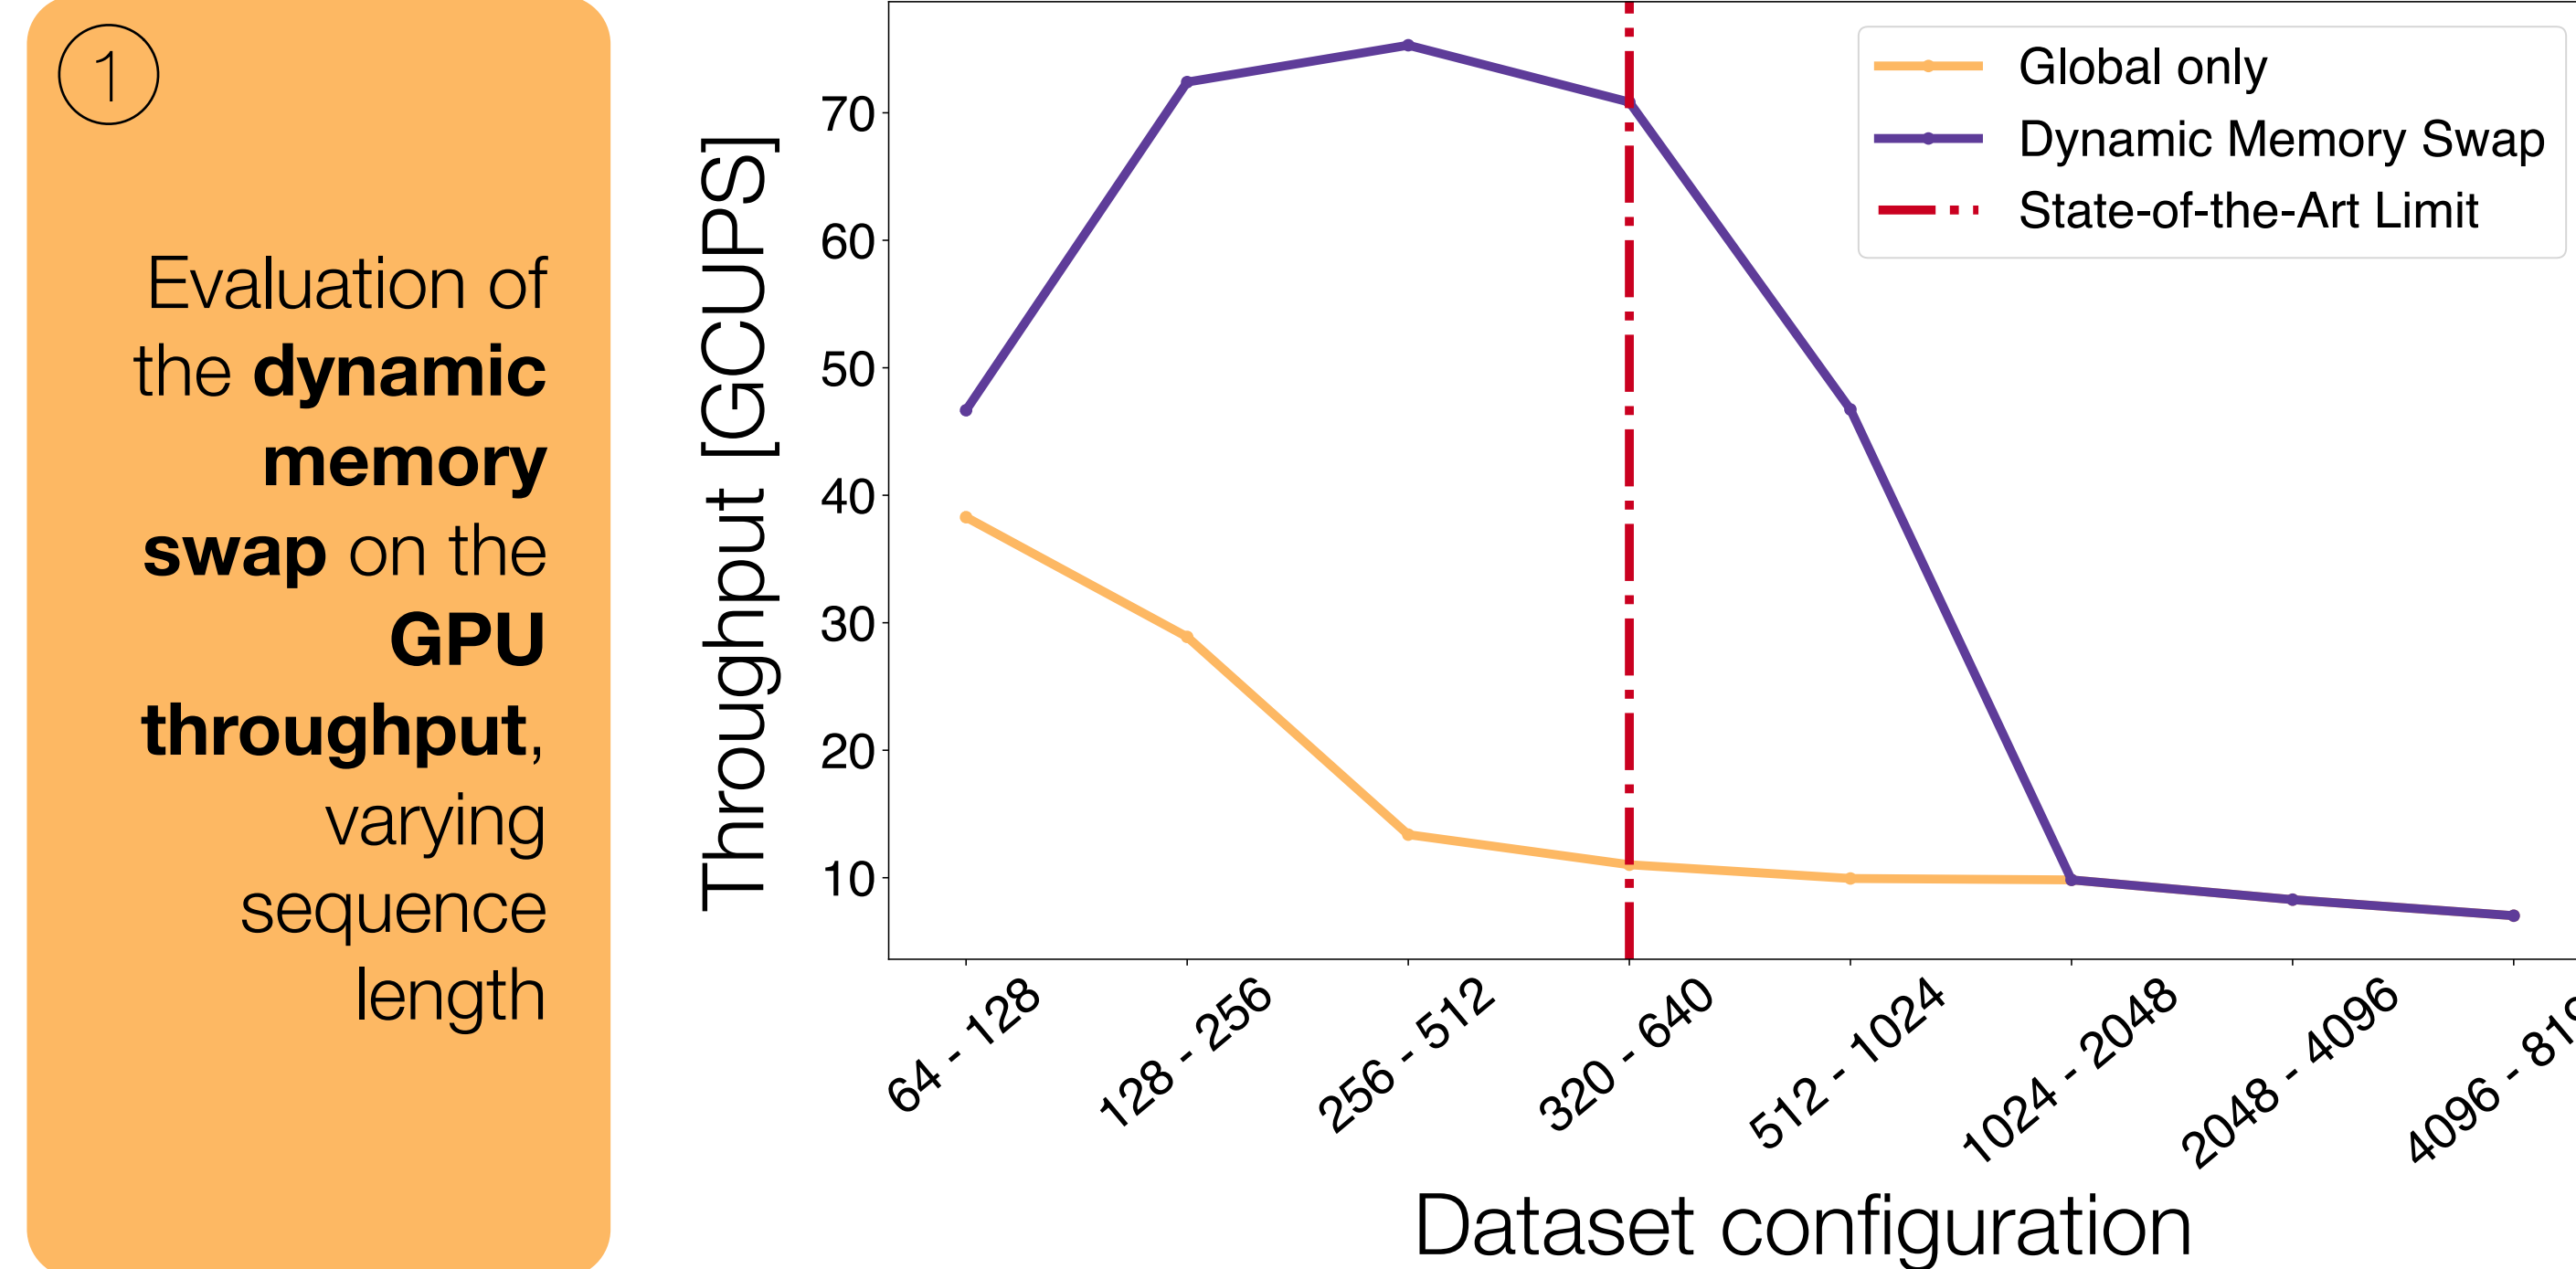


Proposed Solution




Experimental Results

GPU: NVIDIA Tesla V100 connected to a dual-socket Intel Xeon Platinum 8167M



② Evaluation with the **10s dataset^[2]**, comparing the **attained performance** with **State-of-the-Art solutions**



Conclusions

We present an **high-performance GPU accelerator** for the PairHMM algorithm featuring the **dynamic memory swap** to provide support for long sequences. Experimental results show **8154x** and **1.6x performance improvements** against the software baseline and the fastest hardware-accelerated solution, respectively.

Acknowledgments

This work has been partially supported by Compiler Frameworks and Hardware Generators for Innovative US Government Design. The Authors would also like to thank Oracle Research Program for the Oracle Cloud Credits