High-performance Matrix Computations

Prof. **Paolo Bientinesi**

pauldj@aices.rwth-aachen.de

Processor Info

Quad-core AMD Opteron Barcelona 8356 @2.3GHz.

The theoretical peak performance of one node is
27.6GFLOPS = 4cores * 2.3GHz * 3FLOPS.

Cache Information.

L1 Data Cache:

Total size:	64 KB
Line size:	64 B
Number of Lines:	1024
Associativity:	2

L1 Instruction Cache:

Total size:	64 KB
Line size:	64 B
Number of Lines:	1024
Associativity:	2

L2 Unified Cache:

Total size:	512 KB
Line size:	64 B
Number of Lines:	8192
Associativity:	16

L3 Unified Cache:

Total size:	2048 KB
Line size:	64 B
Number of Lines:	32768
Associativity:	32

GER: $A := A + \alpha x y^T$

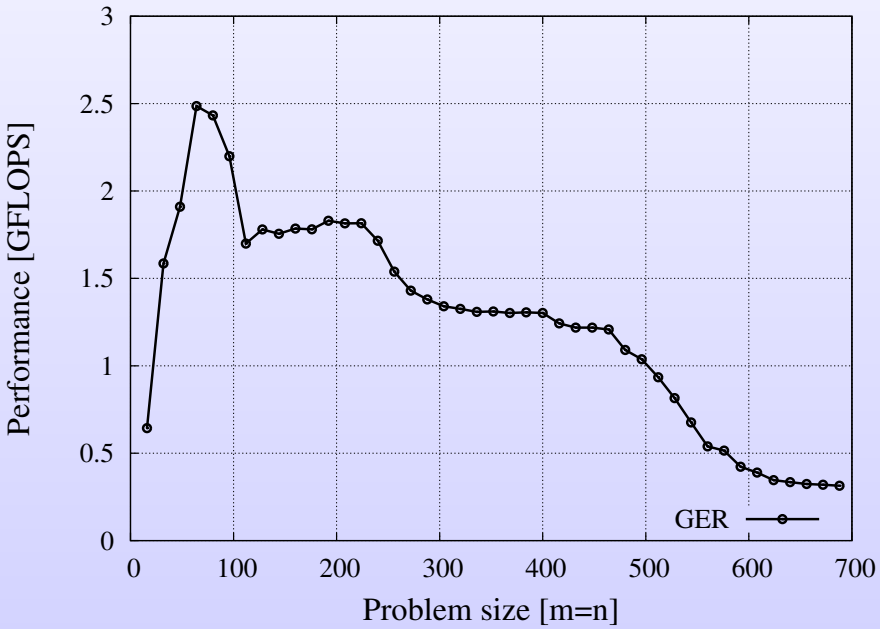

(a) execution 1

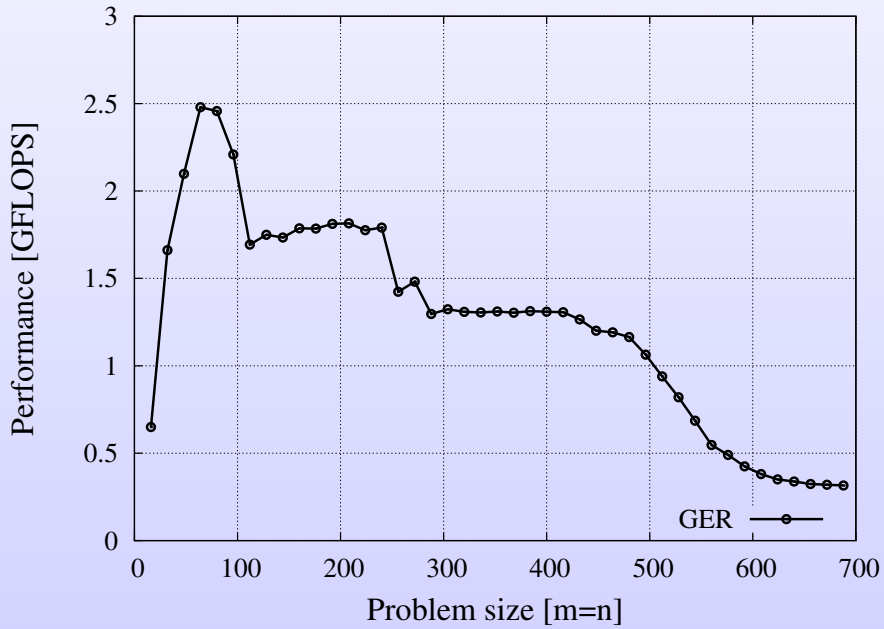

(b) execution 2

GER (2)

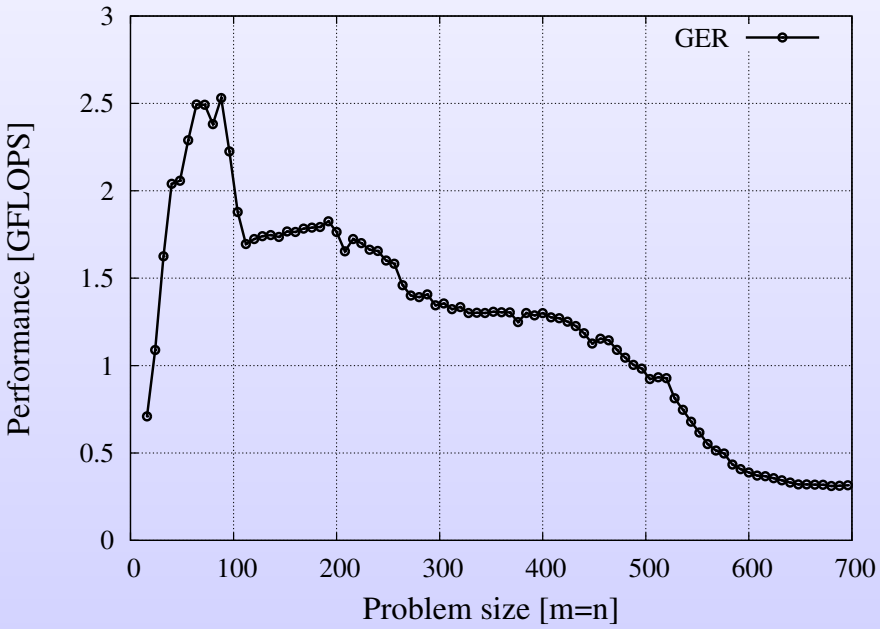

(c) step: 8

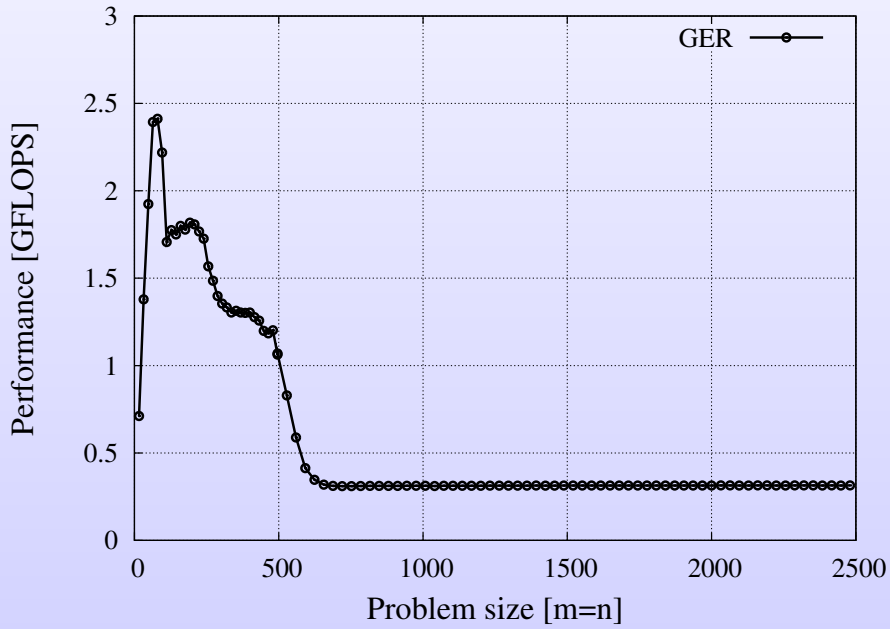

(d) larger size

GEMM

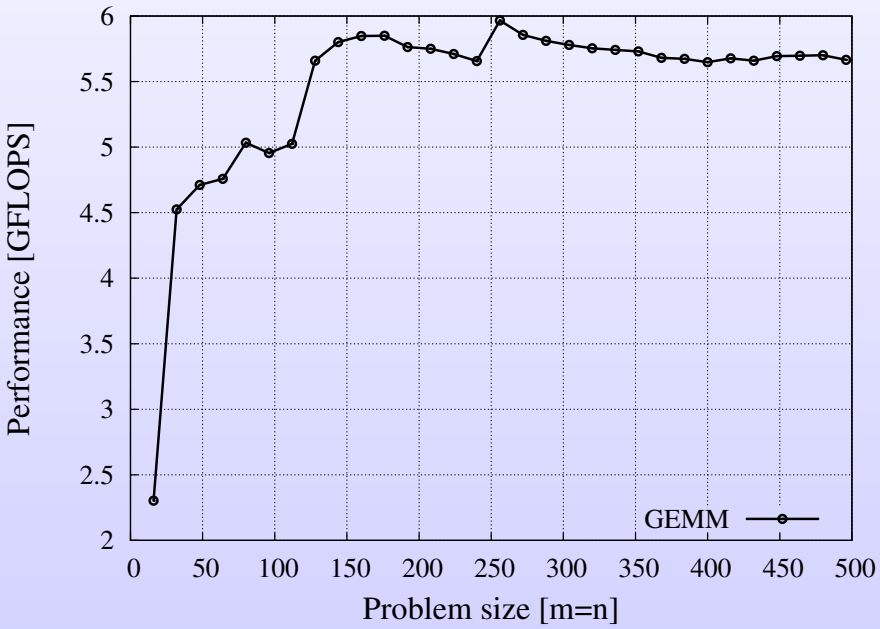

(e) small size

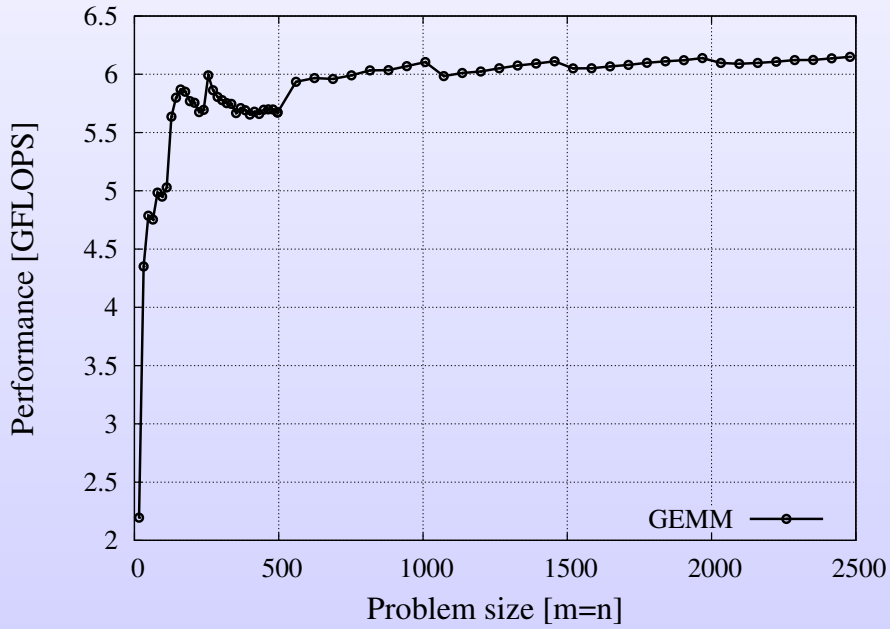

(f) larger size