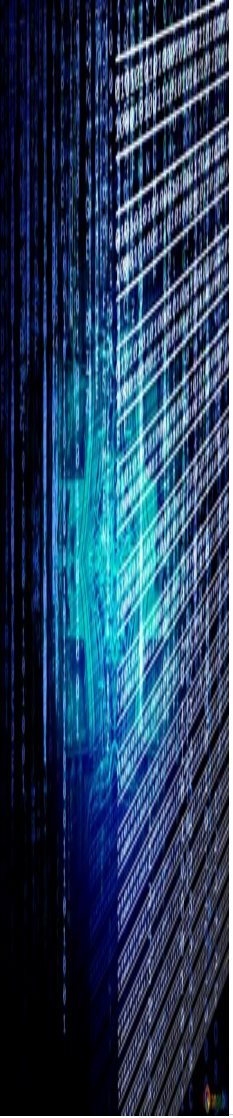

An evaluation and comparison of GPU solver libraries and hardware for OPM acceleration

Tong Dong Qiu
Razvan Nane

Structure of Flow

- Assembly
  - Finite volume method
  - Grid determines sparsity pattern
  - Linear system is assembled
- Linear solve, A*x=b
  - Preconditioned bicgstab solver
  - Sparsity pattern is constant

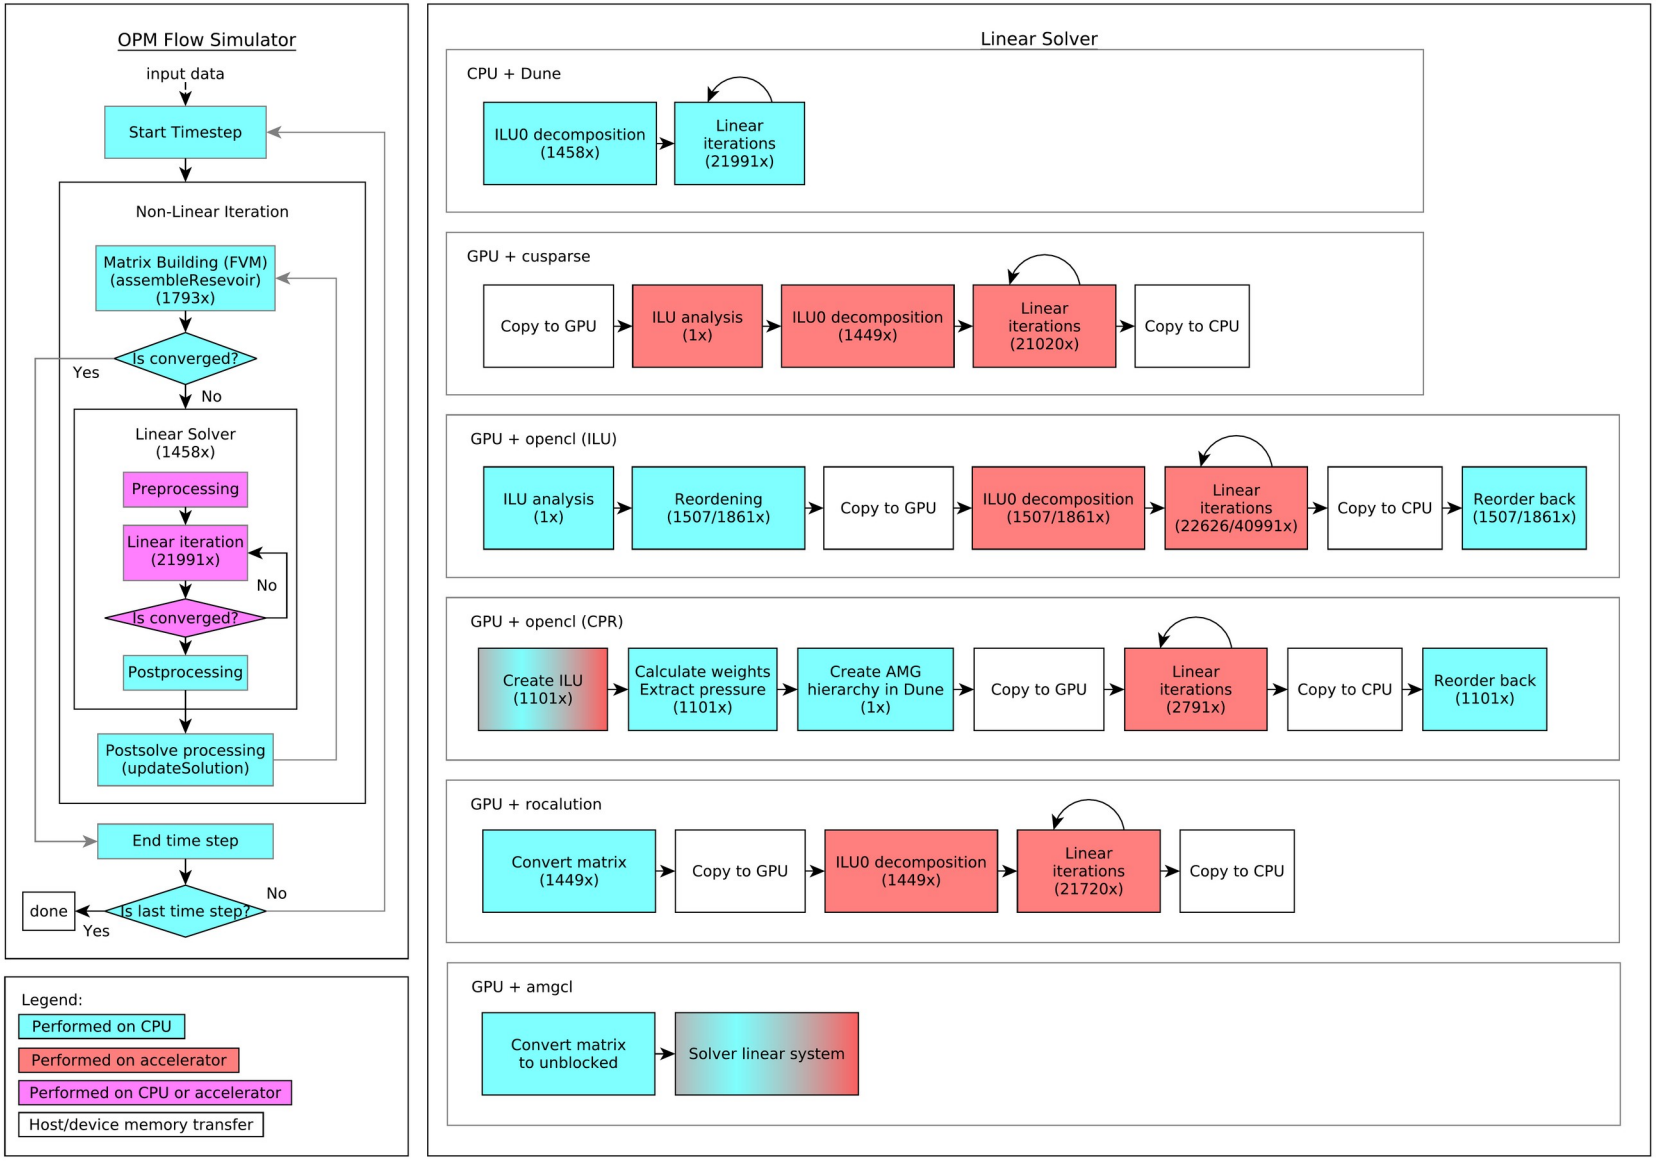
OPM Flow Simulator
input data
Start Timestep
Non-Linear Iteration
Matrix Building (FVM)
(assembleResevoir)
(1793x)
Is converged?
Yes
No
Linear Solver
(1458x)
Preprocessing
Linear iteration
(21991x)
No
Is converged?
Postprocessing
Postsolve processing
(updateSolution)
End time step
Is last time step?
No
done
Yes
Linear Solver
CPU + Dune
ILU0 decomposition
(1458x)
Linear
iterations
(21991x)
GPU + cusparse
Copy to GPU
ILU analysis
(1x)
ILU0 decomposition
(1449x)
Linear
iterations
(21020x)
Copy to CPU
GPU + opencl (ILU)
ILU analysis
(1x)
Reordening
(1507/1861x)
Copy to GPU
ILU0 decomposition
(1507/1861x)
Linear
iterations
(22626/40991x)
Copy to CPU
Reorder back
(1507/1861x)
GPU + opencl (CPR)
Create ILU
(1101x)
Calculate weights
Extract pressure
(1101x)
Create AMG
hierarchy in Dune
(1x)
Copy to GPU
Linear
iterations
(2791x)
Copy to CPU
Reorder back
(1101x)
GPU + rocalution
Convert matrix
(1449x)
Copy to GPU
ILU0 decomposition
(1449x)
Linear
iterations
(21720x)
Copy to CPU
GPU + amgcl
Convert matrix
to unblocked
Solver linear system
Legend:
Performed on CPU
Performed on accelerator
Performed on CPU or accelerator
Host/device memory transfer

Preconditioner: Blocked ILU0

- Blocked ILU0
  - Extract parallelism (level_scheduling or graph_coloring)
  - Reordering
  - Exact or ChowPatel decomposition
  - Application via 2 triangular solves

Preconditioner: Blocked ILU0

- Jacobi-ILU
  - Use MPI partitioning algorithm to divide grid
  - Partitioning stays constant
  - Create single matrix with more parallelism
  - Copy into jacobi matrix
  - Number of colors: 194 → 42
  - Linear solver gets 2 matrices

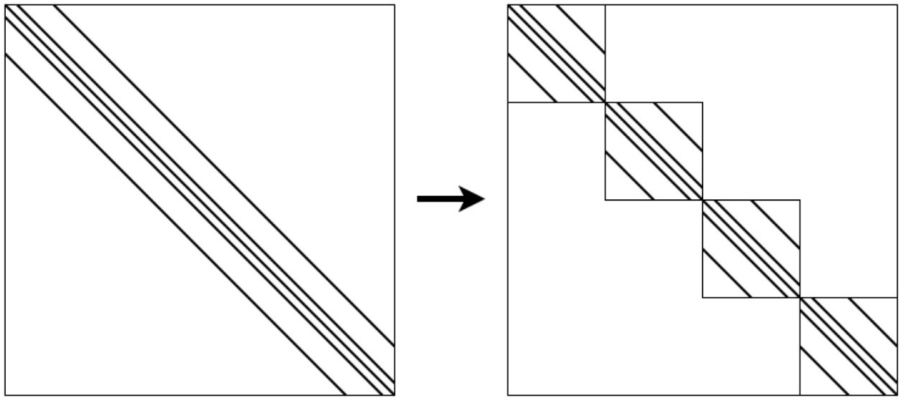

Preconditioner: CPR

- Constrained Pressure Reduction (CPR)
  - Blocked ILU0
  - Quasiimpes weights
  - Extract pressure component
  - Solve scalar system using Algebraic Multigrid (AMG)
  - Jacobi smoothing

WellContributions

- Commandline parameter
- StandardWell on GPU
- MultisegmentWell on CPU with UMFPack

amgcl

- 3rd party library
- Linear solvers, preconditioners, smoothers
- CPU, CUDA, VexCL(OpenCL)
- JSON parameter file

rocalution

- ROCm
- Transpose every block to CSC
- Blackbox ilu0-bicgstab

rocsparse

- ROCm
- Same API as cusparse
- More control → jacobi-ILU possible

Experimental Setup

	Cores	Clock	L3 Cache	Memory	Mem. Bandwidth (GB/s)
AMD EPYC 7753	64	2.45 – 3.5 GHz	256 MB	PCIe 4.0, 3200 MHz	205
AMD EPYC 7601	32	2.2 – 3.2 GHz	64 MB	PCIe 3.0, 2666 MHz	171

Name	Num. cores[1]	Max. FP64 TFLOPS	Mem. Size (GB)	Mem. Bandwidth (GB/s)
NVIDIA Tesla A100	6912	9.7	40, 80	2000
AMD Instinct MI100	7680	11.5	32	1229

[1]: NVIDIA CUDA cores are not the same as AMD stream processors

Benchmarks eX^3, NORNE

- AMD EPYC 7763, NVIDIA A100-SXM-80GB
- ILU0, masters of Jan 2022, coupled wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU	514	158	158	1857	1489	21107
cusparse	462 (1.11)	148	123 (1.28)	1798	1438	20689
opencl LS	565 (0.91)	150	224 (0.71)	1810	1447	20797
opencl LS, 150 blocks	457 (1.12)	142	129 (1.22)	1775	1419	24245
opencl GC	562 (0.91)	189	158 (1.00)	2238	1857	43442
amgcl+vexcl 9	521 (0.99)	157	170 (0.93)	1867	1502	54947

Benchmarks eX^3, NORNE

- AMD EPYC 7763, AMD Instinct MI100
- ILU0, masters of Jan 2022, coupled wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU	502	152	157	1857	1489	21107
rocalution	434 (1.16)	138	119 (1.32)	1810	1449	21720
opencl LS	610 (0.82)	144	277 (0.57)	1810	1447	20797
opencl LS, 150 blocks	455 (1.10)	142	127 (1.24)	1775	1419	24245
opencl GC	587 (0.86)	186	179 (0.88)	2303	1915	43519
amgcl+vexcl 7	518 (0.97)	147	179 (0.88)	1865	1503	62803

Benchmarks eX^3, NORNE

- AMD EPYC 7601, AMD Instinct MI100
- ILU0, masters of Jan 2022, coupled wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU	733	227	237	1857	1489	21107
rocalution	675 (1.09)	244	138 (1.72)	1810	1449	21720

Benchmarks eX^3, NORNE

- AMD EPYC 7601
- ILU0, masters of Jun 2022, separate wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU (4 MPI)	616	173	250	1857	1493	22054
CPU (16 MPI)	412	76	222	1857	1495	22197
CPU (64 MPI)	146	29	45	1888	1522	23416

Benchmarks eX^3, bigger case

- AMD EPYC 7763, AMD Instinct MI100
- ILU0, masters of Feb 2022, coupled wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU (1 MPI)	10591	4562	3885	2250	2008	12418
CPU (4 MPI)	6869 (1.54)	3190	2338 (1.66)	2335	2078	12370
opencl LS	13346 (0.79)	5951	4640 (0.84)	2257	2006	12603
opencl LS, 2 blocks	16394 (0.65)	7381	5481 (0.71)	2320	2077	12841
rocalution	7975 (1.33)	4187	1846 (2.10)	2035	1807	13033

Benchmarks eX^3, bigger case

- AMD EPYC 7763
- ILU0, masters of Jun 2022, separate wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU (1 MPI)	11327	5595	3041	2447	2181	12993
CPU (64 MPI)	1298 (8.73)	439	541 (5.62)	2864	2567	13147
CPU (128 MPI)	1117 (10.14)	348	460 (6.61)	2513	2248	12950

Benchmarks eX^3, NORNE

- AMD EPYC 7763, AMD Instinct MI100
- CPR, masters of Jan 2022, coupled wellcontributions

	Total time	Assembly time	Linear solve time	Linearizations	Newton Iterations	Linear Iterations
CPU	438	139	117	1701	1345	3392
opencl LS, 0 blocks	404 (1.08)	116	120 (0.98)	1446	1101	2791
opencl LS, 150 blocks	393 (1.11)	115	111 (1.05)	1442	1099	3288

Conclusion

- GPU can be faster than single MPI process
- Multiple MPI is still faster

Future work

- Remove reordering for OpenCL GPUs
- Extract pressure on GPU
- Rocsparse with Jacobi-ILU
- Assembly

Questions