







### Recent and upcoming developments in Ginkgo. Pratik Nayak



KIT - The Research University in the Helmholtz Association



## Outline

- Motivation and objectives
- Ginkgo?
- Design philosophy and features
- A quick overview of Release 1.6.0, features and performance
- Features in development and possibly in Release 1.7.0
- Future outlook



3

### **OpenCARP** simulation workflow





# Why Ginkgo ?

- Significant part of the total time spent in linear solver.
- Crucial to optimize this bottleneck, particularly for the cell-by-cell model.
- Move computation to GPU, minimize data movement to and from GPU.



Bi-domain model: Ginkgo (1 GPU) v/s PETSc (32 CPUs)



## Ginkgo: A high performance library





### Ginkgo: Sustainable development





### Ginkgo: Latest release - 1.6.0

- Support for sparse direct factorizations + solvers.
- Improved AMG performance.
- Profiler annotations: NVTX, ROCTX, VTune, TAU
- Stream support: Executors can now take a stream argument
- New distributed Schwarz preconditioner.
- Matrix reorderings: METIS (nested dissection), AMD (Fill-in reducing)
- Mixed precision SpMV with CSR.

More details: <u>https://github.com/ginkgo-project/ginkgo/releases/tag/v1.6.0</u>



### Overview of features: Single executor

	Functionality		CUDA	HIP	DPC++
Basic	SpMV	Ś	S	S	ø
	SpMM	S	S	S	S
	SpGeMM	ø	ø	ø	ø
vers	BICG	ø	S	ø	ø
	BICGSTAB	S	S	ø	ø
sol	CG	ø	S	ø	ø
<sup>o</sup>	CGS	S	S	ø	S
Kry	GMRES	ø	ø	ø	S
	IDR	ø	S	ø	ø
ers	(Block-)Jacobi	Ś	Ś	S	S
ion	ILU/IC		S	ø	S
Precondit	Parallel ILU/IC	ø	ø	ø	S
	Parallel ILUT/ICT	ø	S	ø	ø
	Sparse Approximate Inverse	ø	ø	ø	ø

	Functionality	OMP	CUDA	HIP	DPC++
	AMG preconditioner	ø	S	T	ø
AMG	AMG solver	S	S	ø	ø
	Parallel Graph Match	Ś	S	S	ø
t	Symbolic Cholesky	ø	S	ø	S
ire	Numeric Cholesky	S	Ś	ø	
sed	Symbolic LU	S	Ś	S	ø
Dars	Numeric LU	S	S	S	
S	Sparse TRSV	S	S	ø	
	<b>On-Device Matrix Assembly</b>	S	ø	ø	ø
sa	MC64/RCM reordering	ø			
Utiliti	Wrapping user data		<hr/>		
	Logging		<hr/>		
	PAPI counters			1	



## Overview of features: Distributed

- All crucial solvers and operations supported.
- Local block preconditioners supported through the distributed Schwarz preconditioner.
- Compose matrix formats between diagonal and off-diagonal matrices.





## Ginkgo: Distributed data distribution



Sketch created by Yu-Hsiang Tsai

FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing



### Weak scaling: SpMV on Frontier

Cray MPICH + AMD MI250X on 16k GCDs



Weak scaling: problem size increases with parallel resources

FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing



## Strong scaling: Solvers on Frontier



FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing

Strong scaling: constant problem size



**EXASGD** 

### GPU resident sparse direct solvers

- Power Grid Simulations
- All GPU solvers outperform CPU solvers
- Ginkgo first GPU-resident solver
- Enables factorization on AMD and also on Intel GPUs (No vendor alternatives available)



© Slaven Peles

#### FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing



### Some utilities: Profiler annotations

NVTX





ROCTX

FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing



### Some utilities: Profiler annotations





generate(gko::solver::Cg <double< td=""><td>apply(gko::solver::Cg<double>)</double></td></double<>	apply(gko::solver::Cg <double>)</double>
generate(gko::preconditioner::)	
sort_by generate#9	



# What's in development and maybe in 1.7.0

- Batched iterative methods.
- More distributed preconditioners:
  - Distributed AMG (Coming soon)
  - Distributed BDDC (See Fritz's talk)
  - Two-level Schwarz methods.
- Easy solver configuration with JSON (already in use in openCARP)
- Local-only distributed matrix format: Abstract matrix format to describe and operate on overlapping partitions.
- Custom allocator support: Get around large allocation overheads with memory pool allocators (in applications).



## Batched methods?

- Batching: <u>Independent</u> computations that can be <u>scheduled in parallel</u>.
- Are highly suitable for GPUs and processors with many parallel computing units.
- Can maximize utilization of the GPU, due to excellent scalability.



## Batched methods in Combustion

- Low mach number, non-compressible flow: leads to very stiff ODEs
- Linear multi-step time stepping with BDF

$$(I - h_n \beta_{n,0} \frac{df}{d\phi}) [\phi^{n(m+1)} - \phi^{n(m)}] = -G(\phi^{n(m)})$$

- Sparsity pattern,  $\mathbf{A} \equiv (I h_n \beta_{n,0} \frac{df}{d\phi})$ same for each cell.
- Hence independent systems to be solved at each cell of the order of the number of cells 200k ~ 2e6 on each GPU

FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing









## Batched methods in Fusion Plasma

- Similar idea as before, but for Particle in cell code (PIC) XGC simulating gyrokinetic plasma in a tokamak.
- Reduction in linear solve time by about 90%.
- Scales extremely well.





## Batched methods in Fusion Plasma

- Similar idea as before, but for Particle in cell code (PIC) XGC simulating gyrokinetic plasma in a tokamak.
- Reduction in linear solve time by about 90%.
- Scales extremely well, due to the embarrassing parallelism.



ime/step (sec)



# Easy configuration for Ginkgo

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

- User provides a simple
   JSON config file at runtime.
- Solver and interface within OpenCARP configured automatically from these settings.

FiNE – Fixed point Numerics for Exascale SCC – Steinbuch Centre for Computing

{		
	"name":	"solver",
	"base":	"CgFactory",
	"criter:	ia": [
	{	
		"name": "AbsTol",
		"base": "ResidualNormFactory",
		"baseline": "absolute",
		"reduction_factor": 1e-8
	},	
	{	
		"name": "Iter".
		"base": "IterationFactory".
		"max iters": 100
	}	
	1.	
	"exec":	"executor"
}.		
{		
	"name":	"precond".
	"base":	"JacobiFactory"
	"max bl	ock size": 32
	"exec":	"executor"
7		
,		



20 21

22 23

24 25

26

27 28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

56

```
"name": "precond",
"base": "MultigridFactory",
"criteria": [
        "base": "IterationFactory",
        "max iters": 1
],
"max_levels": 9,
"min_coarse_rows": 64,
"post_uses_pre": true,
"zero_guess": true,
"pre_smoother": [
        "base": "IrFactory",
        "criteria": [
                "base": "IterationFactory",
                "max_iters": 1
        ],
        "solver": {
            "base": "JacobiFactory",
            "max block size": 1
        }.
        "relaxation_factor": 0.9
],
"mg_level": [
        "base": "AmgxPgmFactory",
        "deterministic": false
],
"coarsest solver": {
    "base": "IrFactory",
    . .. . . . .
```



### Easy configuration for Ginkgo

83 template <typename T, typename S> void ginkgo\_solver<T, S>::setup\_solver(abstract\_matrix<T, S>& mat, double tol, int max\_it, short norm, 84 85 const char\* name, bool has\_nullspace, void\* void\_logger, 86 const char\* solver\_opts\_file, const char\* default\_opts) 87 88 using RM = gko::extension::resource\_manager::ResourceManager; auto &gko\_mat = dynamic\_cast<ginkgo\_matrix<T,S>&>(mat); 89 auto exec = qko\_mat.data->get\_executor(); 90 91 auto non\_const\_exec = std::const\_pointer\_cast<qko::Executor>(exec); auto str\_name = std::string{name} + " (Ginkgo)"; 92 93 this->name = str\_name.c\_str(); 94 this->options\_file = solver\_opts\_file; auto logger = reinterpret\_cast<opencarp::FILE\_SPEC>(void\_logger); 95 103 RM resource\_manager; std::ifstream options\_json(this->options\_file); 104 105 rapidjson::IStreamWrapper options\_in(options\_json); rapidjson::Document f\_options; 106 107 f\_options.ParseStream(options\_in); 108 resource\_manager.insert\_data<gko::Executor>("executor", non\_const\_exec); 109 resource\_manager.read(f\_options); 110 115 116 auto solver\_factory = resource\_manager.search\_data<gko::LinOpFactory>("solver"); this->solver = solver\_factory->generate(gko\_mat.data); 117 118 dynamic\_cast<gko::Preconditionable \*>(this->solver.get())->set\_preconditioner(this->precond); this->set\_stopping\_criterion(normtype, tol, max\_it, verbose, logger); 119 120 }



### Conclusion

- Various new features available in 1.6.0
  - Distributed block preconditioners.
  - Profiling annotations.
  - Improved on on GPU AMG performance.
  - Scaling on Frontier.
- New upcoming features (1.7.0 and later)
  - Batched iterative solvers.
  - Distributed preconditioners (BDDC, AMG)
  - Easy configuration with JSON files
- Happy to discuss any additional feature requests.

### Acknowledgements



This project has received funding from the European High–Performance Computing Joint Undertaking EuroHPC (JU) under grant agreement No 955495.

The JU receives support from the European Union's Horizon 2020 research and innovation

programme and France, Italy, Germany, Austria, Norway, and Switzerland.

The project was co-funded by the French National Research Agency ANR, the German Federal Ministry of Education and Research, the Italian ministry of economic development, the Swiss State Secretariat for Education, Research and Innovation, the Austrian Research Promotion Agency FFG, and the Research Council of Norway.



https://github.com/ginkgo-project/ginkgo