

# Sustaining Simulation Performance in the US Exascale Computing Project – A tale from the trenches



Approved for public release

Hartwig Anzt, University of Tennessee



# Designing an ECP library for sustaining simulation performance

## MAGMA SPARSE

MAGMA-sparse as a “child” of MAGMA explores the development of sparse linear algebra for NVIDIA GPUs.

### *Limitations:*

- *C code with hand-written build system*
- *Sparse unit testing*
- *Focus on NVIDIA GPUs*
- *Design-specific limitations (flexibility/extensibility)*

# Designing an ECP library for sustaining simulation performance

## MAGMA SPARSE

MAGMA-sparse as a “child” of MAGMA explores the development of sparse linear algebra for NVIDIA GPUs.

### Design considerations for Ginkgo

- Platform Portability
- Performance
- Rapid integration of new algorithms
- xSDK / E4S Community Policies
- BSSw expertise / experience
- Modern C++
- CI/CD and unit testing
- Open source & permissive licensing



**Building Trusted Scientific Software**

**Software Verification**

**Think Locally, Act Globally: Outreach for Better Scientific Software**

I have worked in the scientific software field for more than 30 years. Verification is doing things right, and vigilance to it gives us memory in order to avoid confusion when the bit

**Facing internal and external concerns**

Verification focuses on internal concerns of a good software

**An ambitious goal**

The ECP needs to deliver a software environment and applications ready to run on modern computers, which are scheduled to be deployed starting in 2021. Achieving this goal entails a large, large-scale software development effort. Recognizing the challenges development teams will face, the ECP is supporting the IDEAS Productivity project to help scientific researchers improve their development practices.



# Designing an ECP library for sustaining simulation performance

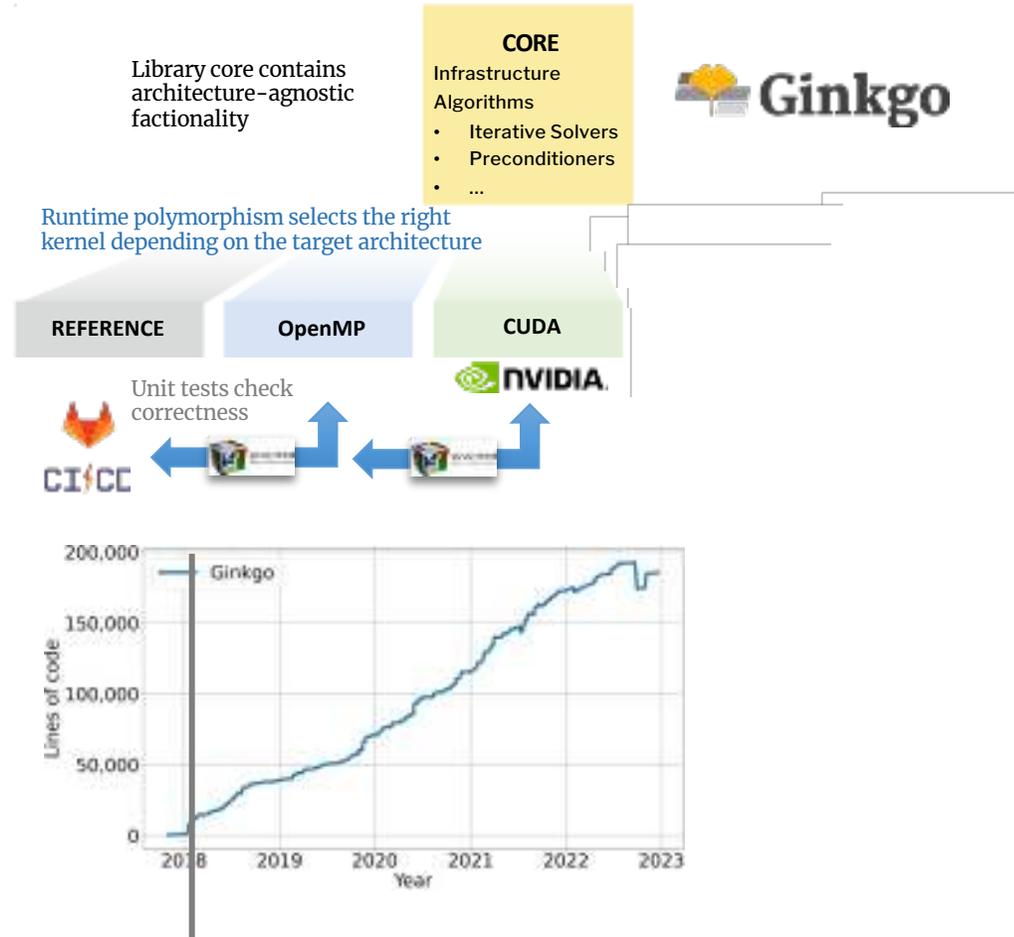
## MAGMA SPARSE

MAGMA-sparse as a “child” of MAGMA explores the development of sparse linear algebra for NVIDIA GPUs.

## Design considerations for Ginkgo

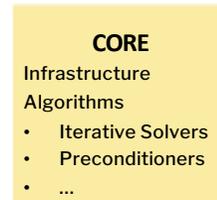
- Platform Portability
- Performance
- Rapid integration of new algorithms
- xSDK / E4S Community Policies
- BSSw expertise / experience
- Modern C++
- CI/CD and unit testing
- Open source & permissive licensing

*Before the first line of code is written, we spend a year on whiteboard discussions.*



# The LinOp abstraction

Library core contains architecture-agnostic functionality



Runtime polymorphism selects the right kernel depending on the target architecture



	Functionality	OMP	CUDA
Basic	SpMV	☑	☑
	SpMM	☑	☑
	SpGeMM	☑	☑
Krylov solvers	BiCG	☑	☑
	BiCGSTAB	☑	☑
	CG	☑	☑
	CGS	☑	☑
Preconditioners	GMRES	☑	☑
	IDR	☑	☑
	(Block-1)Jacobi	☑	☑
Preconditioners	ILU/IC	☑	☑
	Parallel ILU/IC	☑	☑
	Parallel ILUT/ICT	☑	☑
	Sparse Approximate Inverse	☑	☑

**Linear Operator Interface**

We express everything as Linear Operator

- Technically we leverage C++ class inheritance
- Applications can apply any functionality as a linear operator

Matrix Vector Product      Preconditioner (for matrix A)      Solver (for system Ax = b)

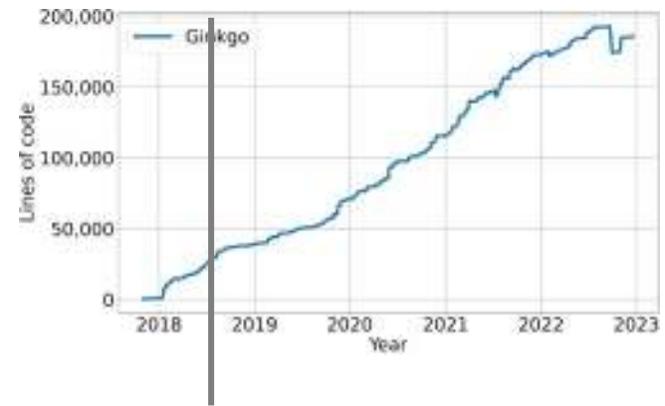
$x = A \cdot b$        $x = M^{-1} \cdot b$        $x = M \cdot b$

$M^{-1} = A^{-1}$        $S = A^{-1}$

$M^{-1} = \text{inv}(A)$        $S = \text{inv}(A)$

All of them can be expressed as

Application of a linear operator (LinOp)  $A \cdot x = b$



# Extending to AMD GPUs



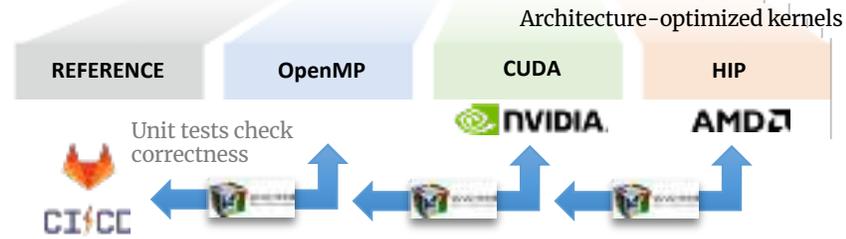
Library core contains architecture-agnostic functionality

**CORE**  
Infrastructure Algorithms

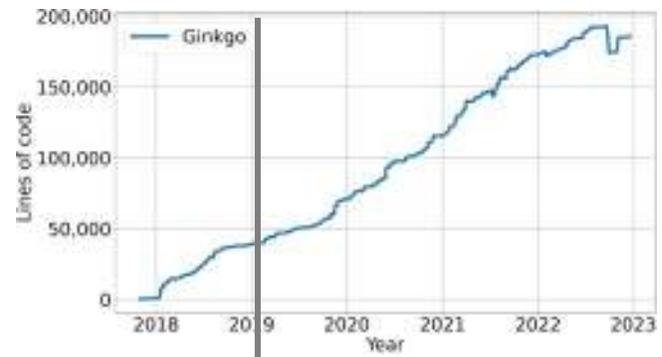
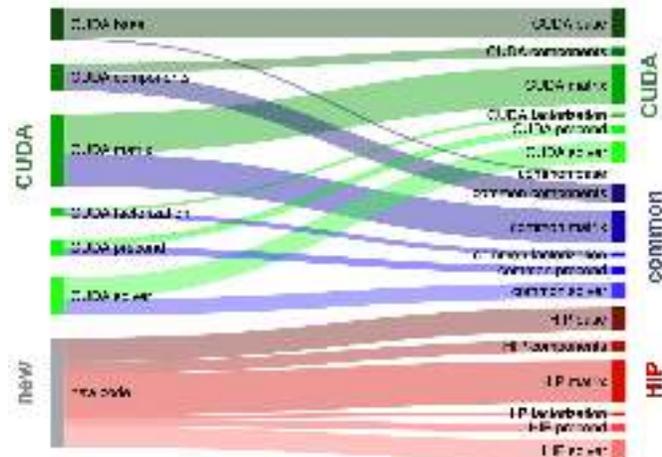
- Iterative Solvers
- Preconditioners
- ...



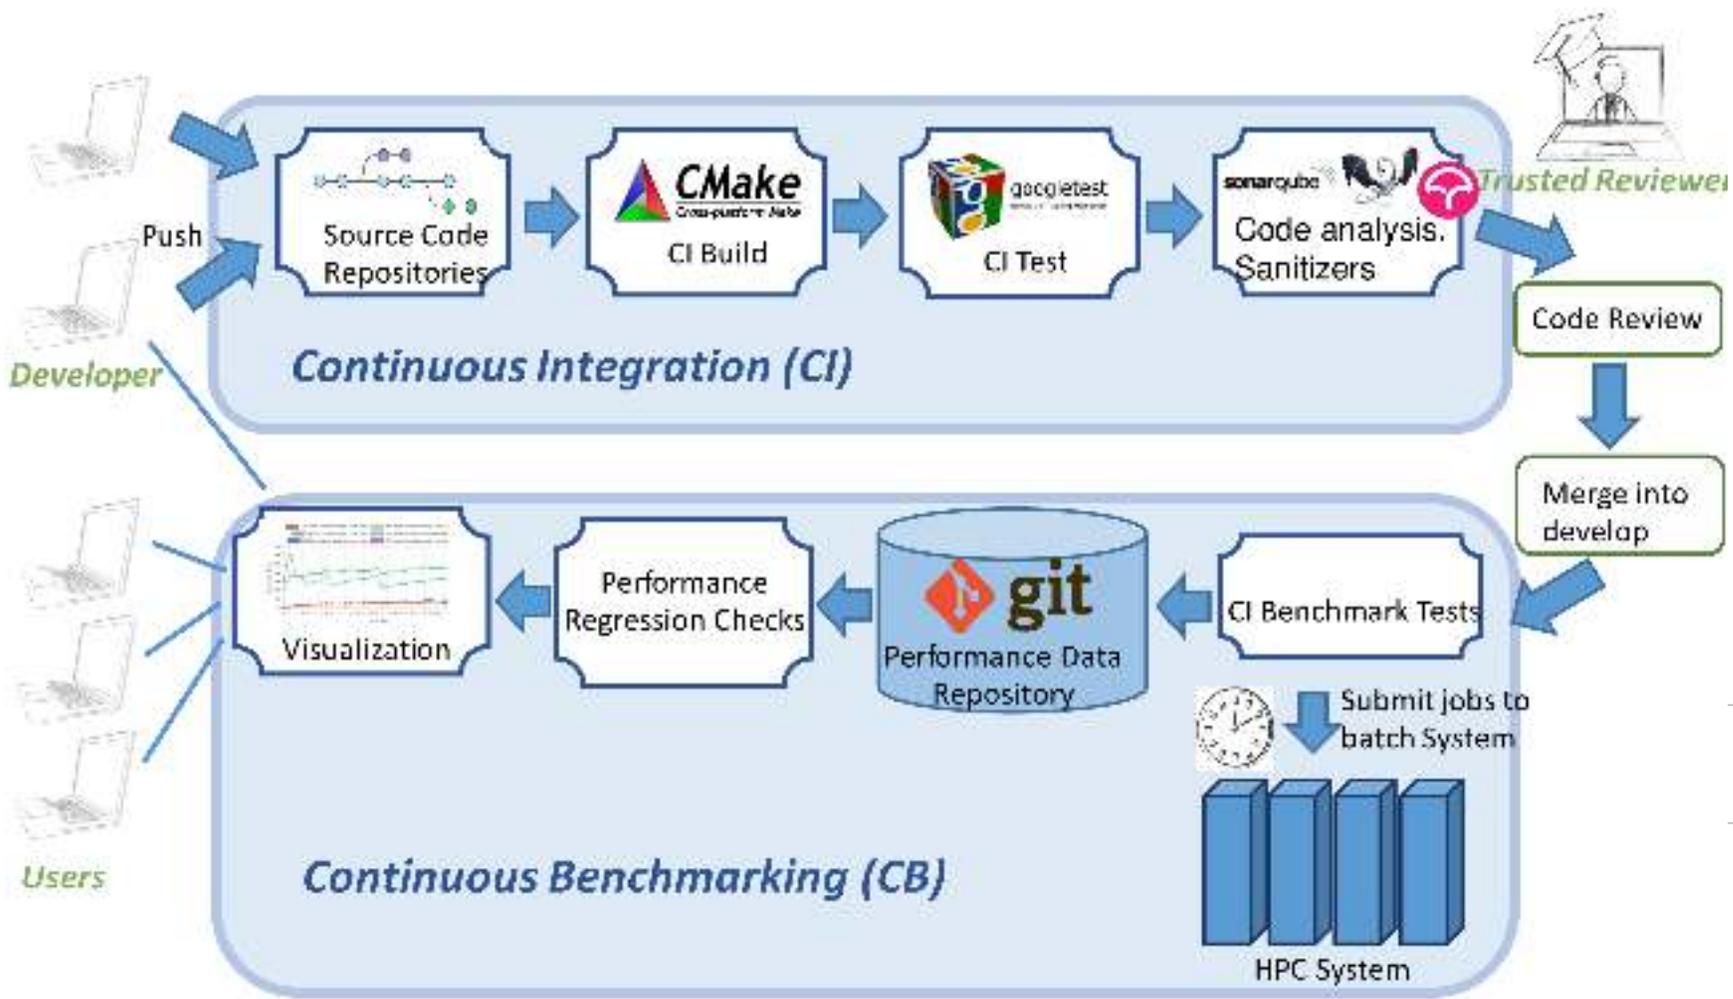
Runtime polymorphism selects the right kernel depending on the target architecture



Functionality	OMP	CUDA	HIP
<b>Basic</b>			
SpMV	✓	✓	✓
SpMM	✓	✓	✓
SpGeMM	✓	✓	✓
<b>Krylov solvers</b>			
BICGSTAB	✓	✓	✓
BICGSTAB-L	✓	✓	✓
CG	✓	✓	✓
CGS	✓	✓	✓
GMRES	✓	✓	✓
IDR	✓	✓	✓
(Block-Jacobi)	✓	✓	✓
<b>Preconditioners</b>			
ILU/IC	✓	✓	✓
Parallel ILU/IC	✓	✓	✓
Parallel ILUT/ICT	✓	✓	✓
Sparse Approximate Inverse	✓	✓	✓



# Sustainable software development & CI/CD

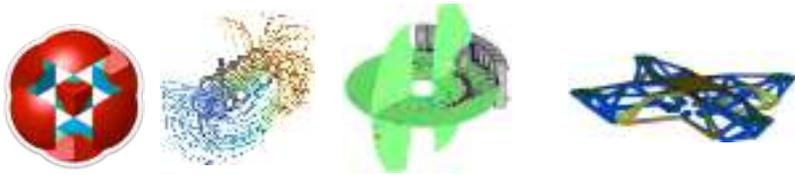




# Sustainable software development & CI/CD

	Build	Test	Deploy
NVIDIA	✓ build/cuda90/openmpi/gnu5/llvm39	✓ test/cuda90/openmpi/gnu5/llvm39	✓ deploy/cuda90/openmpi/gnu5/llvm39
	✓ build/cuda91/mvapich2/gnu6/llvm40	✓ test/cuda91/mvapich2/gnu6/llvm40	✓ deploy/cuda91/mvapich2/gnu6/llvm40
	✓ build/cuda92/mvapich2/gnu7/llvm50/intel2017	✓ test/cuda92/mvapich2/gnu7/llvm50/intel2017	✓ deploy/cuda92/mvapich2/gnu7/llvm50/intel2017
	✓ build/cuda100/mvapich2/gnu7/llvm60/intel2018	✓ test/cuda100/mvapich2/gnu7/llvm60/intel2018	✓ deploy/cuda100/mvapich2/gnu7/llvm60/intel2...
	✓ build/cuda101/openmpi/gnu8/llvm7/intel2019	✓ test/cuda101/openmpi/gnu8/llvm7/intel2019	✓ deploy/cuda101/openmpi/gnu8/llvm7/intel2019
	✓ build/cuda101/openmpi/gnu8/llvm11/intel2019	✓ test/cuda101/openmpi/gnu8/llvm11/intel2019	✓ deploy/cuda101/openmpi/gnu8/llvm11/intel2019
	✓ build/cuda102/ncmpi/gnu6/llvm8/intel2019	✓ test/cuda102/ncmpi/gnu6/llvm8/intel2019	✓ deploy/cuda102/ncmpi/gnu6/llvm8/intel2019
	✓ build/cuda110/mvapich2/gnu9/llvm9/intel2020	✓ test/cuda110/mvapich2/gnu9/llvm9/intel2020	✓ deploy/cuda110/mvapich2/gnu9/llvm9/intel2020
	✓ build/cuda114/openmpi/gnu11/llvm12	✓ test/cuda114/openmpi/gnu11/llvm12	✓ deploy/cuda114/openmpi/gnu11/llvm12
CPU only	✓ build/nocuda/mvapich2/gnu5/llvm39/intel2018	✓ test/nocuda/mvapich2/gnu5/llvm39/intel2018	✓ deploy/nocuda/mvapich2/gnu5/llvm39/intel2018
	✓ build/nocuda/openmpi/gnu9/llvm8	✓ test/nocuda/openmpi/gnu9/llvm8	✓ deploy/nocuda/openmpi/gnu9/llvm8
Intel	✓ build/oneapi/openmpi	✓ test/oneapi/openmpi	✓ deploy/oneapi/openmpi
AMD	✓ build/rocm40/openmpi/gnu5/llvm50	✓ test/rocm40/openmpi/gnu5/llvm50	✓ deploy/rocm40/openmpi/gnu5/llvm50
	✓ build/rocm45/mvapich2/gnu8/llvm8	✓ test/rocm45/mvapich2/gnu8/llvm8	✓ deploy/rocm45/mvapich2/gnu8/llvm8
	✓ build/rocm502/openmpi/gnu11/llvm11	✓ test/rocm502/openmpi/gnu11/llvm11	✓ deploy/rocm502/openmpi/gnu11/llvm11

# Input from the "first customer"



MFEM is a *free, lightweight, scalable* C++ library for finite element methods.

### Speeding up MFEM's "example 22" on GPUs

Example 22 of the MFEM finite element library solves harmonic oscillation problems, with a forced oscillation imposed at the boundary. In this test, we use variant 1:

$$-\nabla \cdot (a \nabla u) - \omega^2 b u + i \omega c u = 0$$

with  $a = 1, b = 1, \omega = 10, c = 20$

Speedup of Ginkgo's Compressed Basis-GMRES solver vs MFEM's GMRES solver for three different orders of basis functions (p), using MFEM matrix-free operators and the Ginkgo MFEM integration wrappers in MFEM. CUDA 10.1/V100 and ROCm 4.0/MI50.

Real part of solution (top), imaginary part of solution

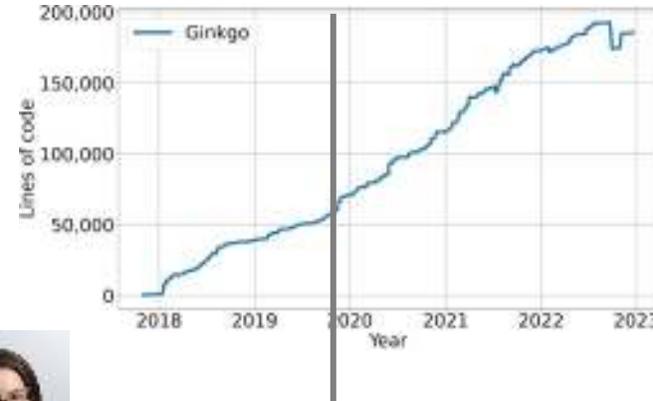
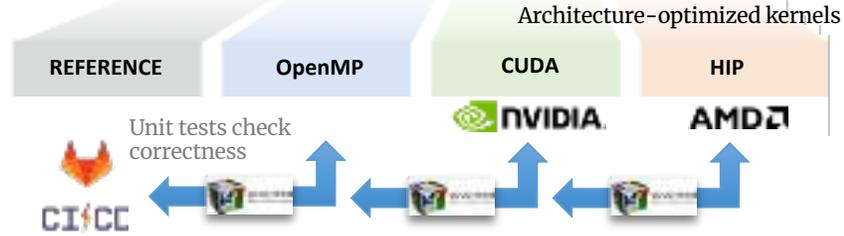
Library core contains architecture-agnostic functionality

**CORE**  
Infrastructure Algorithms

- Iterative Solvers
- Preconditioners
- ...



Runtime polymorphism selects the right kernel depending on the target architecture



	Functionality	OMP	CUDA	HIP
Basic	SpMV	☑	☑	☑
	SpMM	☑	☑	☑
	SpGeMM	☑	☑	☑
Krylov solvers	BiCG	☑	☑	☑
	BiCGSTAB	☑	☑	☑
	CG	☑	☑	☑
	CGS	☑	☑	☑
	GMRES	☑	☑	☑
Preconditioners	IDR	☑	☑	☑
	(Block-)Jacobi	☑	☑	☑
	ILU/IC	☑	☑	☑
	Parallel ILU/IC	☑	☑	☑
	Parallel ILUT/ICT	☑	☑	☑
	Sparse Approximate Inverse	☑	☑	☑

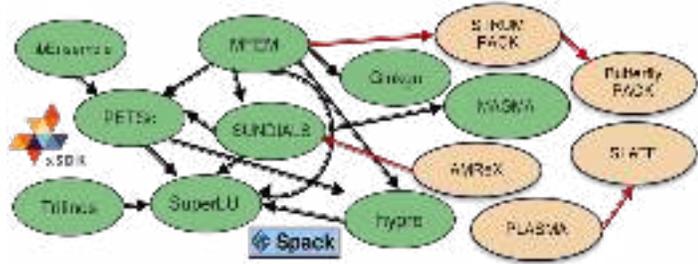
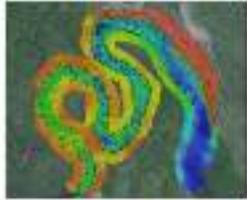


	OMP	CUDA	HIP
On-Device Matrix Assembly	☑	☑	☑
MC64/RCM-reordering	☑	☑	☑
Wrapping user data	☑	☑	☑



# Part of the xSDK effort

## xSDK: Extreme-scale Scientific Software Development Kit



xsdk-examples v.0.3.0

The xSDK provides infrastructure for and interoperability of a **collection of related and complementary software elements**—developed by diverse, independent teams throughout the high-performance computing (HPC) community—that provide the building blocks, tools, models, processes, and related artifacts for rapid and efficient development of high-quality applications.

### November 2022

- 26 math libraries
- 2 domain components
- 16 mandatory xSDK community policies
- Spack xSDK installer

### xSDK community policies:

- 16 mandatory policies,
- 8 recommended policies,
- 4 Spack variant guidelines
- Available on Github

<https://xsdk.info/policies/>

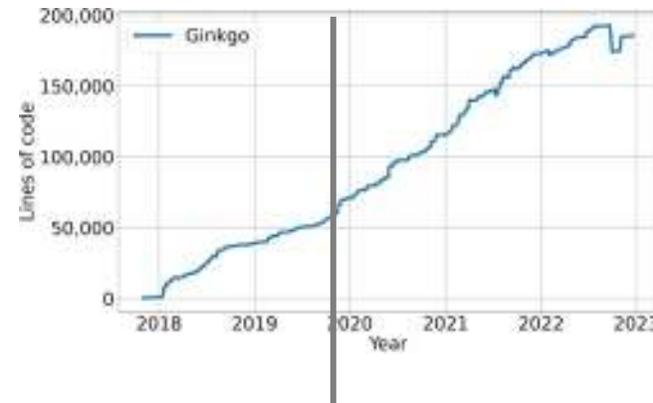
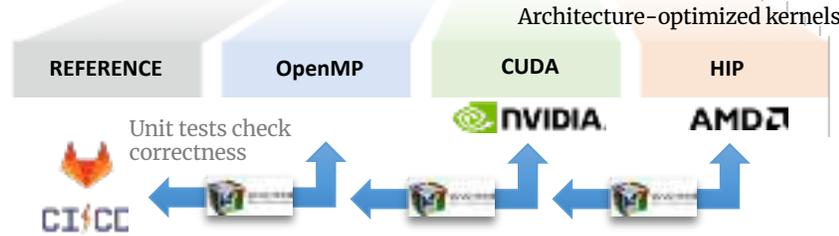
Library core contains architecture-agnostic functionality

**CORE**  
Infrastructure Algorithms

- Iterative Solvers
- Preconditioners
- ...



Runtime polymorphism selects the right kernel depending on the target architecture



	Functionality	OMP	CUDA	HIP
Basic	SpMV	☑	☑	☑
	SpMM	☑	☑	☑
	SpGeMM	☑	☑	☑
Krylov solvers	BiCG	☑	☑	☑
	BiCGSTAB	☑	☑	☑
	CG	☑	☑	☑
	CGS	☑	☑	☑
Preconditioners	GMRES	☑	☑	☑
	IDR	☑	☑	☑
	(Block-)Jacobi	☑	☑	☑
Utilities	ILU/IC	☑	☑	☑
	Parallel ILU/IC	☑	☑	☑
	Parallel ILUT/ICT	☑	☑	☑
	Sparse Approximate Inverse	☑	☑	☑

	Functionality	OMP	CUDA	HIP
Utilities	On-Device Matrix Assembly	☑	☑	☑
	MC54/RCM-reordering	☑	☑	☑
	Wrapping user data	☑	☑	☑





# Extending to Intel GPUs



Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

- Home
- Technologies
- Sectors
- COVID-19
- AI/ML/DL



March 23, 2021

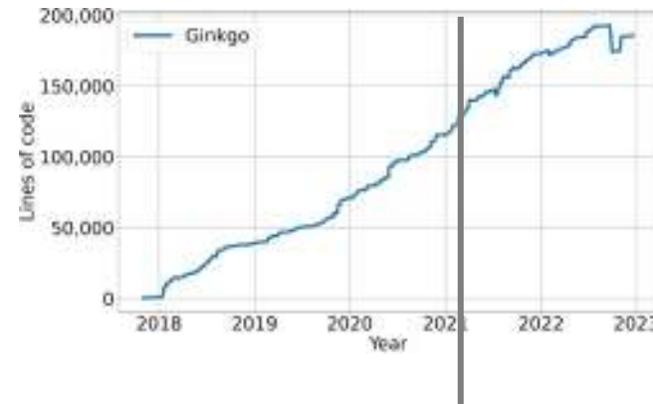
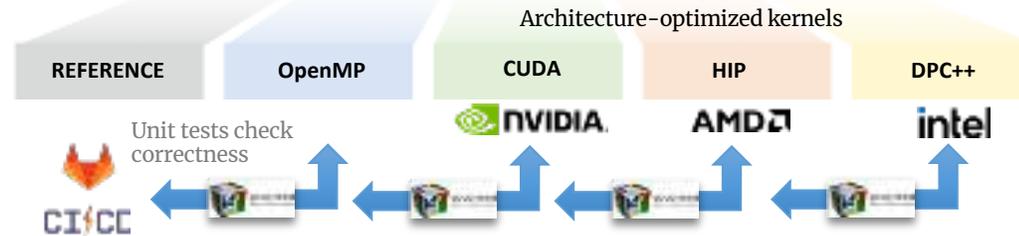


Library core contains architecture-agnostic functionality

- CORE**
- Infrastructure
  - Algorithms
    - Iterative Solvers
    - Preconditioners
    - ...



Runtime polymorphism selects the right kernel depending on the target architecture



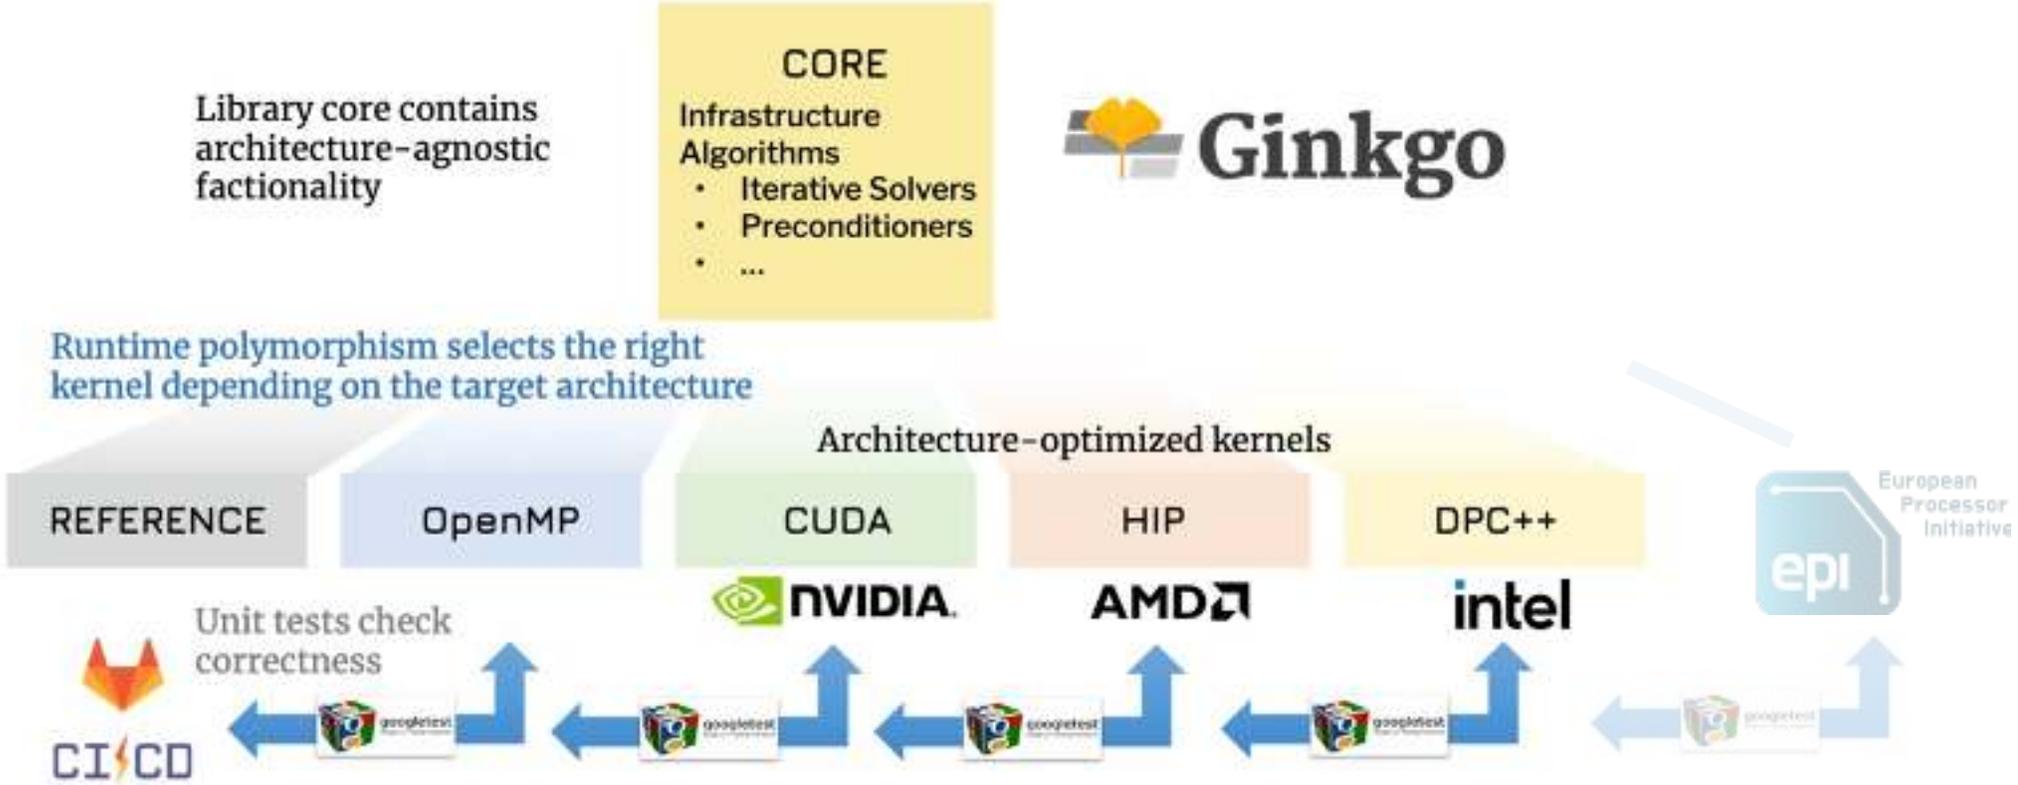
	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	✓	✓	✓	✓
	SpMM	✓	✓	✓	✓
	SpBeMM	✓	✓	✓	✓
Krylov solvers	BICG	✓	✓	✓	✓
	BICGSTAB	✓	✓	✓	✓
	CG	✓	✓	✓	✓
	CGS	✓	✓	✓	✓
	GMRES	✓	✓	✓	✓
Preconditioners	IDR	✓	✓	✓	✓
	(Block-1)Jacobi	✓	✓	✓	✓
	ILU/IC	✓	✓	✓	✓
	Parallel ILU/IC	✓	✓	✓	✓
	Parallel ILU/ICT	✓	✓	✓	✓
	Sparse Approximate Inverse	✓	✓	✓	✓

Utilities	On-Device Matrix Assembly	✓	✓	✓	✓
	MC64/RCM-reordering	✓			
	Wrapping user data		✓		
	Logging		✓		
	PAPI counters		✓		





# Portability as central design principle



# Focus efforts as lightweight tool in ECP to address challenges



Focus efforts

- Mixed precision
- batched



- Address recent hardware trends (tensor cores, etc.)
- Address hardware requirements

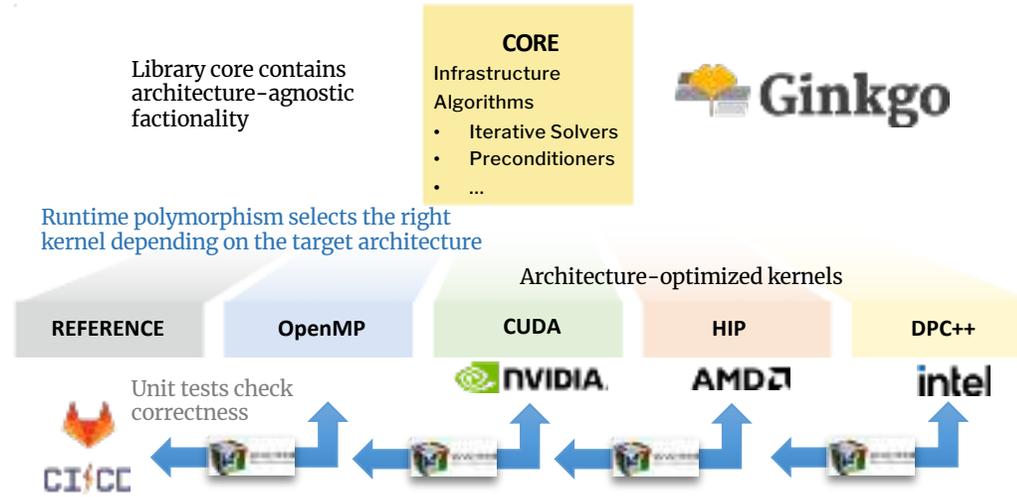
Library core contains architecture-agnostic functionality

**CORE**  
Infrastructure Algorithms

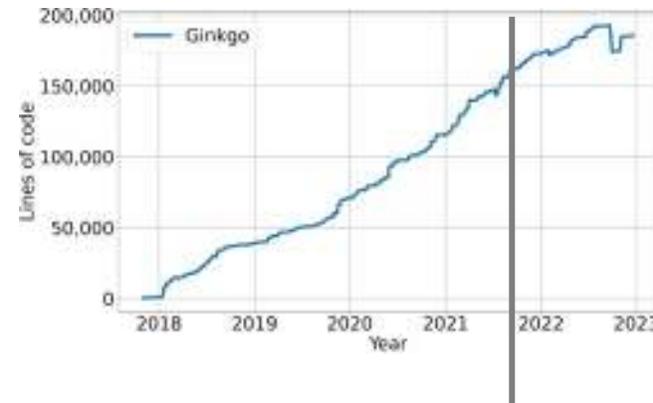
- Iterative Solvers
- Preconditioners
- ...



Runtime polymorphism selects the right kernel depending on the target architecture



Functionality	OMP	CUDA	HIP	DPC++
<b>Basic</b>				
SpMV	☑	☑	☑	☑
SpMM	☑	☑	☑	☑
SpGeMM	☑	☑	☑	☑
<b>Krylov solvers</b>				
BCG	☑	☑	☑	☑
BCGSTAB	☑	☑	☑	☑
CG	☑	☑	☑	☑
CGS	☑	☑	☑	☑
GMRES	☑	☑	☑	☑
IDR	☑	☑	☑	☑
(Block-)Jacobi	☑	☑	☑	☑
<b>Preconditioners</b>				
ILU/IC	☑	☑	☑	☑
Parallel ILU/IC	☑	☑	☑	☑
Parallel ILUT/ICT	☑	☑	☑	☑
Sparse Approximate Inverse	☑	☑	☑	☑



Industry Collaboration with bi-weekly meetings

<b>Utilities</b>				
On-Device Matrix Assembly	☑	☑	☑	☑
MC64/RCM-reordering	☑			
Wrapping user data			<input type="checkbox"/>	
Logging			<input type="checkbox"/>	
PAPI counters			<input type="checkbox"/>	



# Mixed precision focus effort



Focus efforts

- Mixed precision
- batched



- Address recent hardware trends (tensor cores, etc.)
- Address hardware requirements

## Advances in Mixed Precision Algorithms: 2021 Edition

by the ECP Multiprecision Effort Team (Lead: Hartwig Anzt)

Ahmad Abdelfattah, Hartwig Anzt, Alan Ayala, Erik G. Boman, Erin Carson, Sebastien Cayrols, Terry Cojean, Jack Dongarra, Rob Falgout, Mark Gates, Thomas Grützmacher, Nicholas J. Higham, Scott B. Kruger, Sherry Li, Neil Lindquist, Yang Liu, Jennifer Luo, Piotr Luszczek, Pratik Nayak, Daniel Osei-Kuffuor, Sri Pranesh, Sivasankaran Rajamanikam, Tobias Ribizel, Barry Smith, Kasia Swirydowicz, Stephen Thomas, Stanimire Tomov, Yaohung M. Tsai, Ichi Yamazaki, Urike Meier Yang

Available access Research article First published online March 18, 2021

A survey of numerical linear algebra methods utilizing mixed-precision arithmetic

Ahmad Abdelfattah, Hartwig Anzt, Pratik Nayak, and Urike Meier Yang View all authors and affiliations

Volume 35, Issue 1 | https://doi.org/10.1177/10984140211009331

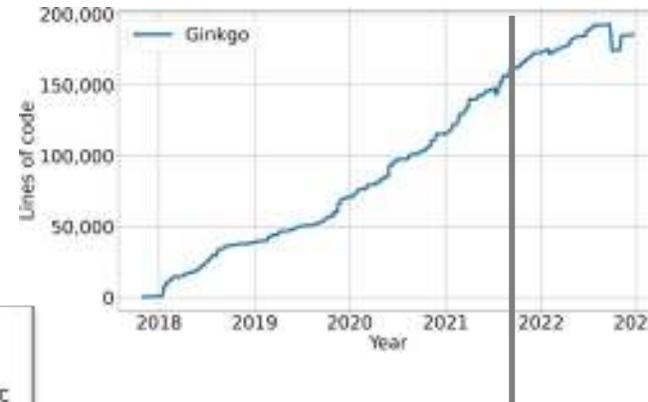
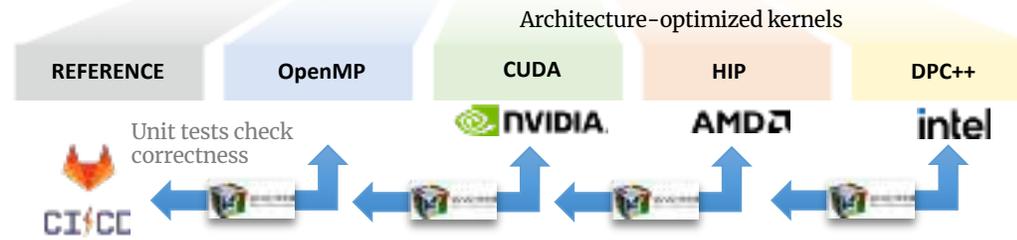
Library core contains architecture-agnostic factuality

**CORE**  
Infrastructure Algorithms

- Iterative Solvers
- Preconditioners
- ...



Runtime polymorphism selects the right kernel depending on the target architecture



Industry Collaboration with bi-weekly meetings

	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	✓	✓	✓	✓
	SpMM	✓	✓	✓	✓
	SpGeMM	✓	✓	✓	✓
Krylov solvers	BiCG	✓	✓	✓	✓
	BiCGSTAB	✓	✓	✓	✓
	CG	✓	✓	✓	✓
	CGS	✓	✓	✓	✓
	GMRES	✓	✓	✓	✓
Preconditioners	IDR	✓	✓	✓	✓
	(Block-)Jacobi	✓	✓	✓	✓
	ILU/IC	✓	✓	✓	✓
	Parallel ILU/IC	✓	✓	✓	✓
	Parallel ILU/ICT	✓	✓	✓	✓
	Sparse Approximate Inverse	✓	✓	✓	✓

Utilities	On-Device Matrix Assembly	MC64/RCM reordering	Wrapping user data	Logging	PAPI counters
	✓	✓	✓	✓	✓



# Mixed precision AMG on GPUs

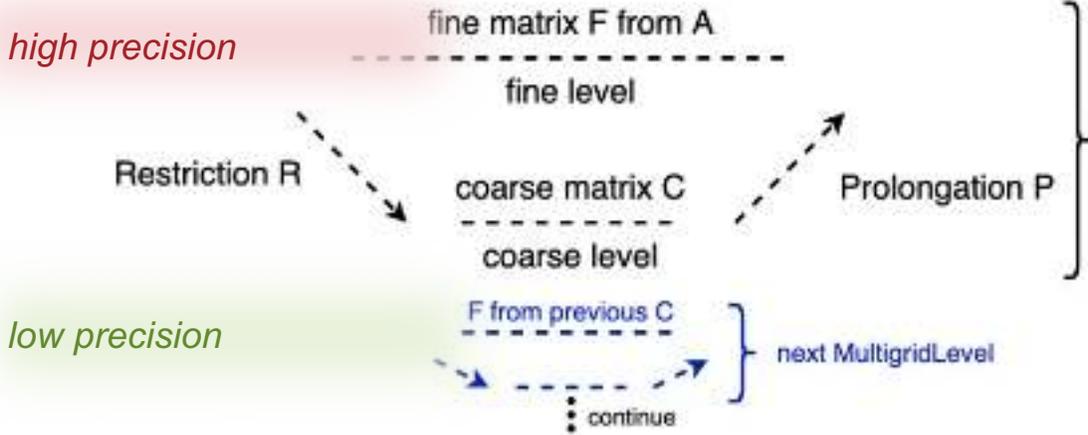
- Preconditioning iterative solvers

- Idea: Approximate inverse of system matrix to make the system “easier to solve”:  $P^{-1} \approx A^{-1}$   
and solve  $Ax = b \Leftrightarrow P^{-1}Ax = P^{-1}b \Leftrightarrow \tilde{A}x = \tilde{b}$

- Mixed Precision Multigrid Preconditioner

```

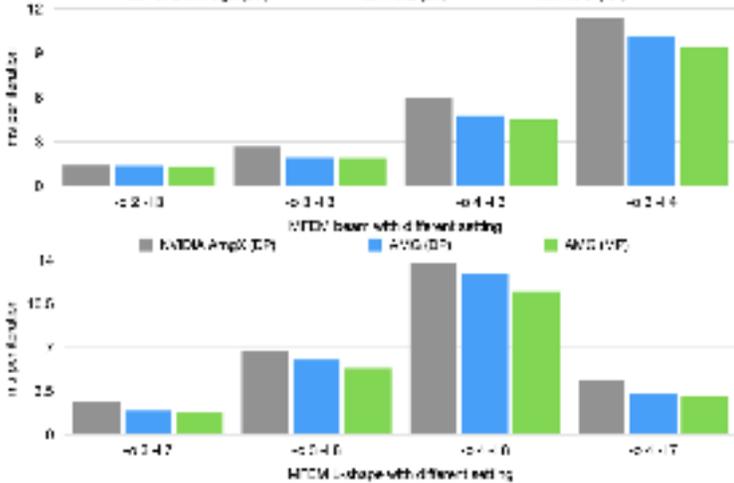
1 multigrid::build()
2   .with_max_levels(10u) // equal to NVIDIA/AMGX 11 max levels
3   .with_min_coarse_row(64u)
4   .with_pre_smoother(sm, sm_f)
5   .with_mg_level(pgm, pgm_f)
6   .with_level_selector(
7     [](const size_type level, const LinOp*) -> size_type {
8       // Only the first level is generated by MultigridLevel(double).
9       // The subsequent levels are generated by MultigridLevel(float)
10      return level >= 1 ? 1 : 0;
11    })
12   .with_coarest_solver(coarest_solver_f)
  
```



MultigridLevel



Mike Tsai



# Mixed precision AMG on GPUs



- Focus efforts
- Mixed precision
  - batched



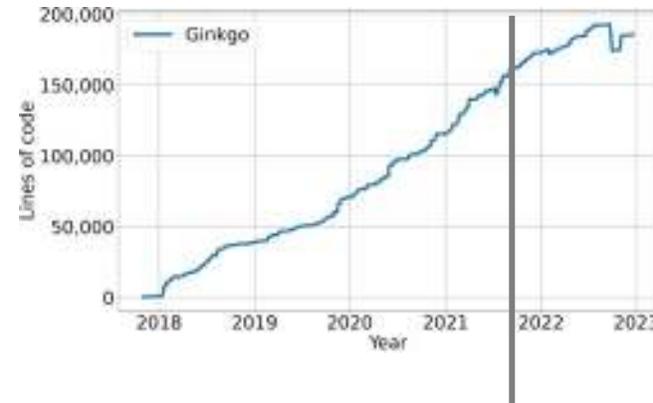
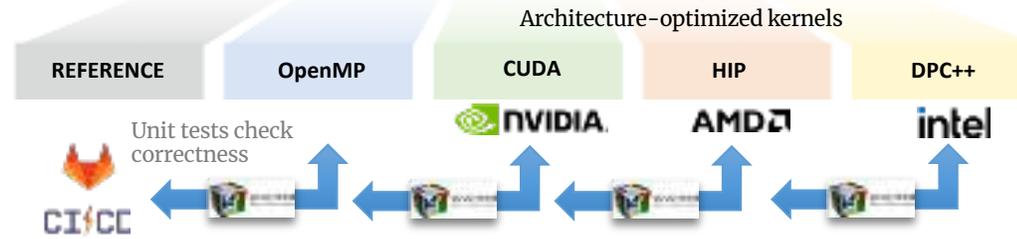
Library core contains architecture-agnostic functionality

**CORE**  
Infrastructure Algorithms

- Iterative Solvers
- Preconditioners
- ...



Runtime polymorphism selects the right kernel depending on the target architecture



Industry Collaboration with bi-weekly meetings

	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	☑	☑	☑	☑
	SpMM	☑	☑	☑	☑
	SpGeMM	☑	☑	☑	☑
Krylov solvers	BiCG	☑	☑	☑	☑
	BiCGSTAB	☑	☑	☑	☑
	CG	☑	☑	☑	☑
	CGS	☑	☑	☑	☑
	GMRES	☑	☑	☑	☑
Preconditioners	IDR	☑	☑	☑	☑
	(Block-1)Jacobi	☑	☑	☑	☑
	ILU/IC	☑	☑	☑	☑
	Parallel ILU/IC	☑	☑	☑	☑
	Parallel ILU/ICT	☑	☑	☑	☑
	Sparse Approximate Inverse	☑	☑	☑	☑

	OMP	CUDA	HIP	DPC++
AMG preconditioner	☑	☑	☑	☑
AMS solver	☑	☑	☑	☑
Parallel Graph Match	☑	☑	☑	☑

	OMP	CUDA	HIP	DPC++
On-Device Matrix Assembly	☑	☑	☑	☑
MC64/RCM-reordering	☑			
Wrapping user data		☑		
Logging		☑		
PAPI counters		☑		

ICL UTK @ICL\_UTK · Sep 13  
 Congratulations to Yu-Hsiang Mike Tsai from @KITKarlsruhe, in collaboration with ICL's Natalie Beams and @HartwigArenz! Their paper "Mixed Precision Algebraic Multigrid on GPUs" took home a best paper award at PPAM2022.  
[ppam.edu.pl](http://ppam.edu.pl)

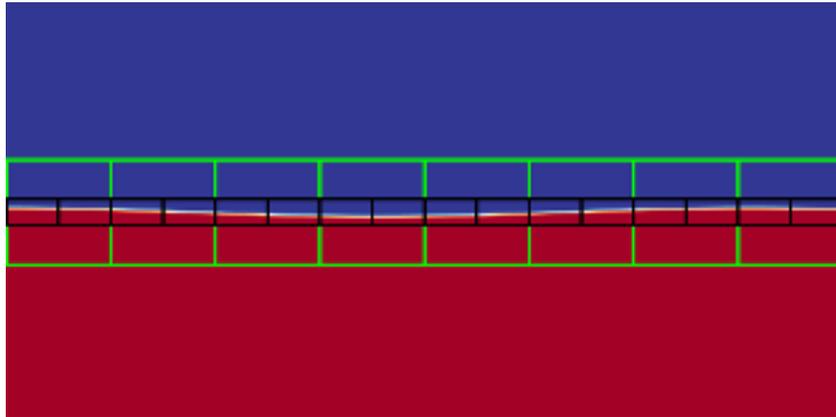


# Batched focus effort – Combustion Simulations

## Batched iterative solvers for SUNDIALS / PeleLM

PeleLM is a parallel, adaptive mesh refinement (AMR) code that solves the reacting Navier-Stokes equations in the low Mach number regime. The core libraries for managing the subcycling AMR grids and communication are found in the [AMReX source code](https://amrex-combustion.github.io/AMReX).

<https://amrex-combustion.github.io/PeleLM/overview.html>



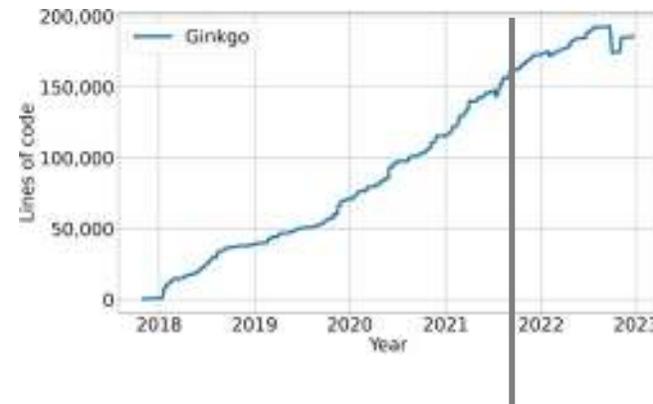
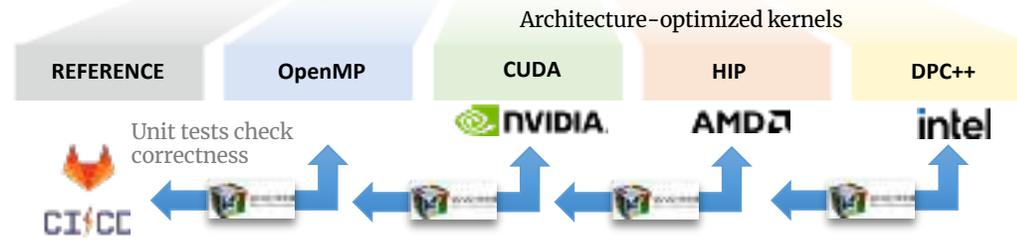
Problem	Size	Non-zeroes (A)	Non-zeroes (L+U)
draco2016_ju	54	2,332 (90%)	2,754 (94%)
draco19	22	438 (90%)	442 (91%)
gr12	33	978 (90%)	1,018 (93%)
gr30	54	2,560 (88%)	2,660 (98%)
bonnacore	144	6,135 (90%)	20,307 (98%)
Takaya	10	91 (91%)	91 (91%)

Library core contains architecture-agnostic factuality

- CORE**  
Infrastructure Algorithms
- Iterative Solvers
  - Preconditioners
  - ...



Runtime polymorphism selects the right kernel depending on the target architecture



	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	✓	✓	✓	✓
	SpMM	✓	✓	✓	✓
	SpGeMM	✓	✓	✓	✓
Krylov solvers	BICG	✓	✓	✓	✓
	BICGSTAB	✓	✓	✓	✓
	CG	✓	✓	✓	✓
	CGS	✓	✓	✓	✓
	GMRES	✓	✓	✓	✓
Preconditioners	IDR	✓	✓	✓	✓
	(Block-1)Jacobi	✓	✓	✓	✓
	ILU/IC	✓	✓	✓	✓
	Parallel ILU/IC	✓	✓	✓	✓
	Parallel ILU/ICT	✓	✓	✓	✓
	Sparse Approximate Inverse	✓	✓	✓	✓

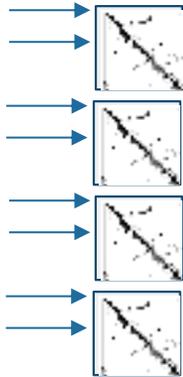
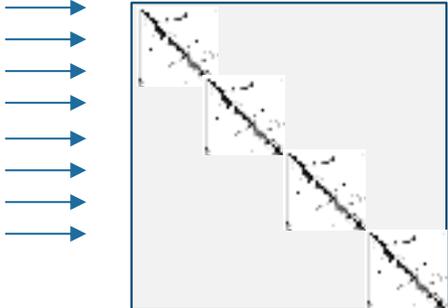
	OMP	CUDA	HIP	DPC++
AMG preconditioner	✓	✓	✓	✓
AMG solver	✓	✓	✓	✓
Parallel Graph Match	✓	✓	✓	✓

	OMP	CUDA	HIP	DPC++
On-Device Matrix Assembly	✓	✓	✓	✓
MC64/RCM reordering	✓			
Wrapping user data		✓		
Logging		✓		
PAPI counters		✓		



# Batched focus effort – Combustion Simulations

- Many sparse problems of medium size have to be solved concurrently.
  - ~ 50 – 2,000 unknowns, < 50% dense;
  - All sparse systems may share the same sparsity pattern;
  - An approximate solution may be acceptable (e.g., inside a non-linear solver);
- One solution is to arrange the individual systems on the main diagonal of one large system.
  - Convergence determined by the “hardest” problem;
  - No reuse of sparsity pattern information;
  - Global synchronization points;
- Better approach: design batched iterative solve functionality that solves all problems concurrently.
  - Problem-dependent convergence accounted for;
  - No global synchronization;
  - Reuse of sparsity pattern information;

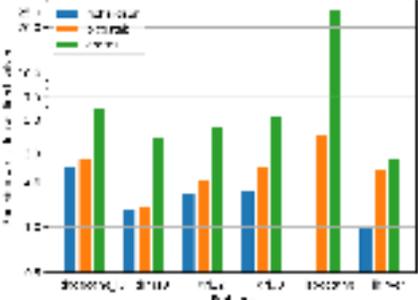




**Batched Sparse Iterative Solvers for Computational Chemistry Simulations on GPUs**

Publisher: IEEE [Cite This](#) [PDF](#)

Isha Aggarwal; Achiya Kozli; Pratik Nayak; Cody J. Balco; Carol S. Woodward; Harshil Anzi; All Authors

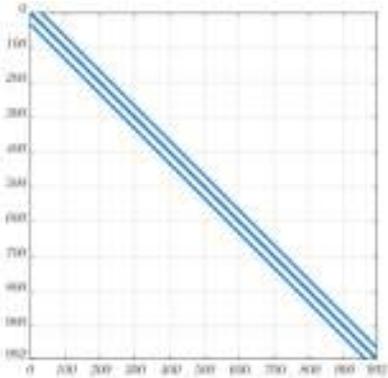
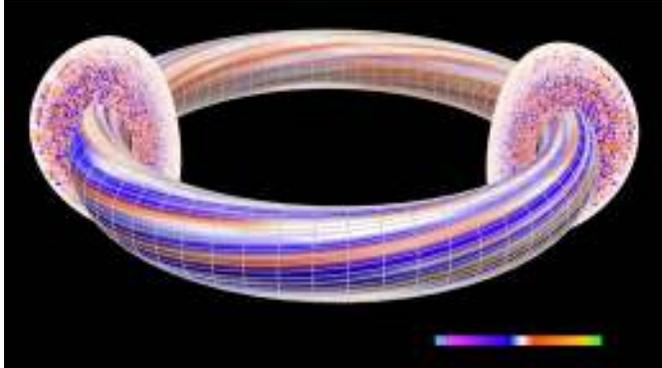
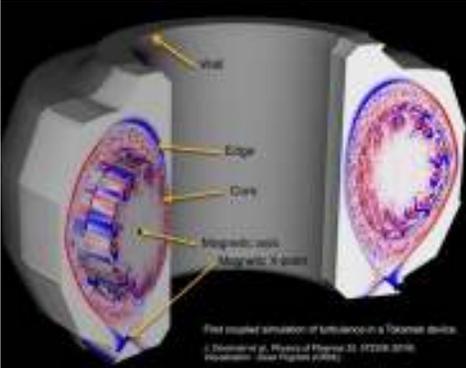


Problem Size	Individual	Block	Batch
1000x1000	~0.45	~0.55	~0.75
2000x2000	~0.35	~0.45	~0.65
4000x4000	~0.30	~0.40	~0.60
8000x8000	~0.25	~0.35	~0.55
16000x16000	~0.20	~0.30	~0.50
32000x32000	~0.15	~0.25	~0.45

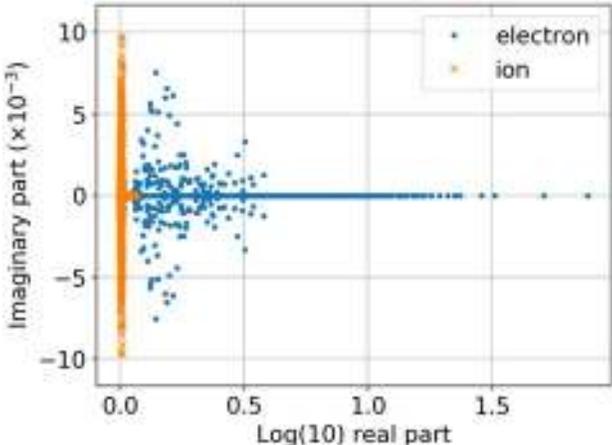
# Batched focus effort – Fusion Plasma Simulations

*XGC is a gyrokinetic particle-in-cell code, which specializes in the simulation of the edge region of magnetically confined thermonuclear fusion plasma. The simulation domain can include the magnetic separatrix, magnetic axis and the biased material wall. XGC can run in total-delta-f, and conventional delta-f mode. The ion species are always gyrokinetic except for ETG simulation. Electrons can be adiabatic, massless fluid, driftkinetic, or gyrokinetic.*

Source: [https://xgc.pppl.gov/html/general\\_info.html](https://xgc.pppl.gov/html/general_info.html)

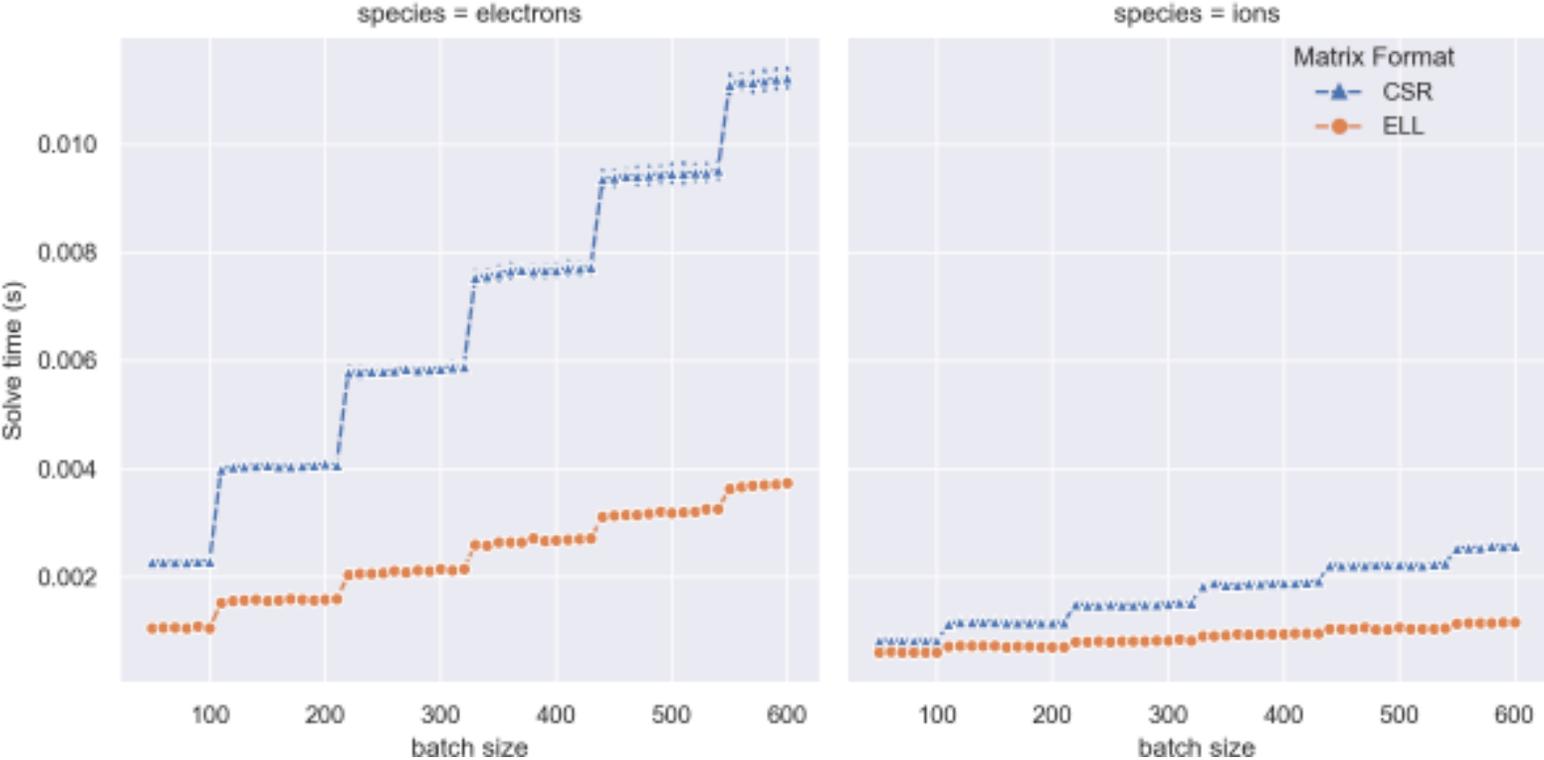


- Two species
- Ions easy to solve
- Electrons hard to solve
- Banded matrix structure
- Non-symmetric, need BiCGSTAB
- $n = \sim 1,000$
- $nz = \sim 9,000$



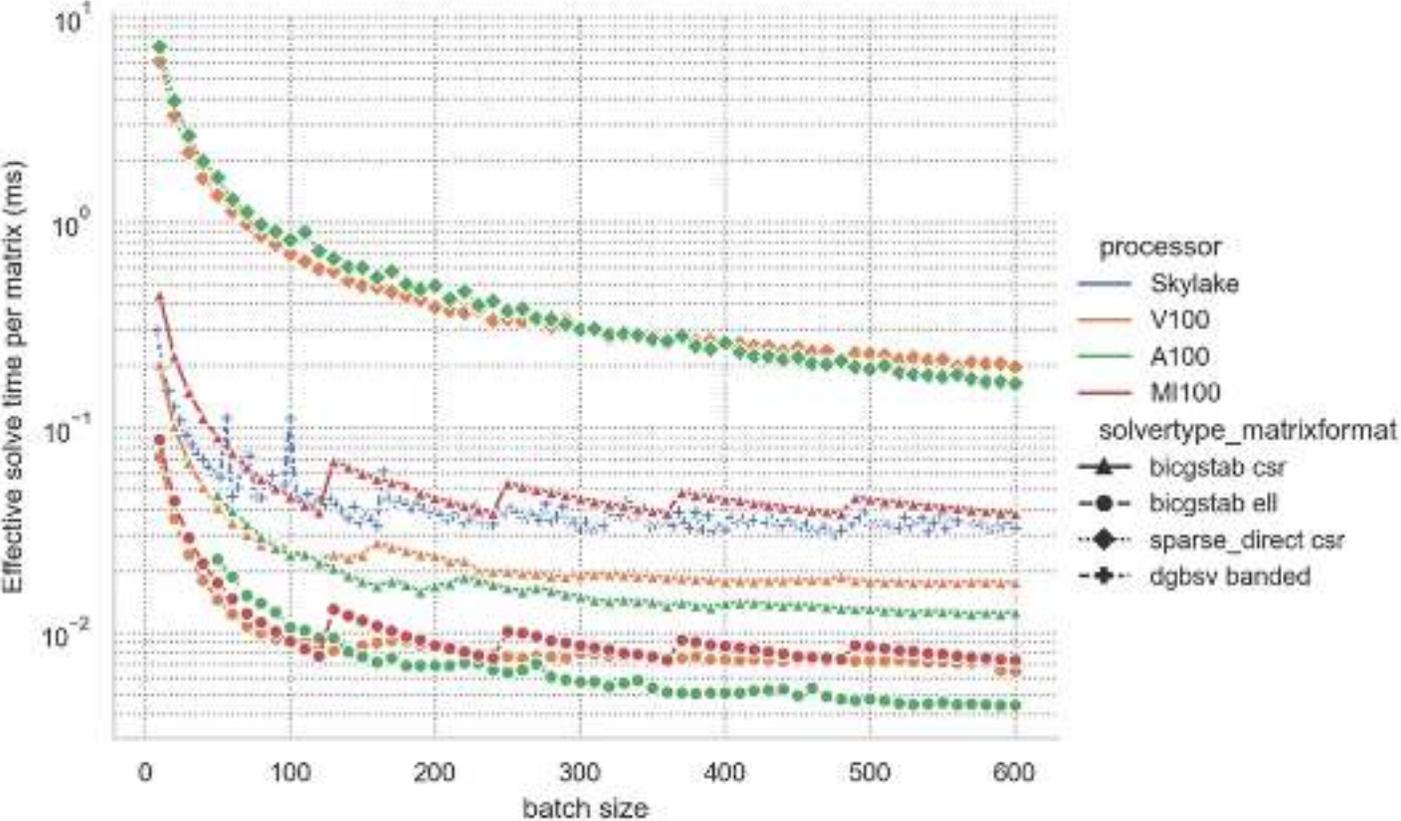
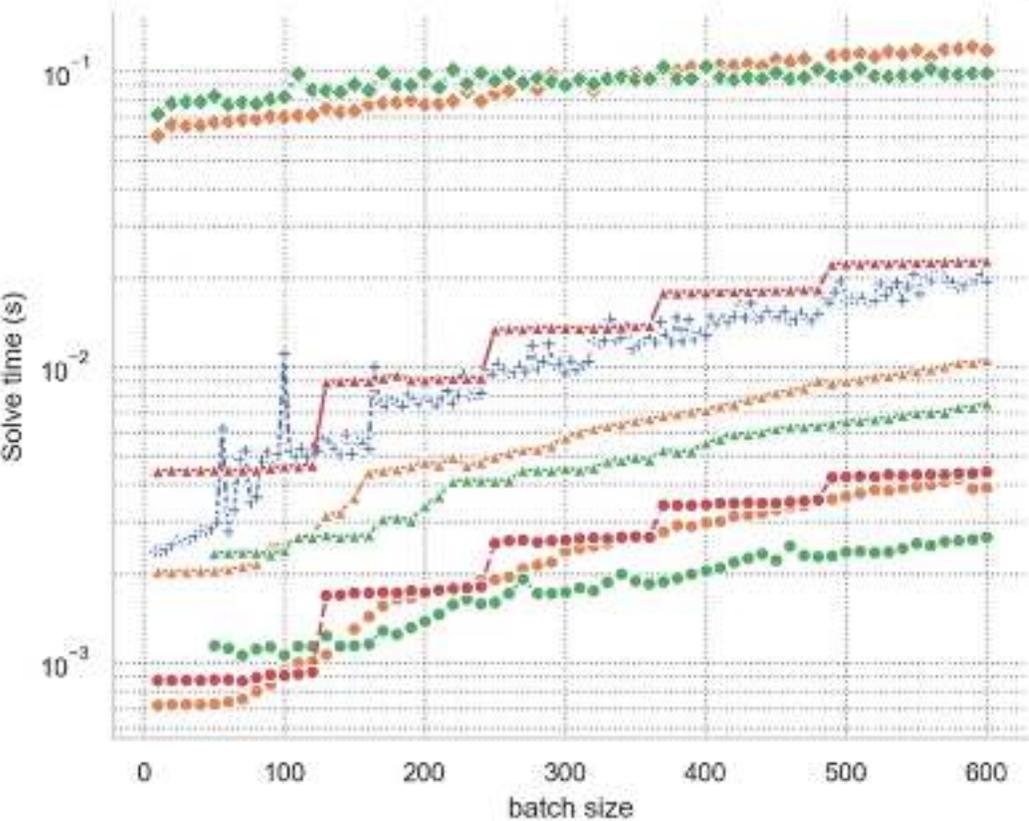
# Batched focus effort – Fusion Plasma Simulations

NVIDIA A100 GPU



- Ions easy to solve
- Electrons hard to solve
- ELL format more suitable
- “step pattern” for runtime filling up multiprocessors

# Batched focus effort – Fusion Plasma Simulations



Aditya Kashi, Pratik Nayak, Dhruva Kulkarni, Aaron Scheinberg, Paul Lin, and Hartwig Anzt. **Batched sparse iterative solvers on gpu for the collision operator for fusion plasma simulations.** In *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 157–167. IEEE, 2022.

# Batched focus effort – Fusion Plasma Simulations

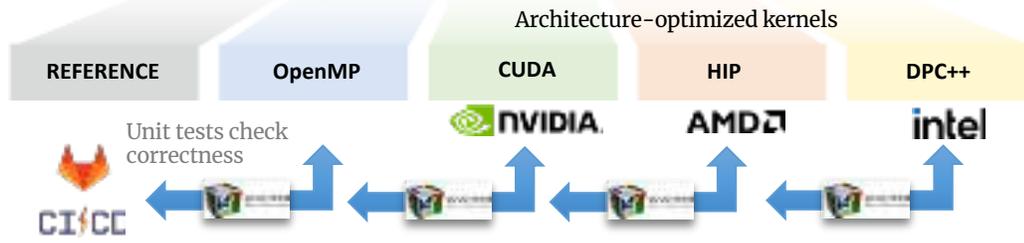


Library core contains architecture-agnostic factuality

- CORE**  
Infrastructure Algorithms
- Iterative Solvers
  - Preconditioners
  - ...



Runtime polymorphism selects the right kernel depending on the target architecture



**XGC collision operator solve LAPACK vs. Ginkgo: XGC pe459\_d3d\_EM\_healload test case**

- XGC pe459\_d3d\_EM\_healload (Aaron's test case; used for Summit, Perlmutter and Crusher scaling studies)
- Preliminary study on 32 nodes of Perlmutter (128 A100s)
  - 2 poloidal planes (216k nodes per plane); 22.4M p1/GPU, 89.6M p1/node (p1\_num=700k)
  - Ran 20 time steps; collisions calculated every other time step

	Time step 1	Time step 2
MAIN_LOOP	75.02	1.028
COLLISION_OPERATOR	3.842	5.10
FLUX_CALCULATION	1.12	0.21
FLUX_CALCULATION_COPY	1.12	0.18
COLLISION_OPERATOR_COPY	1.12	0.05
COLLISION_OPERATOR_COPY_COPY	1.12	0.05
COLLISION_OPERATOR_COPY_COPY_COPY	1.12	0.05

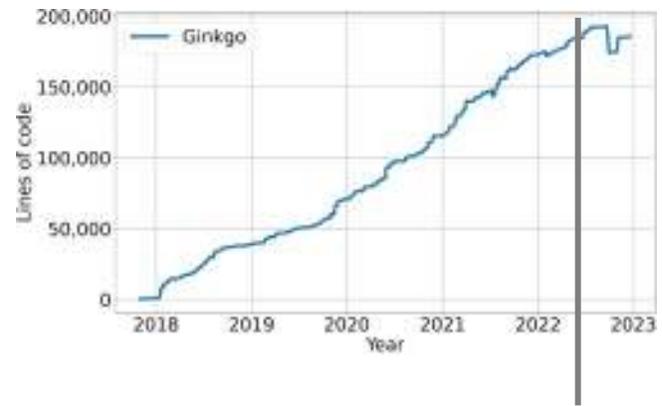
**With CPU LAPACK dgbsv**

- COLLISION time is 17% of MAIN\_LOOP time
- COLLISION\_OPERATOR\_SOLVE is 24% of COLLISION time and 4.3% of MAIN\_LOOP time
- COLLISION\_OPERATOR\_DGBSV is 67% of COLLISION\_OPERATOR\_SOLVE
- COLLISION\_OPERATOR\_DGBSV\_BANDWIDTH is 10% of COLLISION\_OPERATOR\_SOLVE

**Replacing CPU LAPACK dgbsv by GPU Ginkgo**

- COLLISION\_OPERATOR\_SOLVE reduced from 0.026s to 0.002s per step (reduction of 92%)
- COLLISION\_OPERATOR\_COPY\_COPY\_COPY reduced from 0.016s to 0.002s per step (reduction of 88%)
- MAIN\_LOOP time reduced by 4.1%

**XGC collision operator solve performed on GPU**



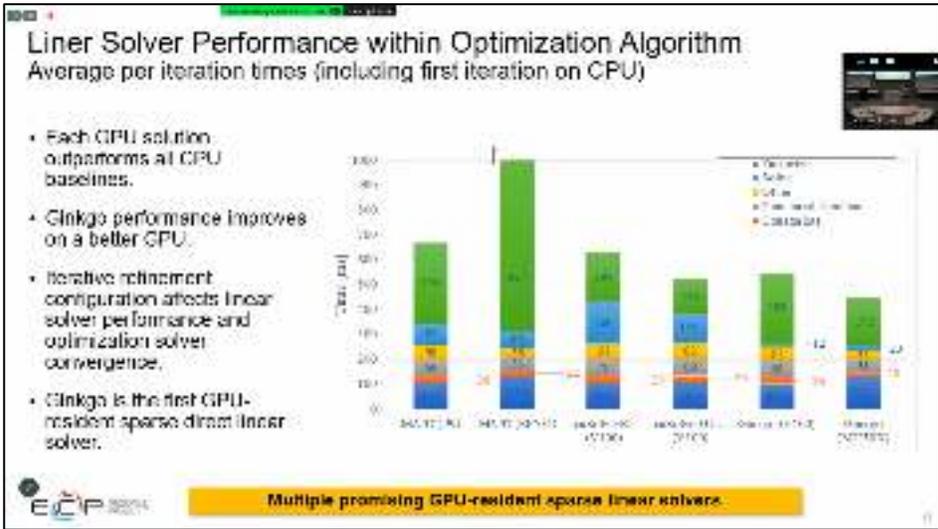
	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	✓	✓	✓	✓
	SpMM	✓	✓	✓	✓
	SpBEMM	✓	✓	✓	✓
Krylov solvers	BICGSTAB	✓	✓	✓	✓
	BICGSTAB_L	✓	✓	✓	✓
	CG	✓	✓	✓	✓
	CGS	✓	✓	✓	✓
	GMRES	✓	✓	✓	✓
Preconditioners	IDR	✓	✓	✓	✓
	(Block-1)Jacobi	✓	✓	✓	✓
	ILU/IC	✓	✓	✓	✓
	Parallel ILU/IC	✓	✓	✓	✓
Batched	Sparse Approximate Inverse	✓	✓	✓	✓
	Batched BICGSTAB	✓	✓	✓	✓
	Batched CG	✓	✓	✓	✓
	Batched GMRES	✓	✓	✓	✓
	Batched ILU	✓	✓	✓	✓
	Batched Jacobi	✓	✓	✓	✓
AMG	AMG preconditioner	✓	✓	✓	✓
	Parallel Graph Match	✓	✓	✓	✓
Utilities	On-Device Matrix Assembly	✓	✓	✓	✓
	MC64/RCM reordering	✓			
	Wrapping user data		✓		
	Logging		✓		
	PAPI counters		✓		



© Doug Kothe



# Sparse direct solvers for power grid simulations



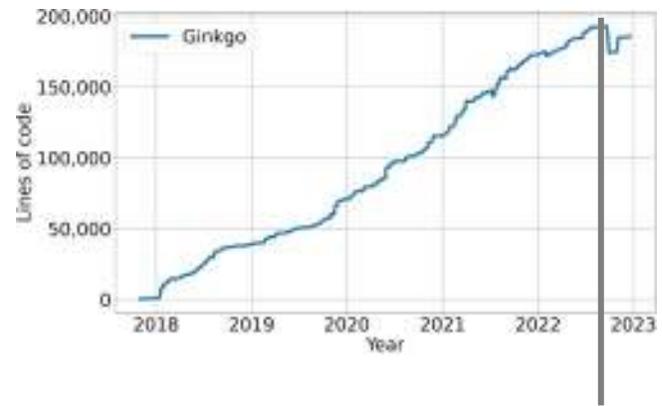
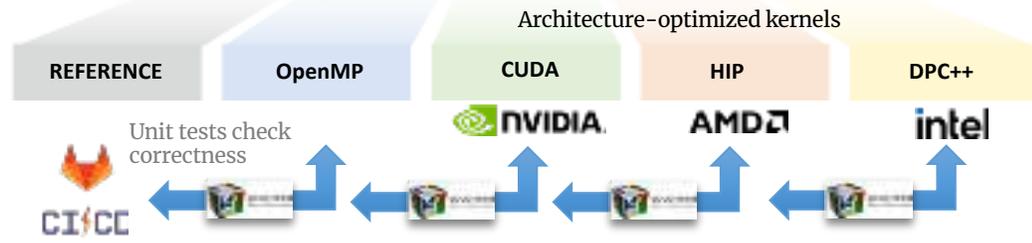
**EXASGD** © Slaven Peles

Library core contains architecture-agnostic factuality

- CORE**  
Infrastructure Algorithms
- Iterative Solvers
  - Preconditioners
  - ...



Runtime polymorphism selects the right kernel depending on the target architecture



**EXASGD**

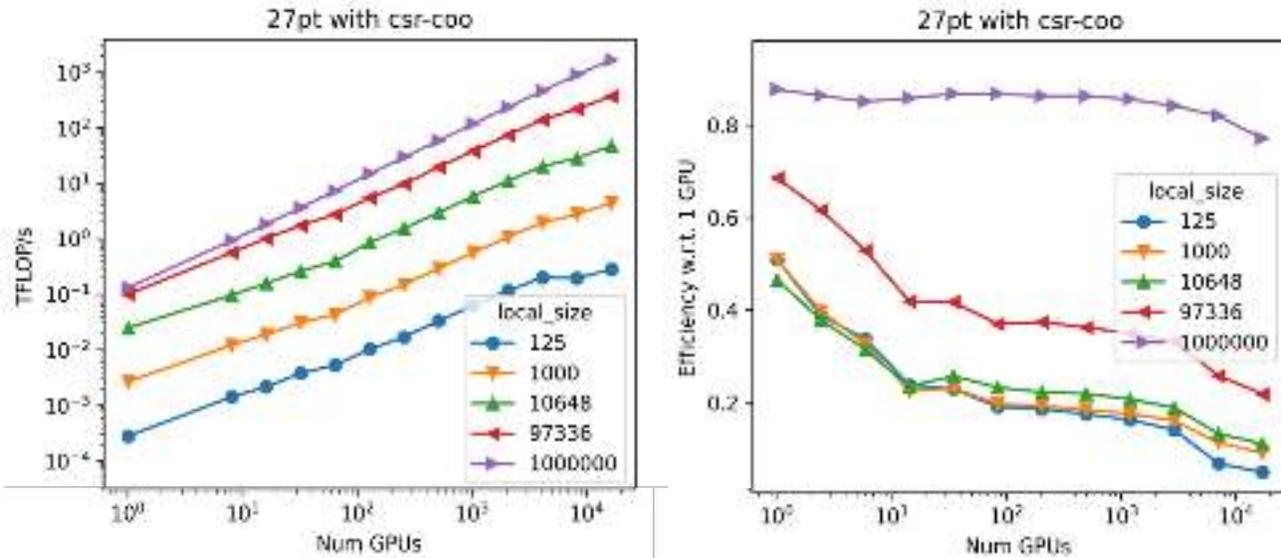
	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	✓	✓	✓	✓
	SpMM	✓	✓	✓	✓
	SpGeMM	✓	✓	✓	✓
Krylov solvers	BiCG	✓	✓	✓	✓
	BiCGSTAB	✓	✓	✓	✓
	CG	✓	✓	✓	✓
	CGS	✓	✓	✓	✓
	GMRES	✓	✓	✓	✓
Preconditioners	IDR	✓	✓	✓	✓
	(Block-)Jacobi	✓	✓	✓	✓
	ILU/IC	✓	✓	✓	✓
	Parallel ILU/IC	✓	✓	✓	✓
Batched	Sparse Approximate Inverse	✓	✓	✓	✓
	Batched BiCGSTAB	✓	✓	✓	✓
	Batched CG	✓	✓	✓	✓
	Batched GMRES	✓	✓	✓	✓
	Batched ILU	✓	✓	✓	✓
AMG	Batched ISA1	✓	✓	✓	✓
	Batched Jacobi	✓	✓	✓	✓
	AMG preconditioner	✓	✓	✓	✓
Sparse direct	AMG solver	✓	✓	✓	✓
	Parallel Graph Match	✓	✓	✓	✓
	Symbolic Cholesky	✓	✓	✓	✓
	Numeric Cholesky	✓	✓	✓	✓
	Symbolic LU	✓	✓	✓	✓
Utilities	Numeric LU	✓	✓	✓	✓
	Sparse TRSV	✓	✓	✓	✓
	On-Device Matrix Assembly	✓	✓	✓	✓
	MCs4/RCM-reordering	✓	✓	✓	✓
	Wrapping user data	✓	✓	✓	✓
	Logging	✓	✓	✓	✓
	PAPI counters	✓	✓	✓	✓



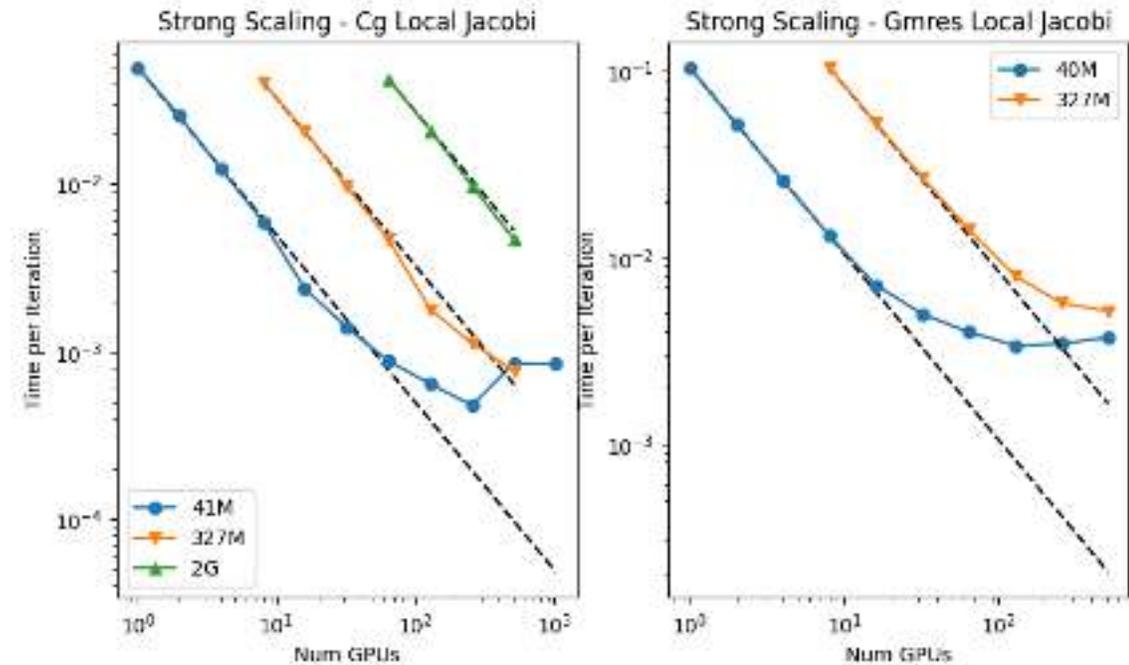
# Distributed runs on Frontier (Cray + AMD MI250 GPUs)

*Weak scaling: problem size increases with parallel resources*

Weak scaling up to 16k GCDs (8k GPUs)

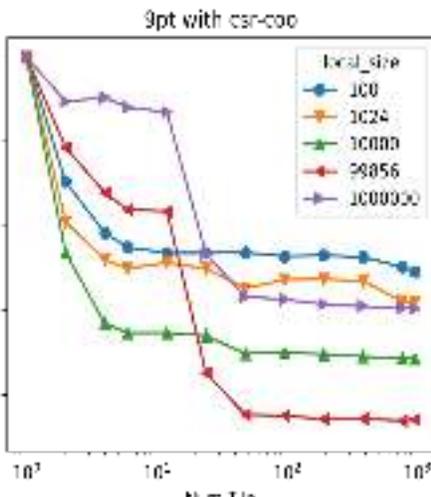
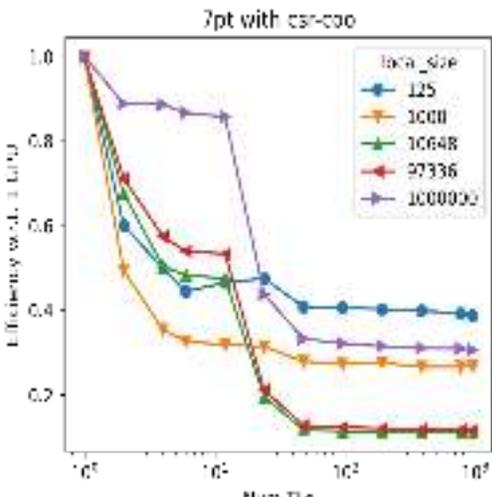
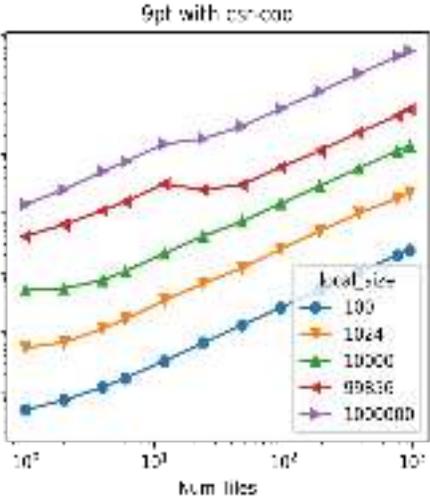
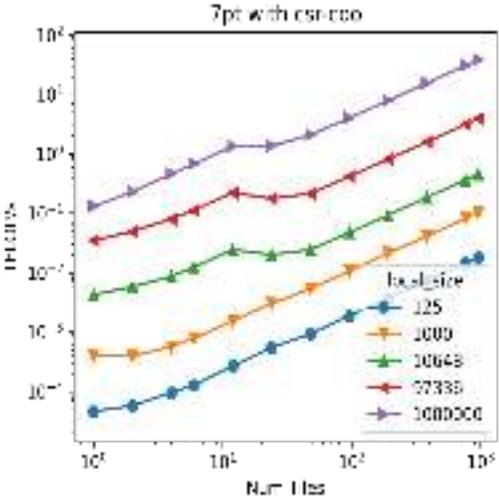
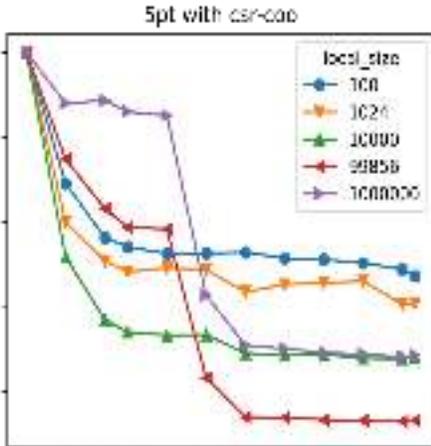
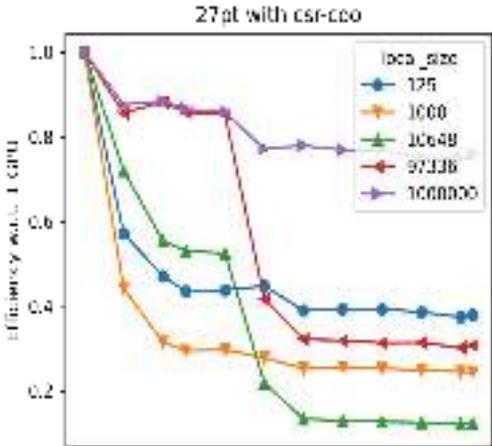
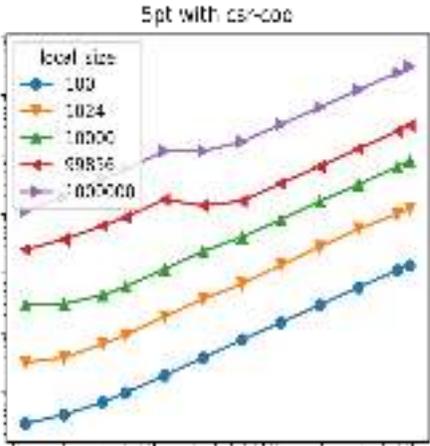
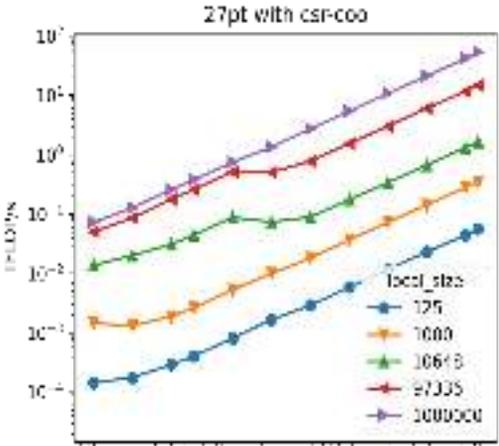


*Strong scaling: problem size constant*



# Distributed runs on Sunspot (Intel PVCA GPUs)

*Weak scaling: problem size increases with parallel resources*



# “Now” – Near completion of ECP

- Sustainable software design ready for the addition of new backends.

- EuroHPC Project MICROCARD uses Ginkgo



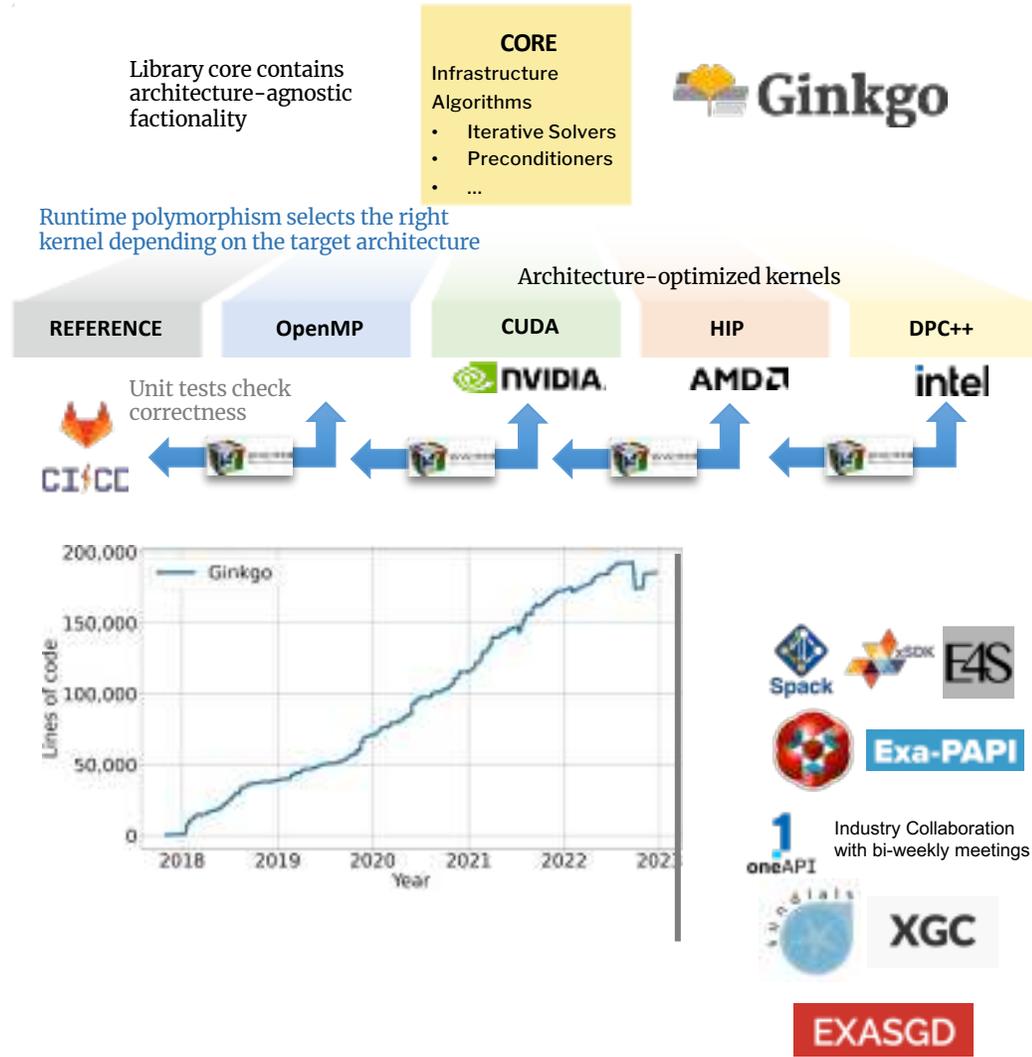
- BMBF PDExa project uses Ginkgo



- BMBF ExaSIM project uses Ginkgo



<https://exasim-project.com>



	Functionality	OMP	CUDA	HIP	DPC++
Basic	SpMV	☑	☑	☑	☑
	SpMM	☑	☑	☑	☑
	SpGeMM	☑	☑	☑	☑
Krylov solvers	BICG	☑	☑	☑	☑
	BICGSTAB	☑	☑	☑	☑
	CG	☑	☑	☑	☑
	CGS	☑	☑	☑	☑
	GMRES	☑	☑	☑	☑
Preconditioners	IDR	☑	☑	☑	☑
	(Block-1)Jacobi	☑	☑	☑	☑
	ILU/IC	☑	☑	☑	☑
	Parallel ILU/IC	☑	☑	☑	☑
Batched	Sparse Approximate Inverse	☑	☑	☑	☑
	Batched BICGSTAB	☑	☑	☑	☑
	Batched CG	☑	☑	☑	☑
	Batched GMRES	☑	☑	☑	☑
	Batched ILU	☑	☑	☑	☑
	Batched ISAI	☑	☑	☑	☑
AMG	Batched Jacobi	☑	☑	☑	☑
	AMG preconditioner	☑	☑	☑	☑
	AMG solver	☑	☑	☑	☑
	Parallel Graph Match	☑	☑	☑	☑
Sparse direct	Symbolic Cholesky	☑	☑	☑	☑
	Numeric Cholesky	☑	☑	☑	☑
	Symbolic LU	☑	☑	☑	☑
	Numeric LU	☑	☑	☑	☑
Utilities	Sparse TRSV	☑	☑	☑	☑
	On-Device Matrix Assembly	☑	☑	☑	☑
	MC64/RCM/reordering	☑	☑	☑	☑
	Wrapping user data	☑	☑	☑	☑
	Logging	☑	☑	☑	☑
	PAPI counters	☑	☑	☑	☑



# Lessons learnt from the Ginkgo development process

- **ECP earmarking roughly half the budget to Software & App development is a game changer.**
  - **Central component for the success of ECP.**
  - This concept needs to – and does become - the blueprint for other nations and projects.
- **Workforce recruitment and workforce retention are the key to success in software development.**
  - Money does not write software. RSEs do. **We need to create attractive career plans.**
  - We need to make research software development attractive to students. **Academic recognition.**
- **Anticipating the future in hardware development accelerates the porting process.**
  - **Blueprints** and **early access systems** both useful.
  - **Interaction with industry** is mutually beneficial.
- **Management, tools, and strategic initiatives, interaction and collegial behavior are important.**
  - Jira/Notion/[...] milestones and deliverables give projects and collaborative interactions a structure and timeline.
  - **Strategic focus groups, conferences, and meetings** bring experts together and **create collaboration.**
  - **Listen to the application needs. Value input and acknowledge collaborators.**