

# Porting Machine Learning and Modeling from a Laptop to a Supercomputer: Numerical Libraries

Matthew Knepley



P

Portable

E

Extensible

T

Toolkit for

S

Scientific

c

computing

<https://petsc.org>

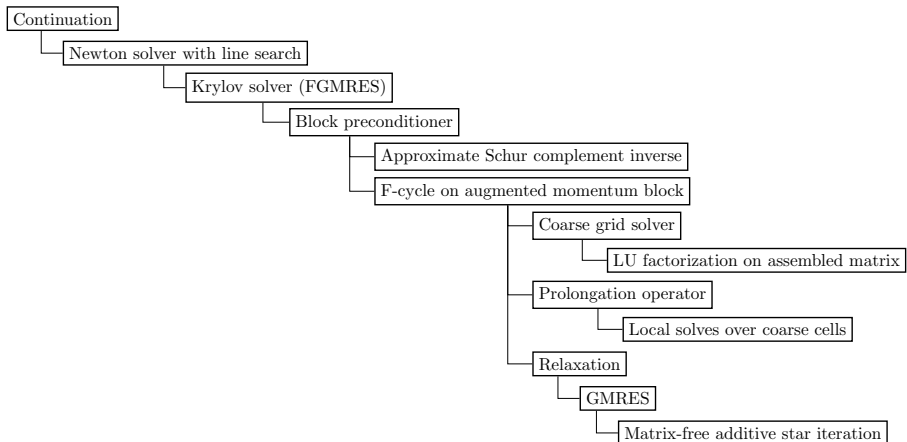
- [PyLith](#) (CIG)
- [Underworld](#) (Monash)
- [Salvus](#) (ETHZ)
- [TerraFERMA](#) (Columbia)
- [PFLOTRAN](#) (DOE)
- [STOMP](#) (DOE)
- [pTatin3d](#) (UCSD)
- [Rhea](#) (NYU)
- [Waiwera](#) (U Auckland)

- Vec/Mat bound to a compute resource at runtime  
*“Toward Performance-Portable PETSc for GPU-based Exascale Systems”*  
[arXiv:2011.00715](https://arxiv.org/abs/2011.00715)
- Communication happens directly between GPUs  
*“The PetscSF Scalable Communication Layer”*  
[arXiv:2102.13018](https://arxiv.org/abs/2102.13018)

- Allow accuracy limits to supercede stability
  - CFL limits
  - Geometric limits
- Scale separation
  - Parabolization
  - Localization
- Inverse problems
  - Bayesian inference
  - Parameter estimation

# Composable Solvers


## Linear



Farrell, Mitchell, and Wechsung (2018), [arXiv:1810.03315](https://arxiv.org/abs/1810.03315)


# Preconditioning NRichardson with Newton

```
./ex19 -lidvelocity 100 -grashof 1.3373e2  
-da_grid_x 16 -da_grid_y 16 -da_refine 2  
-snes_type nrichardson -snes_max_it 200  
-npc_snes_type newtonls -npc_snes_max_it 3 -npc_pc_type lu
```



# Preconditioning NRichardson with Newton


```
./ex19 -lidvelocity 100 -grashof 1.3373e2  
-da_grid_x 16 -da_grid_y 16 -da_refine 2  
-snes_type nrichardson -snes_max_it 200  
-npc_snes_type newtonls -npc_snes_max_it 4 -npc_pc_type lu
```





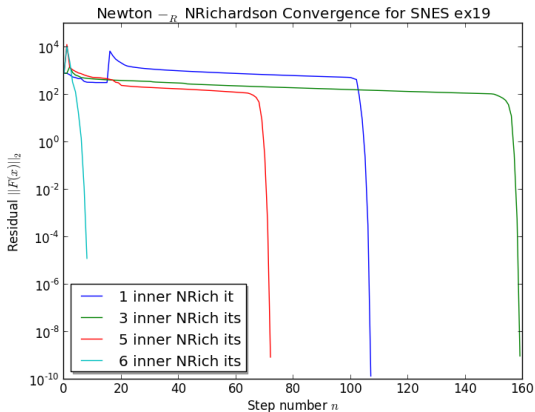
# Preconditioning Newton with NRichardson

```
./ex19 -lidvelocity 100 -grashof 1.3373e2  
-da_grid_x 16 -da_grid_y 16 -da_refine 2  
-snes_type newtonls -snes_max_it 1000 -pc_type lu  
-npc_snes_type nrichardson -npc_snes_max_it 1
```




# Preconditioning Newton with NRichardson

```
./ex19 -lidvelocity 100 -grashof 1.3373e2  
-da_grid_x 16 -da_grid_y 16 -da_refine 2  
-snes_type newtonls -snes_max_it 1000 -pc_type lu  
-npc_snes_type nrichardson -npc_snes_max_it 3
```




# Preconditioning Newton with NRichardson

```
./ex19 -lidvelocity 100 -grashof 1.3373e2  
-da_grid_x 16 -da_grid_y 16 -da_refine 2  
-snes_type newtonls -snes_max_it 1000 -pc_type lu  
-npc_snes_type nrichardson -npc_snes_max_it 5
```




# Preconditioning Newton with NRichardson

```
./ex19 -lidvelocity 100 -grashof 1.3373e2  
-da_grid_x 16 -da_grid_y 16 -da_refine 2  
-snes_type newtonls -snes_max_it 1000 -pc_type lu  
-npc_snes_type nrichardson -npc_snes_max_it 6
```




# Adaptive Mesh Refinement

- DM interface with [p4est](#) package from Burstedde and Isaac
- PETSc solvers can be used seamlessly
- 2015 Gordon Bell Winner for Mantle Convection Simulation




# Adaptive Mesh Refinement

- DM interface with [p4est](#) package from Burstedde and Isaac
- PETSc solvers can be used seamlessly
- 2015 Gordon Bell Winner for Mantle Convection Simulation
- Inversion for basal traction on the full Antarctic ice sheet



Observed surface flow velocity (Rignot et al., 2011)




Antarctic ice sheet inversion for the basal friction parameter field using InSAR surface velocity measurements

*a*


# Adaptive Mesh Refinement

- DM interface with [ParMmg](#) package from Froehly and Cirrottola
- PETSc solvers can be used seamlessly
- Wallwork, Barral, Ham, Piggot (2021) [eartharXiv:4205](#)



# Adaptive Mesh Refinement

- DM interface with [ParMmg](#) package from Froehly and Cirrottola
- PETSc solvers can be used seamlessly
- Wallwork, Barral, Ham, Piggot (2021) [eartharXiv:4205](#)





- Low-level AD

- PyTorch, Tensorflow

- Enzyme

“Performance-Portable Solid Mechanics via Matrix-Free  $p$ -Multigrid”

[arXiv:2204.01722](https://arxiv.org/abs/2204.01722)

- High-level AD

- pyadjoint (Firedrake, FEniCS)

[github:dolfin-adjoint/pyadjoint](https://github.com/dolfin-adjoint/pyadjoint)

- libadjoint

[bitbucket:dolfin-adjoint/libadjoint](https://bitbucket.org/dolfin-adjoint/libadjoint)

# Libraries Hide Hardware Details

# Libraries Hide Hardware Details



# Libraries Hide Hardware Details

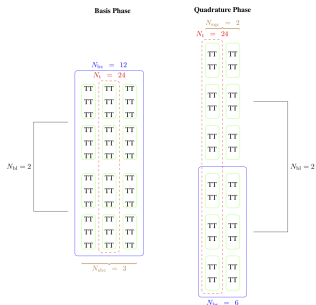


## Libraries Hide Hardware Details


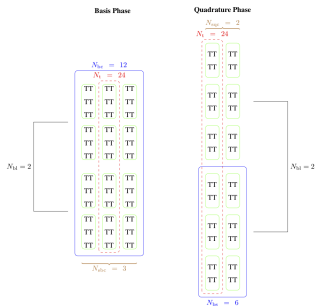


# Libraries Hide Implementation Complexity

## Libraries Hide Implementation Complexity




## Libraries Hide Implementation Complexity





## Libraries Hide Implementation Complexity



# Libraries Accumulate Best Practices

# Libraries Accumulate Best Practices

Classical Gram-Schmidt orthogonalization with  
selective reorthogonalization

# Libraries Accumulate Best Practices

Classical Gram-Schmidt orthogonalization with selective reorthogonalization

Eigen-estimation for AMG with Krylov bootstrap

# Libraries Accumulate Best Practices

Classical Gram-Schmidt orthogonalization with selective reorthogonalization

Eigen-estimation for AMG with Krylov bootstrap

Improvement without code changes

# Libraries Simplfy Determining Provenance

# Libraries Simplify Determining Provenance

Rather than making build-time choices,

# Libraries Simplify Determining Provenance

Rather than making build-time choices,  
that must be plumbed through all levels



# Libraries Simplify Determining Provenance

Rather than making build-time choices,  
that must be plumbed through all levels  
with another layer of workflow scripts

# Libraries Simplify Determining Provenance

Rather than making build-time choices,  
that must be plumbed through all levels  
with another layer of workflow scripts  
and brittle top-level interfaces,

# Libraries Simplfy Determining Provenance

we use packages without modification,

# Libraries Simplify Determining Provenance

we use packages without modification,  
compiled in a standard way

# Libraries Simplify Determining Provenance

we use packages without modification,  
compiled in a standard way  
and controlled entirely via runtime options.

# Why PETSc?

- **Dependability & Maintainability**
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- **Dependability & Maintainability**
  - Nearly 30 years of continuous development
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility



# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
  - Tested at every supercomputer installation and 10,000+ users
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
  - Many Gordon Bell Winners
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
  - State-of-the-Art Linear and Nonlinear Solvers, Eigensolvers, Optimization Solvers, Timesteppers
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility

# Why PETSc?

- Dependability & Maintainability
- Portability & Robustness
- Performance & Scalability
- Optimality & Robustness
- Flexibility & Extensibility
  - More than 8000 citations and hundreds of application packages
  - Aerodynamics, Arterial Flow, Corrosion, Combustion, Data Mining, Earthquake Mechanics, . . .