



# Massive Scaling of MASSIF: Algorithm Development for Hooke's Law Simulations on Distributed GPU systems

Anuva Kulkarni<sup>1</sup>, Jelena Kovačević<sup>2</sup>, Franz Franchetti<sup>1</sup>

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>NYU Tandon School of Engineering

## GPUs and the need for algorithm re-design

### Motivation


Common characteristics of scientific codes:

- FFT-based simulations involve all-to-all communication
- High memory requirement


Incompatibility with GPUs:

- GPUs have small on-chip memory (~16GB)
- Various communication latencies

MASSIF simulates Composite microstructure made up of grains



Stress field at grain boundary



### Algorithm 1 MASSIF Inner loop

```

1: Initialize:  $\epsilon^0 \leftarrow E, \sigma_{mn}^0(x) \leftarrow C_{mnlk}(x) : \epsilon_{kl}^0(x)$ 
2: while  $\epsilon_k > \epsilon_{tol}$  do
3:    $\hat{\epsilon}_{mn}^{(i)} \leftarrow \text{FFT}(\sigma_{mn}^{(i)}(x))$ 
4:   Check convergence
5:    $\Delta \hat{\epsilon}_{kl}^{(i+1)}(\xi) \leftarrow \hat{\Gamma}_{klmn}(\xi) : \hat{\sigma}_{mn}^{(i)}(\xi)$ 
6:   Update strain:  $\hat{\epsilon}_{kl}^{(i+1)}(\xi) \leftarrow \hat{\epsilon}_{kl}^{(i)}(\xi) + \Delta \hat{\epsilon}_{kl}^{(i+1)}(\xi)$ 
7:    $\hat{\epsilon}_{kl}^{(i+1)}(x) \leftarrow \text{iFFT}(\hat{\epsilon}_{kl}^{(i+1)}(\xi))$ 
8:   Update stress:  $\sigma_{mn}^{(i+1)}(x) \leftarrow C_{mnlk}(x) : \hat{\epsilon}_{kl}^{(i+1)}(x)$ 

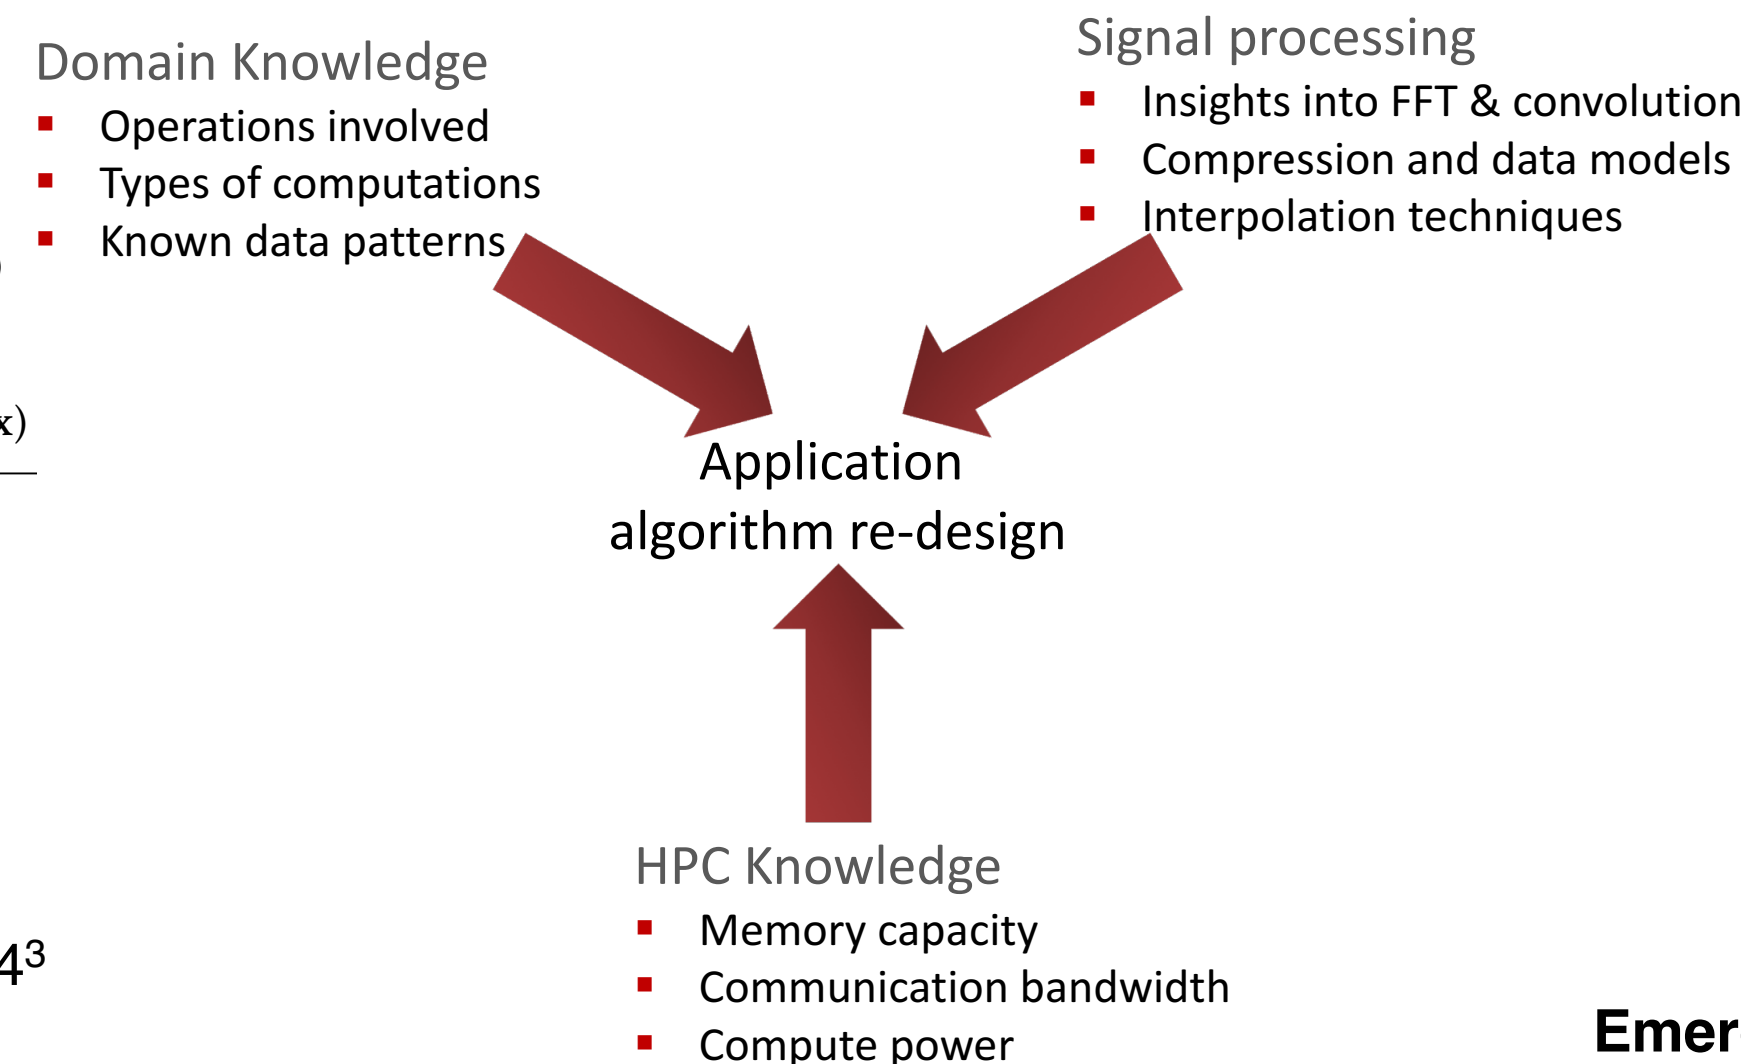
```

### Main issues faced:

As problem size increases for high resolution simulations, storage and communication requirements increase. Problems larger than  $1024^3$  grids not currently simulated.

### Proposed Solution:

Re-design the application algorithm using domain expertise and knowledge about high performance computing platforms.

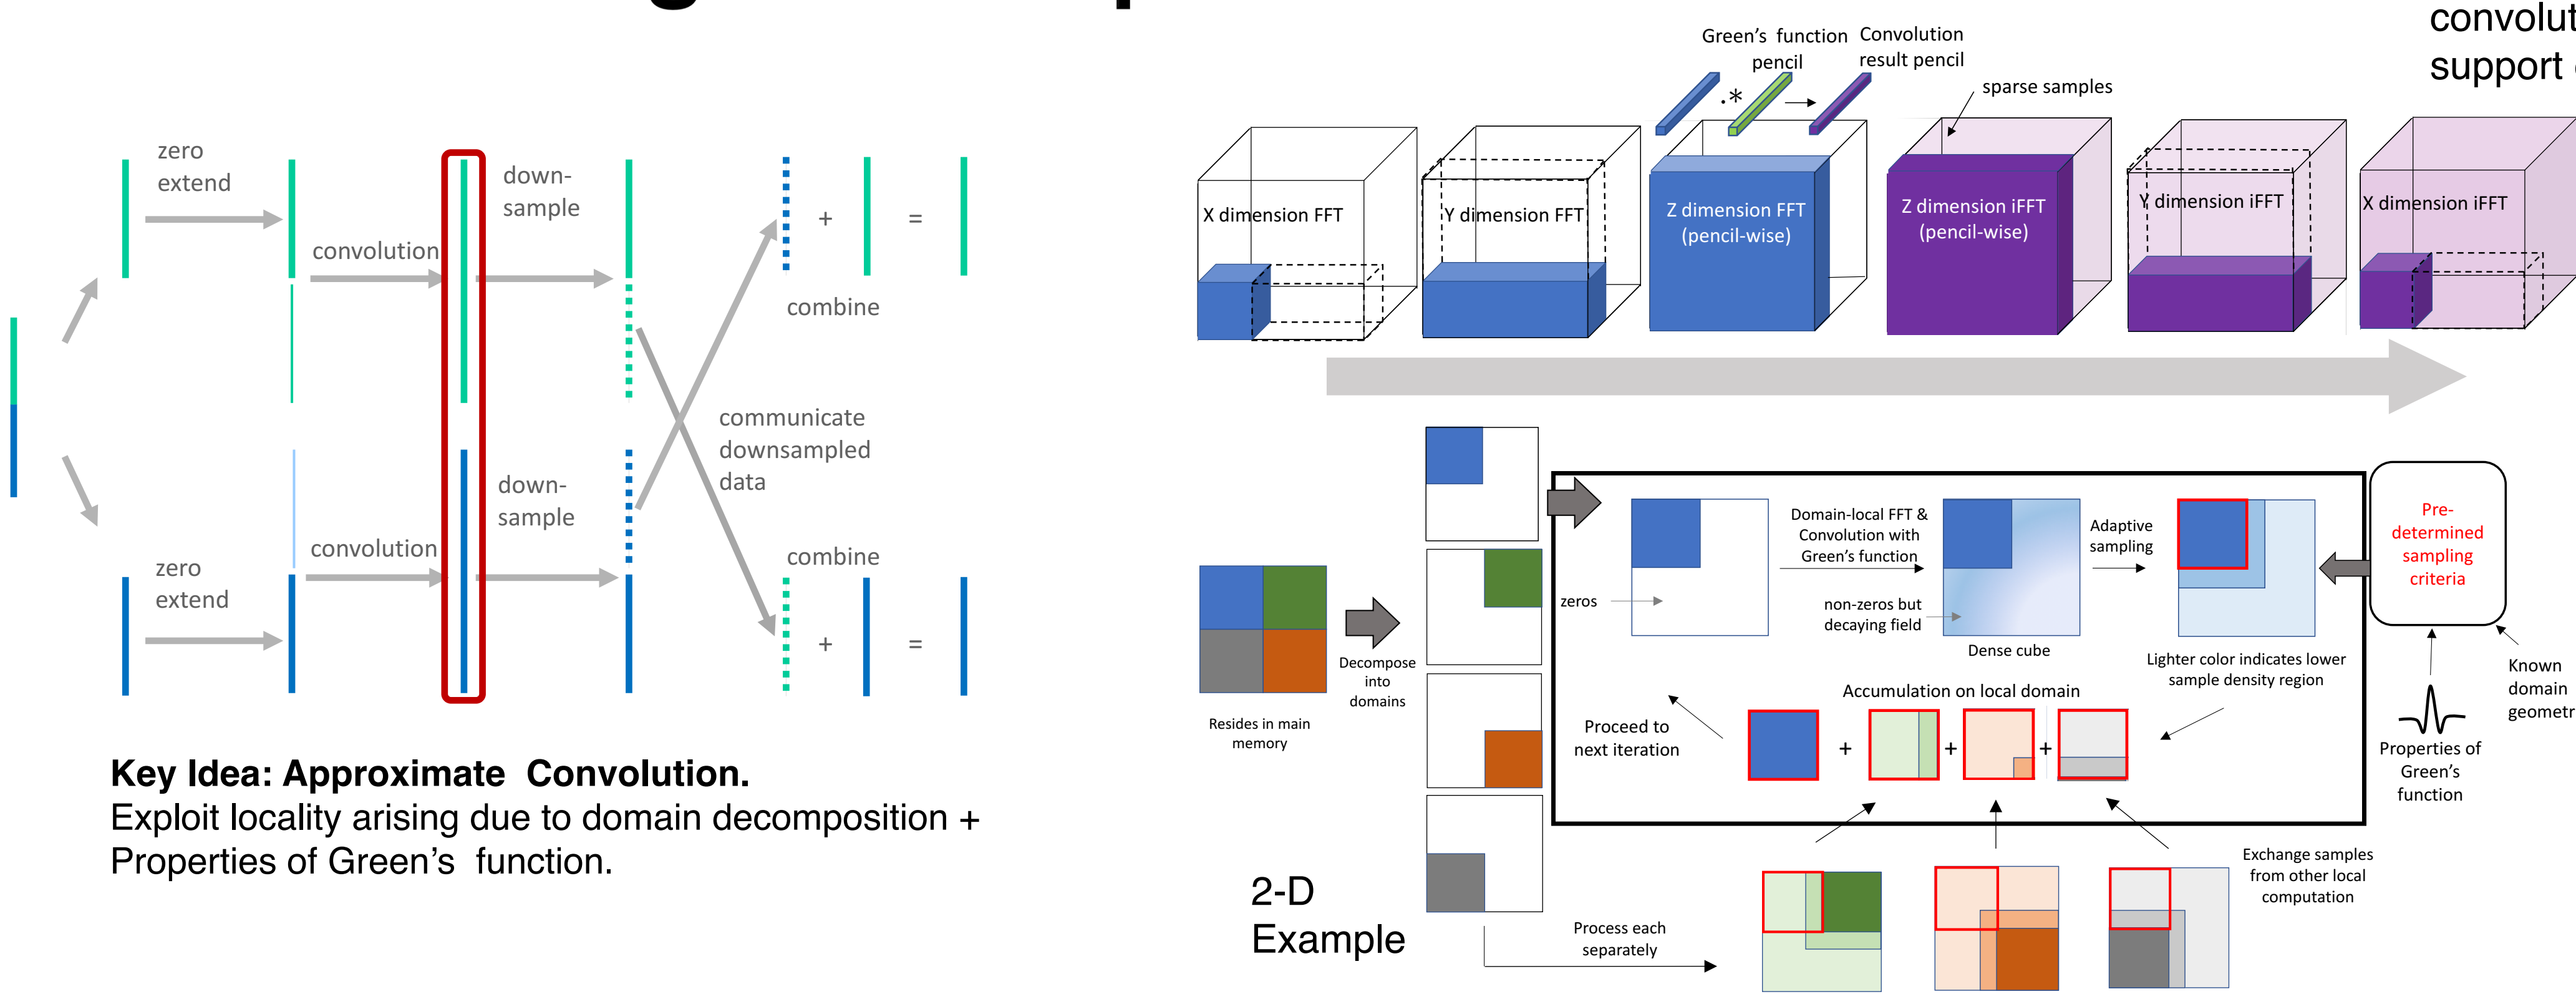


### Challenges with FFTW

- FFTW is de-facto standard interface for FFT
- Vendor libraries support the FFTW 3.X interface: Intel MKL, IBM ESSL, AMD ACML (end-of-life), Nvidia cuFFT, Cray LibSci/CRAFFT
- Some Issues:
  - No native support for accelerators (GPUs, Xeon PHI, FPGAs) and SIMT
  - Parallel/MPI version does not scale beyond 32 nodes
  - No analogue to LAPACK for spectral method
  - Complex data patterns may need to be expressed, FFTW currently falls short. But, extensions like FFTX could add new descriptors.

Emerging interfaces like FFTX, extension of FFTW, enables algorithm specification as composition of sub-plans

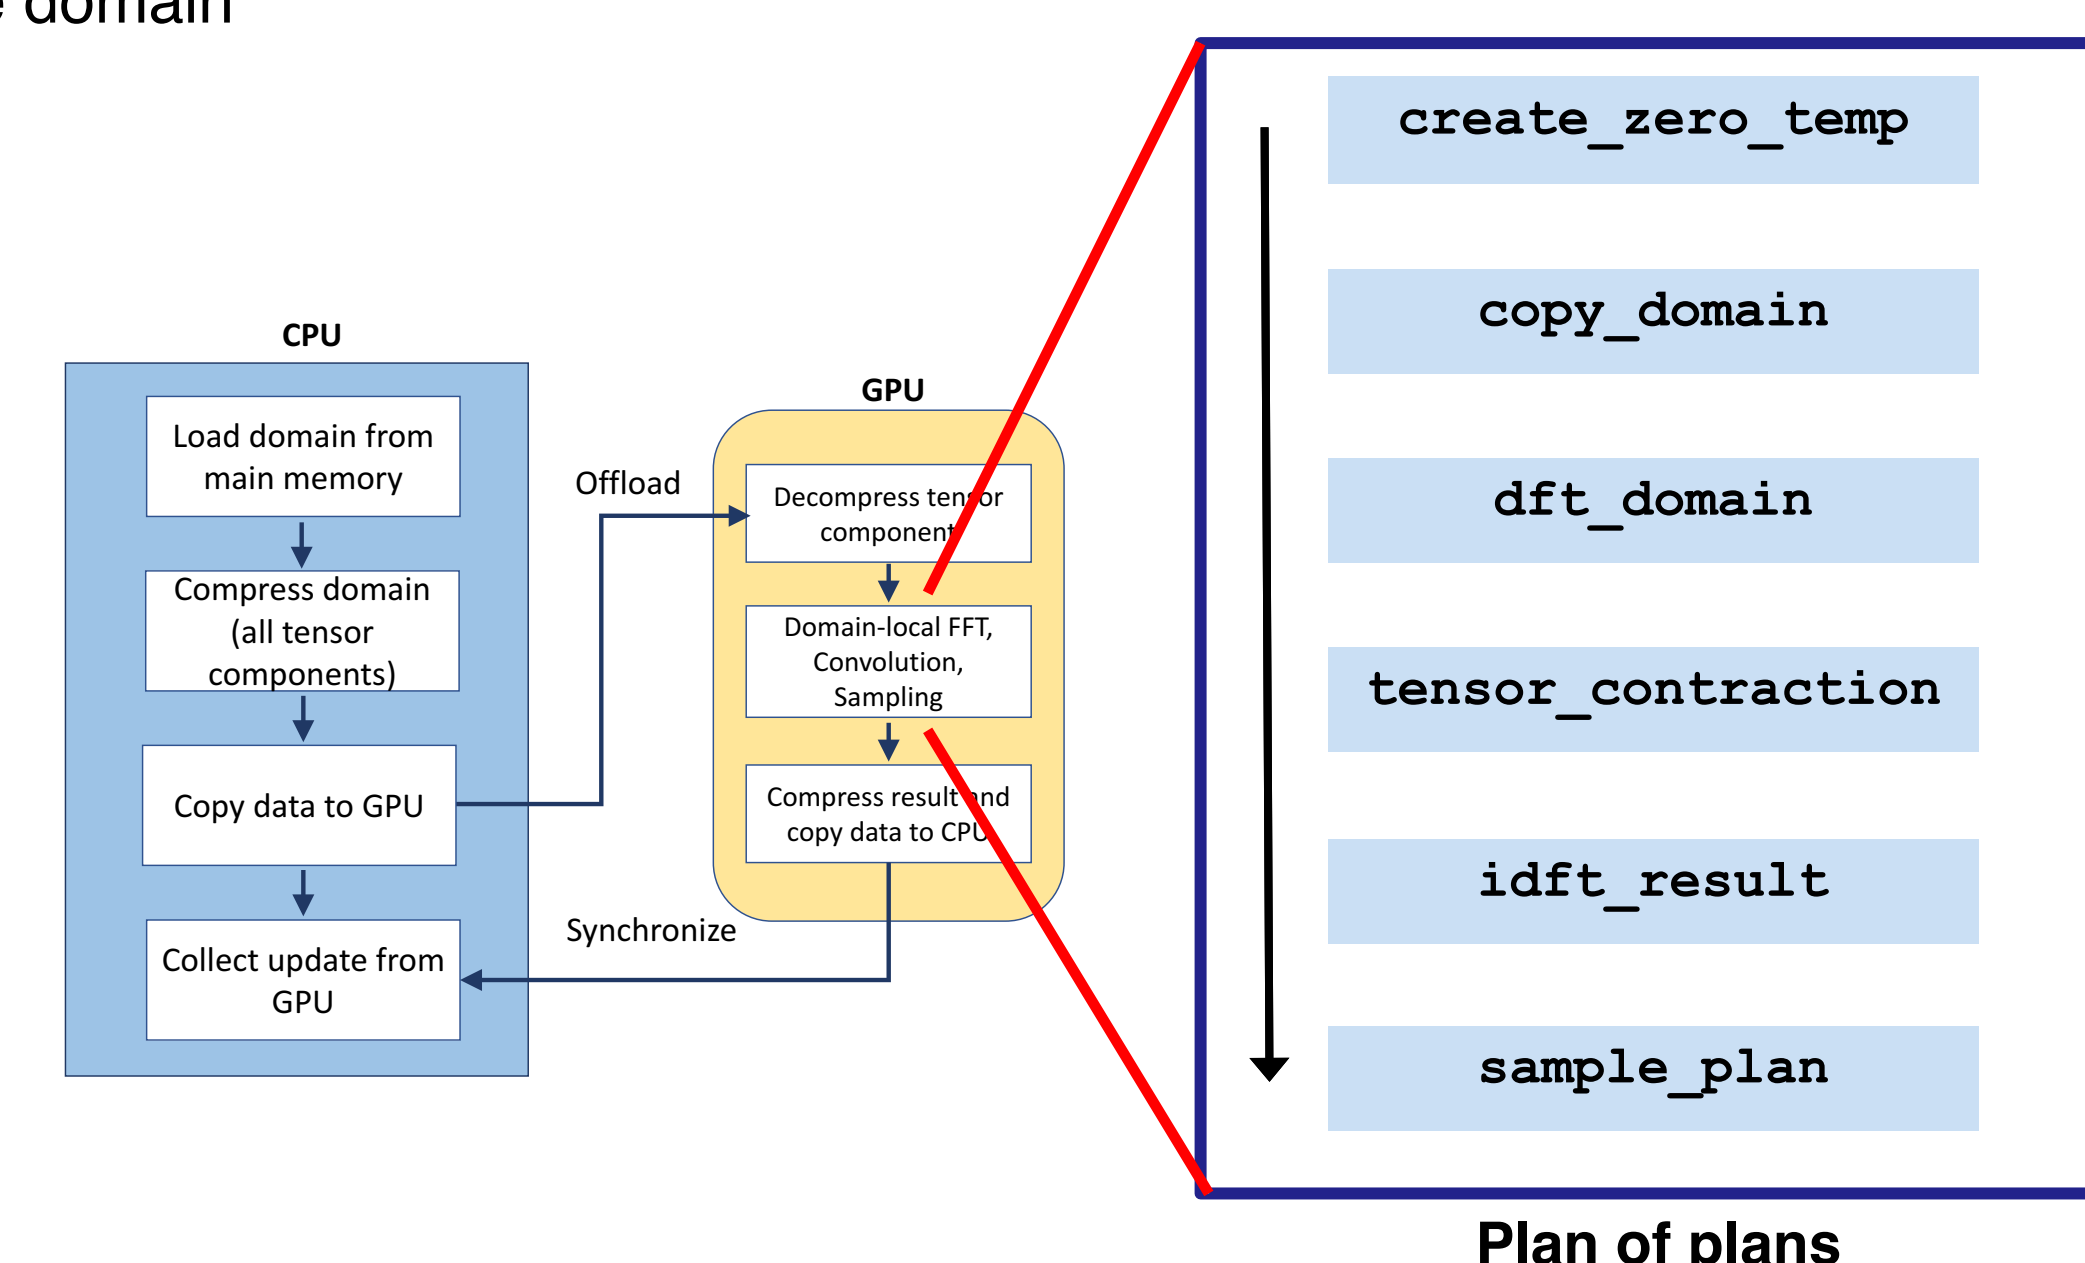
## Front end: Algorithm Specification



**Key Idea: Approximate Convolution.**  
Exploit locality arising due to domain decomposition + Properties of Green's function.

Sparse sampling of convolution result outside the support of the domain

### Making algorithm specification easier for the user



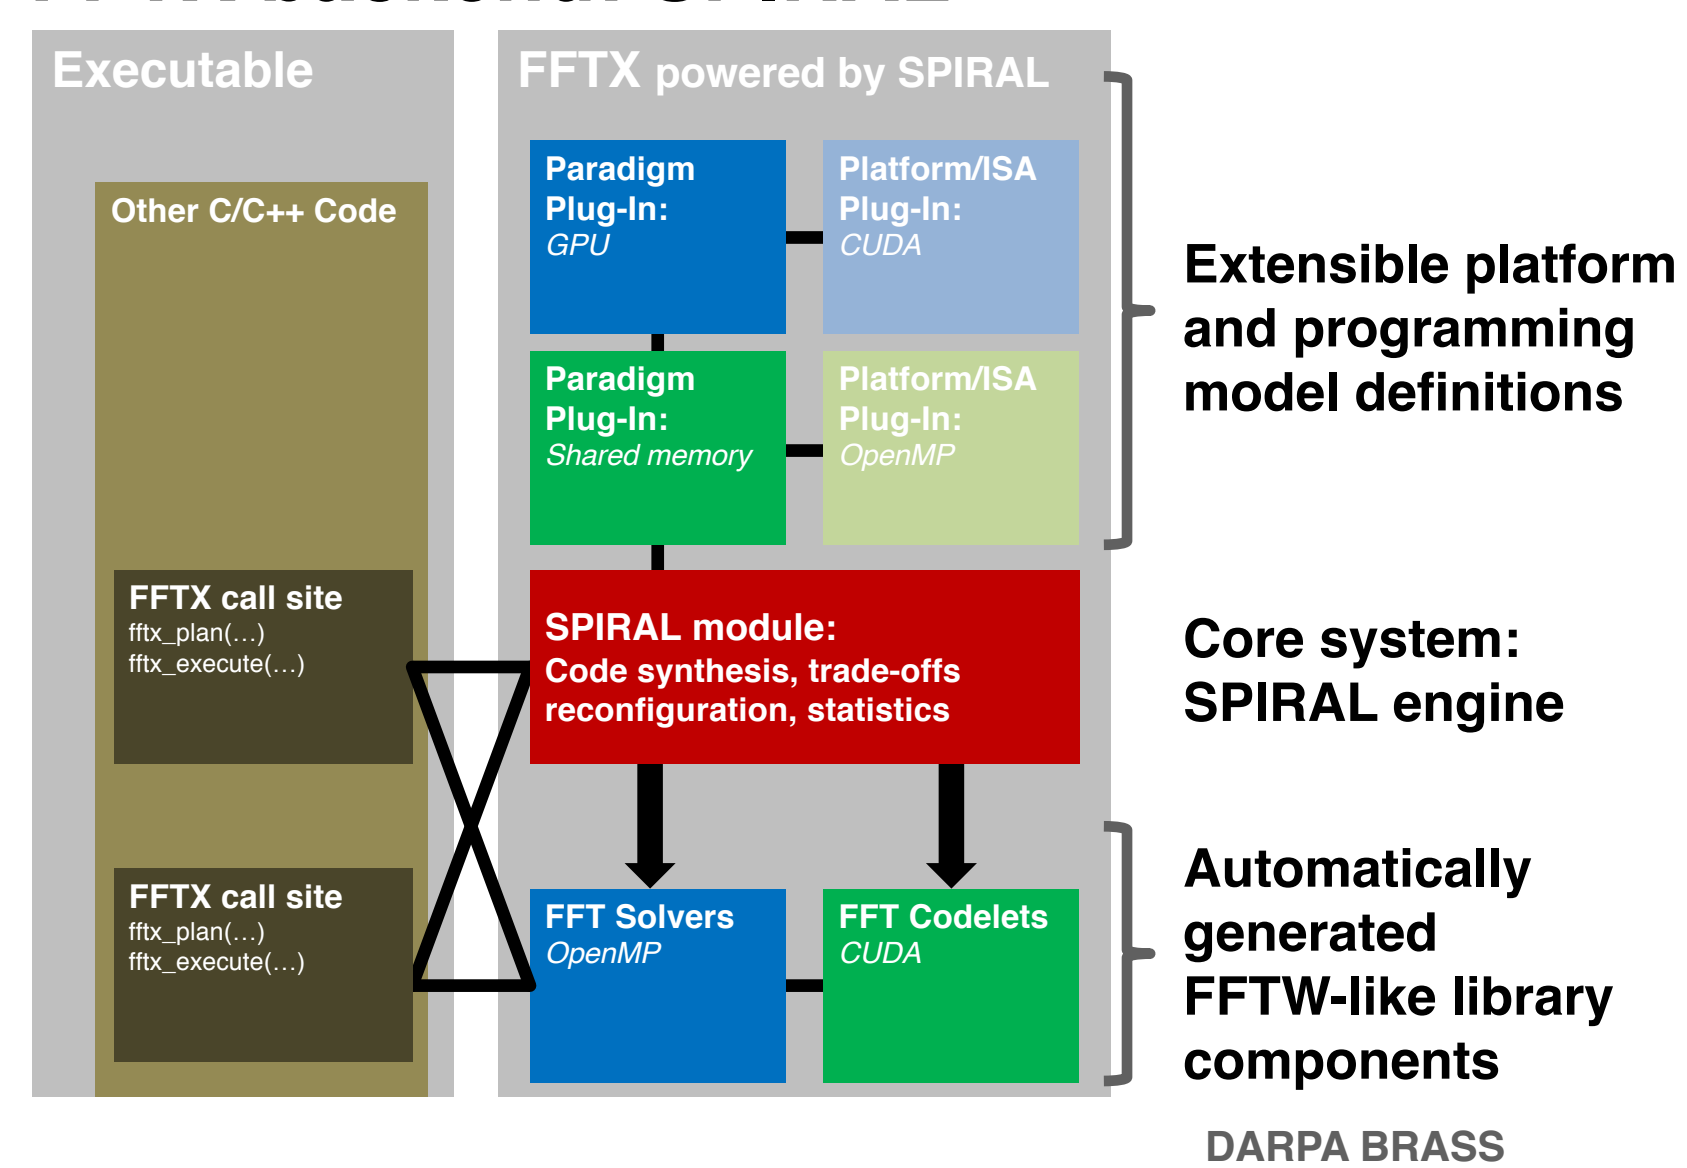
Goal: Domain decomposition and sampling expressed easily using an FFTW-like interface

## Back end: Code Optimization

### FFTX is..

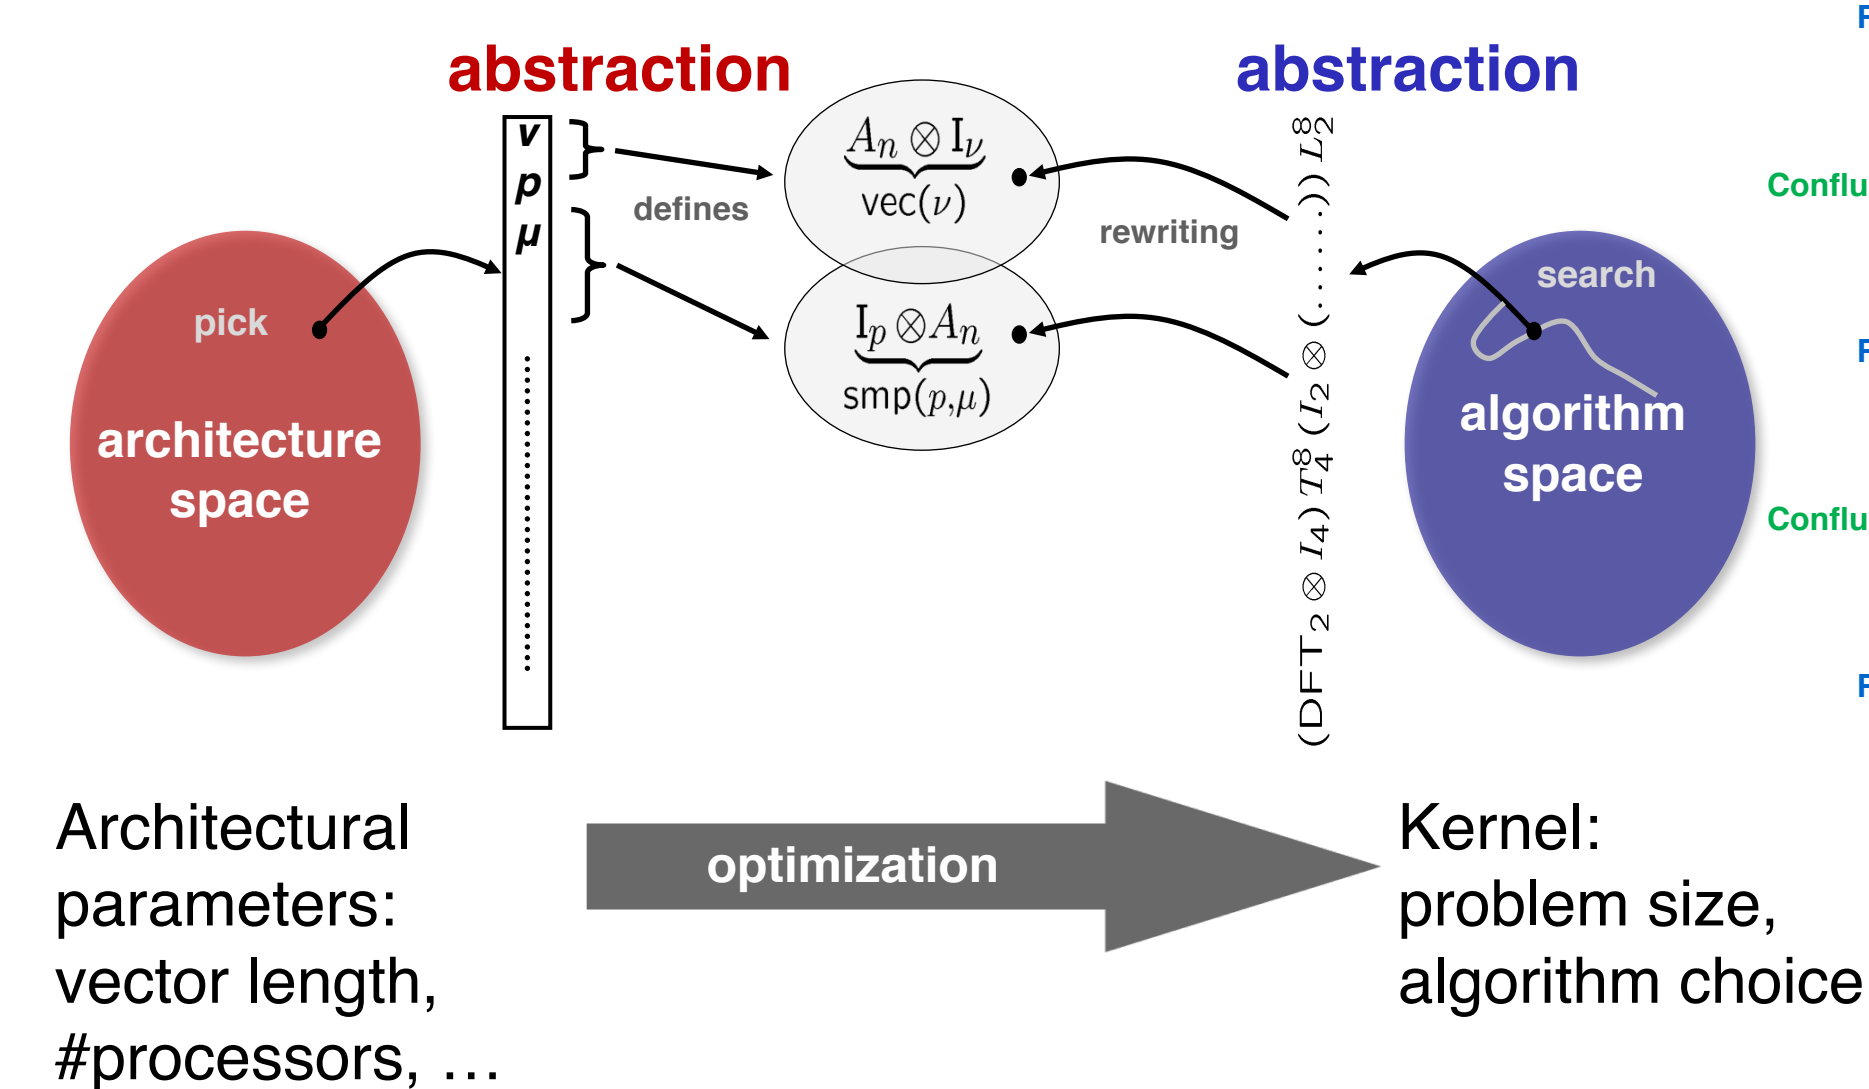
- Modernized FFTW-style interface
- Backwards compatible to FFTW 2.X and 3.X
- Small number of new features, familiar interface
- Code generation backend using SPIRAL
- Compilation and advanced performance optimization cross-call and cross library optimization, accelerator off-loading,...

### FFTX backend: SPIRAL

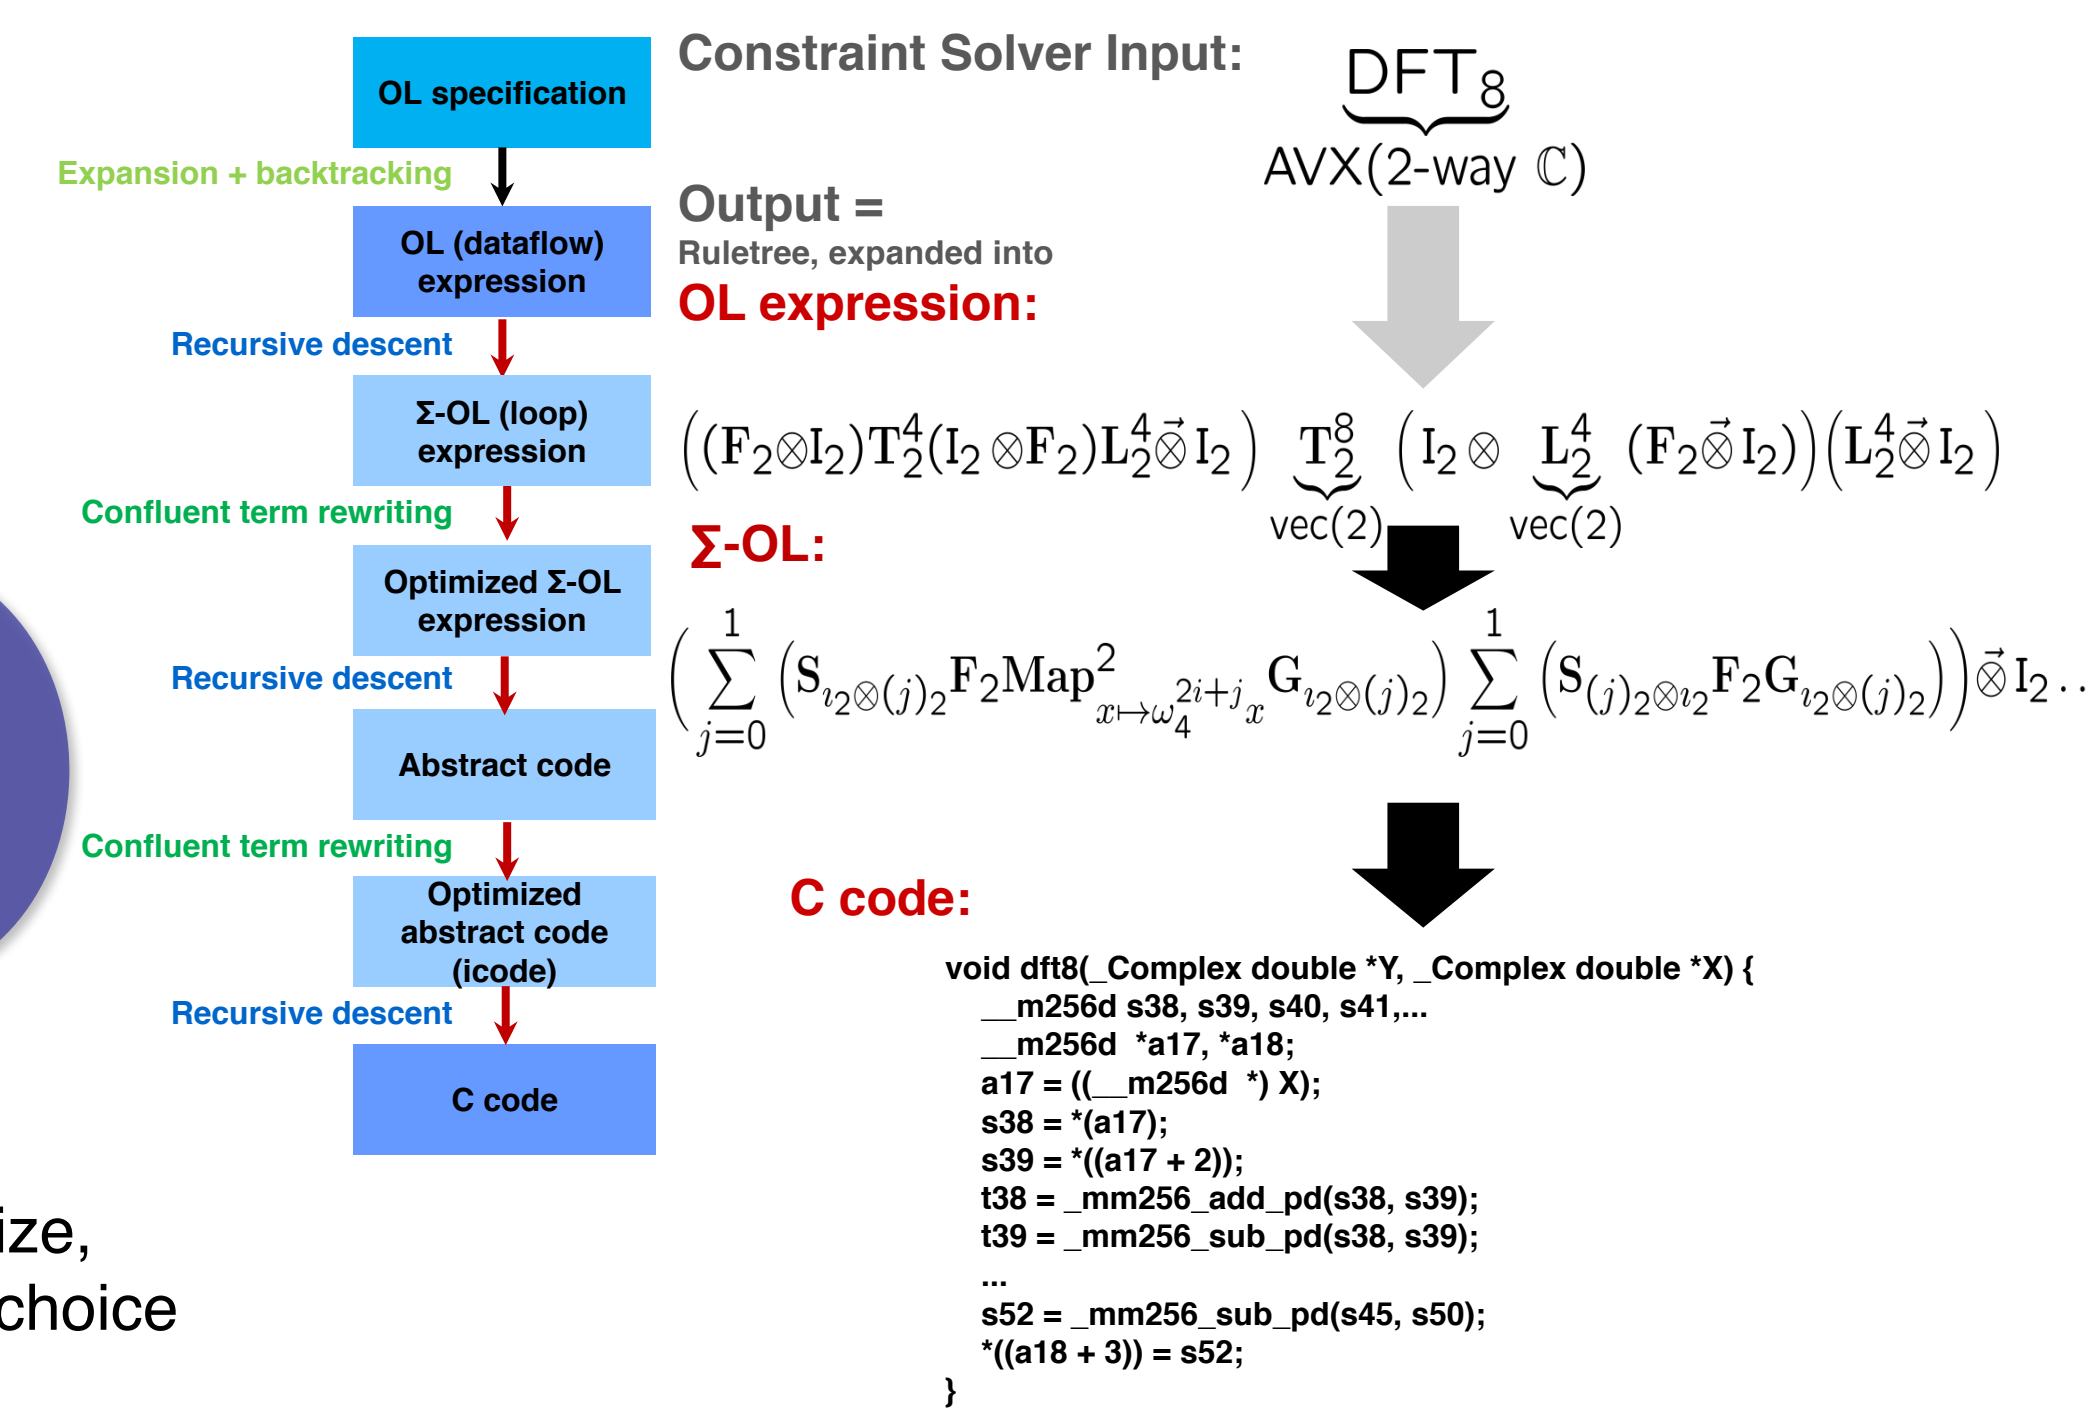


### Platform-aware formal program synthesis

Model: common abstraction = spaces of matching formulas



### Translating an OL expression into code



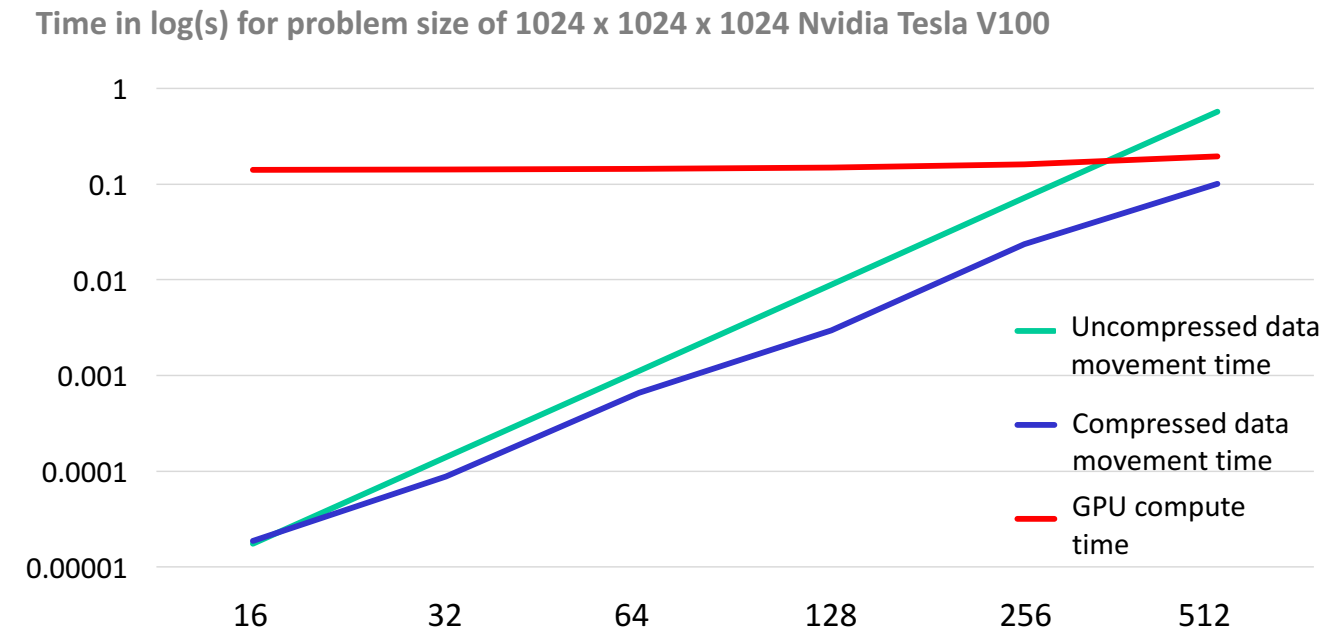
## First-order Performance Analysis

Single node, multi-GPU (HPE Apollo 6500 node)

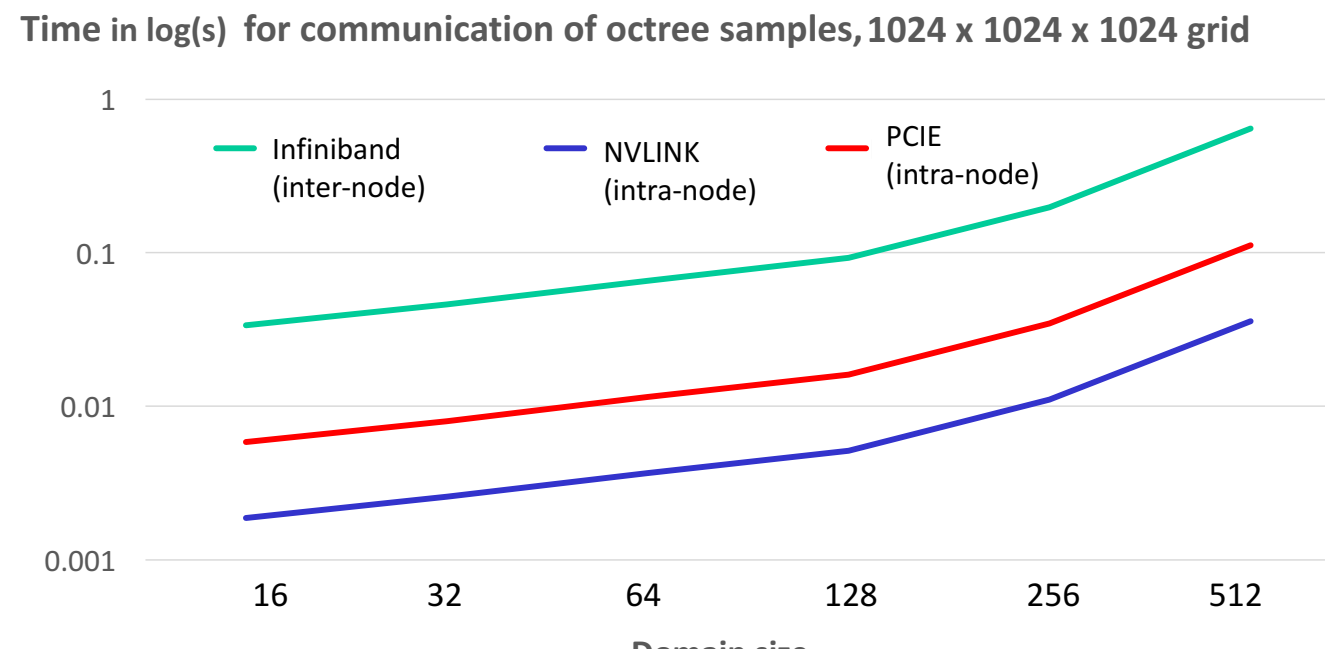
DGX2 Workstation with 16 GPUs connected by NVLINK

Summit supercomputer: many-node, with multiple GPUs/node.

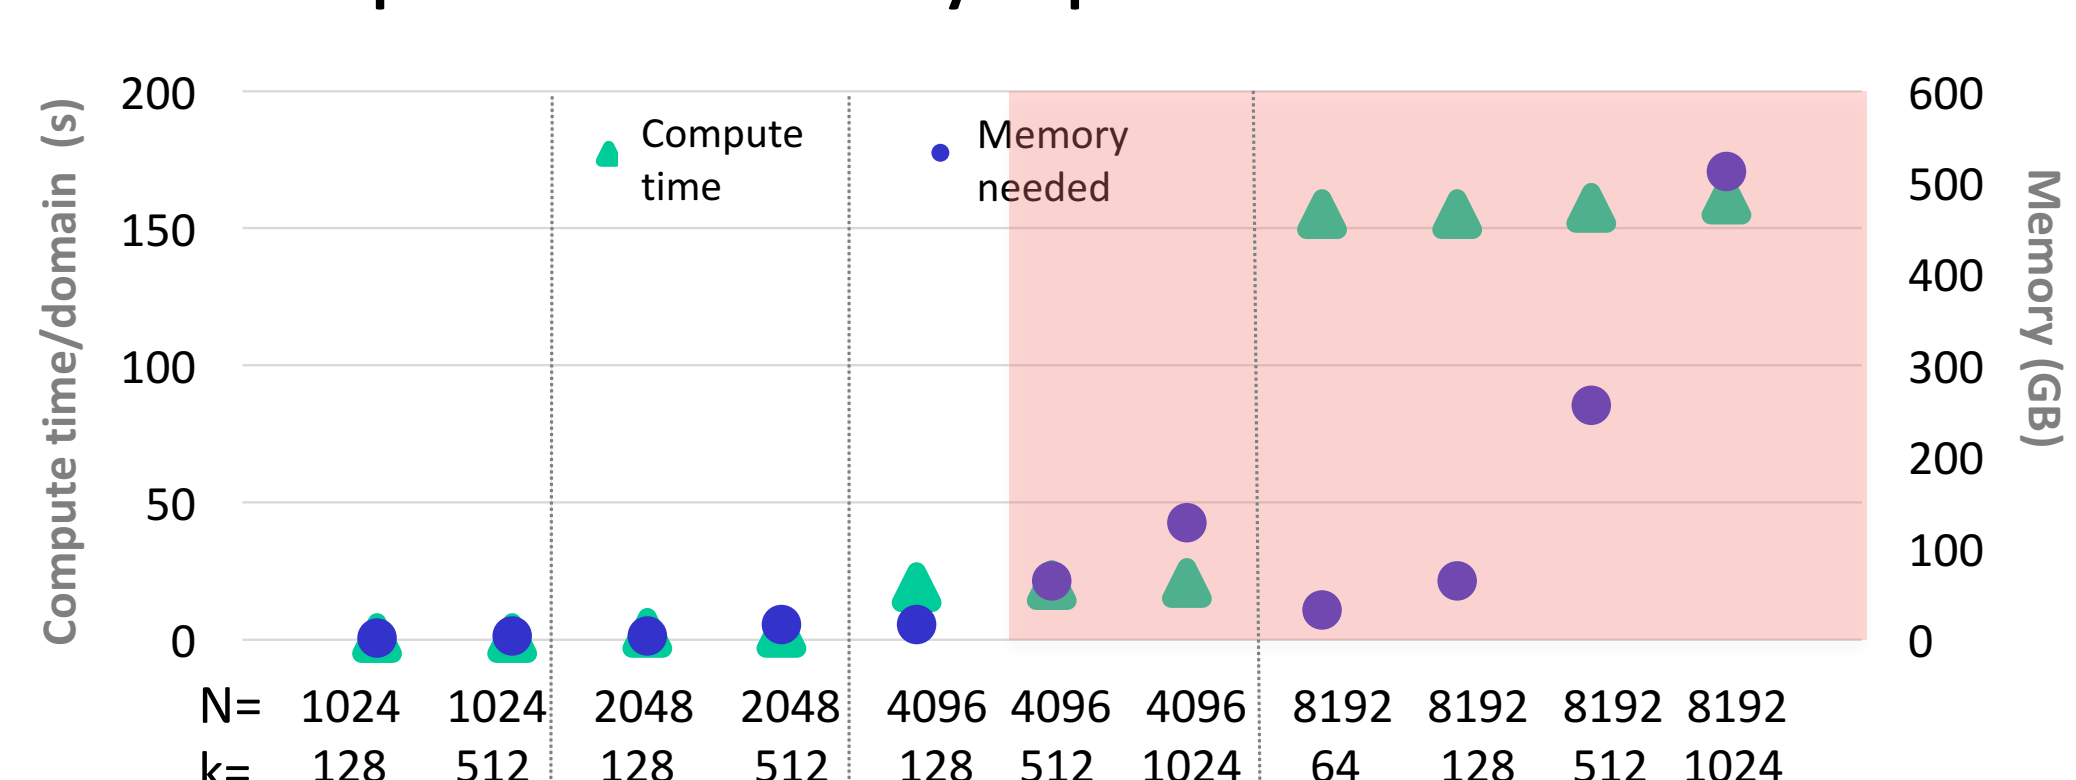
### Comparison of communication time and compute time



### Intra-node and Inter-node communication times



### GPU compute time & Memory requirement

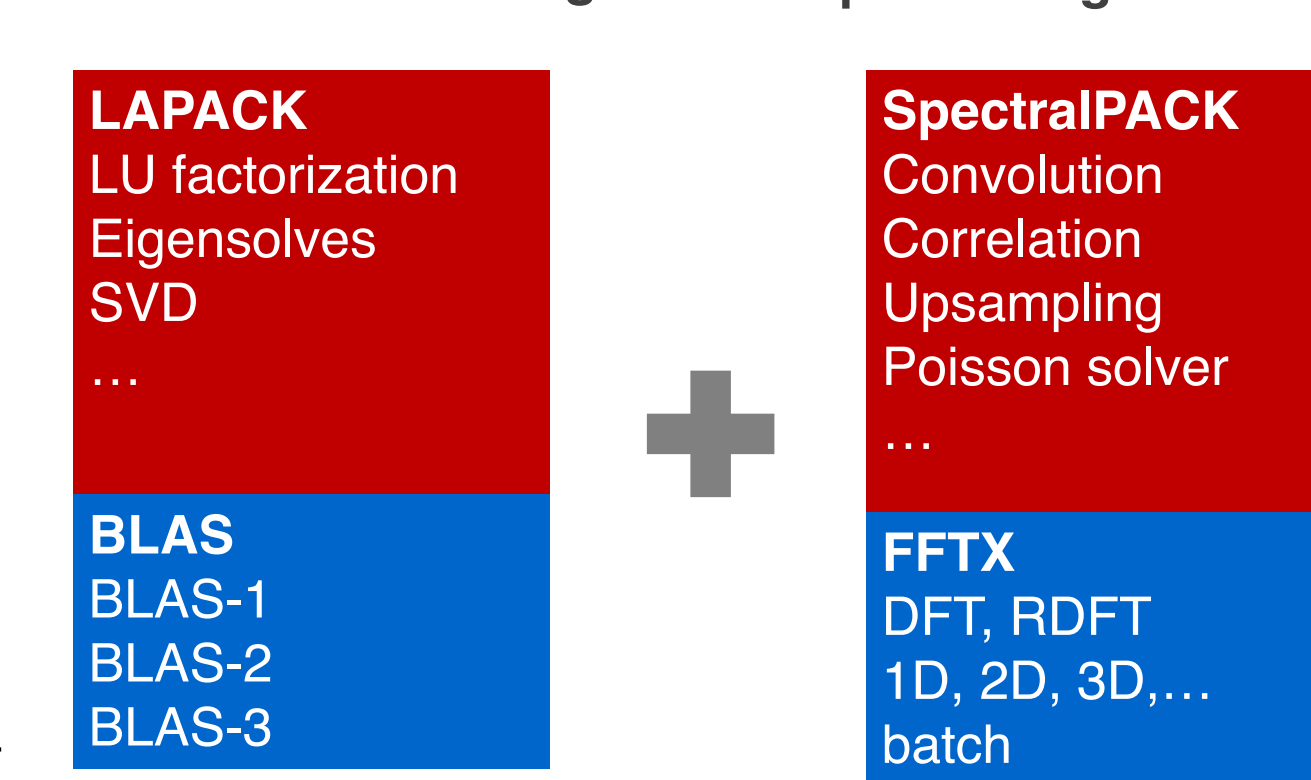


Machine	k →		N = 1024		N = 2048		N = 4096		N = 8192	
	128	512	128	512	128	512	64	128	512	1024
HPE Apollo 6500 node	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
DGX2	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
Summit (ORNL)	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗

Memory requirement exceeds available GPU memory

### Future work: FFTX and SpectralPACK

Numerical Linear Algebra Spectral Algorithms



- LAPACK for spectral algorithms
- Define FFTX as the analogue to BLAS
- Define class of numerical algorithms to be supported by SpectralPACK
- Define SpectralPACK functions

FFTX project:



### References

- [1] V. Tari, R. A. Lebensohn, R. Pokharel, T. J. Turner, P. A. Shade, J. V. Bernier, and A. D. Rollett. 2018. Validation of micro-mechanical FFT-based simulations using High Energy Di reaction Microscopy on Ti-7Al. Acta Materialia 154 (8 2018). https://doi.org/10.1016/j.actamat.2018.05.036
- [2] M. Frigo and S. G. Johnson. 2005. The Design and Implementation of FFTW3. Proc. IEEE 93, 2 (Feb 2005), 216–231. https://doi.org/10.1109/JPROC.2004.840301
- [3] A. Kulkarni, F. Franchetti, and J. Kovačević. 2018. Large-Scale Algorithm Design for Parallel FFT-based Simulations on GPUs. In 2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP). 301–305. https://doi.org/10.1109/GlobalSIP.2018.8646675
- [4] F. Franchetti, et al. 2018. FFTX and SpectralPACK: A First Look. In IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC).
- [5] P. McCorquodale, P. Colella, G. T. Bails, and S. B. Baden. 2006. A Local Corrections Algorithm for Solving Poisson's Equation in Three Dimensions. 2 (10 2006).

Images credit: NVIDIA, ORNL website, SpiralGen.