Global Computing Lab:
Performance Portable Plasma Simulations for the Exascale Era

Nigel Tan

THE UNIVERSITY OF
TENNESSEE
KNOXVILLE
BIG ORANGE. BIG IDEAS.®

Performance Portable Plasma Simulations for the Exascale Era

- Collaboration between GCLab and Los Alamos National Laboratory
- Nigel Tan (UTK), Bob Bird (LANL), Michela Taufer(UTK)

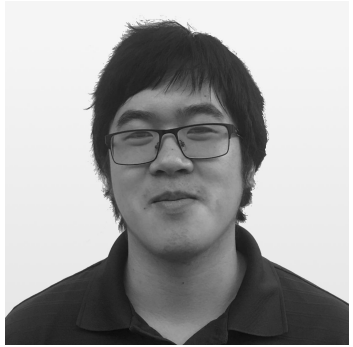

Vector Particle-In-Cell (VPIC)

- State of the art plasma simulation code solving the Vaslov-Maxwell equations
- Performance focused PIC code
  - 32 bit floating point arithmetic
  - Heavy use of vector intrinsics
- Only CPUs and KNL supported
  - Would require full rewrite to run on GPU
  - Want portable, performant, and modern code

Supercomputer	CPU	Accelerator
Summit/Sierra	Power9	Tesla V100
Trinity	Xeon	Xeon Phi Knight Landing
Perlmutter	Epyc	Tesla
Aurora	Xeon	Intel Xe
Frontier	Epyc	Radeon Instinct

*VPIC would require at least 3 ports just to run on the major US supercomputers

Challenges!

- Legacy codebase
  - Lots of macros
  - Few comments
  - Already parallelized kernels are difficult to understand
- Accuracy
  - Are our answers too wrong?

Porting Methodologies

- Data layout & movement
  - 3 Copies of data
  - Array of Structs to Struct of Arrays
  - Use subviews to reduce data movement

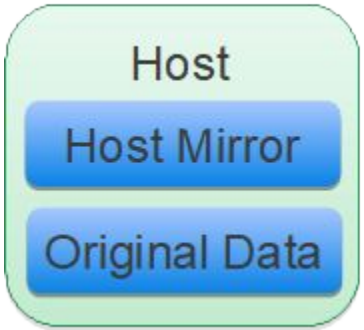

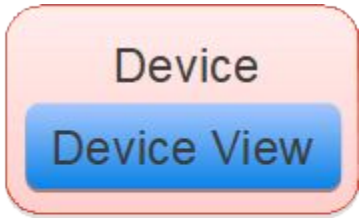

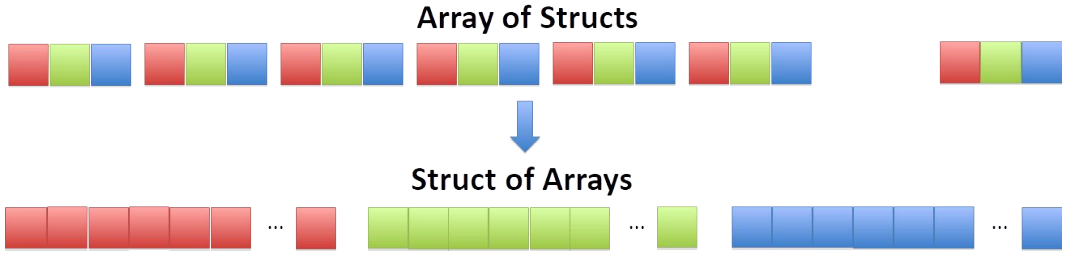

Porting Methodologies

- Communication

Original VPIC

Kokkos VPIC

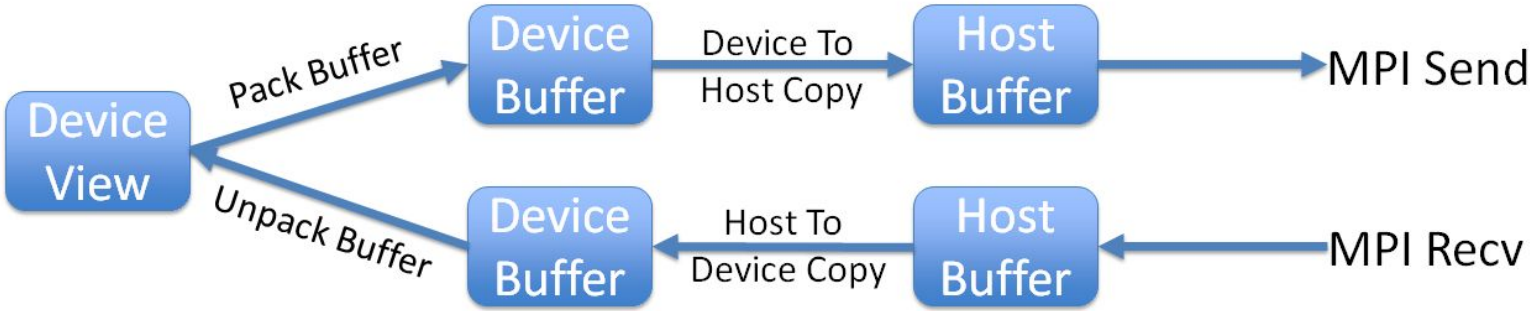

Weak Scaling

- Near ideal scaling on 2048 Summit nodes with 12288 GPUs
- At most ~10% loss using nearly half of Summit

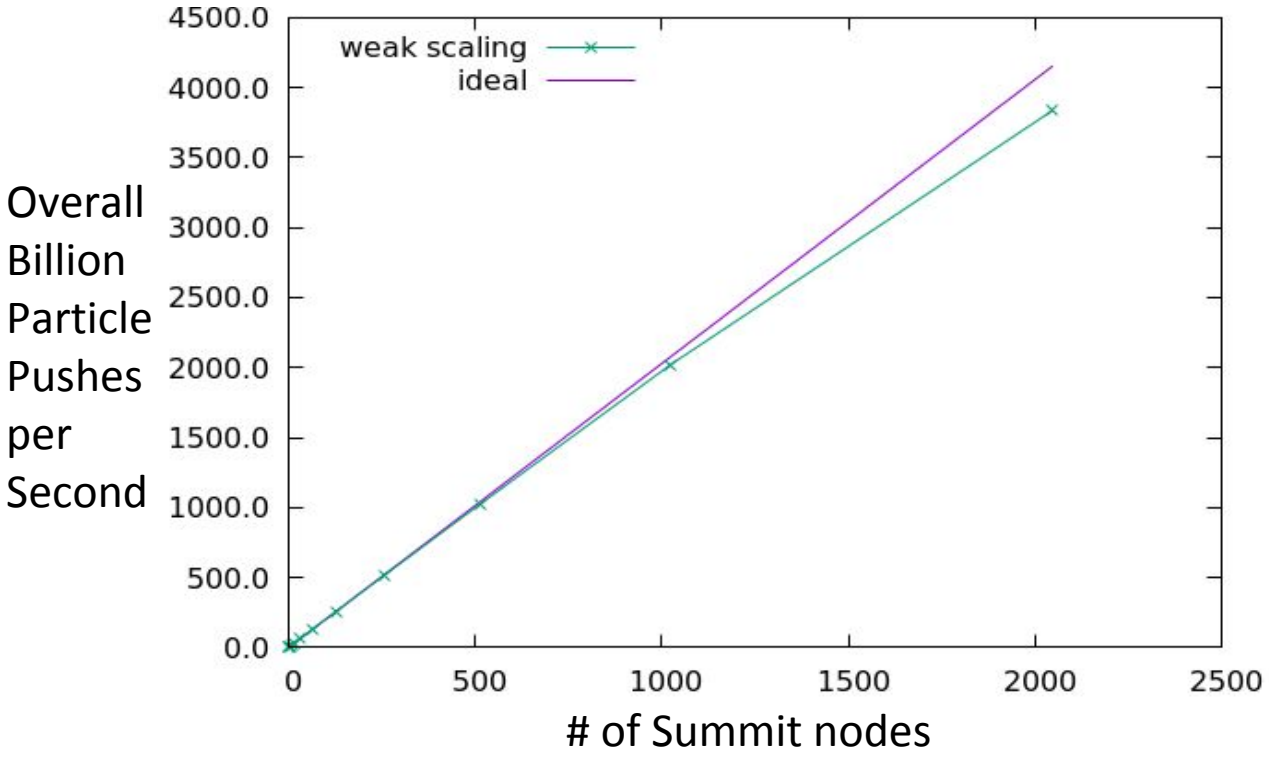

Strong Scaling

$$\Phi(a,p,H) = \begin{cases} \dfrac{|H|}{\sum_{i \in H} \dfrac{1}{e_i(a,p)}} & \text{if } i \text{ is supported } \forall i \in H \\ 0 & \text{otherwise} \end{cases}$$

1. Pennycook, Simon J. et al. “A Metric for Performance Portability.” ArXiv abs/1611.07409 (2016): n. pag.

VPIC Version	Φ(CPU) App. Eff.	Φ(GPU) App. Eff.	Φ(All) App. Eff.
Base	61.26%	0%	0%
SIMD	100%	0%	0%
Kokkos	48.89%	100%	52.74%