

APPLICATION STUDY OF GYROKINETIC PIC CODES ON INTEL KNL ARCHITECTURE



BRIAN MACKIE-MASON

TIMOTHY J. WILLIAMS

Argonne Leadership Computing
Facility

PAULIUS VELESKO

Intel Corporation

ROBERT HAGER

C-S CHANG

Princeton Plasma Physics
Laboratory



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

September 26, 2018
Hillsboro, OR

XGC TEAM

Contributing Organizations

■ DOE National Labs

- Oak Ridge National Laboratory
- Argonne National Laboratory
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory
- Princeton Plasma Physics Laboratory (home base)
- Lawrence Berkeley National Laboratory

■ Universities

- Rennsselaer Polytechnic Institute
- University of Colorado-Boulder
- University of Colorado-Denver

- National Institute for Fusion Science (Japan)
- National Fusion Research Institute (Korea)

■ Industry Partners

- NVIDIA
- Intel
- International Business Machines

XGC INTRODUCTION




U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.


Argonne 
NATIONAL LABORATORY

MOTIVATION

Fusion Energy Sciences



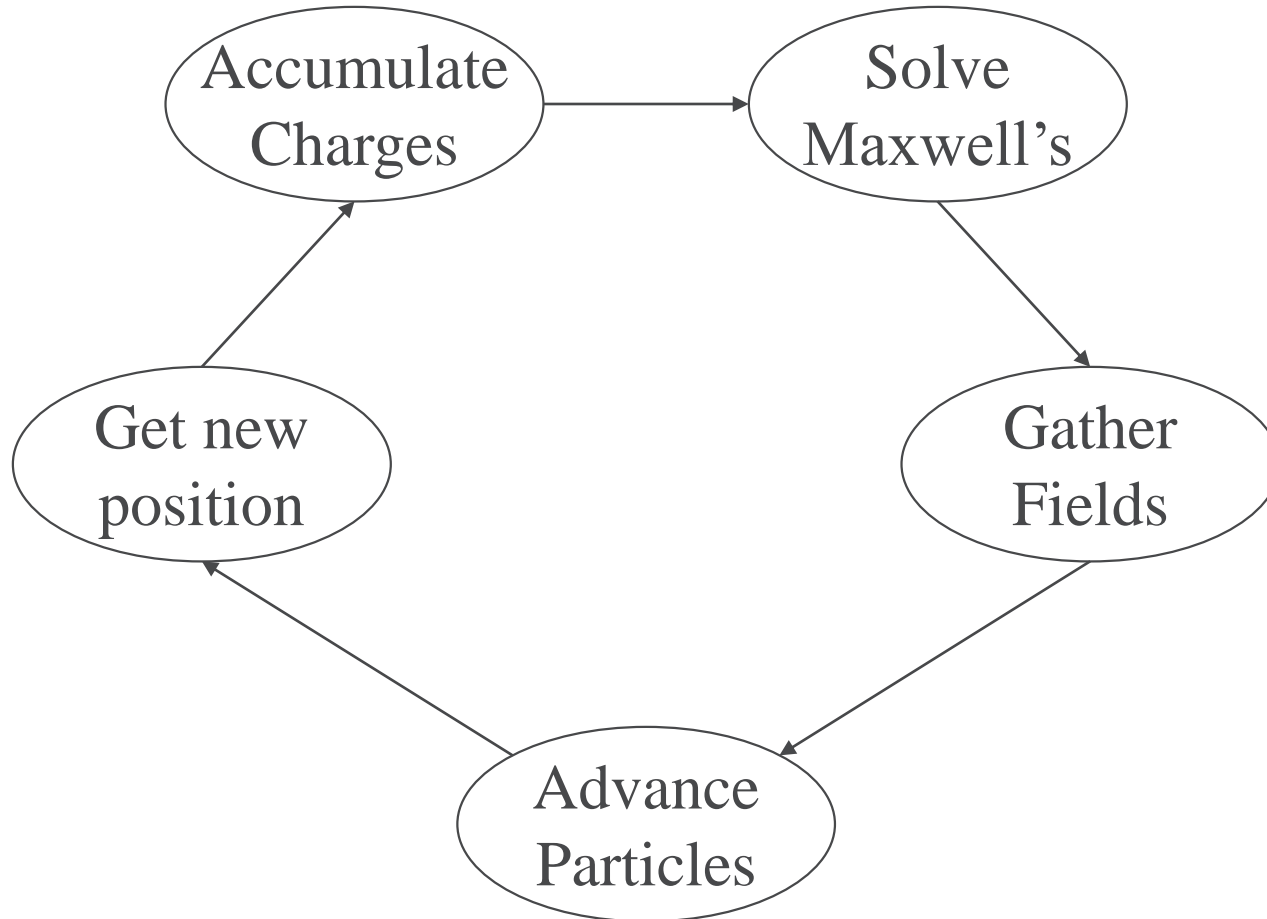
Wikimedia Commons



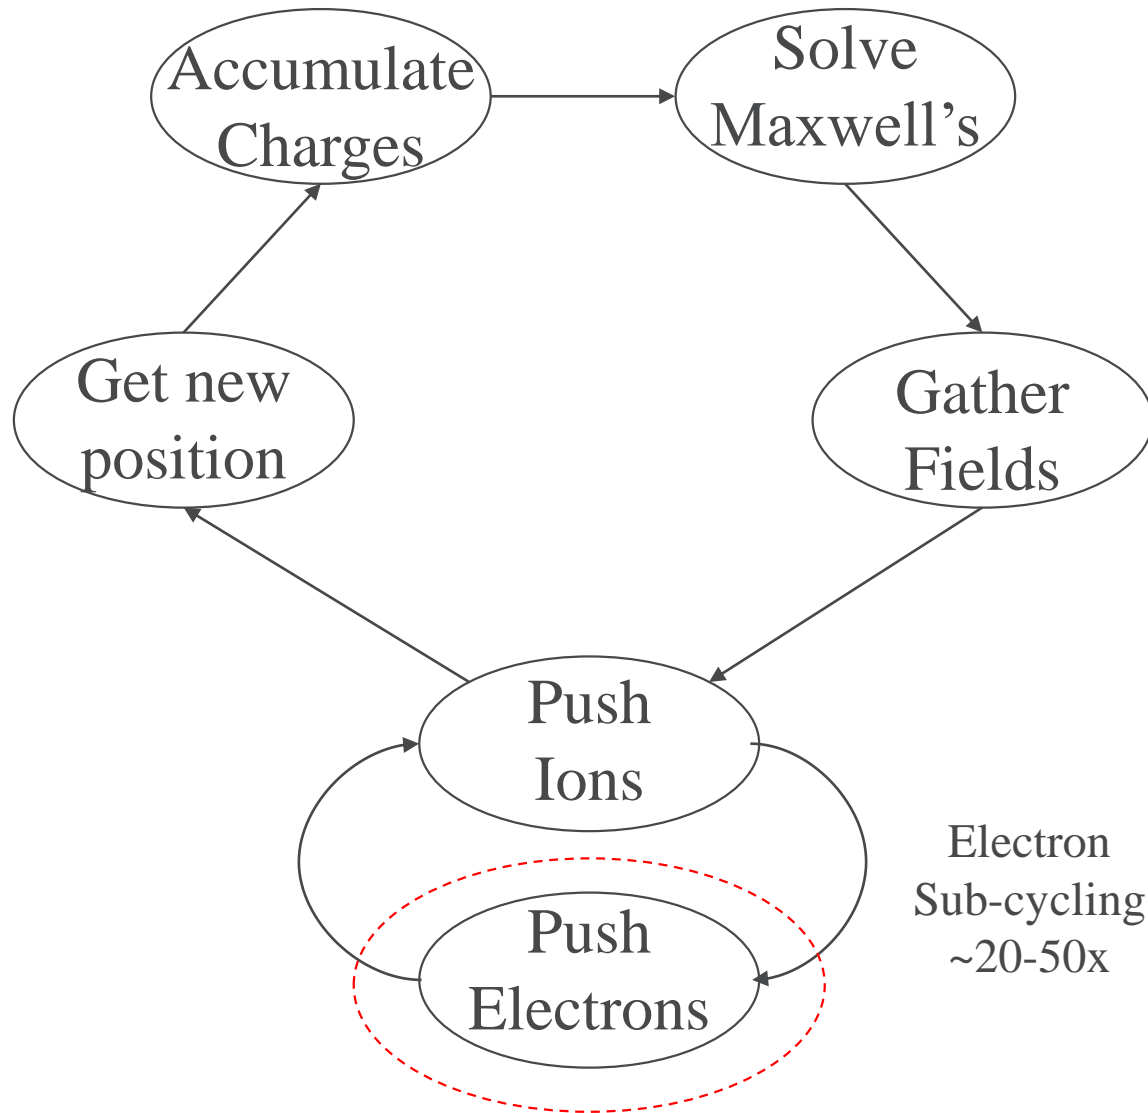
C-S Chang, et. al.

- Nuclear fusion binds nuclei and uses excess energy
- Can use water as its primary source of fuel (extract D and T)
- Turbulent flow along edge is major challenge

GYROKINETIC PIC ALGORITHMIC FLOWCHART




GYROKINETIC PIC ALGORITHMIC FLOWCHART



PIC CYCLE

Graphical look



IMPLEMENTATION DETAILS

- Fortran 90 w/MPI + OpenMP®
- Different libraries: pspline (kernel), PETSc*, ADIOS*, Intel® Math Kernel Library
- Timing done with camtimers library
- Extensive codebase, multiple version with different physics assumptions
- This work considers most computationally intensive part (electron push)

OPTIMIZATION APPROACH - SOA



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

Argonne 
NATIONAL LABORATORY

PARTICLE AND FIELD DATA STRUCTURES

AoS vs. SoA

- SoA: fld%r(veclen)
- AoS: fld(veclen)%r
- SoA: ptl%ph(6)
- AoS: ptl(6)%ph
- Expectation is that SoA will provide a significant performance improvements the data structures will be chunked into groups of veclen

AoS

```
type fld_struct
  real(8) r
  real(8) z
  real(8) phi
  real(8) br
  real(8) bz
  real(8) bphi
  real(8) dbrdr
  real(8) dbrdz
  real(8) dbrdp
  real(8) dbzdr
  real(8) dbzdz
```

SoA

```
real(8), dimension(:), allocatable :: r, z, phi !> NOT USED. R,Z coordinate variable
real(8), dimension(:), allocatable :: br, bz, bphi !> B_r, B_z, B_phi
real(8), dimension(:), allocatable :: dbrdr, dbrdz, dbrdp, dbzdr, dbzdz, &
    dbzdp, dbpdr, dbpdz, dbpdp !> dB_x / dx , x=R,Z,phi
real(8), dimension(:), allocatable :: Er, Ez, Ephi !> E_r, E_z, E_phi
real(8), dimension(:), allocatable :: Er00, Ez00 !> only for electron
real(8), dimension(:), allocatable :: dpsidr, dpsidz !> dpsi / dr, dpsi / dz
real(8), dimension(:), allocatable :: tdr, tdbz, tdbphi !> RMP variables
```

BENCHMARK PROBLEM DESCRIPTION

- Key parameters
 - Physics
 - Number of particles: 50k
 - No collisions/impurities
 - How many sub-cycles? 20
 - How much particle sorting? Once
 - Computational
 - How many nodes? One
 - How many MPI ranks per node? 16
 - How many threads per rank? 16
 - Cache-quad memory mode
- 4 simulations launched
- Camtimer libraries used to collect data

OPTIMIZATION RESULTS



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

Argonne 
NATIONAL LABORATORY

AOS->SOA OPTIMIZATION RESULTS

By key subroutine

Category	FLD	PTL	FLD	PTL	FLD	PTL	FLD	PTL
	SOA	SOA	SOA	AOS	AOS	SOA	AOS	AOS
Sort	0.7454 ± 0.0074		0.7387 ± 0.0095		0.7344 ± 0.022		0.7356 ± 0.0217	
Scatter	0.2710 ± 0.037		0.2726 ± 0.071		0.2913 ± 0.0541		0.2986 ± 0.0645	
Gather	0.3610 ± 0.0095		0.3624 ± 0.0064		0.3593 ± 0.0104		0.3606 ± 0.0099	
Time-step	9.9207 ± 0.036		9.9258 ± 0.0423		10.6269 ± 0.0608		10.5743 ± 0.0359	
Total	11.2982 ± 0.090		11.2994 ± 0.129		12.0119 ± 0.147		11.9690 ± 0.132	

Key Results:

1. SoA is ~5-8% faster
2. Almost all of speed up occurs in time-stepping
3. The switch from AoS->SoA *appears* to be more significant on the fields

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system. For more complete information [click here](#).

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Configuration Details: These optimizations were performed on Intel® Xeon Phi™ x200 using a Cray XC40 system with Dragonfly Aries interconnect.

Performance results are based on testing as of Sep. 12, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Note: Date doesn't need to be red and bold; it's that way to stick out so people know to change it.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Intel, the Intel logo, Xeon, and others are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© Intel Corporation.

AOS->SOA OPTIMIZATION RESULTS

By physical operation

	FLD/PTL	FLD/PTL	FLD/PTL	FLD/PTL
Category	SoA/SoA	AoS/AoS	SoA/AoS	AoS/SoA
Field	6.890	7.549	6.961	7.534
Particle	0.645	0.661	0.631	0.664
Field-Particle	1.063	1.117	1.091	1.116

FLD SoA		PTL SoA	
PTL AoS->SoA	-2.13%	FLD AoS-> SoA	8.54%
FLD AoS		PTL AoS	
PTL AoS->SoA	0.43%	FLD AoS->SoA	7.79%

Key Results:

1. Majority of time touches field quantities
2. Little to no impact, sometimes negative, when particle flag is changed
3. Significant impact when field flag is changed

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system. For more complete information [click here](#).

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Configuration Details: These optimizations were performed on Intel® Xeon Phi™ x200 using a Cray XC40 system with Dragonfly Aries interconnect.

Performance results are based on testing as of Sep. 12, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Note: Date doesn't need to be red and bold; it's that way to stick out so people know to change it.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Intel, the Intel logo, Xeon, and others are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© Intel Corporation.


WHY?

Vectorized Snippets & Roofline Analysis

- Snippet from `derivs_with_e_elec_vec`:

```
do iv = 1, veclen
#ifdef FIELD_SOA
  yprime_r(iv) = D(iv) * ( drift_on * (fld%bz(iv) * Fp(iv) - fld%Bphi(iv) * Fz(iv)) * over_B2(iv) &
    + cmrho(iv) * (fld%br(iv) + fld%tdbr(iv)) &
    + drift_on * cmrho2(iv) * (fld%dbzdp(iv) * inv_r(iv) - fld%dbpdz(iv)) ) &
  yprime_z(iv) = D(iv) * ( drift_on * (fld%bphi(iv) * fr(iv) - fld%br(iv) * fp(iv)) * over_B2(iv) &
    + cmrho(iv) * (fld%bz(iv) + fld%tdbz(iv)) &
    + drift_on * cmrho2(iv) * (fld%bphi(iv) * inv_r(iv) + fld%dbpdr(iv) - fld%dbdrp(iv) * inv_r(iv)) ) &
  yprime_p(iv) = D(iv) * ( drift_on * (fld%br(iv) * fz(iv) - fld%bz(iv) * fr(iv)) * over_B2(iv) &
    + cmrho(iv) * (fld%bphi(iv) + fld%tdbphi(iv)) &
    + drift_on * cmrho2(iv) * ( fld%dbrdz(iv) - fld%dbzdr(iv)) ) * inv_r(iv)
```

- Intel® Advisor Roofline analysis:




NEXT STEPS FOR INTEL® XEON PHI™ X200 (KNIGHTS LANDING)

- More advanced data structure tiling (i.e. SoSoA)
- Tuning of additional openMP parameters
- Scaling up particle density
- Vector implementation of index searching
- Context of electron push kernel optimization within whole code
- More work on the potential derivatives

NOTICES AND DISCLAIMERS

THIS SLIDE MUST BE USED WITH ANY SLIDES REMOVED FROM THIS PRESENTATION

- Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.
- No computer system can be absolutely secure.
- Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.
- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/benchmarks>.
- Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.
- Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
- Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.
- Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.
- Intel, the Intel logo, and Intel Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.
- *Other names and brands may be claimed as property of others.
- © 2018 Intel Corporation.



THANK YOU!
QUESTIONS?



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

Argonne 
NATIONAL LABORATORY

BACKUP SLIDES



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

Argonne 
NATIONAL LABORATORY

MEMORY ALLOCATION

Static vs. Dynamic

- Static:
 - `real (kind=8), dimension(vecLEN) :: Er,Ez,Ep`
- Dynamic:
 - `real (kind=8), dimension(:), allocatable :: Er,Ep,Ez`
 - call `efield_gk_vec_init(vecLEN)`
- Expectation is that static is a small improvement

MEMORY ALLOCATION RESULTS

Table showing no performance improvement

Call-stack top				Call-stack bottom		
No alignment				No alignment		
Sub-cycles	Static [s]	Dynamic [s]		Sub-cycles	Static [s]	Dynamic [s]
20	0.598	0.509		20	0.603	0.515
40	0.589	0.485		40	0.592	0.495
60	0.593	0.492		60	0.607	0.485
80	0.599	0.487		80	0.607	0.492
100	0.605	0.488		100	0.611	0.486
Call-stack top				Call-stack bottom		
Alignment				Alignment		
Sub-cycles	Static [s]	Dynamic [s]		Sub-cycles	Static [s]	Dynamic [s]
20	0.583	0.531		20	0.589	0.529
40	0.595	0.505		40	0.590	0.505
60	0.592	0.508		60	0.593	0.512
80	0.598	0.512		80	0.600	0.501
100	0.605	0.501		100	0.607	0.503

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system. For more complete information [click here](#).

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Configuration Details: These optimizations were performed on Intel® Xeon Phi™ x200 using a Cray XC40 system with Dragonfly Aries interconnect.

Performance results are based on testing as of Sep. 12, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Note: Date doesn't need to be red and bold; it's that way to stick out so people know to change it.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Intel, the Intel logo, Xeon, and others are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© Intel Corporation.