NVIDIA DEVELOPMENTS FOR EARTH SYSTEM MODELING

Stan Posey, HPC Program Manager, ESM Domain, NVIDIA (HQ), Santa Clara, CA, USA

NVIDIA

TOPICS OF DISCUSSION

- NVIDIA HPC AND ESM UPDATE
- GPU TECHNOLOGY FOR ESM

NVIDIA Company Introduction

HQ in Santa Clara, CA, USA

FY16 Revenue ~$5B USD

Employs ~9,500 in 20 Countries

NVIDIA Core Markets

GPU Computing

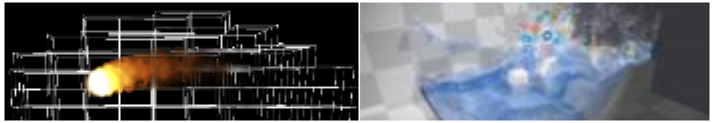

Computer Graphics

Artificial Intelligence

NVIDIA Growth from Advancement of HPC and AI

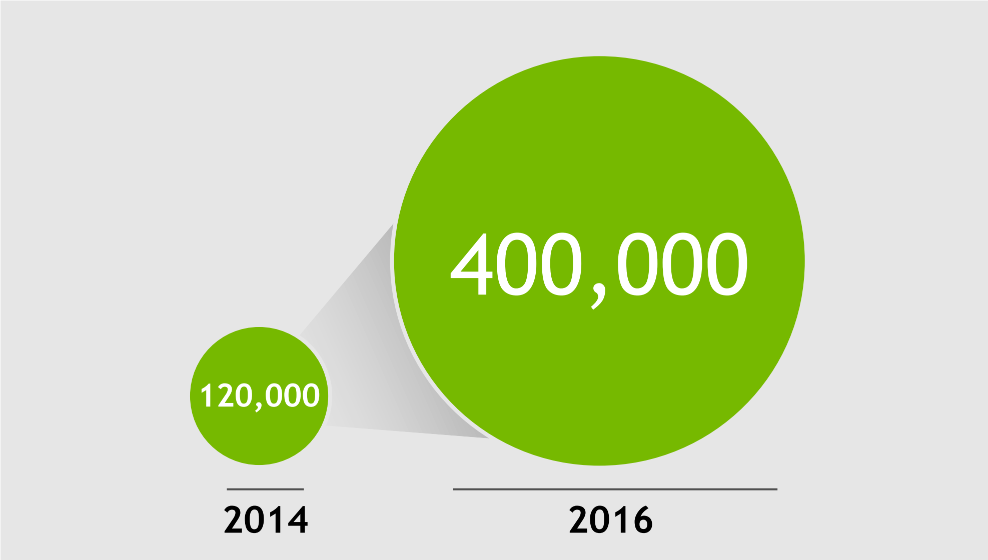

3X HPC Developers

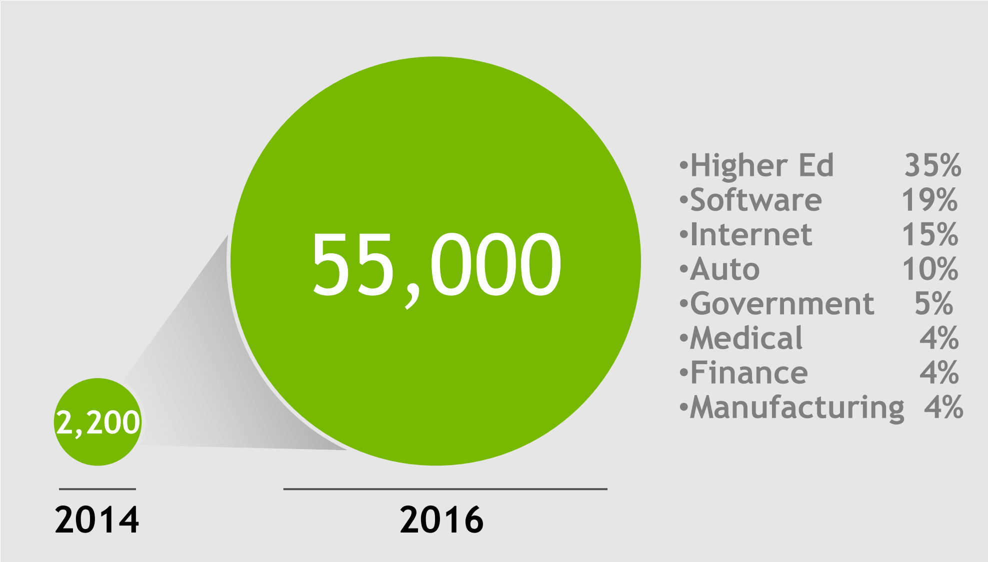

25x Deep Learning Developers

GPUs Power World's Leading Data Centers for HPC and AI:

amazon web services · facebook · Google · IBM · Lawrence Livermore National Laboratory · Microsoft · Oak Ridge National Laboratory · twitter

HPC Advancements: US DOE Supercomputers

Pre-Exascale Systems Powered by NVIDIA Tesla

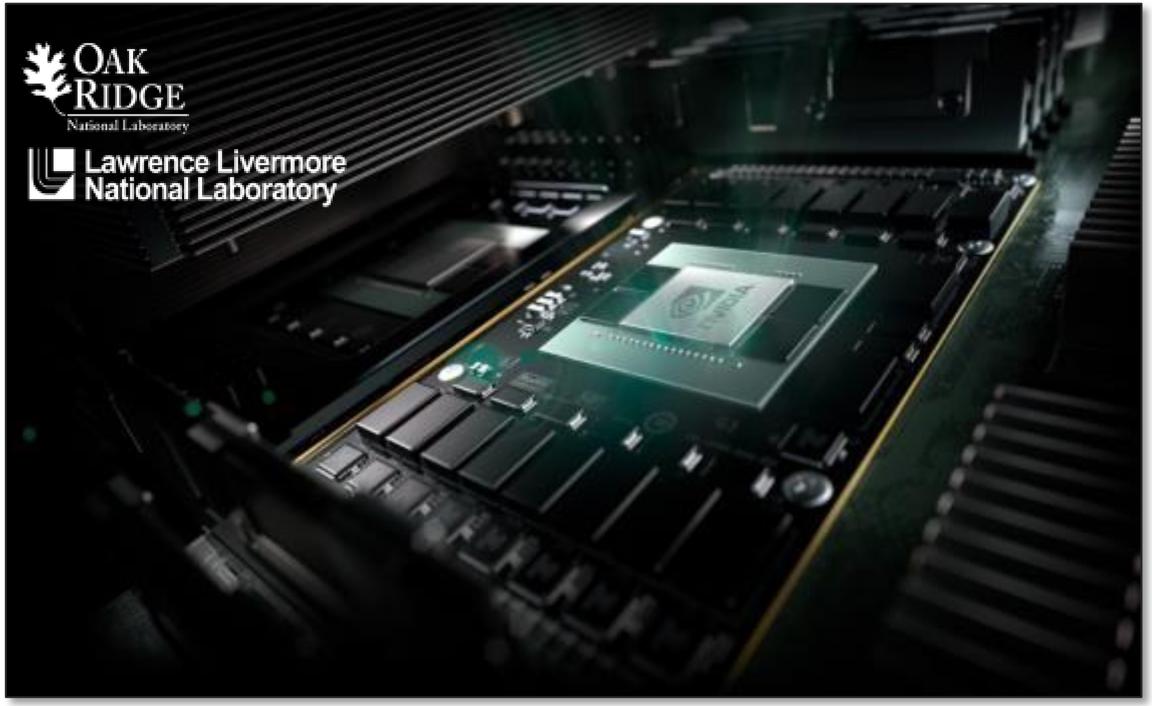

Summit & Sierra Supercomputers

150-300 PFLOPS Peak

Featuring NVIDIA Volta GPU

NVLink High Speed Interconnect

40 TFLOPS per Node, >3,400 Nodes

2018

RECENT NWP DEPLOYMENTS WITH GPU SYSTEMS

Cray CS-Storm, 192 x K80, 8 GPUs per node

MeteoSwiss Deploys World's 1st Operational NWP on GPUs

2-3x higher resolution for daily forecasts

14x more simulation with ensemble approach for medium-range forecasts

NWP Model: COSMO

Cray CS-Storm, 760 x P100, 8 GPUs per node

NOAA To Improve NWP and Climate Research with GPUs

Develop global model with 3km resolution, five-fold increase from today's resolution

NWP Model: FV3/GFS (also climate research)

Deep Learning Growth in NWP and Observations

Monitoring Effects of Carbon
and Greenhouse Gas Emissions

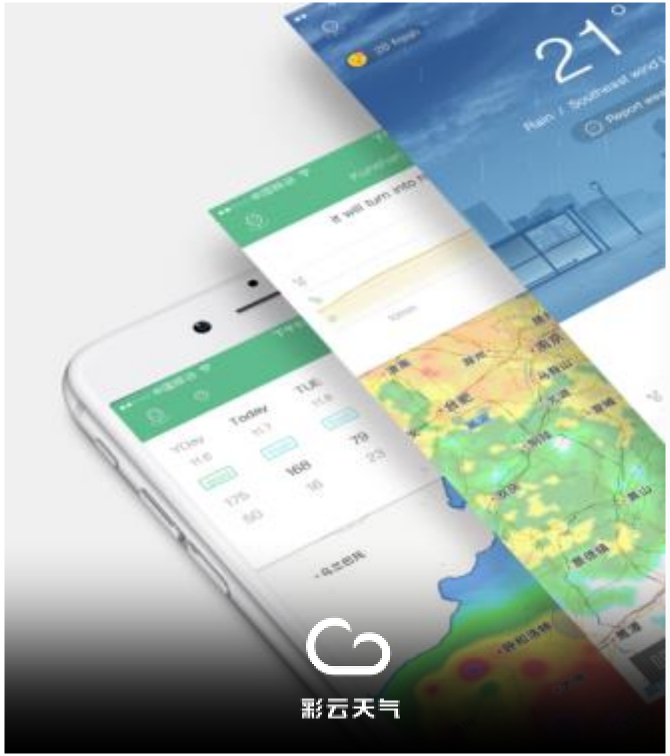

Colorful Clouds (CN)
AI Weather Forecasting

Yandex Meteum (RU)
AI Hyperlocal Nowasting

NVIDIA Volta - Latest GPU Announced at GTC 2017

TESLA V100 SPECIFICATIONS

Compute	7.5 TF DP · 15 TF SP · 120 TF DL
Memory	HBM2: **900 GB/s** · 16 GB
Interconnect	NVLink (up to 300 GB/s) + PCIe Gen3 (up to 32 GB/s)
Availability	DGX-1: Q3 2017; OEM : Q4 2017 DELL EMC · Hewlett Packard Enterprise · IBM · SUPERMICRO

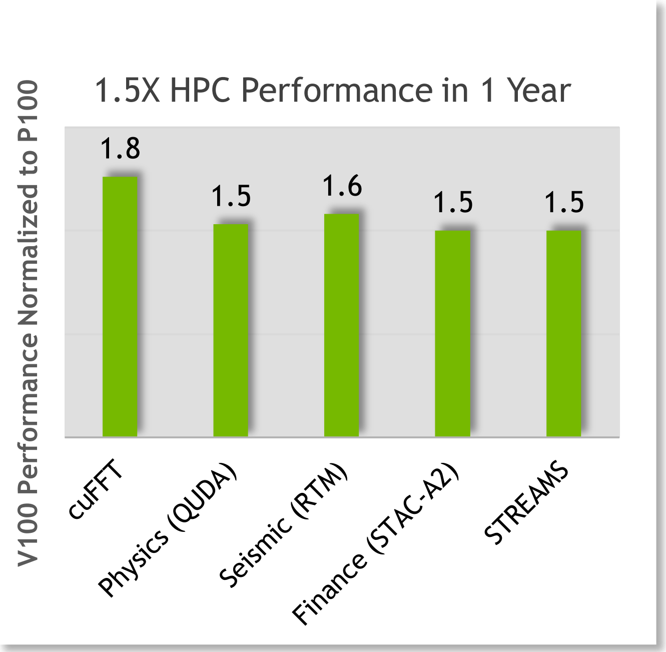

System Config Info: 2X Xeon E5-2690 v4, 2.6GHz, w/ 1X Tesla P100 or V100. V100 measured on pre-production hardware.

Update on DOE Pre-Exascale CORAL Systems

- **LLNL Sierra 150PF in 2018**
- **ORNL Summit 200PF in 2018**
  - **CAAR support from IBM and NVIDIA**

TITAN VS SUMMIT

Compute System Comparison

FEATURE	TITAN	SUMMIT	
Application Performance	Baseline	5-10x Titan	
Number of Nodes	18,688	~4,600	~1/4x
Node performance	1.4 TF	> 40 TF	~29x
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + HBM	
NV memory per Node	0	1600 GB	
Total System Memory	710 TB	>10 PB DDR4 + HBM + Non-volatile	
System Interconnect (node injection bandwidth)	Gemini (6.4 GB/s)	Dual Rail EDR-IB (23 GB/s)	
Interconnect Topology	3d Torus	Non-blocking Fat Tree	
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™	27,600 GPUs
File System	32 PB, 1 TB/s, Lustre©	250 PB, 2.5 TB/s, GPFS™	
Peak power consumption	9 MW	15 MW	~1.7x

ACME Fully GPU-Accelerated Coupled Climate Model

- ACME: Accelerated Climate Modeling for Energy
  - First fully accelerated climate model (GPU and MIC)
  - Consolidation of DOE ESM projects from 7 into 1
  - DOE Labs: Argonne, LANL, LBL, LLNL, ORNL, PNNL, Sandia
  - Towards NH global Atm 12 km, Ocn 15 km, 80 year

ACME Accelerated Climate Modeling for Energy

- ACME component models and GPU progress
  - Atm – ACME-Atmosphere (NCAR CAM-SE fork)
    - Dycore now in trunk, CAM physics started with OpenACC
  - Ocn – MPAS-O (LANL)
    - LANL team at ORNL OpenACC Hackathon during 2015
  - Others – published OpenACC progress
    - Sea-Ice – ACME-CICE (LANL)
    - Land – CLM (ORNL, NCAR)
    - Cloud Superparameterization – SAM (SBU, CSU)
    - Land-Ice – PISCEES (Multi-lab – LLNL, Sandia)

PROGRAMMING STRATEGIES FOR GPU ACCELERATION

Applications

GPU Libraries	OpenACC Directives	Programming in CUDA
Provides Fast "Drop-In" Acceleration	GPU-acceleration in Standard Language (Fortran, C, C++)	Maximum Flexibility with GPU Architecture and Software Features

Increasing Development Effort →

NOTE: Many application developments include a combination of these strategies

Examples

GPU Libraries:
- IFS (FFT, DGEMM)
- COSMO (Tridiag Solve)

OpenACC Directives:
- FV3
- MPAS
- IFS
- ICON
- COSMO (Physics)
- WRF -NCAR
- ACME
- CAM-SE
- NICAM
- ICON
- NEMO
- UM/Gungho

Programming in CUDA:
- COSMO (Dycore)
- NUMA
- ICON
- WRF -TQI
- NICAM (Dycore)

TOPICS OF DISCUSSION

- NVIDIA HPC AND ESM UPDATE
- **GPU TECHNOLOGY FOR NWP**
- DISCUSSION AND Q&A

Select NVIDIA ESM Highlights During 2016

- **WW ESM growth in GPU funded-development:** NOAA, NCAR, ECMWF, DOE, DoD
- **New ESM-driven GPU systems (K80/P100):** NOAA, ECMWF, CSCS, NIES, Others
- **First ever GPU-based operational NWP:** MeteoSwiss with COSMO (Mar 2016)
  - ~4x speedup with ~5x less energy vs. conventional CPU-only; New COSMO evaluations by Met's in DE, RU, IT
- **DOE climate model ACME-Atm v1 production on TITAN using PGI OpenACC**
- **New NCAR collaboration launched with GPU Hands-on Workshops**
  - https://www2.cisl.ucar.edu/news/summertime-hackathon (focus on GPU development of MPAS-A)
- **ECWMF selected NVIDIA as partner for ESCAPE Exascale weather project**
  - https://software.ecmwf.int/wiki/display/OPTR/NVIDIA+Basic+GPU+Training+with+emphasis+on+Fortran+and+OpenACC
- **NEMO Systems Team invited NVIDIA as member of HPC working group**
  - Following successful NVIDIA OpenACC scalability of NEMO for ORCA025 configuration (NEMO UGM 2014)
- **New ESM opportunities developing in new solution focus areas**
  - DL in climate and weather; BI for Ag and Actuary; Air quality monitoring (CN, KR); Commercial WRF start-up TQI
- **ESM model teams have consistent participation in GPU Hackathons**
  - DOE/ACME, NRL/COAMPS, MPI-M/ECHAM6, ODU/FVCOM, NRL/HYCOM, NOAA GFDL radiation models

GPU Funded-Development Growing for ESM

HPC Programs with Funding Specifically Targeted for GPU Development of Various ESMs

- SENA – NOAA funding for accelerator development of WRF, NGGPS (FV3), GFDL climate, NMMB
- ESCAPE – ECMWF-led EUC Horizon 2020 program for IFS; NVIDIA 1 of 11 funded partners
- ACME – US DOE accelerated climate model: CAM-SE, MPAS-O, CICE, CLM, SAM, PISCEES, others
- AIMES – Govt's from DE, FR, and JP for HPC (and GPU) developments of ICON, DYNAMICO, NICAM
- SIParCS – NCAR academia funding for HPC (and GPU) developments of MPAS, CESM, DART, Fields
- AOLI – US DoD accelerator development of operational models HYCOM, NUMA, CICE, RRTMG
- GridTools – Swiss gov funding MCH/CSCS/ETH for accelerator-based DSL in COSMO, ICON, others

NOTE: Follow each program LINK for details; Programs listed from top-down in rough order of newest to oldest start date

GPUs Deployed for ESM and NWP Modeling (Apr 17)

Organization	Location	GPUs	System	
DOE ORNL	Oak Ridge, TN, US	Volta – 200 PF	IBM Power9 – *Summit*	New P100
DOE LLNL	Livermore, CA, US	Volta – 150 PF	IBM Power9 – *Sierra*	
CSCS	Lugano, CH	P100 – 4,500	Cray XC-40	
NOAA	Fairmont, WV, US	P100 – 760	Cray CS-Storm	
NIES	Tsukuba, JP	P100 – 266	HP/SGI ICE	
DoD MHPCC	Maui, HI, US	P100 – 64	IBM Power8	
MCH/CSCS	Lugano, CH	K80 – 192	Cray CS-Storm – *Piz Kesch*	K80/K40
ECMWF	Reading, UK	K80 – 68	Dell	
DoD AFRL	Dayton, OH, US	K40 – 356	HP/SGI ICE X – *Thunder*	
NASA GSFC	Greenbelt, MD, US	K40 – 36	HP/SGI, IBM x86	
NASA ARC	Mtn View, CA, US	K40 – 64	HP/SGI – *Pleiades*	
DOE ORNL	Oak Ridge, TN, US	K20X – 18,688	Cray – *TITAN*	K20X
CSCS	Lugano, CH	K20X – 5,272	Cray – *Piz Daint, Piz Dora*	
TiTech	Tokyo, JP	K20X – 4,224	NEC/HP – *TSUBAME 2.0*	
NCSA	Urb-Ch, IL, US	K20X – 3,072	Cray – *Blue Waters*	
NCAR	Cheyenne, WY, US	K20X – 30	IBM x86 – *Yellowstone*	

GPU Developments for Atmospheric Models

	Organization	Location	Model	GPU Approach
Global 	ORNL, SNL	US	ACME-Atmosphere	OpenACC (migration from CUDA-F)
	ORNL, PNNL, UCI, SBU	US	SAM	OpenACC
	NCAR; THU	US	CAM-SE	OpenACC (migration from CUDA-F)
	NCAR, KISTI	US	MPAS-A	OpenACC
	NOAA GFDL, ESRL	US	FV3/GFS	OpenACC
	NASA GSFC	US	GEOS-5	OpenACC (migration from CUDA-F)
	US Naval Res Lab, NPS	US	NUMA/NEPTUNE	DSL – dycore only
	ECMWF	UK	IFS	Libs + OpenACC
	MetOffice , STFC	UK	UM/GungHo	OpenACC back-end to PSyKAI
	DWD, MPI-M, CSCS	DE, CH	ICON	DSL – dycore, OpenACC – physics
	JAMSTEC, UT, RIKEN	JP	NICAM	OpenACC
Regional 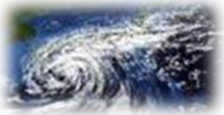	NCAR; TQI/SSEC	US	WRF-ARW	(i) OpenACC, (ii) CUDA
	DWD, MCH, CSCS	DE, CH	COSMO	DSL – dycore, OpenACC – physics
	Bull, MF	FR	HARMONIE	OpenACC
	TiTech, JMA	JP	ASUCA	Hybrid-Fortran, OpenACC

GPU Developments for Ocean Models

Model	Organization(s)	GPU Development Status
MPAS-O	DOE LANL (US)	Ocean component model for ACME, OpenACC GPU development committed
POP	DOE LANL, NCAR (US)	No GPU plans at LANL; Paper published on GPU development by NUDT (CN)
HYCOM	NOPP, DoD NRL, NOAA NCEP, FSU, Others (US)	Funded OpenACC GPU development through AOLI; Participation in 2015 Hackathon
MOM	NOAA GFDL (US)	Funded GPU development through SENA program, initial GPU profiling underway
NEMO	CNRS, Mercator, UKMO, NERC, CMCC, INGV (EU, UK)	NVIDIA-led OpenACC development results at GTC 2013 and 2014 NEMO Users meeting; HPC working group (NVIDIA as a member) will evaluate GPU potential
GOcean	UKMO, STFC (UK)	Use of UKMO GungHo framework, GPU development led by NVIDIA Devtech
POM	Princeton, ODU (US)	GPU development led by Tsinghua University (CN), results in 2015 publication
MITgcm	MIT, NASA (US)	Early GPU investigations by MIT (C. Hill); nothing currently in development
LICOM	CAS-IAP/LASG (CN)	GPU development ongoing, led by NVIDIA Devtech
ROMS	UCLA, Rutgers (US)	Early CUDA GPU investigations by Cal Poly (C. Lupo), results at GTC 2013
ADCIRC	UNC, UND, Others (US)	GPU investigations ongoing for DGM formulations, potential use of HPCX
FVCOM	UMass (US)	Participation in 2015 Hackathon, evaluation of OpenACC and AmgX GPU solver library

PGI® COMPILERS & TOOLS

PGI COMPILERS & TOOLS UPDATE

PGI Compilers for Heterogeneous Supercomputing, June 2017

NVIDIA.

AGENDA

PGI Overview

OpenACC Applications & Performance

OpenPOWER & Volta V100

OpenACC and CUDA Unified Memory

PGI — THE NVIDIA HPC SDK

- Fortran, C & C++ Compilers
  - Optimizing, SIMD Vectorizing, OpenMP
- Accelerated Computing Features
  - OpenACC Directives, CUDA Fortran
- Multi-Platform Solution
  - X86-64 and OpenPOWER Multicore CPUs
  - NVIDIA Tesla GPUs
  - Supported on Linux, macOS, Windows
- MPI/OpenMP/OpenACC Tools
  - Debugger
  - Performance Profiler
  - Interoperable with DDT, TotalView

PGI®

The Compilers & Tools
for Supercomputing

OPENACC IS FOR MULTICORE, MANYCORE & GPUS

```
 98 !$ACC KERNELS
 99 !$ACC LOOP INDEPENDENT
100      DO k=y_min-depth,y_max+depth
101 !$ACC LOOP INDEPENDENT
102        DO j=1,depth
103          density0(x_min-j,k)=left_density0(left_xmax+1-j,k)
104        ENDDO
105      ENDDO
106 !$ACC END KERNELS
```

CPU

```
% pgfortran -ta=multicore -fast -Minfo=acc -c \
  update_tile_halo_kernel.f90
  . . .
  100, Loop is parallelizable
       Generating Multicore code
       100, !$acc loop gang
  102, Loop is parallelizable
```

GPU

```
% pgfortran -ta=tesla,cc35,cc60 -fast -Minfo=acc -c \
  update_tile_halo_kernel.f90
  . . .
  100, Loop is parallelizable
  102, Loop is parallelizable
       Accelerator kernel generated
       Generating Tesla code
       100, !$acc loop gang, vector(4) ! blockidx%y threadidx%y
       102, !$acc loop gang, vector(32) ! blockidx%x threadidx%x
```

CLOVERLEAF V1.3

AWE Hydrodynamics mini-app

6500+ lines, !$acc kernels

OpenACC or OpenMP

Source on GitHub

http://uk-mac.github.io/CloverLeaf

OPENACC PERFORMANCE PORTABILITY

AWE Hydrodynamics CloverLeaf mini-App, bm32 data set

http://uk-mac.github.io/CloverLeaf

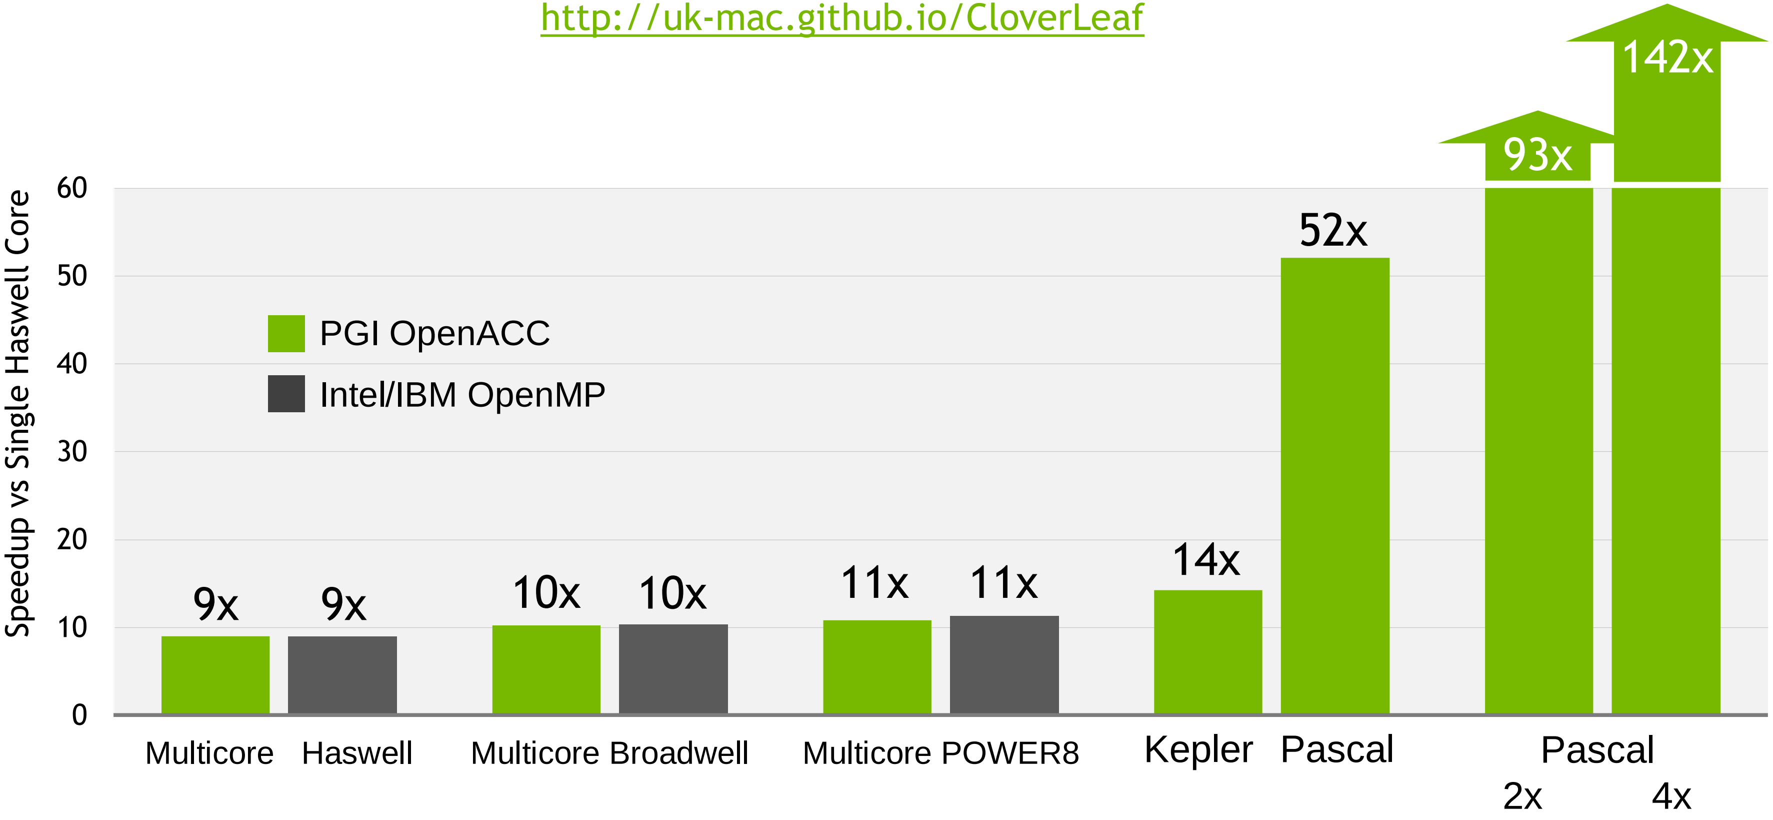

Systems: Haswell: 2x16 core Haswell server, four K80s, CentOS 7.2 (perf-hsw10), Broadwell: 2x20 core Broadwell server, eight P100s (dgx1-prd-01), Minsky: POWER8+NVLINK, four P100s, RHEL 7.3 (gsn1.
Compilers: Intel 17.0.1, IBM XL 13.1.3, PGI 16.10.
Benchmark: CloverLeaf v1.3 downloaded from http://uk-mac.github.io/CloverLeaf the week of November 7 2016; CloverlLeaf_Serial; CloverLeaf_ref (MPI+OpenMP); CloverLeaf_OpenACC (MPI+OpenACC)
Data compiled by PGI November 2016.

PGI COMPILERS & TOOLS FOR EVERYONE

PGI Community Edition Now Available

	FREE PGI Community Edition	PGI Professional Edition	PGI Enterprise Edition
PROGRAMMING MODELS OpenACC, CUDA Fortran, OpenMP, C/C++/Fortran Compilers and Tools	✓	✓	✓
PLATFORMS X86, OpenPOWER, NVIDIA GPU	✓	✓	✓
UPDATES	1-2 times a year	6-9 times a year	6-9 times a year
SUPPORT	User Forums	PGI Support	PGI Premier Services
LICENSE	Annual	Perpetual	Volume/Site

PGI 2017 NEW FEATURES

Fortran/C/C++ OpenACC 2.5 on Tesla GPUs and multicore CPUs

- CUDA 8.0/P100 GPU accelerator support
- OpenACC for multicore Haswell, Broadwell and OpenPOWER CPUs
- User-driven OpenACC performance optimizations

Full C++ 14 language support, interoperable with GCC/g++ 5.1 thru 6.2

OpenMP 4.5 Fortran for OpenPOWER multicore CPUs (no target offload, yet)

Multicore CPU Performance Optimizations

- OpenPOWER performance enhancements averaging 5-10%
- New tuned numerical intrinsics for Haswell/Broadwell
- Improved inlining

Updated OpenMPI, NetCDF, ESMF libraries; new OS distro's support; debugger enhancements

Full list of new features at www.pgicompilers.com/whats-new

… COMING SOON

OpenACC enhancements

- Deep copy of aggregate data structures
- CUDA Unified Memory enhancements
- GPU loop scheduling optimizations

New processors and platforms

- NVIDIA Volta V100 GPUs
- CUDA 9.0 support
- POWER9, Xeon Skylake, Xeon Phi KNL (Beta)

OpenMP 4.5 Fortran/C/C++ for Multicore CPUs (no target offload)

C++ atomics/abstraction performance, initial C++17, GNU 6.3 interop

Tools – OpenACC profiling, PGDBG/AVX-512, OpenPOWER DWARF gen

AGENDA

PGI Overview

OpenACC Applications & Performance

OpenPOWER & Volta V100

OpenACC and CUDA Unified Memory

OPENACC ADOPTION

NEW PLATFORMS

OpenPOWER™

Sunway TaihuLight
#1 Top 500, Nov. 2016

GROWING COMMUNITY

- 6,000+ enabled developers
- 4,500+ course registrants
- 350+ Hackathon attendees
- 150+ User Group members

PORTING SUCCESS

- Gaussian 16 ported to Tesla with OpenACC
- Five of 13 ORNL CAAR codes using OpenACC
- ANSYS Fluent R18 production release is GPU accelerated using OpenACC

GAUSSIAN 16

A Leading Computation Chemistry Code

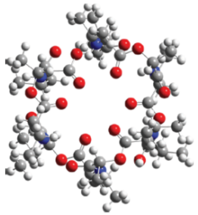

Valinomycin
wB97xD/6-311+(2d,p) Freq
2.25X speedup

Hardware: HPE server with dual Intel Xeon E5-2698 v3 CPUs (2.30GHz ; 16 cores/chip), 256GB memory and 4 Tesla K80 dual GPU boards (boost clocks: MEM 2505 and SM 875). Gaussian source code compiled with PGI Accelerator Compilers (16.5) with OpenACC (2.5 standard).

PGI

Mike Frisch, Ph.D.
President and CEO
Gaussian, Inc.

"Using OpenACC allowed us to continue development of our fundamental algorithms and software capabilities simultaneously with the GPU-related work. In the end, we could use the same code base for SMP, cluster/network and GPU parallelism. PGI's compilers were essential to the success of our efforts."

NUMECA FINE/OPEN

Unstructured CDF Solver

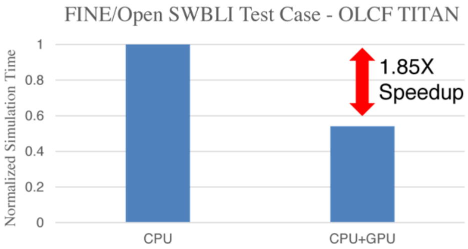

System: Four nodes on the OLCF Titan supercomputer (16 cores and 1 GPU per node).

PGI

David Gutzwiller
Lead Software Developer
NUMECA

"Porting our unstructured C++ CFD solver FINE/Open to GPUs using OpenACC would have been impossible two or three years ago, but OpenACC has developed enough that we're now getting some really good results."

ANSYS FLUENT

Discrete Ordinate (DO) Radiation Solver

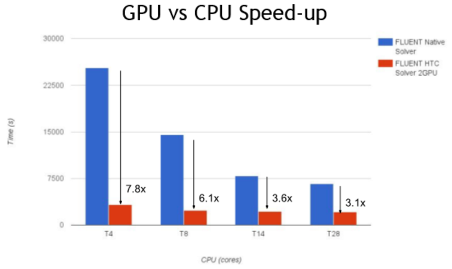

CPU Hardware:
(Haswell EP) Intel(R) Xeon(R) CPU E5-2695 v3 @ 2.30GHz, 2 socket x 14 = 28 cores
GPGPU Hardware:
Tesla K80 12+12 GB, Driver 346.46

PGI

Sunil Sathe
Lead Software Developer
ANSYS Fluent

"We've effectively used OpenACC for heterogeneous computing in ANSYS Fluent with impressive performance. We're now applying this work to more of our models and new platforms."

MPAS ATMOSPHERE

Climate and Weather Modeling

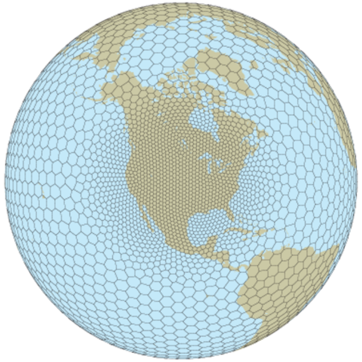

A variable resolution MPAS Voronoi mesh

Source: mpas github
https://mpas-dev.github.io/atmosphere/atmosphere.html

PGI

Richard Loft
Director, Technology Development
NCAR

"Our team has been evaluating OpenACC as a pathway to performance portability for the Model for Prediction (MPAS) atmospheric model. Using this approach on the MPAS dynamical core, we have achieved performance on a single P100 GPU equivalent to 2.7 dual socketed Intel Xeon nodes on our new Cheyenne supercomputer."

COSMO

Regional Atmospheric Model

Image courtesy Meteoswiss.

Dr. Oliver Fuhrer
Senior Scientist
Meteoswiss

"OpenACC made it practical to develop for GPU-based hardware while retaining a single source for almost all the COSMO physics code."

VASP

The Vienna Ab Initio Simulation Package

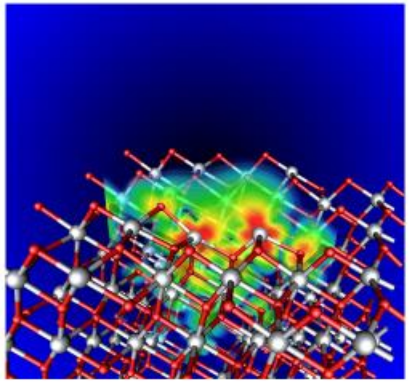

Dr. Martijn Marsman
Computational Materials Physics
University of Vienna

"Early indications are that we can nearly match the performance of CUDA using OpenACC on GPUs. This will enable our domain scientists to work on a uniform GPU accelerated Fortran source code base."

FV3 WEATHER MODEL

Global Weather Forecast Model

OpenACC Performance

- PGI compiler V16.10
- 2X faster performance on GPU
  - Dual-socket Haswell CPU and NVIDIA Pascal (P100) GPU

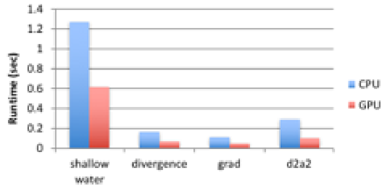

Slide courtesy of Mark Govett, NOAA / Earth System Research Laboratory

Mark Govett
Chief, HPC Section
NOAA

"Lessons learned in the development of NIM and F2C-ACC have proven invaluable in our current efforts to create a single, performance portable version of the FV3 weather model using OpenACC."

QUANTUM ESPRESSO

Quantum Chemistry Suite

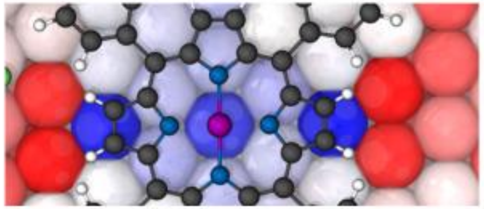

QUANTUMESPRESSO

www.quantum-espresso.org

Filippo Spiga
Head of Research Software Engineering
University of Cambridge

"CUDA Fortran gives us the full performance potential of the CUDA programming model and NVIDIA GPUs. !$CUF KERNELS directives give us productivity and source code maintainability. It's the best of both worlds."

AGENDA

PGI Overview

OpenACC Applications & Performance

OpenPOWER & Volta V100

OpenACC and CUDA Unified Memory

The PGI Compiler Infrastructure

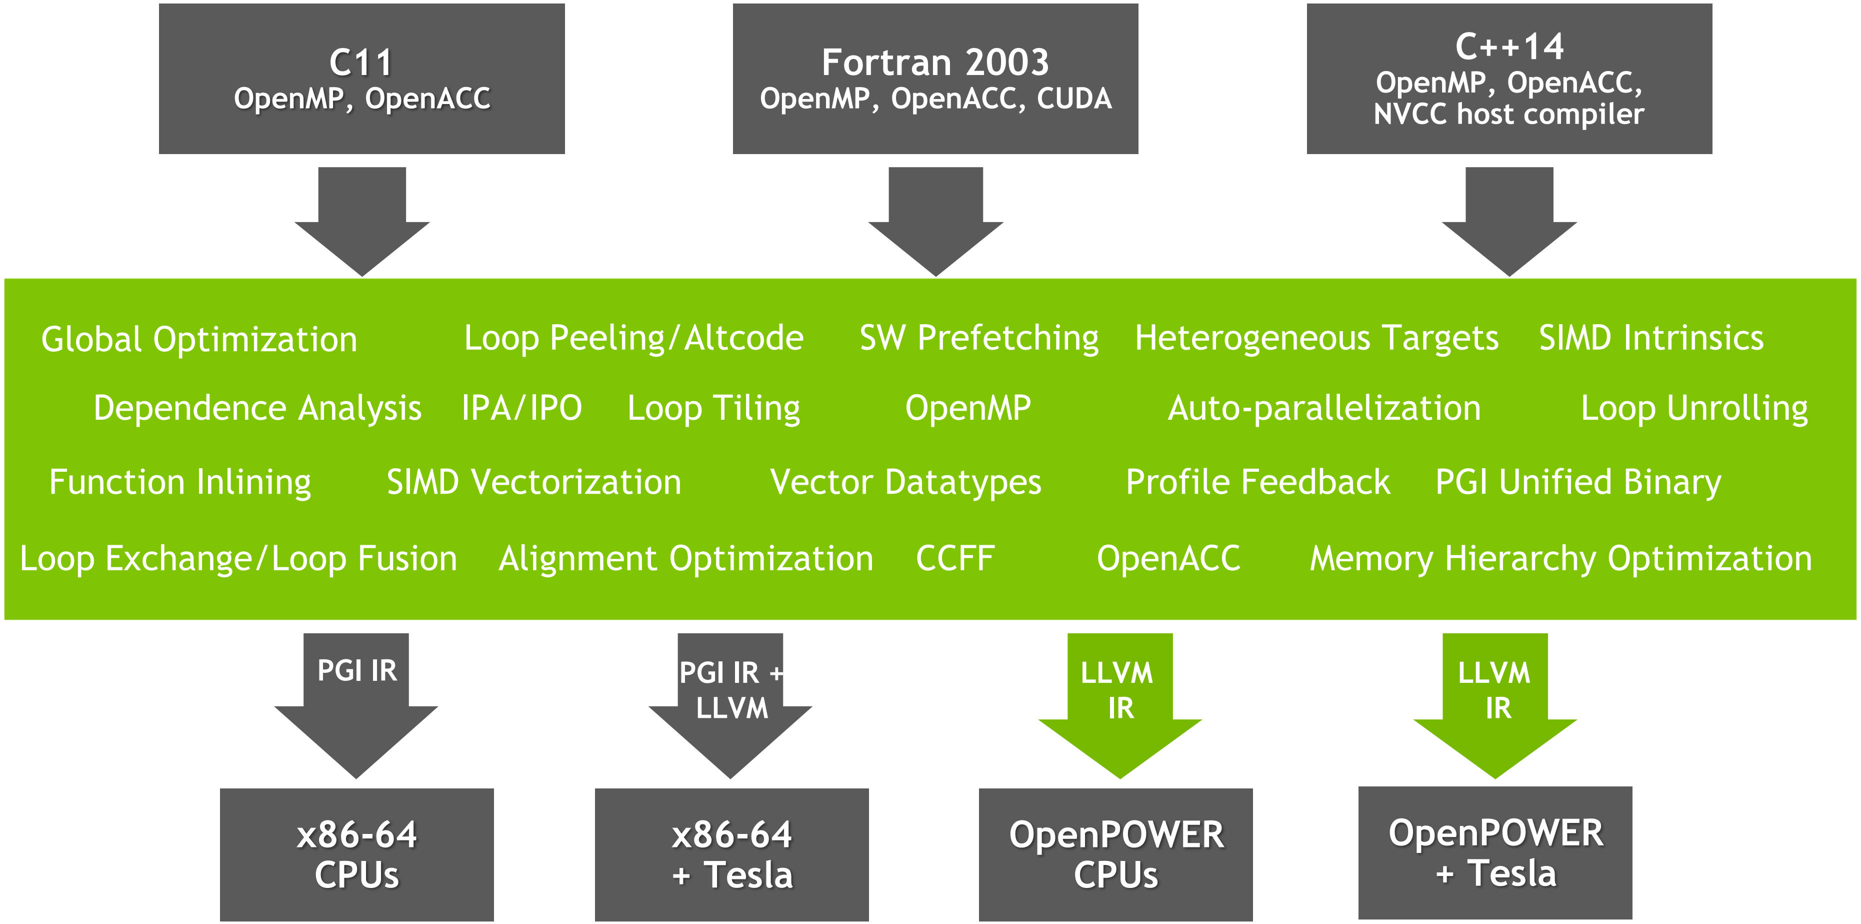

PGI FOR OpenPOWER+TESLA

Fortran 2003, C11, C++14 compilers, performance profiler

CUDA Fortran, OpenACC, OpenMP, NVCC host compiler

Integrated IBM-optimized LLVM OpenPOWER code generator

Available now at:

pgicompilers.com/openpower

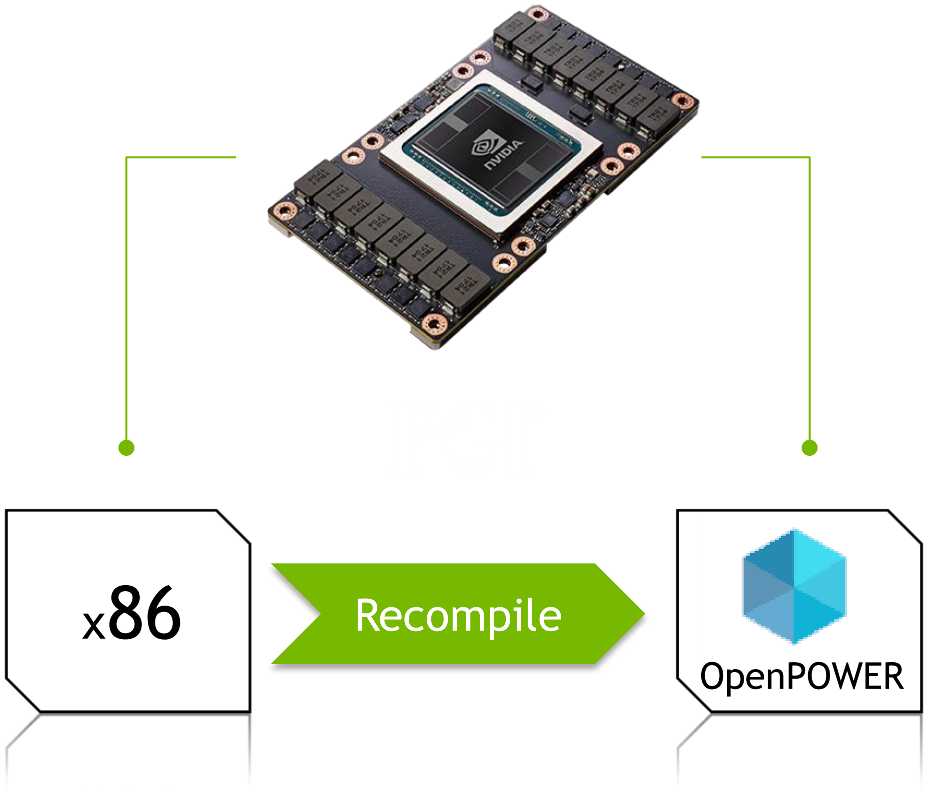

PGI FOR OpenPOWER+TESLA

Fortran 2003, C11, C++14 compilers, performance profiler

CUDA Fortran, OpenACC, OpenMP, NVCC host compiler

Integrated IBM-optimized LLVM OpenPOWER code generator

Available now at:

pgicompilers.com/openpower

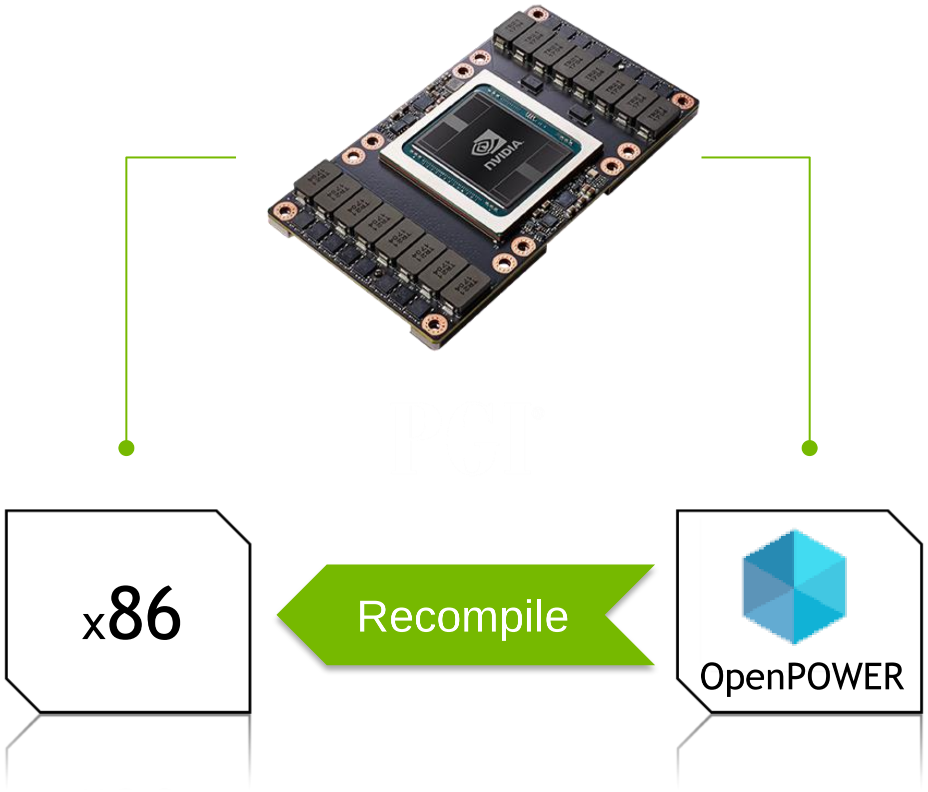

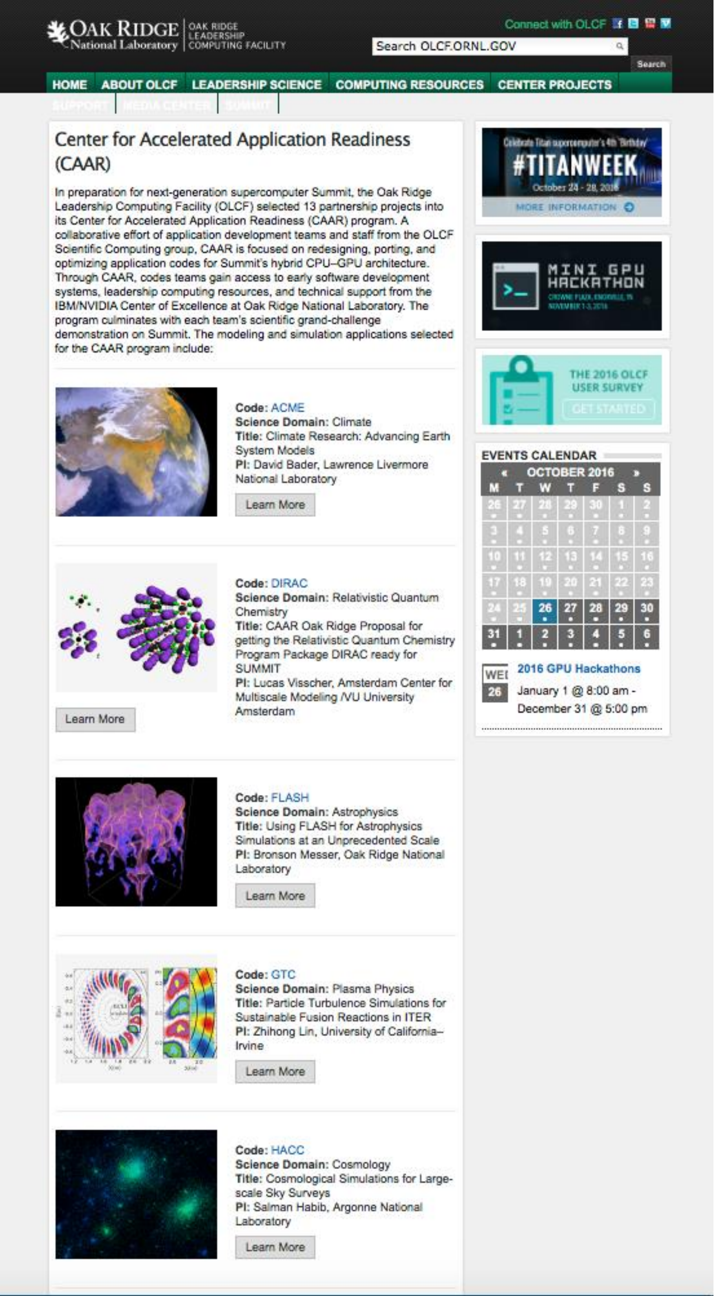

CENTER FOR ACCELERATED APPLICATION READINESS (CAAR)

Oak Ridge Leadership Computing Facility

Collaboration of application development teams and OLCF staff focused on redesigning, porting, and optimizing 13 application codes for Summit's hybrid POWER9+Volta node architecture.

www.olcf.ornl.gov/caar

X86-64 TO OPENPOWER PORTING

C/C++ ABI DIFFERENCES

- signed vs unsigned default char
- long double

NUMERICAL DIFFERENCES

- Intrinsics accuracy may differ across targets
- FMA vs. no FMA

PLATFORM DIFFERENCES

- Large memory model
- C varargs

X86-SPECIFIC FEATURES

- Inline *asm* statements
- SSE/AVX intrinsics

NVIDIA TESLA V100 GPU

GIANT LEAP FOR AI & HPC

5,120 CUDA cores | 640 Tensor cores

7.5 FP64 TFLOPS | 15 FP32 TFLOPS

NEW 120 Tensor TFLOPS

20MB SM RF | 16MB Cache

16GB HBM2 @ 900 GB/s

300 GB/s NVLink

IBM POWER NVLINK SYSTEMS

APPROVED DESIGNS FOR OPENPOWER ECOSYSTEM

4 GPU POWER8 FOR PASCAL

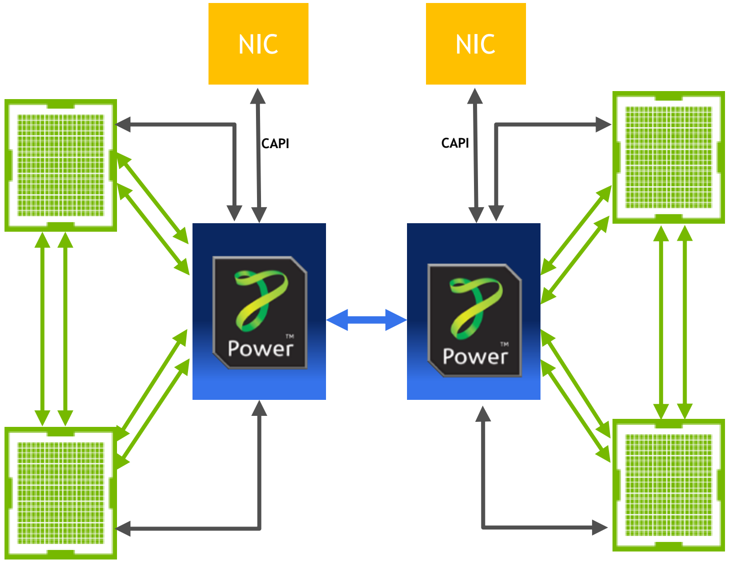

Page Migration Engine

4 GPU POWER9 FOR VOLTA

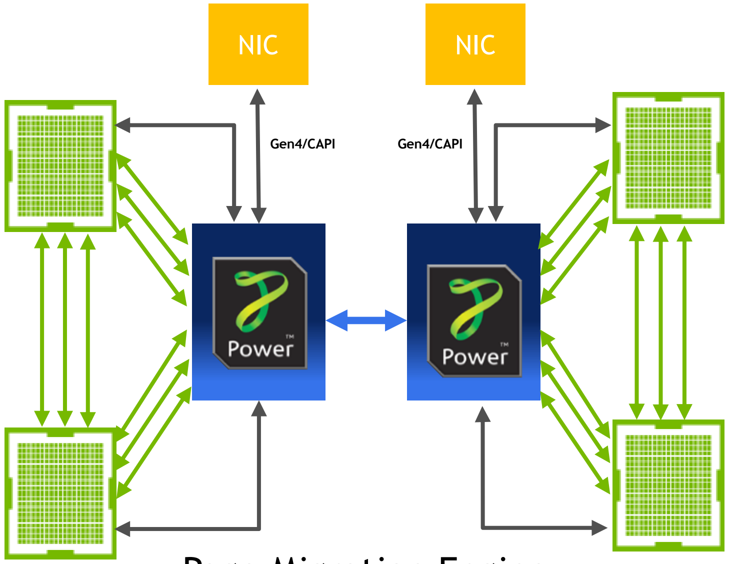

Page Migration Engine
Coherence, Atomics, ATS

P100 SXM2
40 GB/s full duplex each CPU:GPU and GPU:GPU link
IBM POWER8+ CPU

VOLTA SXM2
75 GB/s full duplex each CPU:GPU and GPU:GPU link
IBM POWER9 CPU

THE FLANG PROJECT

An open source Fortran front-end for LLVM

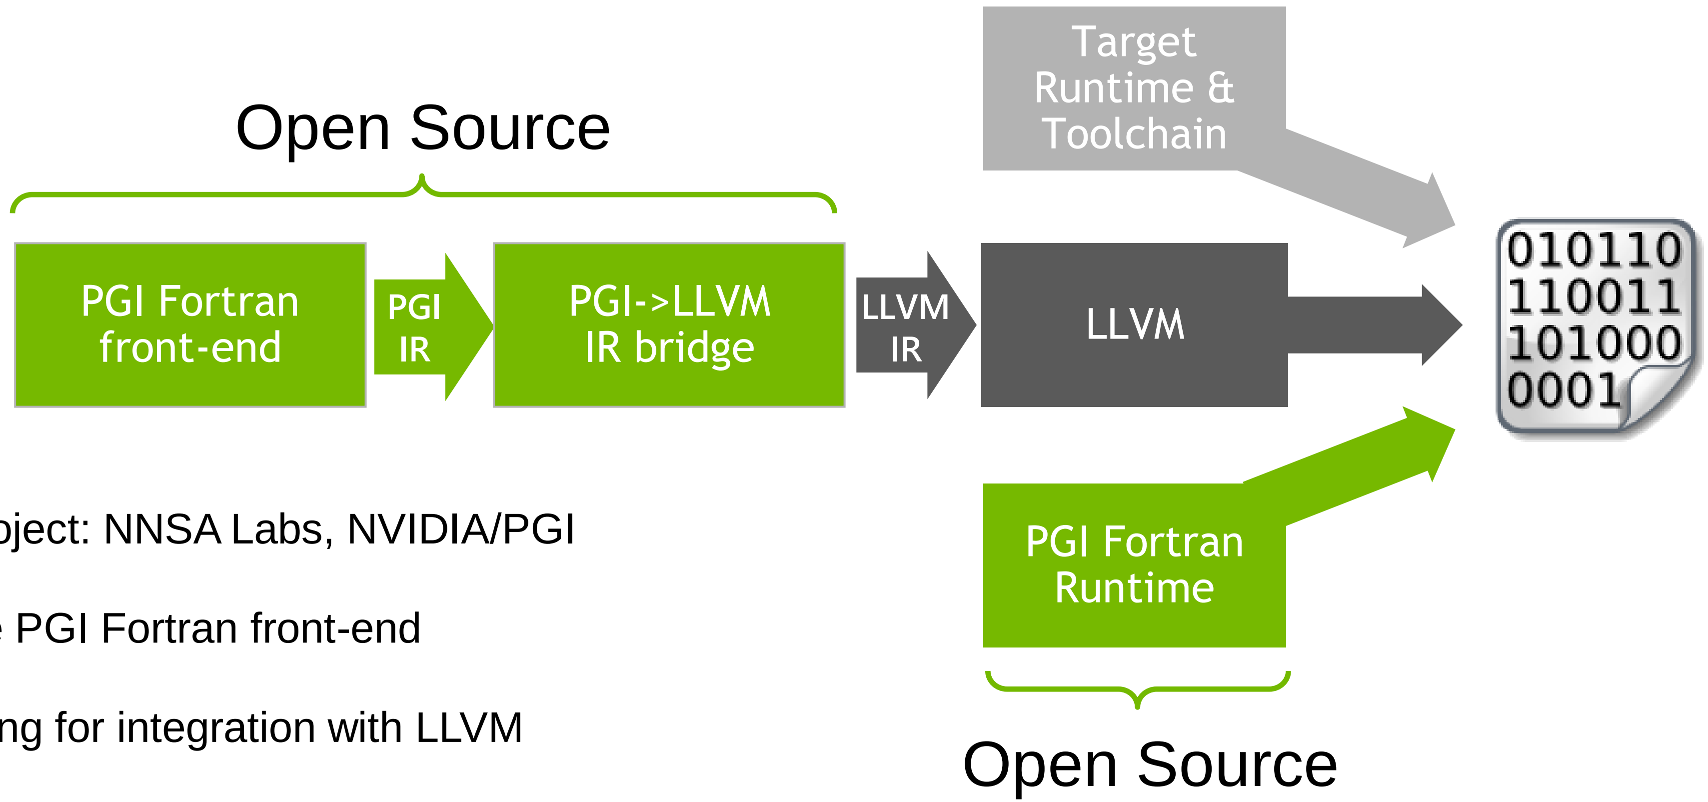

Multi-year project: NNSA Labs, NVIDIA/PGI

Based on the PGI Fortran front-end

Re-engineering for integration with LLVM

Develop CLANG-quality Fortran msg facility

LLVM: COMMUNITY POWER

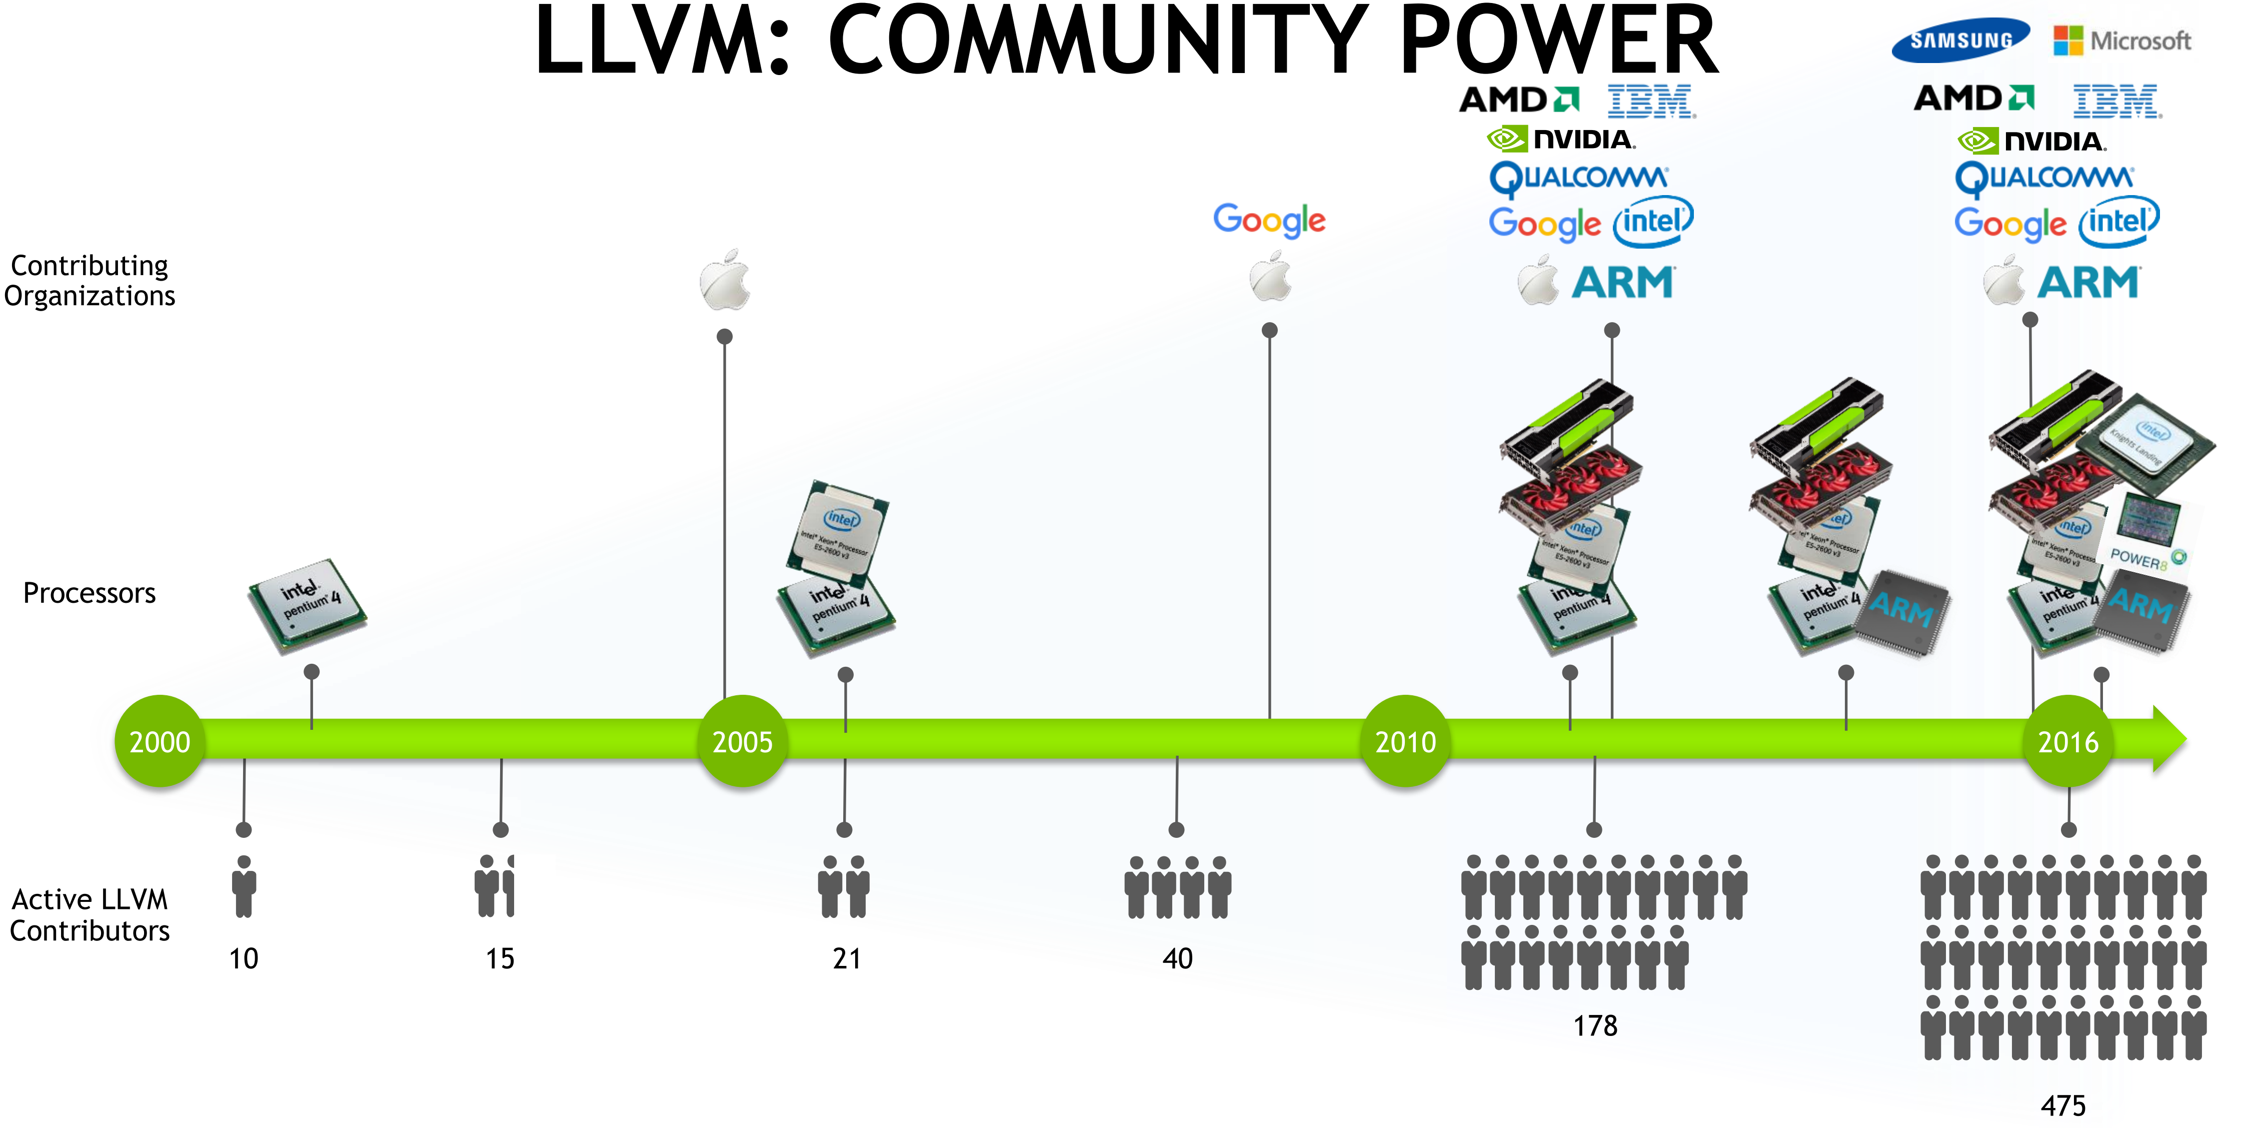

AGENDA

PGI Overview

OpenACC Applications & Performance

OpenPOWER & Volta V100

OpenACC and CUDA Unified Memory

CUDA Unified Memory

Dramatically Lower Developer Effort

Developer View Today

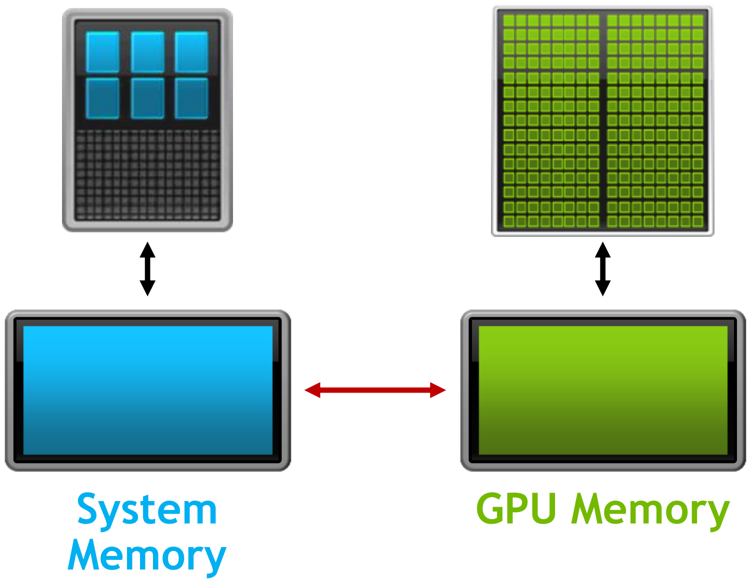

Developer View With CUDA Unified Memory

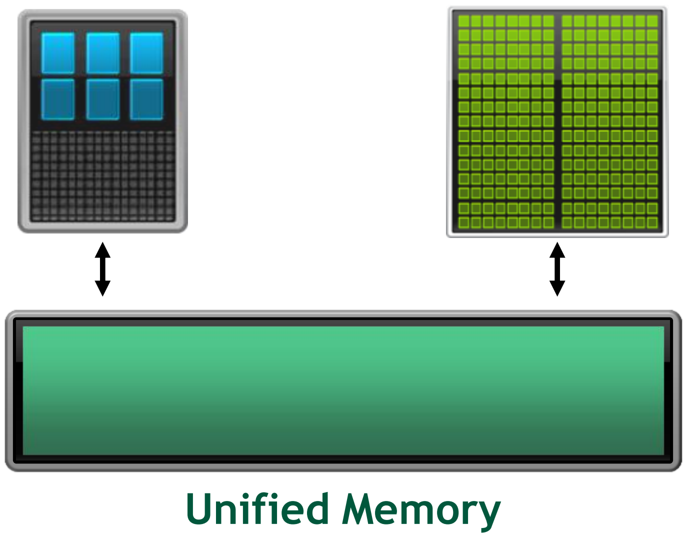

CUDA UNIFIED MEMORY ON P100

Servicing CPU *and* GPU Page Faults

GPU Code

```
__global__
Void setValue(char *ptr, int index, char val)
{
  ptr[index] = val;
}
```

CPU Code

```
cudaMallocManaged(&array, size);
memset(array, size);
setValue<<<...>>>(array, size/2, 5);
```

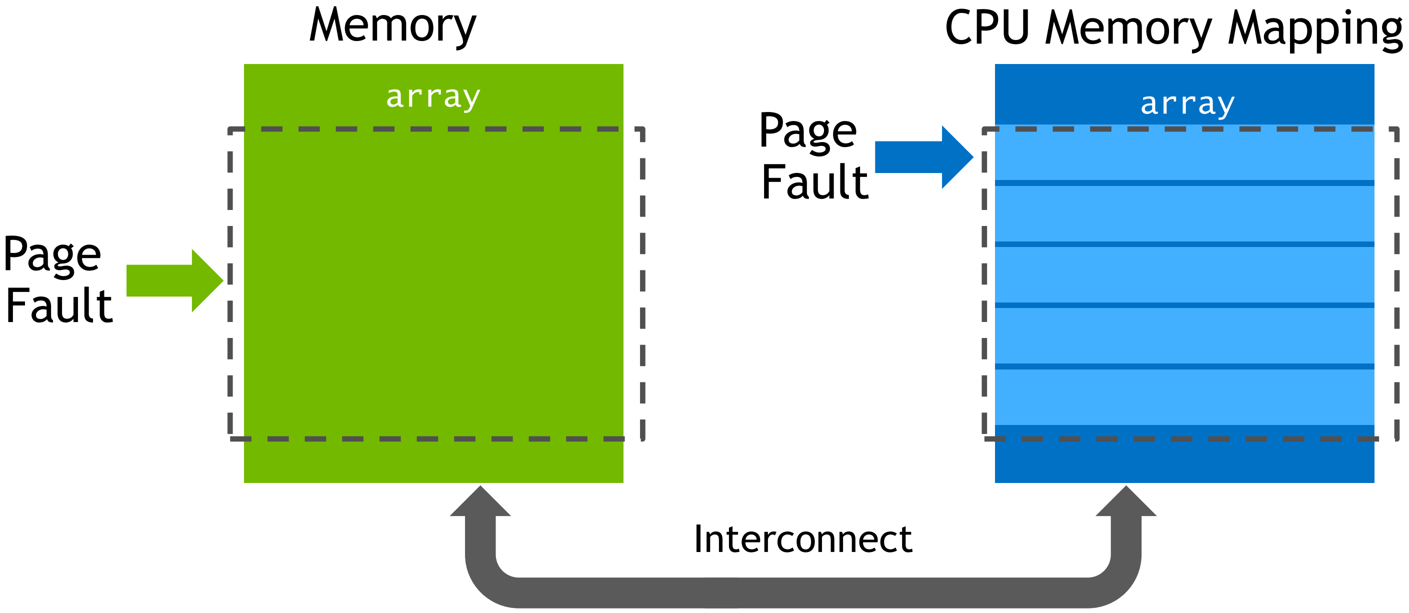

OPENACC 2.6 MANUAL DEEP COPY

```
typedef struct points {
    float* x;  float* y;  float* z;
    int n;
    float coef, direction;
} points;


void sub ( int n, float* y ) {
    points p;
    #pragma acc data create (p)
    {
        p.n = n;
        p.x = ( float*) malloc ( sizeof ( float )*n );
        p.y = ( float*) malloc ( sizeof ( float )*n );
        p.z = ( float*) malloc ( sizeof ( float )*n );
        #pragma acc update device (p.n)
        #pragma acc data copyin (p.x[0:n], p.y[0: n])
        {
            #pragma acc parallel loop
            for ( i =0; i<p.n; ++I ) p.x[i] += p.y[i];
            . . .
```

DRAFT OPENACC 3.0 TRUE DEEP COPY

```
typedef struct points {
    float* x;  float* y;  float* z;
    int n;
    float coef, direction;
    #pragma acc policy inout(x[0:n],y[0:n])
} points;

void sub ( int n, float* y ) {
    points p;


        p.n = n;
        p.x = ( float*) malloc ( sizeof ( float )*n );
        p.y = ( float*) malloc ( sizeof ( float )*n );
        p.z = ( float*) malloc ( sizeof ( float )*n );

        #pragma acc data copy (p)
        {
            #pragma acc parallel loop
            for ( i =0; i<p.n; ++I ) p.x[i] += p.y[i];
            . . .
```

OPENACC 2.6 MANUAL DEEP COPY

A real-world example: managing one aggregate data structure

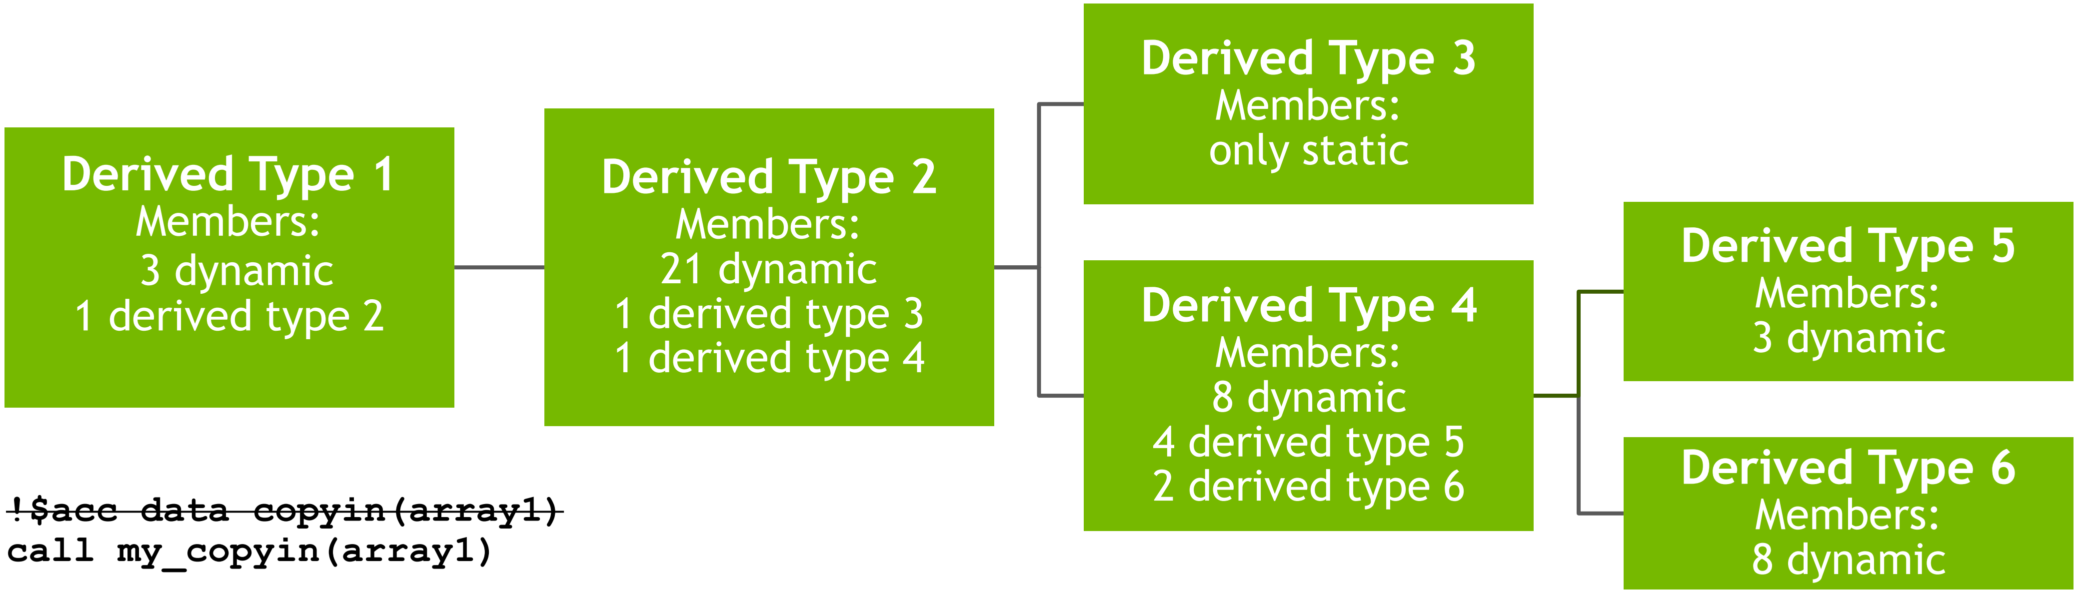


```
~~!$acc data copyin(array1)~~
call my_copyin(array1)
```

-> 12 lines of code

-> 48 lines of code

-> 26 lines of code

-> 8 lines of code

-> 13 lines of code

-> 107 lines of code just for COPYIN

Plus additional lines of code for COPYOUT, CREATE, UPDATE

OPENACC WITH UNIFIED MEMORY

A real-world example: managing one aggregate data structure

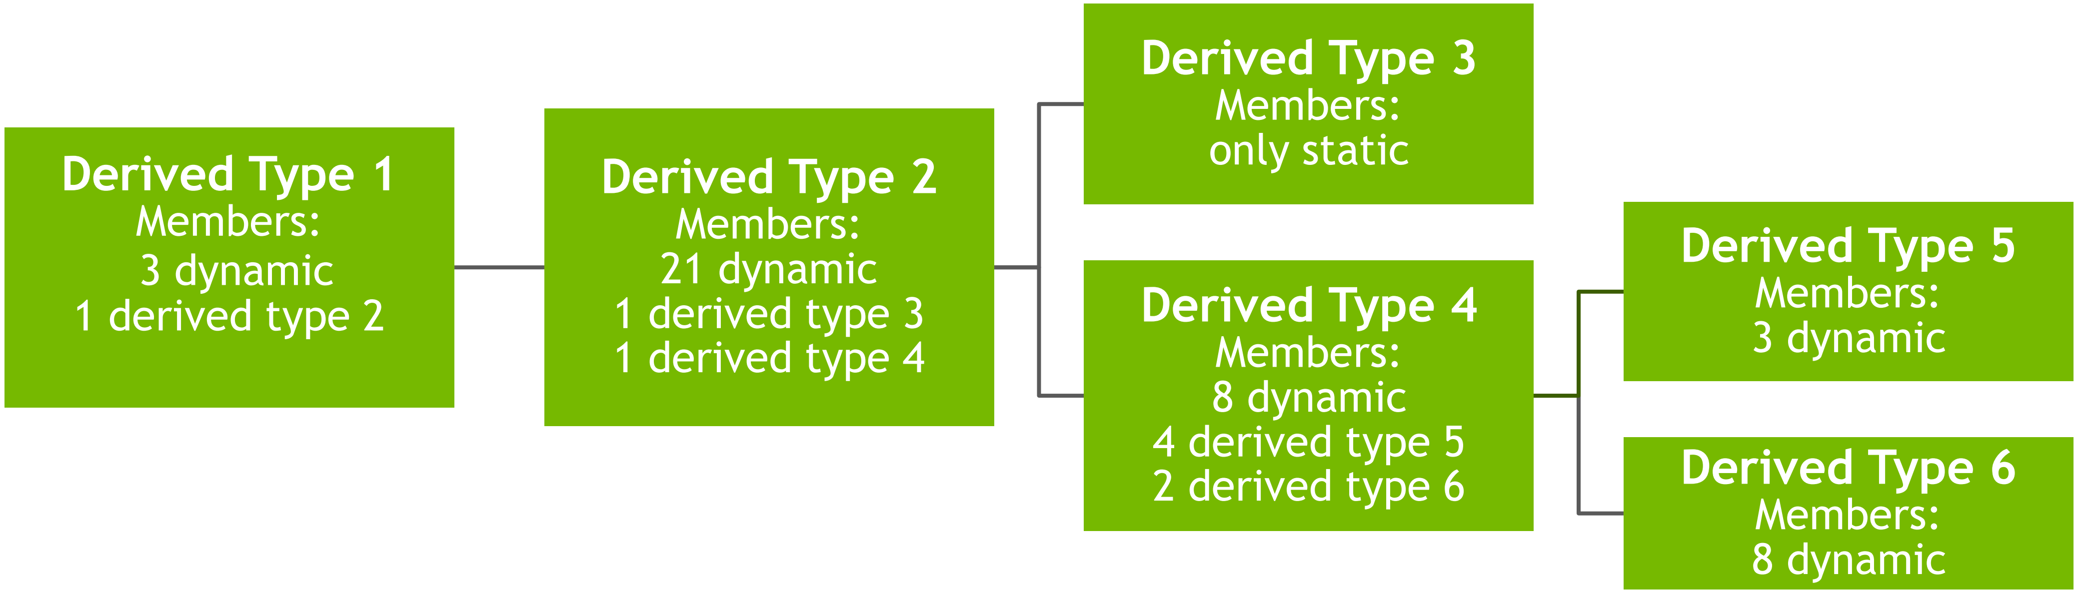

0 lines of code! It just works.

NVIDIA