

# Directive-based Accelerator Programming With OpenACC

Mathew Colgrove

[mathew.colgrove@pgroup.com](mailto:mathew.colgrove@pgroup.com)

July 2013

**PGI**

- C99, C++, Fortran 2003 Compilers
  - Optimizing
  - Vectorizing
  - Parallelizing
- Graphical parallel tools
  - PGDBG® debugger
  - PGPROF® profiler
- AMD, Intel, NVIDIA
- PGI Unified Binary™
- Linux, OS X, Windows
- Visual Studio & Eclipse integration
- GPGPU Features
  - PGI Accelerator™, OpenACC
  - CUDA Fortran/C/C++
  - CUDA-x86
  - OpenCL (ARM CPUs only)


The Portland Group

Technology Products Services Support Download Resources User Forums Purchase About

Site Map Contact Log In Search

# PGI 2013


Now available for [download](#). This release includes OpenMP performance 1.5x faster than GCC, new GNU-compatible OpenACC C++ compilers, support for CUDA 5 and Tesla K20 and more. Read [what's new](#).



**OpenACC**  
 DIRECTIVES FOR ACCELERATORS


PGI Accelerator™ compilers now include comprehensive support for the OpenACC specification for directive-based programming of GPUs and accelerators. [Learn more](#).

## PGI® Optimizing Fortran, C and C++ Compilers & Tools




**PGI Workstation™ and PGI Server™ for x64**

PGI optimizing multi-core x64 compilers for Linux, MacOS & Windows with support for debugging and profiling of local MPI processes. A complete OpenMP/MPI SDK for high performance computing on the latest Intel and AMD CPUs. [Try | Buy](#)




**PGI CUDA Fortran and CUDA-x86 Compilers**

CUDA Fortran enables GPU acceleration of HPC applications using the NVIDIA CUDA parallel programming model in a native optimizing Fortran 2003 compiler. Use CUDA-x86 to compile your CUDA C or CUDA Fortran program to run on x64 targets. Both products are compatible and interoperable with NVIDIA's CUDA C. [Try | Buy](#)




**PGI Accelerator™ with OpenACC**

PGI Accelerator C99 & Fortran enable high level programming of HPC applications for x64+accelerators using OpenACC compiler directives. Portable, incremental, and easy to use for application domain experts. [Try | Buy](#)



**The PGI CDK® Cluster Development Kit®**


The PGI CDK includes optimizing Fortran/C/C++ compilers configured to build, debug and profile MPI and hybrid MPI/OpenMP HPC applications for Linux or Windows Clusters using the major open source MPI implementations or MSMPI. [Try | Buy](#)



**PGI Visual Fortran® for Microsoft Windows**

PGI Visual Fortran brings optimizing multi-core x64 Fortran with integrated OpenMP/MPI debugging to scientists & engineers on Microsoft Windows within Microsoft Visual Studio. [Try | Buy](#)

## PGI Compilers for Mobile and Embedded Applications



**PGI OpenCL Compiler for Multi-core ARM**

PGCL™ is PGI's optimizing OpenCL™ compiler for the ST-Ericsson NovaThor highly integrated mobile platform. NovaThor is based on ARM® processing cores and includes advanced graphics cores, powerful multimedia engines, and the latest mobile broadband and connectivity technologies. [Download](#)

**The New PGC++ Compiler for Linux**

The new PGC++ compiler uses the system C++ libraries, avoiding problems that arise from dealing with another C++ and Standard Template Library.

**Targeting AVX-Enabled Processors**

PGI support for the new processors from AMD and Intel.

**First Look: PGI CUDA C/C++ for x86**

How it works, what's included in this first functional release and what's in store in the coming months.

**Object Oriented Programming with Fortran 2003 Part 1**

Procedure polymorphism can operate on a variety of data types and values.

**Debugging CUDA-x86 Applications**

Using PGDBG to debug CUDA C and C++ programs on x86 targets.

**Introduction to CUDA Fortran**

Step by step example of writing a basic CUDA Fortran program using PGI Fortran.

**The PGI Accelerator Programming Model on NVIDIA GPUs Part 1**

The first in a series of articles on PGI's accelerator-enabled compilers. Part 1 takes an in-depth look at PGI's Accelerator Programming Model.

**Understanding the CUDA Data Parallel Threading Model—A Primer**

This article describe the data parallelism model supported in CUDA on NVIDIA GPUs.

**Tuning a Monte Carlo Algorithm on GPUs**

A step-by-step example demonstrating many useful CUDA Fortran techniques including device resident data, sum reductions, data transfer optimization and calling a CUDA C kernel from CUDA Fortran.

**Parallel Random Number Generation Using OpenMP, OpenCL and PGI Accelerator Directives**

www.pgroup.com

# Talk Roadmap

- Introduction to OpenACC
- Compute Constructs
- Loop Directives
- Demo – Himeno
- Data Region
- Porting Example – Seismic CPML
- Future OpenACC features

# OpenACC

## Open Programming Standard for Parallel Computing

“PGI OpenACC will enable programmers to easily develop portable applications that maximize the performance and power efficiency benefits of the hybrid CPU/GPU architecture of Titan.”

--Buddy Bland, Titan Project Director, Oak Ridge National Lab



“OpenACC is a technically impressive initiative brought together by members of the OpenMP Working Group on Accelerators, as well as many others. We look forward to releasing a version of this proposal in the next release of OpenMP.”

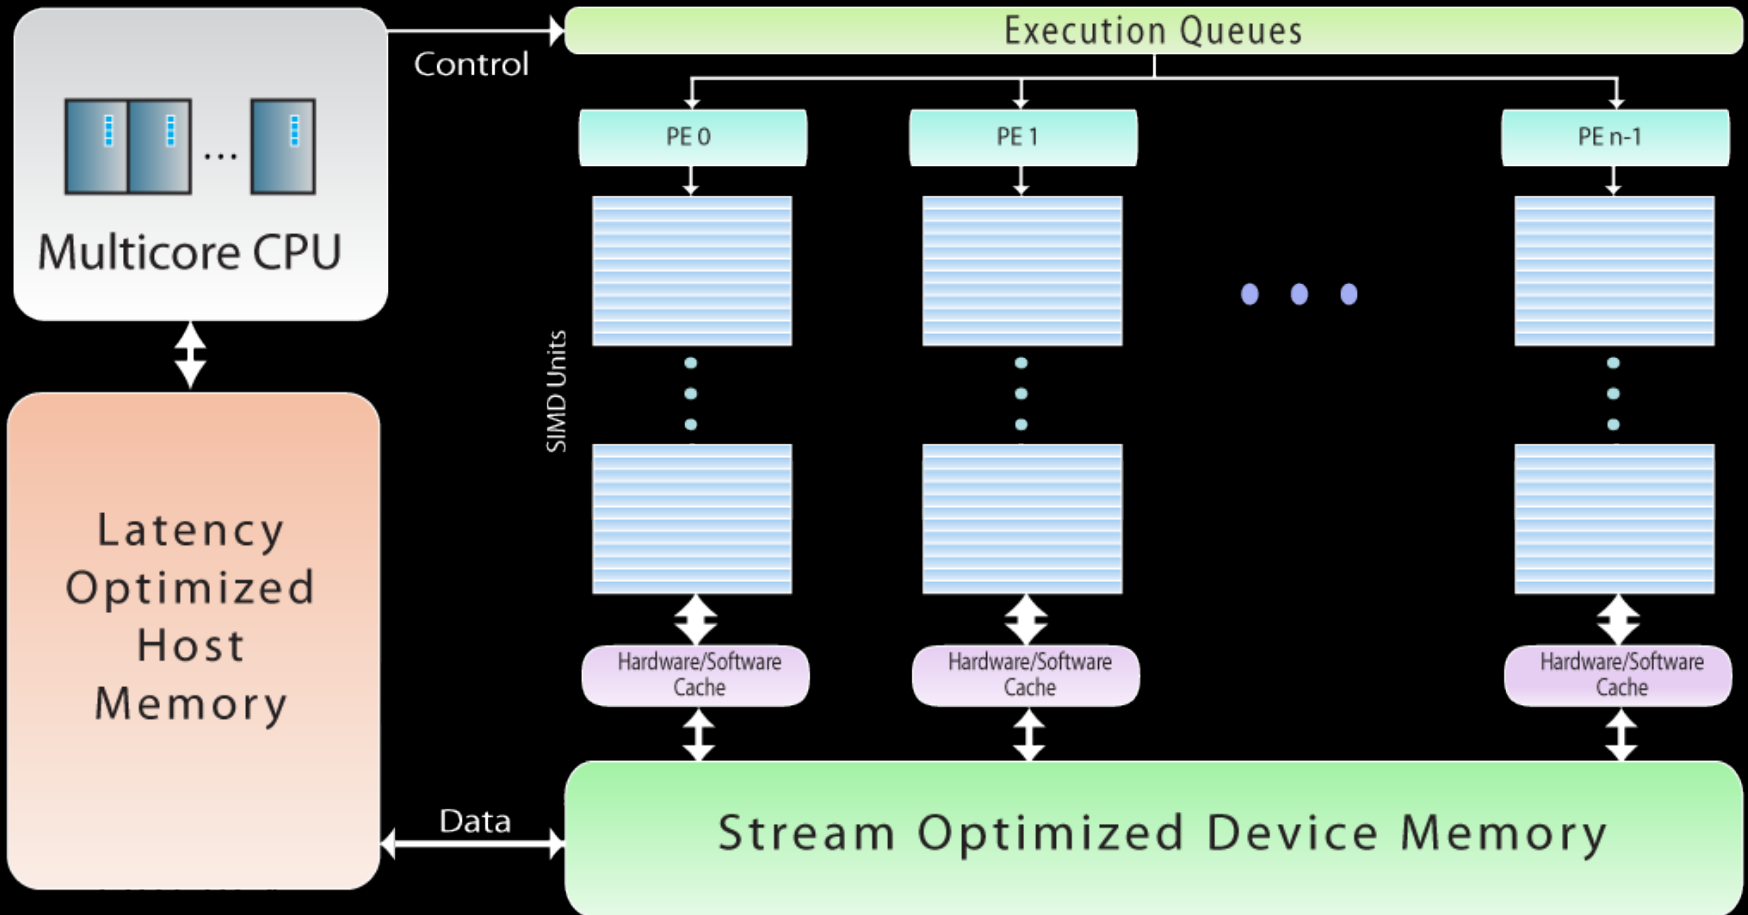
--Michael Wong, CEO OpenMP Directives Board



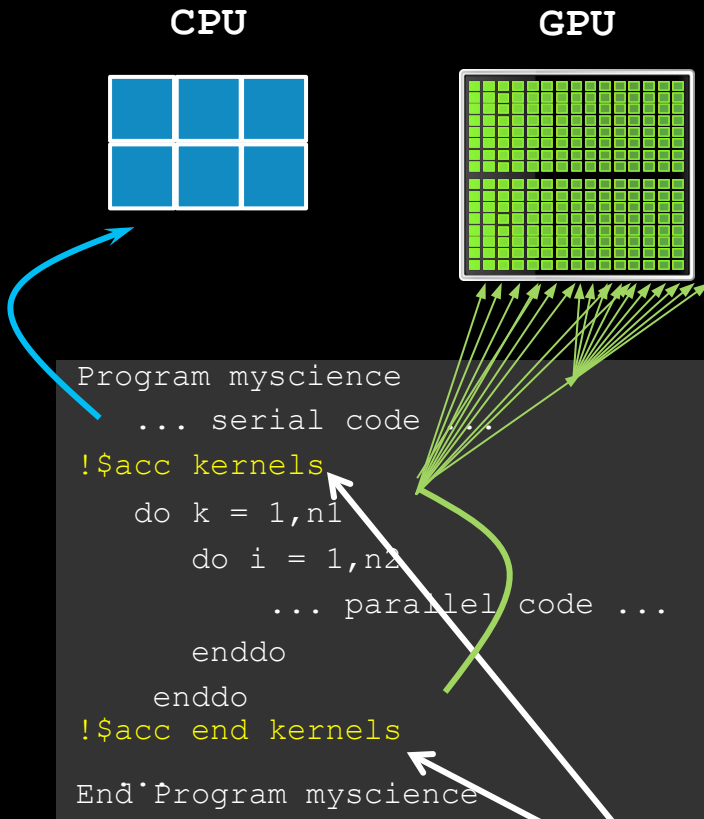
## OpenACC Standard



# OpenACC Abstract Machine Architecture



# OpenACC Directives



Portable compiler hints

Compiler parallelizes code

Designed for multicore CPUs & many core GPUs / Accelerators

**Your original  
Fortran or C code**

OpenACC  
Compiler  
Directive

# PGI Accelerator Directive-based Compilers

```
#pragma acc kernels loop
for( i = 0; i < nrows; ++i ){
  float val = 0.0f;
  for( d = 0; d < nzeros; ++d ){
    j = i + offset[d];
    if( j >= 0 && j < nrows )
      val += m[i+nrows*d] * v[j];
  }
  x[i] = val;
}
```

compile



Link



```
matvec:
  subq  $328, %rsp
  ...
  call  __pgi_cu_alloc
  ...
  call  __pgi_cu_uploadx
  ...
  call  __pgi_cu_launch2
  ...
  call  __pgi_cu_downloadx
  ...
  call  __pgi_cu_free
  ...
```

+

```
.entry matvec_14_gpu( ...
.reg .u32 %r<70> ...
cvt.s32.u32 %r1, %tid.x;
mov.s32 %r2, 0;
setp.ne.s32 $p1, %r1, %r2
cvt.s32.u32 %r3, %ctaid.x;
cvt.s32.u32 %r4, %ntid.x;
mul.lo.s32 %r5, %r3, %r4;
@%p1 bra $!t_0_258;
st.shared.s32 [__i2s], %r5
$!t_0_258:
  bar.sync 0;
  ...
```

Unified  
Object

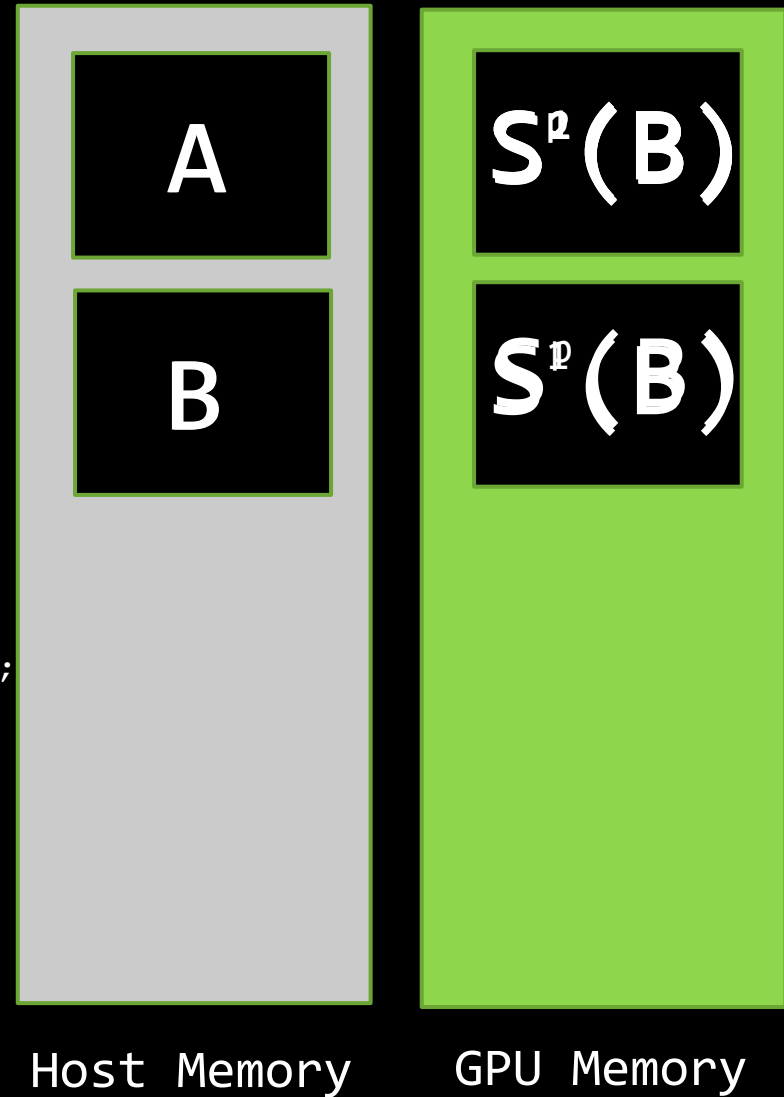
execute



... no change to existing makefiles, scripts, IDEs,  
programming environment, etc.

# PGI Accelerator OpenACC example

```
#pragma acc data \
    copy(b[0:n][0:m]) \
    create(a[0:n][0:m])
{
for (iter = 1; iter <= p; ++iter){
    #pragma acc kernels
    {
        for (i = 1; i < n-1; ++i){
            for (j = 1; j < m-1; ++j){
                a[i][j]=w0*b[i][j]+
                    w1*(b[i-1][j]+b[i+1][j]+
                        b[i][j-1]+b[i][j+1])+
                    w2*(b[i-1][j-1]+b[i-1][j+1]+
                        b[i+1][j-1]+b[i+1][j+1]);
            } }
        for( i = 1; i < n-1; ++i )
            for( j = 1; j < m-1; ++j )
                b[i][j] = a[i][j];
    }
}
}
```





# Why use OpenACC Directives?

- Productivity
  - Higher level programming model
  - a la OpenMP
- Portability
  - ignore directives, portable to the host
  - portable across different accelerators
  - *performance portability*
- Performance feedback

# Matrix Multiply Source Code Size Comparison:

```
1 void matrixMulGPU(cl_uint clDeviceCount, cl_mem h_A, float* h_B_data,
2                unsigned int mem_size_b, float* d_C)
3 {
4     cl_mem d_A[MAX_GPU_COUNT];
5     cl_mem d_C[MAX_GPU_COUNT];
6     cl_mem d_B[MAX_GPU_COUNT];
7     cl_event GPUDone[MAX_GPU_COUNT];
8     cl_event GPUExecution[MAX_GPU_COUNT];
9
10    // Create buffers for each GPU
11    // Each GPU will compute sizePerGPU rows of the result
12    int sizePerGPU = NA / clDeviceCount;
13
14    int workOffset[MAX_GPU_COUNT];
15    int workSize[MAX_GPU_COUNT];
16
17    workOffset[0] = 0;
18    for(unsigned int i=0; i < clDeviceCount; ++i)
19    {
20        // Input buffer
21        workSize[i] = (i != clDeviceCount - 1) ? sizePerGPU : (NA - workOffset[i]);
22
23        d_A[i] = clCreateBuffer(clGetContext, CL_MEM_READ_ONLY, workSize[i] * sizeof(float) * NA, NULL, NULL);
24
25        // Copy only assigned rows from host to device
26        clEnqueueCopyBuffer(commandQueue[i], h_A, d_A[i], workOffset[i] * sizeof(float) * NA,
27                          0, workSize[i] * sizeof(float) * NA, 0, NULL, NULL);
28
29        // create OpenCL buffer on device that will be initialized from the host memory on first use
30        // on device
31        d_B[i] = clCreateBuffer(clGetContext, CL_MEM_READ_ONLY | CL_MEM_COPY_HOST_PTR,
32                              mem_size_b, h_B_data, NULL);
33
34        // Output buffer
35        d_C[i] = clCreateBuffer(clGetContext, CL_MEM_WRITE_ONLY, workSize[i] * WC * sizeof(float), NULL, NULL);
36
37        // set the args values
38
39        clSetKernelArg(multiplicationKernel[i], 0, sizeof(cl_mem), (void *) d_C[i]);
40        clSetKernelArg(multiplicationKernel[i], 1, sizeof(cl_mem), (void *) d_A[i]);
41        clSetKernelArg(multiplicationKernel[i], 2, sizeof(cl_mem), (void *) d_B[i]);
42        clSetKernelArg(multiplicationKernel[i], 3, sizeof(float) * BLOCK_SIZE *BLOCK_SIZE, 0);
43        clSetKernelArg(multiplicationKernel[i], 4, sizeof(float) * BLOCK_SIZE *BLOCK_SIZE, 0);
44
45        if(i+1 < clDeviceCount)
46            workOffset[i + 1] = workOffset[i] + workSize[i];
47    }
48
49    // Execute Multiplication on all GPUs in parallel
50    size_t localWorkSize[] = {BLOCK_SIZE, BLOCK_SIZE};
51    size_t globalWorkSize[] = {shRoundUp(BLOCK_SIZE, WC), shRoundUp(BLOCK_SIZE, workSize[0])};
52
53    // Launch kernels on devices
54    for(unsigned int i = 0; i < clDeviceCount; i++)
55    {
56        // Multiplication - non-blocking execution
57        globalWorkSize[i] = shRoundUp(BLOCK_SIZE, workSize[i]);
58        clEnqueueNDRangeKernel(commandQueue[i], multiplicationKernel[i], 2, 0, globalWorkSize, localWorkSize,
59                              0, NULL, &GPUExecution[i]);
60    }
61
62    for(unsigned int i = 0; i < clDeviceCount; i++)
63    {
64        clFinish(commandQueue[i]);
65    }
66
67    for(unsigned int i = 0; i < clDeviceCount; i++)
68    {
69        // Non-blocking copy of result from device to host
70        clEnqueueReadBuffer(commandQueue[i], d_C[i], CL_FALSE, 0, WC * sizeof(float) * workSize[i],
71                            h_C + workOffset[i] * WC, 0, NULL, &GPUDone[i]);
72    }
73
74    // CPU sync with GPU
75    clWaitForEvents(clDeviceCount, GPUDone);
76
77    // Release mem and event objects
78    for(unsigned int i = 0; i < clDeviceCount; i++)
79    {
80        clReleaseMemObject(d_A[i]);
81        clReleaseMemObject(d_C[i]);
82        clReleaseMemObject(d_B[i]);
83        clReleaseEvent(GPUExecution[i]);
84        clReleaseEvent(GPUDone[i]);
85    }
86
87    kernel void
88    matrixMul( __global float* C, __global float* A, __global float* B,
89              __local float* Aa, __local float* Bb)
90    {
91        int bx = get_group_id(0), tx = get_local_id(0);
92        int by = get_group_id(1), ty = get_local_id(1);
93        int aEnd = NA * BLOCK_SIZE * by + NA - 1;
94        float Csub = 0.0f;
95
96        for (int a = WA*BLOCK_SIZE*by, b = BLOCK_SIZE * bx;
97             a <= aEnd; a += BLOCK_SIZE, b += BLOCK_SIZE*WB) {
98            Aa[tx + ty * BLOCK_SIZE] = A[a + WA * ty + tx];
99            Bb[ty + ty * BLOCK_SIZE] = B[b + WB * ty + tx];
100            barriers(CLK_LOCAL_MEM_FENCE);
101            for (int k = 0; k < BLOCK_SIZE; ++k)
102                Csub += Aa[k + ty * BLOCK_SIZE]*Bb[tx + k * BLOCK_SIZE];
103            barriers(CLK_LOCAL_MEM_FENCE);
104        }
105        C[get_global_id(1) * get_global_id(0) + get_global_id(0)] = Csub;
106    }
107 }
```

Directives

CUDA C

OpenCL

# Talk Roadmap

- Introduction to OpenACC
- Compute Constructs
- Loop Directives
- Demo – Himeno
- Data Region
- Porting Example – Seismic CPML
- Future OpenACC features

# Kernels Construct

- C

```
#pragma acc kernels clause...  
{  
    for( i = 0; i < n; ++i ) r[i] = a[i]*2.0f;  
}
```

- Fortran

```
!$acc kernels clause...  
    do i = 1,n  
        r(i) = a(i) * 2.0  
    enddo  
!$acc end kernels
```

# Parallel Construct

- C

```
#pragma acc parallel clause...  
{  
    #pragma acc loop gang vector  
    for( i = 0; i < n; ++i ) r[i] = a[i]*2.0f;  
}
```

- Fortran

```
!$acc parallel clause...  
!$acc loop gang vector  
    do i = 1,n  
        r(i) = a(i) * 2.0  
    enddo  
!$acc end parallel
```

# Kernels vs. Parallel

- Kernels Construct
  - Derived from PGI Accelerator Model
  - More implicit giving the compiler more freedom to create optimal code for a given accelerator
  - Works best for tightly nested loops
  - May require some additional 'hints' to the compiler
    - i.e. C99 restrict keyword
- Parallel Construct
  - Based on OpenMP "workshare"
  - Create parallel gangs that execute redundantly
  - Each gang executes a portion of a work-sharing loop
  - More explicit requiring some user intervention

<http://www.pgroup.com/lit/articles/insider/v4n2a1.htm>

# Talk Roadmap

- Introduction to OpenACC
- Compute Constructs
- Loop Directives
- Demo – Himeno
- Data Region
- Porting Example – Seismic CPML
- Future OpenACC features

# Loop Directive

- C

```
#pragma acc loop clause...  
for( i = 0; i < n; ++i ){  
    ....  
}
```

- Fortran

```
!$acc loop clause...  
do i = 1, n
```

Note: Compute Constructs and the Loop directive may be combined



# Loop Directive Clauses

- `independent`
  - use with care, overrides compiler analysis for dependence, private variables (kernels only, implied with `parallel`)
- `private( list )`
  - private data for each iteration of the loop
- `reduction( red:var )`
  - reduction across the loop
- **Scheduling Clauses**
  - `vector or vector(width)`
  - `gang or gang(width)`
  - `worker or worker(width)`
  - `seq`

# Loop Scheduling Clauses

- `!$acc loop gang`
  - runs in 'gang' mode only (`blockIdx`)
  - does not declare that the loop is in fact parallel (use independent)
- `!$acc loop gang(32)`
  - runs in 'parallel' mode only with `gridDim == 32` (32 blocks)
- `!$acc loop vector(128)`
  - runs in 'vector' mode (`threadIdx`) with `blockDim == 128`  
(128 threads)
  - vector size, if present, must be compile-time constant
- `!$acc loop gang vector(128)`
  - strip mines loop
  - inner loop runs in vector mode, 128 threads (`threadIdx`)
  - outer loop runs across thread blocks (`blockIdx`)

# Time for a Demo

- Himeno
  - Compute constructs
  - Loop schedules
  - Compiler Feed-back messages
  - Profiling

# Talk Roadmap

- Introduction to OpenACC
- Compute Constructs
- Loop Directives
- Demo – Himeno
- Data Region
- Porting Example – Seismic CPML
- Future OpenACC features

# Data Region

- C

```
#pragma acc data  
{  
    . . . .  
}
```

- Fortran

```
!$acc data  
    . . . .  
!$acc end data
```

- May span across host code and multiple compute regions
- May be nested
- May not be nested within a compute region
- Data is not implicitly synchronized between the host and device

# Data Clauses

- Data clauses

- `copy( list )`
- `copyin( list )`
- `copyout( list )`
- `create( list )`
- `present( list )`
- `present_or_copy(list)`      `pcopy(list)`
- `present_or_copyin(list)`      `pcopyin(list)`
- `present_or_copyout(list)`      `pcopyout(list)`
- `present_or_create(list)`      `pcreate(list)`
- `deviceptr( list )`

# Data Attributes

- predetermined data attributes
  - loop variables are private
- implicit data attributes
  - `array, struct – present_or_copy`
  - `scalar – firstprivate`
- explicit data attributes
  - in a data clause

# Data Regions Across Procedures

```
subroutine sub( a, b )  
  real :: a(:), b(:)  
  !$acc kernels pcopyin(b)  
    do i = 1,n  
      a(i) = a(i) * b(i)  
    enddo  
  !$acc end kernels  
  ...  
end subroutine
```

```
subroutine bus(x, y)  
  real :: x(:), y(:)  
  !$acc data copy(x)  
  call sub( x, y )  
  !$acc end data
```



# Data Regions Across Procedures

```
void sub( float* a, float* b, int n ){
    int i;
    #pragma acc kernels pcopyin(b[0:n])
        for( i = 0; i < n; ++i )
            a[i] *= b[i];
    ...
}
```

```
void bus( float* x, float* y, int n ){
    #pragma acc data copy(x[0:n])
    {
        sub( x, y, n );
    }
    ...
}
```

# Update Directives

- `update host( list )`
- `update device( list )`
  - data must be in a data allocate clause for an enclosing data region
  - both may be on a single line
    - `update host(list) device(list)`
- **Data update clauses on data construct (PGI)**
  - `updatein( list )` or `update device( list )`
  - `updateout( list )` or `update host( list )`
  - shorthand for update directive just inside data construct

# Additional Concepts not covered today

- Asynchronous Data Movement
- Asynchronous Compute
- Wait Directive
- Cache Directive
- Host\_data Directive
- OpenACC and CUDA interoperability
- OpenACC Runtime Library Calls
- OpenACC Environment Variables
- Multi-device programming
- Multi-target accelerators (Unified Binary)
- Performance Optimization
- Data Layout

# Talk Roadmap

- Introduction to OpenACC
- Compute Constructs
- Loop Directives
- Demo – Himeno
- Data Region
- **Porting Example – Seismic CPML**
- Future OpenACC features

# Porting SEISMIC CPML

- Set of ten open-source Fortran programs
- Solves two-dimensional or three-dimensional isotropic or anisotropic elastic, viscoelastic or poroelastic wave equation.
- Uses finite-difference method with Convolutional or Auxiliary Perfectly Matched Layer (C-PML or ADE-PML) conditions
- Developed by Dimitri Komatitsch and Roland Martin from University of Pau, France.
- Accelerated source used is taken from the 3D elastic finite-difference code in velocity and stress formulation with Convolutional-PML (C-PML) absorbing conditions.

[http://www.geodynamics.org/cig/software/seismic\\_cpml](http://www.geodynamics.org/cig/software/seismic_cpml)

Full Article: <http://www.pgroup.com/lit/articles/insider/v4n1a3.htm>

# Step 1: Evaluation

- Is my algorithm right for a GPU?
  - SEISMIC\_CPML models seismic waves through the earth. Has an outer time step loop with 9 inner parallel loops. Uses MPI and OpenMP parallelization.
- Good candidate for the GPU, but not ideal.

# Step 2: Add Compute Regions

## !\$acc kernels

```
do k = kmin,kmax
do j = NPOINTS_PML+1, NY-NPOINTS_PML
do i = NPOINTS_PML+1, NX-NPOINTS_PML
total_energy_kinetic = total_energy_kinetic + 0.5d0 * rho*(vx(i,j,k)**2 + vy(i,j,k)**2 + vz(i,j,k)**2)
epsilon_xx = ((lambda + 2.d0*mu) * sigmaxx(i,j,k) - lambda * sigmayy(i,j,k) - lambda*sigmazz(i,j,k)) / (4.d0 * mu * (lambda + mu))
epsilon_yy = ((lambda + 2.d0*mu) * sigmayy(i,j,k) - lambda * sigmaxx(i,j,k) - lambda*sigmazz(i,j,k)) / (4.d0 * mu * (lambda + mu))
epsilon_zz = ((lambda + 2.d0*mu) * sigmazz(i,j,k) - lambda * sigmaxx(i,j,k) - lambda*sigmayy(i,j,k)) / (4.d0 * mu * (lambda + mu))
epsilon_xy = sigmaxy(i,j,k) / (2.d0 * mu)
epsilon_xz = sigmaxz(i,j,k) / (2.d0 * mu)
epsilon_yz = sigmayz(i,j,k) / (2.d0 * mu)
total_energy_potential = total_energy_potential + 0.5d0 * (epsilon_xx * sigmaxx(i,j,k) + epsilon_yy * sigmayy(i,j,k) + &
epsilon_yy * sigmayy(i,j,k)+ 2.d0 * epsilon_xy * sigmaxy(i,j,k) + 2.d0*epsilon_xz * sigmaxz(i,j,k)+2.d0*epsilon_yz * sigmayz(i,j,k))
enddo
enddo
enddo
```

## !\$acc end kernels

# Compiler Feedback

```
% pgfortran -Mmpi=mpich2 -fast -acc -Minfo=accel
seismic_CPML_3D_isotropic_MPI_OACC_1.F90 -o gpu1.out
seismic_cpml_3d_iso_mpi_openmp:
  1107, Generating copyin(vz(11:91,11:631,kmin:kmax))
      Generating copyin(vy(11:91,11:631,kmin:kmax))
      Generating copyin(vx(11:91,11:631,kmin:kmax))
      Generating copyin(sigmmaxx(11:91,11:631,kmin:kmax))
      Generating copyin(sigmayy(11:91,11:631,kmin:kmax))
      Generating copyin(sigmazz(11:91,11:631,kmin:kmax))
      Generating copyin(sigmaxy(11:91,11:631,kmin:kmax))
      Generating copyin(sigmazx(11:91,11:631,kmin:kmax))
      Generating copyin(sigmayz(11:91,11:631,kmin:kmax))
      Generating compute capability 1.3 binary
      Generating compute capability 2.0 binary
```



```
1108, Loop is parallelizable
1109, Loop is parallelizable
1110, Loop is parallelizable
    Accelerator kernel generated
    1108, !$acc do gang, vector(4)
    1109, !$acc do gang, vector(4)
    1110, !$acc do vector(16)
1116, Sum reduction generated for total_energy_kinetic
1134, Sum reduction generated for total_energy_potential
```

# Initial Timings

Version	MPI Processes	OpenMP Threads	GPUs	Time (sec)	Approx. Programming Time (min)
Original MPI/OMP	2	4	0	951	
ACC Step 1	2	0	2	3031	10

System Info:  
4 Core Intel Core-i7 920 Running at 2.67Ghz  
Includes 2 Tesla C2070 GPUs  
Problem Size: 101x64x1x128

Why the slowdown?

# Step 3: Optimize Data Movement

```
!$acc data &
!$acc      copyin(a_x_half,b_x_half,k_x_half, &
!$acc          a_y_half,b_y_half,k_y_half, &
!$acc          a_z_half,b_z_half,k_z_half, &
!$acc          a_x,a_y,a_z,b_x,b_y,b_z,k_x,k_y,k_z, &
!$acc          sigmaxx,sigmaxz,sigmaxy,sigmayy,sigmayz,sigmazz, &
!$acc          memory_dvx_dx,memory_dvy_dx,memory_dvz_dx, &
!$acc          memory_dvx_dy,memory_dvy_dy,memory_dvz_dy, &
!$acc          memory_dvx_dz,memory_dvy_dz,memory_dvz_dz, &
!$acc          memory_dsigmaxx_dx, memory_dsigmaxy_dy, &
!$acc          memory_dsigmaxz_dz, memory_dsigmaxy_dx, &
!$acc          memory_dsigmaxz_dx, memory_dsigmayz_dy, &
!$acc          memory_dsigmayy_dy, memory_dsigmayz_dz, &
!$acc          memory_dsigmazz_dz)

      do it = 1,NSTEP
.... Cut ....
      enddo
!$acc end data
```

# Timings Continued

Version	MPI Processes	OpenMP Threads	GPUs	Time (sec)	Approx. Programming Time (min)
Original MPI/OMP	2	4	0	951	
ACC Step 1	2	0	2	3031	10
ACC Step 2	2	0	2	124	180

System Info:  
4 Core Intel Core-i7 920 Running at 2.67Ghz  
Includes 2 Tesla C2070 GPUs

Data movement  
time now only 5  
seconds!

# Step 4: Fine Tune Schedule

```
!$acc do vector(4)
  do k=k2begin,NZ_LOCAL
    kglobal = k + offset_k
!$acc do gang, vector(4)
    do j=2,NY
!$acc do gang, vector(16)
      do i=1,NX-1
        value_dvx_dx = (vx(i+1,j,k)-vx(i,j,k)) * ONE_OVER_DELTAX
        value_dvy_dy = (vy(i,j,k)-vy(i,j-1,k)) * ONE_OVER_DELTAY
        value_dvz_dz = (vz(i,j,k)-vz(i,j,k-1)) * ONE_OVER_DELTAZ
```

# Final Timings

Version	MPI Processes	OpenMP Threads	GPUs	Time (sec)	Approx. Programming Time (min)
Original MPI/OMP	2	4	0	951	
ACC Step 1	2	0	2	3031	10
ACC Step 2	2	0	2	124	180
ACC Step 3	2	0	2	120	120

7x in 5 Hours!

# Cluster Timings

---

Version	Size	MPI Processes	OpenMP Threads	GPUs	Time (sec)
MPI/OMP	101x641x3072	24	96	0	2081
MPI/ACC	101x641x3072	24	0	24	446

---

System Info: 8 Nodes

2 socket, 6 Core Intel X5675 Running at 3.06Ghz with 3 Tesla C2070 GPU

Compiler: PGI 2012 version 12.5

Just  
under 5x!

# Talk Roadmap

- Introduction to OpenACC
- Compute Constructs
- Loop Directives
- Demo – Himeno
- Data Region
- Porting Example – Seismic CPML
- Future OpenACC features



# OpenACC 2.0 Highlights

- Procedure calls, separate compilation
- Nested parallelism
- Data management features and global data
- atomic operations

# Currently Procedure Calls in Compute Regions Must Be Inlined

```
#pragma acc parallel loop num_gangs(200)..  
for( int i = 0; i < n; ++i ){  
    v[i] += rhs[i];  
    matvec( v, x, a, i, n );  
    // must inline matvec  
}
```

```
pgcc -Minline a.c  
pgcc -Mextract=lib:mylib b.c  
pgcc -Minline=lib:mylib a.c  
pgcc -Minline=levels:10 a.c  
pgcc -Minline=levels:2,foo,phoo,bar,lib:mylib a.c
```

# OpenACC 2.0 Function Calling

```
#pragma acc routine worker
extern void matvec(float* v, float* x, ... );
...
#pragma acc parallel loop num_gangs(200)...
for( int i = 0; i < n; ++i ){
    v[i] += rhs[i];
    matvec( v, x, a, i, n );
    // procedure call on the device
}

#pragma acc routine worker
void matvec( float* v, float* x,
            float* a, int i,
            int n ){
    float xx = 0;
    #pragma acc loop reduction(+:xx)
    for( int j = 0; j < n; ++j )
        xx += a[i*n+j]*v[j];
    x[i] = xx;
}
```

# OpenACC 2.0 routine bind

```
#pragma acc routine worker bind(mvdev)
extern void matvec(float* v, float* x, ... );
...
#pragma acc parallel loop num_gangs(200)...
for( int i = 0; i < n; ++i ){
    v[i] += rhs[i];
    matvec( v, x, a, i, n );
}
```

# OpenACC 2.0 routine bind cont.

```
void matvec( float* v, float* x,  
            float* a, int i, int n ){  
    float xx=0.0;  
  
    for( int j = 0; j < n; ++j )  
        xx += a[i*n+j]*v[j];  
    x[i] = xx;  
}
```

```
#pragma acc routine worker nohost  
void mvdev( float* v, float* x,  
           float* a, int i, int n ){  
    float xx = 0.0;  
    #pragma acc loop reduction(+:xx)  
    for( int j = 0; j < n; ++j )  
        xx += a[i*n+j]*v[j];  
    x[i] = xx;  
}
```

# Nested Parallelism

```
#pragma acc routine  
extern void matvec(float* v, float* x, ... );  
  
...  
#pragma acc parallel loop ...  
for( int i = 0; i < n; ++i )  
    matvec( v, x, i, n );
```

```
#pragma acc routine  
matvec(...){  
    #pragma acc parallel loop  
    for( int i = 0; i < n; ++i ){...}
```

# Nested Parallelism cont.

```
#pragma acc routine
extern void matvec(float* v, float* x, ... );
...
#pragma acc parallel num_gangs(1)
{
    matvec( v0, x0, i, n );
    matvec( v1, x1, i, n );
    matvec( v2, x2, i, n );
}

#pragma acc routine
    matvec(...) {
#pragma acc parallel loop
    for( int i = 0; i < n; ++i ){...}
```

# Dynamic Data Lifetimes

```
void init( int n ){  
...  
#pragma acc enter data copyin( x[0:100] )  
...  
#pragma acc enter data present_or_create( y[0:n] )  
...  
}
```

```
void fini( int n ){  
...  
#pragma acc exit data delete( x[0:100] )  
...  
#pragma acc exit data copyout( y[0:n] )  
...  
}
```



# Global Data

```
float x[1000];  
#pragma acc declare create(x)  
// static allocation, host + device
```

```
float y[1000];  
#pragma acc declare device_resident(y)  
// static allocation, device-only
```

```
float z[10000];  
#pragma acc declare link(z)  
// static allocation on host, dynamic allocation on device
```

# atomic operations

```
#pragma acc parallel loop
  for( j = 1; j < n-1; ++j ){
    y = x[j];
    i = y & 0xf;
    #pragma acc atomic update
      ++bin[i];
  }
```

// essentially the OpenMP atomic operations  
// atomic update, read, write, capture  
// only native data lengths supported

# Other OpenACC 2.0 Features Not Covered Today

- Device-specific tuning, multiple devices
- Multiple host thread support
- Loop directive additions
- Asynchronous behavior additions
- New API routines

# Key Value-add of the PGI Accelerator Compilers

PGI compilers are designed to enable performance-portable programming for all many-core and GPU Accelerator-based systems *without having to re-program for each new successive hardware advancement.*

# Copyright Notice

© Contents copyright 2013, The Portland Group, Inc. This material may not be reproduced in any manner without the expressed written permission of The Portland Group.

PGFORTRAN, PGF95, PGI Accelerator and PGI Unified Binary are trademarks; and PGI, PGCC, PGC++, PGI Visual Fortran, PVF, PGI CDK, Cluster Development Kit, PGPROF, PGDBG, and The Portland Group are registered trademarks of The Portland Group Incorporated. Other brands and names are the property of their respective owners.