

# OpenACC Fundamentals

Steve Abbott <[sabbott@nvidia.com](mailto:sabbott@nvidia.com)>, November 13, 2016



# Who Am I?

2005 - B.S. Physics - Beloit College

2007 - M.S. Physics - University of Florida

2015 - Ph.D. Physics - University of New Hampshire

2015 - 2017 - Postdoctoral Associate - Oak Ridge National Lab

Ported and optimized fusion simulation code for Summit

August - Present - NVIDIA Corp.

Support deployment and use of Summit supercomputer

Problem solving, training, and optimizing science for GPUs

# AGENDA

What is OpenACC?

OpenACC by Example

# 3 Ways to Program GPUs

## Applications

**Libraries**

“Drop-in”  
Acceleration

**Compiler  
Directives**

Easily Accelerate  
Applications

**Programming  
Languages**

Maximum  
Flexibility

# OpenACC Directives

Manage  
Data  
Movement → `#pragma acc data copyin(x,y) copyout(z)`  
{  
...  
Initiate  
Parallel  
Execution → `#pragma acc parallel`  
{  
Optimize  
Loop  
Mappings → `#pragma acc loop gang vector`  
for (i = 0; i < n; ++i) {  
z[i] = x[i] + y[i];  
...  
}  
}  
...  
}

**OpenACC**  
Directives for Accelerators


- Incremental
- Single source
- Interoperable
- Performance portable
- CPU, GPU, MIC

# Accelerated Computing

*10x Performance & 5x Energy Efficiency for HPC*




**CPU**

Optimized for  
Serial Tasks

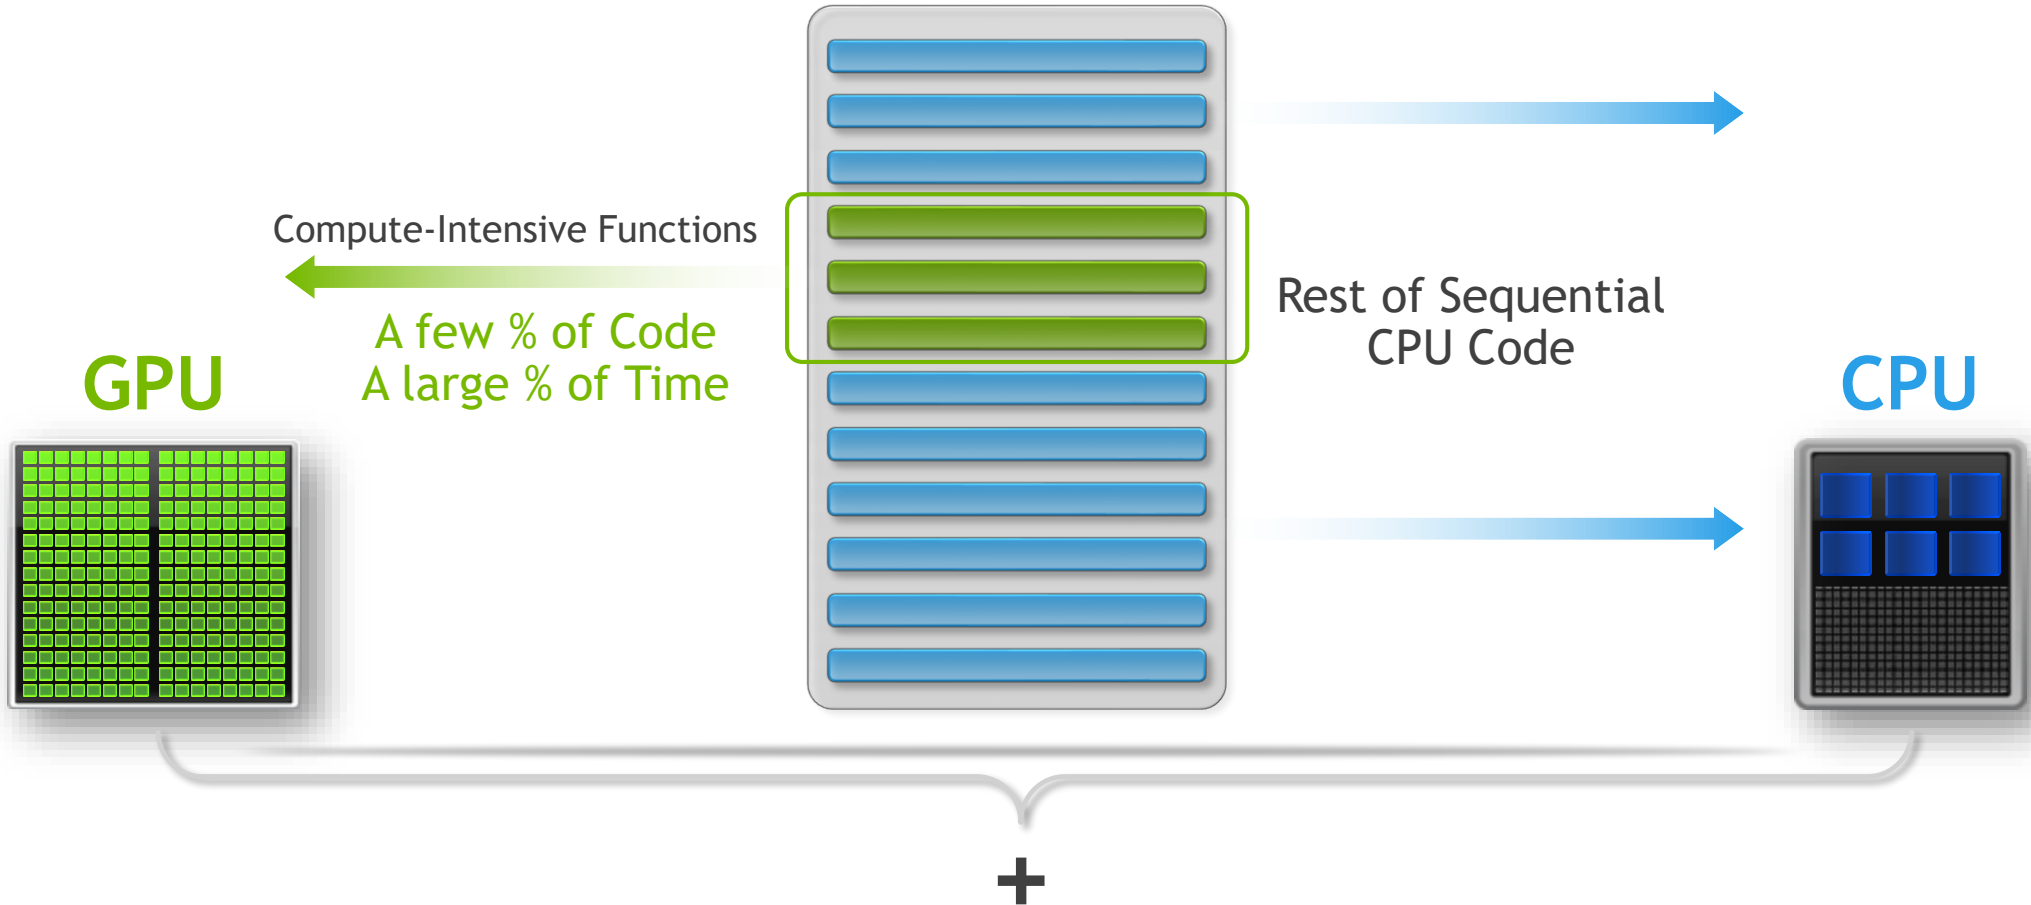


**GPU Accelerator**

Optimized for  
Parallel Tasks





# What is Accelerated Computing?



# 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];
    }
  }
}
```






# Example: Jacobi Iteration

Iteratively converges to correct value (e.g. Temperature), by computing new values at each point from the average of neighboring points.

Common, useful algorithm

Example: Solve Laplace equation in 2D:  $\nabla^2 f(x, y) = 0$



$$A_{k+1}(i, j) = \frac{A_k(i-1, j) + A_k(i+1, j) + A_k(i, j-1) + A_k(i, j+1)}{4}$$

# Jacobi Iteration: C Code

```
while ( err > tol && iter < iter_max ) {  
    err=0.0;  
  
    for( int j = 1; j < n-1; j++) {  
        for(int i = 1; i < m-1; i++) {  
  
            Anew[j][i] = 0.25 * (A[j][i+1] + A[j][i-1] +  
                                A[j-1][i] + A[j+1][i]);  
  
            err = max(err, abs(Anew[j][i] - A[j][i]));  
        }  
    }  
  
    for( int j = 1; j < n-1; j++) {  
        for( int i = 1; i < m-1; i++ ) {  
            A[j][i] = Anew[j][i];  
        }  
    }  
  
    iter++;  
}
```



Iterate until converged



Iterate across matrix elements



Calculate new value from neighbors



Compute max error for convergence



Swap input/output arrays

# Look For Parallelism

```
while ( err > tol && iter < iter_max ) {  
    err=0.0;  
  
    for( int j = 1; j < n-1; j++) {  
        for(int i = 1; i < m-1; i++) {  
  
            Anew[j][i] = 0.25 * (A[j][i+1] + A[j][i-1] +  
                                A[j-1][i] + A[j+1][i]);  
  
            err = max(err, abs(Anew[j][i] - A[j][i]));  
        }  
    }  
  
    for( int j = 1; j < n-1; j++) {  
        for( int i = 1; i < m-1; i++ ) {  
            A[j][i] = Anew[j][i];  
        }  
    }  
  
    iter++;  
}
```



Data dependency  
between iterations.



Independent loop  
iterations



Max Reduction required



Independent loop  
iterations

# OPENACC DIRECTIVE SYNTAX

C/C++

```
#pragma acc directive [clause [,] clause] ...]
```

...often followed by a structured code block

Fortran

```
!$acc directive [clause [,] clause] ...]
```

...often paired with a matching end directive surrounding a structured code block:

```
!$acc end directive
```



Don't forget acc

# OpenACC Parallel Directive

Generates parallelism

```
#pragma acc parallel
```

```
{
```

When encountering the *parallel* directive, the compiler will generate *1 or more parallel gangs*, which execute redundantly.

```
}
```

# OpenACC Parallel Directive

Generates parallelism


```
#pragma acc parallel
```

```
{
```



```
}
```

When encountering the *parallel* directive, the compiler will generate *1 or more parallel gangs*, which execute redundantly.



# OpenACC Loop Directive

Identifies loops to run in parallel

```
#pragma acc parallel
```

```
{
```




```
#pragma acc loop
```

```
for (i=0;i<N;i++)
```

```
{
```

The *loop* directive informs the compiler which loops to parallelize.

```
}
```



# OpenACC Loop Directive

Identifies loops to run in parallel

```
#pragma acc parallel
```

```
{
```


```
#pragma acc loop
```

```
for (i=0;i<N;i++)
```

```
{  
    The loop directive  
    informs the compiler  
    which loops to  
    parallelize.  
}
```



```
}
```





# OpenACC Parallel Loop Directive

Generates parallelism and identifies loop in one directive


```
#pragma acc parallel loop
```

```
for (i=0;i<N;i++)
```

```
{
```

```
}
```

The *parallel* and *loop* directives are frequently combined into one.



# PARALLELIZE WITH OPENACC

```
while ( err > tol && iter < iter_max ) {  
    err=0.0;  
  
    #pragma acc parallel loop reduction(max:err)  
    for( int j = 1; j < n-1; j++) {  
        for(int i = 1; i < m-1; i++) {  
  
            Anew[j][i] = 0.25 * (A[j][i+1] + A[j][i-1] +  
                                A[j-1][i] + A[j+1][i]);  
  
            err = max(err, abs(Anew[j][i] - A[j][i]));  
        }  
    }  
}
```



Parallelize loop on  
accelerator

```
    #pragma acc parallel loop  
    for( int j = 1; j < n-1; j++) {  
        for( int i = 1; i < m-1; i++ ) {  
            A[j][i] = Anew[j][i];  
        }  
    }  
  
    iter++;  
}
```



Parallelize loop on  
accelerator

\* A *reduction* means that all of the N\*M values for err will be reduced to just one, the max.

# OPENACC LOOP DIRECTIVE: PRIVATE & REDUCTION

The **private** and **reduction** clauses are not optimization clauses, they may be required for correctness.

**private** – A copy of the variable is made for each loop iteration

**reduction** – A reduction is performed on the listed variables.

Supports +, \*, max, min, and various logical operations

# BUILDING THE CODE

```
$ pgcc -fast -acc -ta=tesla -Minfo=all laplace2d.c
main:
  40, Loop not fused: function call before adjacent loop
      Generated vector sse code for the loop
  51, Loop not vectorized/parallelized: potential early exits
  55, Accelerator kernel generated
      55, Max reduction generated for error
      56, #pragma acc loop gang /* blockIdx.x */
      58, #pragma acc loop vector(256) /* threadIdx.x */
  55, Generating copyout(Anew[1:4094][1:4094])
      Generating copyin(A[:,:])
      Generating Tesla code
  58, Loop is parallelizable
  66, Accelerator kernel generated
      67, #pragma acc loop gang /* blockIdx.x */
      69, #pragma acc loop vector(256) /* threadIdx.x */
  66, Generating copyin(Anew[1:4094][1:4094])
      Generating copyout(A[1:4094][1:4094])
      Generating Tesla code
  69, Loop is parallelizable
```

# BUILDING THE CODE

```
$ pgcc -fast -acc -ta=tesla -Minfo=all laplace2d.c
main:
  40, Loop not fused: function call before adjacent loop
      Generated vector sse code for the loop
  51, Loop not vectorized/parallelized: potential early exits
  55, Accelerator kernel generated
      55, Max reduction generated for error
      56, #pragma acc loop gang /* blockIdx.x */
      58, #pragma acc loop vector(256) /* threadIdx.x */
  55, Generating copyout(Anew[1:4094][1:4094])
      Generating copyin(A[:,:])
      Generating Tesla code
  58, Loop is parallelizable
  66, Accelerator kernel generated
      67, #pragma acc loop gang /* blockIdx.x */
      69, #pragma acc loop vector(256) /* threadIdx.x */
  66, Generating copyin(Anew[1:4094][1:4094])
      Generating copyout(A[1:4094][1:4094])
      Generating Tesla code
  69, Loop is parallelizable
```

# BUILDING THE CODE


```
$ pgcc -fast -acc -ta=tesla -Minfo=all laplace2d.c
main:
  40, Loop not fused: function call before adjacent loop
      Generated vector sse code for the loop
  51, Loop not vectorized/parallelized: potential early exits
  55, Accelerator kernel generated
      55, Max reduction generated for error
      56, #pragma acc loop gang /* blockIdx.x */
      58, #pragma acc loop vector(256) /* threadIdx.x */
  55, Generating copyout(Anew[1:4094][1:4094])
      Generating copyin(A[:,:])
      Generating Tesla code
  58, Loop is parallelizable
  66, Accelerator kernel generated
      67, #pragma acc loop gang /* blockIdx.x */
      69, #pragma acc loop vector(256) /* threadIdx.x */
  66, Generating copyin(Anew[1:4094][1:4094])
      Generating copyout(A[1:4094][1:4094])
      Generating Tesla code
  69, Loop is parallelizable
```

# BUILDING THE CODE

```
$ pgcc -fast -acc -ta=tesla -Minfo=all laplace2d.c
main:
  40, Loop not fused: function call before adjacent loop
      Generated vector sse code for the loop
  51, Loop not vectorized/parallelized: potential early exits
  55, Accelerator kernel generated
      55, Max reduction generated for error
      56, #pragma acc loop gang /* blockIdx.x */
      58, #pragma acc loop vector(256) /* threadIdx.x */
  55, Generating copyout(Anew[1:4094][1:4094])
      Generating copyin(A[:,:])
      Generating Tesla code
  58, Loop is parallelizable
  66, Accelerator kernel generated
      67, #pragma acc loop gang /* blockIdx.x */
      69, #pragma acc loop vector(256) /* threadIdx.x */
  66, Generating copyin(Anew[1:4094][1:4094])
      Generating copyout(A[1:4094][1:4094])
      Generating Tesla code
  69, Loop is parallelizable
```

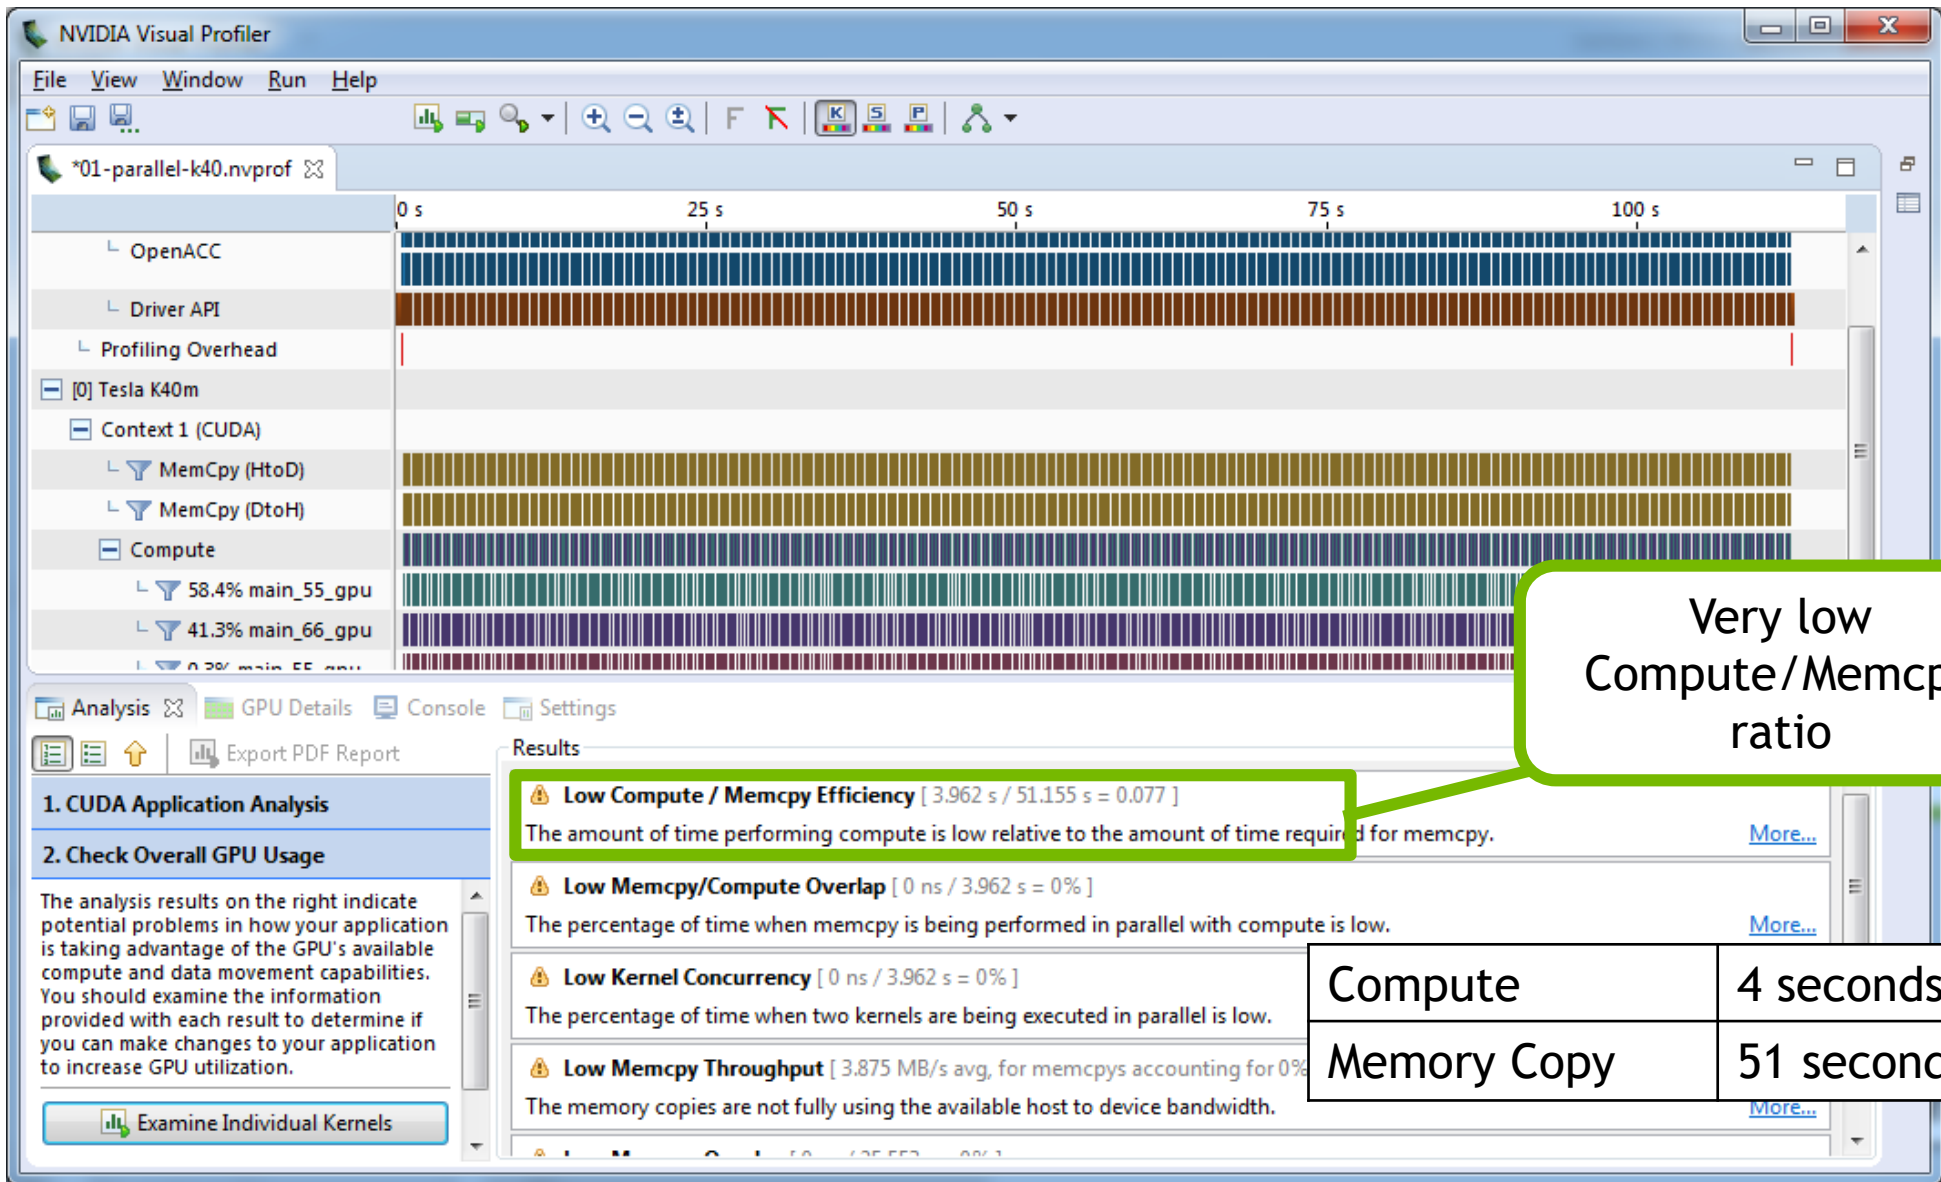
## Speed-up (Higher is Better)

Intel Xeon E5-2698 v3 @ 2.30GHz (Haswell)  
vs.  
NVIDIA Tesla K40 & P100



Why did OpenACC slow down here?





Very low  
Compute/Memcpy  
ratio

**Low Compute / Memcpy Efficiency [ 3.962 s / 51.155 s = 0.077 ]**  
 The amount of time performing compute is low relative to the amount of time required for memcpy.

|             |            |
|-------------|------------|
| Compute     | 4 seconds  |
| Memory Copy | 51 seconds |

# Rest of the Week

Wednesday - Data directives

Friday - Advanced OpenACC and Tuning