

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



#	Question	Answer
1	Could I copy via this directives whole containers from standard library like sets or vectors?	Yes, its possible copy all the data in a container. Complex data structures can be handled also.
2	Can I get somehow online access to some GPUs for practice?	The qwiklab session are being covered in this course. Please refer to the slides.
3	Can it be a good way of acceleration when at first step we use OpenACC and only after that use CUDA for some parts of code that wasn't parallelized so good as it could be?	Yes, that is a good approach, and supported by OpenACC
4	How does number of cores/GPUs influence the behavior of OpenACC directives? In terms of parallelization? Or is it completely opaque?	It should be mostly Opaque. In general, larger more powerful GPUs generally have more cores, so codes will usually run faster on these GPUs.
6	Is there an out of box solution in OpenACC to use multiple GPU nodes in a cluster or do i need to combine it with OpenMPI (that's how I do it right now)	OpenACC can target multiple GPUs in the same node. For distribution of work across nodes, you would typically use MPI
7	Is it still true that doubles slow the code running on GPU heavily or is this an outdated information?	It depends on the GPU and the code.
8	PGPROF is only running on xserver?	there are command line versions also
9	Is the any option to monitor OpenACC code running on GPU using perf?	You can monitor it using available host and GPU monitoring tools, such as top and nvidia-smi
10	Any special reason for choosing PGI over GCC / gnu compilers?	OpenACC will be available with GCC soon, Mentor is working on the implementation as we speak. But for now this version is not available yet.
11	Can we form a cluster of GPU with large memory and use distributed computing? Much like we have clusters with over million cores and have to use MPI for data management and parallelism?	GPUs within a node can be managed from a single program. Multi-node usage will still require a technology like MPI
12	Does OpenACC work with multiple nodes? Does MPI work with GPUs?	You would typically use MPI. MPI and OpenACC can work together
13	Is PGProf open source? Can it be used for python?	it is not open source
14	Can we accelerate with both CPU and GPU?	yes
15	When we use OpenACC with GPU, is multicore CPU idle?	If you target the GPU at compile time, the OpenACC directives will use the GPU exclusively for parallelization
16	Does OpenACC supports acceleration with Xeon Phi?	not today, planned for the near future
17	Does OpenACC optimize the CPU code, to avoid cache misses?	Yes, CPU code is optimized by the PGI compiler
18	Is OpenACC compatible with AMD GPU?	Yes, OpenACC with PGI can target AMD GPUs today
19	How would OpenACC be handle dependencies in scenarios where code contains many loops with outer dependencies?	Data dependencies would normally be handled by the programmer using data directives. Loop-carried dependencies the programmer would normally have

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



		to work to eliminate via code refactoring. Loop-carried dependencies can prevent OpenACC parallelization
20	How big performance differences are between code manually refined with OpenACC vs CUDA? If it covered in webinar no need to respond.	For some codes there will be little difference. Many codes can be refactored to take advantage of OpenACC so that there is little performance difference between OpenACC and hand-optimized CUDA, for example.
21	Is OpenACC free standard? How high costs are involved with PGI compilers and Nvidia compilers, optimization software?	It is a fee and open standard. You can visit OpenACC.org to learn more about the organization and the standard.
22	Is PGI compiler going to replace GCC C++ compiler?	no, both serve various purposes
23	What C compilers are currently optimizing best? Is syntax identical between various compilers?	OpenACC syntax should be identical. PGI and Cray are both good compiler choices today for OpenACC
24	Is the lab usable for all?	the labs are available to everyone
25	About how many people are currently estimated to practice OpenACC? How old is the standard?	the standard has been around since about 2010
26	I am new to mathematical code as a whole however I am slightly experienced in Fortran and I've had to read only C theory. Will I be lost everyone seems ok with being new?	you don't need deep knowledge to join the course
27	CUDA only for GPU's and OPENACC GPUS, CPU's?	OpenACC targets NVIDIA and AMD GPUs, and x86 CPUs.
28	Is assembly level changing of optimization possible?	it's possible, but not typical
29	Say i have 4 cores and a GPU, I use # pragma parallel, is it loaded to 4 CPU or to GPU by compiler???	It depends on how you compile it. If you target the host CPU, it will use the CPU cores. if you target the GPU, it will use the GPU
30	Does OpenACC support INTEL 128 bit SIMD?	Yes, this would be handled when targeting the x86 CPU
31	Would there be any issue with templated functions and OpenACC? Can we use OpenACC on templated functions?	Yes
32	Is there any data on the performance of unified memory versus explicit data statements?	Yes, data, exists, but it won't be covered in this introductory course
33	Any OpenACC for java?	today's OpenACC implementations target C/C++/Fortran
34	Can OpenACC optimized code be run on a cluster machine with multiple x86 sockets?	yes
35	Only C++ and Fortran are use in OpenACC?	correct, for the PGI compiler

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



36	Is CUDA Unified Memory (the fancy one that does not require explicit synchronization) already supported by OS? Last time when I read through the whitepaper, it says that it is not ready yet	CUDA 8 UM with demand paging is only supported by Pascal GPUs, but OS support exists today for it.
37	Is there a version of OpenACC for 32 bits Tegra models? like Tk1?	We are working on an ARM version of OpenACC with PGI, but it is not available yet.
38	Does ifort have options to give compiler feedback, comparable to -Minfo=all for PGI?	ifort doesn't support OpenACC
39	Which Fortran compilers besides PGI support OpenACC?	Cray and others are available.
40	How many ""cores"" (order of magnitude) does a modern graphics card provide? That is, if I'm currently using OpenMP for something like 32 CPU cores, and I use OpenACC instead on a machine with a suitable graphics card, how many ""cores"" would I typically get?	GPUs have thousands of cores comparing to CPUs. The exact number of cores will depend on a GPU model.
41	Is there a document which gives OpenACC equivalents of OMP directives?	There isn't really a 1:1 correspondence.
42	The shared memory is it implicitly managed by OpenACC?	It can be using the OpenACC cache directive
43	Regarding streams, OpenACC does support multiple streams and asynchronous copies?	Yes, these are supported and will be covered in a later session
44	Does OpenACC supports multiple streams and asynchronous copies?	Yes, supported
45	Does OpenACC supports complex C++ classes?	Yes, full C++ support is available.
46	Does the OpenACC manages the copy of class members and decorate the class methods to run on device and host?	Yes, under the hood
47	Does OpenACC support task-parallelism besides data parallelism	You can use task parallelism in a variety of ways using techniques you may be familiar with as well as techniques such as multi-GPU in OpenACC
48	Is there an equivalent to OpenMP parallel tasks or parallel sections in OpenACC?	The parallel section in OpenMP, if the code is running on a GPU, would be closest to the OpenACC parallel directive. But there are important differences.
49	Could I use OpenACC with Thrust?	Yes, interoperability is possible.
50	Is it possible to use iterators in the parallelized loops in OpenACC?	In general, C++ is supported. Some parts of the STL may not be supported, however.
51	Can I run OpenACC parallelized code on NVIDIA GRID cards (Kepler, Maxwell)?	Yes, it's possible.
52	Is there a basic minimum requirement for GPU to run an OpenACC code? Can it be AMD Radeon or Intel Iris or anything else	All current NVIDIA and Radeon GPUs should work
53	Any performance difference between OpenMP and OpenACC for non-GPU implementations?	In many cases, the perf is nearly identical. In some cases OpenACC can even be better.

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



54	Can OpenACC employ integrated graphics processors onboard Intel and AMD CPUs?	It's possible, but I'm not aware of any implementation that do that.
55	How many directives are there to refactor code?	At least 1, but it may require several
56	How many bytes are assigned for float type data in GPU?	4
57	I know somebody has already ask about Xeon Phi. I guess he means the pci-e based Xeon phi. There is also a self-bootable Xeon phi, works just like a regular CPU. Will OpenACC work on that?	Today, Xeon Phi is not supported by PGI compilers. In the future if/when support is announced, then this question can be answered at that point.
58	Does the unified memory access only work on the most recent GPUs?	It works on all Kepler and newer GPUs.
59	What compilers support OpenACC? Is it PGI, only?	PGI, Cray, GCC, and others
60	I have experience with MPI and OpenMP. Question: can we use hybrid MPI-OpenACC?	Yes, you can use MPI with OpenACC. This will be covered in a later session (not this one).
61	Is possible to have like a trial license to use without problems in a test cluster with 6 nodes? Or is necessary to buy the license?	Yes, please send a request to OpenACC@nvidia.com
62	Is it necessary for a pc to have a GPU?	It is not necessary to have a GPU to use OpenACC. OpenACC will work on a CPU as well.
63	Is OpenACC portable?	yes, it's portable
64	Does OpenACC have developer forums?	yes, there are PGI forums and you can also use stackoverflow
65	Does OpenACC have anything to do with NVIDIA Pascal architecture?	They are separate concepts, but OpenACC can certainly run well on Pascal GPUs.
66	Does the platform determine optimal number of gangs per CU on hardware?	The compiler chooses this the number of gangs, but you can adjust the decision with some OpenACC directives/clauses. This will be covered in a later session. The gangs per CU is not normally something you program to or worry about
67	Can you combine ACC and low level CUDA?	yes in several ways.
68	Does the number of Streaming Processors available on a GPU affects on OpenACC performance tuning? or it only depend on the total number of threads supported by the target GPU?	GPUs don't generally have a limit to the number of threads supported. There are various lower-level optimization possibilities with OpenACC that optimize workers per gang or vectors per gang
69	PGPROF can be used to find out area in c# or python code where OpenACC directive can be included?	yes - hotspot profiling
70	I only have some experience with OpenMP, can I understand this lecture?	Possibly. OpenMP and OpenACC are similar in various ways.
71	Would the PGI compiler supporting c++11, c++14, c++17?	c++11 - yes, others are coming
72	Is GCC 6.2 supports OpenACC?	OpenACC support is available to some degree in various branches of the gnu tool chain

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



73	Is it possible to use OpenACC instead OpenCL to accelerate code on FPGA?	not possible today
74	Which one is better OpenMP or OpenACC, in performance point of view?	They have different targets, largely. But if both are targeting the CPU, there may not be much difference in performance
75	Which compiler optimize better OpenACC implementation PGI or GCC?	PGI is farther along in OpenACC support than GCC at this point. support in both will improve over time
76	Can I use intel HD graphics 5500 to accelerate code with OpenACC?	not today
77	So OpenACC targets only Nvidia GPUs??	no it targets NVIDIA and AMD GPUs, x86 and POWER CPUs
78	Is there a way to integrate OpenACC directives into a .NET/C# program?	look at CUDAfy and managedCUDA
79	For beginners, what would you recommend us to use to accelerate large size PDEs? CUDA OpenCL or OpenACC, thanks!	I would start by looking for libraries. If you have a code that is already written in C/C++ or Fortran, it will probably be easier to get started with OpenACC than the other approaches
80	Is it worth attending if I haven't done coding in C++ or fortran at this stage?	It may be hard to follow the programming examples
81	Can OpenACC output a log of CPU/GPU acceleration split in terms of energy usage after a job has completed?	It can output various logs. I don't think they cover energy usage, but there are GPU tools that can monitor energy usage (e.g. nvidia-smi) which can work with oepnACC codes
82	I noticed the OpenACC Toolkit Quick Start Guide states that CUDA 7.5 Driver is required. Will this work with the CUDA 8.0 Driver?	It can be used with CUDA 8 as well.
83	Similarly, when buying the Toolkit license for a group, do I buy one license per computer or one per user?	Steven, could you send your question to OpenACC@nvidia.com? I will forward it to the PGI team.
84	Is the unified memory pool affected by having registered CPU memory and Unregistered GPU Mem (or vice versa)?	UM pool size is determined by device memory for pre-pascal GPUs, or else by available system memory for Pascal GPUs.
85	In fortran I would write this as $A = B + C$, conformable vectors. Would this be treated as a loop and parallelized?	It can be - you still have to decorate it with directives
86	Can OpenACC code run on FPGA	I'm not aware of any implementations that do this.
87	I am a C# developer. How would I best leverage parallel programing and/or GPU programming?	CUDAfy or managedCUDA, or else you can build your own non-managed interface (e.g. via DLL) to OpenACC or other technologies
88	As a C# developer, is the QuantAlea library Alea GPU V3 a good choice for GPU coding from .Net?	possibly but not related to OpenACC
89	Re: LSDALTON - is the GPU version mature? We worked with Patrick Ettenhuber on testing on non-cray and had issues	LSDalton GPU version is about 1 year old and is a part of the production code.

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



90	Does GCC compilers support OpenACC?	Yes, some OpenACC support is in some recent GCC versions
91	What's the promo code for Qwiklabs again?	the promo code for the labs is OPENACC
92	Will the optimization be able to handle switch statements and do while loops or just for loops?	Typically you need a loop that has a discoverable trip count - this will be covered later. It's possible to handle switch in some situations
93	Can I get source code?	Yes, stay tuned
94	Would OpenACC replace CUDA someday, or are these two models useful for differing reasons?	Probably won't replace CUDA. They serve different purposes. OpenMP won't replace C, for example.
95	What if the loops are not divisible equally into gangs. Does OpenACC allow you to have different numbers of threads in different gangs ?	This will be handled by the compiler - all gangs will have the same number of threads.
96	Does the programmer need to ensure that there are no over-lapping arrays before using these pragmas? Or will the pragmas watch for this condition?	typically you want to ensure there is no aliasing in C/C++
97	Is there a difference between OpenMP and OpenACC?	Yes, OpenMP targets CPUs only. OpenACC can target CPUs and accelerators. There are many other differences as well.
98	How to get credits for already registered qwiklab users	send email to OpenACC@nvidia.com
99	Does each gang work on a separate SM and each of the gang's workers run on the SM's CUDA cores?	One SM can host several gangs. Typically each worker would run on a CUDA core.
100	Can I control number of loops in a gang?	It's possible, will be covered later
101	Can you repeat how college students will get their certificates and when are we going to take the tests?	The certification test will be conducted online on your own time. We will be emailing certificates in PDF formats
102	When you say "soon" for GCC, Ms. Levites, do you mean "Agile Development" soon, or "Waterfall Development" soon? ;-)	Soon meaning within a year. A closed beta version is already available.
103	Is it possible to use OpenACC with FORTRAN	Yes, OpenACC works with FORTRAN, C and C++ codes
104	Can you use nvprof to profile CPU only code or does it only do CUDA code?	nvprof can profile CPU codes, GPU codes and OpenACC codes
105	What version of GCC which supports ACC? 5.4 or 6.2?	5.9 or later or 6.x
106	Is there an (easy) implementation for MatLab code?	OpenACC and Matlab are mostly separate. It may be possible to call an OpenACC function from Matlab mex interface
107	What is about a Python implementation of OpenACC?	it's possible to interoperate python and OpenACC codes, but beyond the scope of this course
108	Can OpenACC be integrated with MATLAB code?	Not directly. It may be possible to call an OpenACC function using the mex interface
109	In order to take the official certificate i just have to watch the event live?	In order to receive a certificate, you need to attend live lectures and complete the certification test that will be available after November 9th.

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



110	Linear algebra is essential from what I gather to understand CUDA, right?	linear algebra is not necessary to understanding CUDA or GPU acceleration
111	What are the main differences between CUDA and OpenACC?	OpenACC is a directives-based programming model that is designed to simplify on-ramp to parallel programming. Minimal code modification is required. CUDA is a low-level model that is more involved, but allows for more customizations.
112	Will OpenACC help me understand the main principles behind big data analytics? Can CUDA be used as a tool for data analytics on a large scale?	OpenACC will help you understand parallel programming, same applies to CUDA. And both models can be used for data analytics
113	Does OpenACC support matlab directly?	OpenACC and matlab are mostly separate. It should be possible to call OpenACC code using the mex interface
114	Don't GPUs have a worse performance/watt performance regarding their high watt consumption?	GPUs generally have better perf/watt than CPUs on highly parallel codes
115	Concerning the data that is processed can we apply these techniques on big data?	Yes, you can apply to big data
116	Can you elaborate on the concept of gangs and their mapping on the CPU or GPU cores?	Gangs are groups of vectors and workers. A gang is a level of subdivision of a problem. If you have a for loop of 1 million iterations, it might be broken into 1000 gangs of 1000 workers. Each gang is then responsible for 1000 iterations of the for-loop
117	What happens if we allocate all the graphical processors for computing task, will it run the program or some error may cause.	GPUs can do both graphical and compute tasks at the same time
118	Does pgc++ supports c++14 ? And it is possible to generate kernels from heavily templated functions?	C++14 support is partially available in the latest PGI tools, but does not apply to GPU code yet. Yes, possible to generate kernels from templated code.
119	Are you able to use OpenACC on java?	not directly but there are probably various ways to interoperate
120	Does OpenACC offer parallelism on all levels of multicore CPU, distributed CPU clusters as well as GPUs?	You would use multiple technologies, e.g. OpenACC + OpenMP + MPI
121	Must the PGI compiler be run on the same hardware the compiled code will be run so it can know the hardware configuration and optimize for it? Is it possible to compile on one machine and run the generated code in another machine optimally?	no, the target hardware config is specified by compile switches, not by the platform you are compiling on
122	Is there other languages that are planned to be covered with ACC?	at the moment, with PGI, the supported languages are C/C++ and Fortran
123	Is it possible to use easily any (Nvidia) Graphics-Card?	Nearly all currently supported NVIDIA CUDA GPUs should be usable with OpenACC
124	With CUDA used to compile OpenACC code for Nvidia GPU's, what is used to compile to AMD GPU's?	You can use the PGI compiler to target AMD GPUs for OpenACC codes. The details are beyond the scope of this course.

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



125	Does OpenACC implement any form of memory locking for shared memory access?	OpenACC supports atomics
126	Why do you run code only on one half of the K80?	The K80 is treated as 2 GPUs, you would use multi-accelerator approaches in OpenACC to use both. Covered later in this course.
127	What speed-up can I expect if I use CUDA code instead of OpenACC-generated code for the example used in the presentation (CG)?	Probably not much difference between OpenACC and CUDA for this example, unless you use CUDA libraries.
128	But you can write OpenACC C++ code and call it from python, right?	yes, in the same way you can call a compiled C code from python, for example
129	How many parallel gangs will be instantiated?	this depends on the size of the loops being parallelized - the size of your problem
130	Is global synchronization between the host and device in OpenACC possible?	yes, sync between host and device is possible, there are directives for this
131	Will GCC be available at Qwiklabs hands-on?	no, they will use the PGI toolchain
132	Can I pair OpenMPI with OpenACC? I mean can I send the code to 4 processors and each processor then call the GPU?	Yes, OpenACC can work with OpenMPI.
133	Can I intergrate OpenACC with OpenMPI?	OpenACC and OpenMP can work together in the same code
134	Can I valgrind an OpenACC application for memory leaks?	Yes, it should be possible. It won't cover the GPU code, just the CPU code. A valgrind-like tool is available for GPU code called CUDA-memcheck, it works with OpenACC running on NVIDIA GPUs.
135	Can I transfer multidimensional vectors (say int A[5][5][5]) to the GPU or I need to map A -> B with int B[5*5*5];	You can transfer multidimensional arrays using OpenACC
136	What is the maximum amount of memory I can transfer to the GPU memory? (for an average GPU)	There is no stated limit - it is limited by GPU memory size
137	What free compilers for OpenACC exist?	GCC supports OpenACC
138	As a teacher it's possible to have access to the tools for free?	yes, it's possible
139	Is there an OpenACC version of blas lib?	For host code, use your preferred host blas, for device code, you can interoperate OpenACC with CUBLAS
140	Is it possible to use OpenACC to process a stream data without copying it to RAM? I mean, Is it possible to copy stream data from network to directly GPU memory?	Yes, it's possible. Beyond the scope of this course
141	Is there any profiler to show the cost of each line of code from the perspective of time or processing power?	Yes, depending on the GPU, nvprof and pgprof can do this for both CPU and GPU code.
142	Is it possible to access a specific "thread" index in OpenACC as in CUDA?	typically you wouldnt do this, however it would typically be equivalent to the for-loop index, which you can access of course

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



143	Can OpenACC be integrated with CUDA libraries?	Yes, interoperability is possible
144	Will we be able to download the slides afterwards ?	slides will be available later this week
145	Cisco WebEx question : will we be able to export all questions/answers, I can read all but there are interested	yes, we will be publishing answered questions as well.
146	Is OpenACC compatible for Android, iOS, Mac? it's platform independent right?	The standard is platform independent. The implementations today (e.g. PGI, CRAY) mostly target x86 or OpenPOWER CPUs.
147	Can I use OpenACC directives inside of CUDA kernels?	Not directly, but there are various ways to interoperate CUDA and OpenACC - in both directions
148	Which version of GCC support OpenACC?	In various branches of GCC 5 (very late) and 6
149	How to control overlapping in OpenACC?	overlapping of operations is supported in OpenACC, this will be covered in a later session
150	Can I use Intel® Compiler (ifort)?	no
151	What are some differences between consumer GPUs and the K series?	http://www.nvidia.com/object/why-choose-tesla.html
152	Can you please name the compilers that run OpenACC code in FORTRAN? Does ifort compiler run OpenACC code?	PGI and Cray compilers can be used. There are others as well. ifort does not support OpenACC
153	Can you suggest any book that gives OpenACC directives that the instructor talked about?	We are publishing a new book on OpenACC that will be available starting November 1st - Parallel Programming with OpenACC: http://store.elsevier.com/Parallel-Programming-with-OpenACC/Rob-Farber/isbn-9780124103979/
154	Can I use R to call OpenACC as matlab does??	It should be possible
155	Does OpenACC support Unified memory from CUDA?	Yes, it does, and this will be covered in this course (maybe not today's session)
156	Is it possible to generate random numbers on GPU with OpenACC?	Yes, there are a variety of approaches. For example, OpenACC can interoperate with the CURAND library.
157	pgprog, PGI compiler, what is this "pg" prefix about?	pg = portland group
158	Do I need to have NVIDIA GPUs installed on my laptop for this session? I have not installed CUDA because I run into a mess on Ubuntu 14.04 whenever I try to do so	No you don't need a NVIDIA GPU.
159	Is there a repo containing the source code samples so we can replicate it?	yes, there will be, stay tuned
160	Does OpenACC support nested loops?	Yes, generally two levels of nesting can be directly parallelized, and additional levels can be handled with for example the collapse directive
161	What does it mean to handle unified memory in hardware? Through device drivers?	Starting from Pascal generation of GPUs Unified Memory is implemented not only in software - CUDA drivers, but also in hardware itself

INTRODUCTION TO OPENACC COURSE

Lecture #1 Q&A



162	You mentioned OpenACC is a Parallel programming model. Briefly could you explain, what is the difference between GPU programming model vs Parallel programming model?	GPU programming model is a type of parallel programming model
163	Does OpenACC have directives for Python or Perl languages?	not today
164	I think, you don't need NVIDIA GPU with OpenACC. This meant to be portable API, but will find it out	You are correct. OpenACC works on multiple platforms that includes CPUs and GPUs
165	How much is the performance different with OpenACC as compared to program written with CUDA?	This will vary. Some programs will be approximately the same performance.
166	Why did you make OpenACC when OpenCL does the same thing?	They don't do the same things - stay tuned to the presentation. OpenACC is more like OpenMP in that it can be used to decorate ordinary code to accelerate.
167	I am not expert with OpenCL, but I saw that OpenCL offers the same portability option with other platforms as well like DSP for imaging. Then, how OpenACC is different from OpenCL?	OpenCL is a low-level programming language, while OpenACC is a directives-based approach. With OpenACC just a few directives will accelerate the code and one code will work on multiple platforms.
168	What application can be targeted with OpenACC? Will it work with computer visions/ML/CNN?	Yes, it's possible
169	Exist OpenACC in Fortran???	Yes - there are Fortran directives also , e.g \$!acc kernels ...