PGI 2018

Latest CPUs
NVIDIA Volta
OpenACC 2.6

- Latest CPUs Support
  - Intel Skylake
  - AMD Zen/Epyc
  - IBM POWER9
- AVX-512 code generation
- Full OpenACC 2.6
- OpenMP 4.5 for multicore CPUs
- Integrated CUDA 9.2 toolkit/libraries
- New PGI fastmath intrinsics library
- Partial C++17 support
- Optional LLVM-based code generator

www.pgicompilers.com/whats-new

SPEC CPU 2017 FP SPEED BENCHMARKS

OpenMP 3.1 Performance on Multicore CPUs – smaller is better

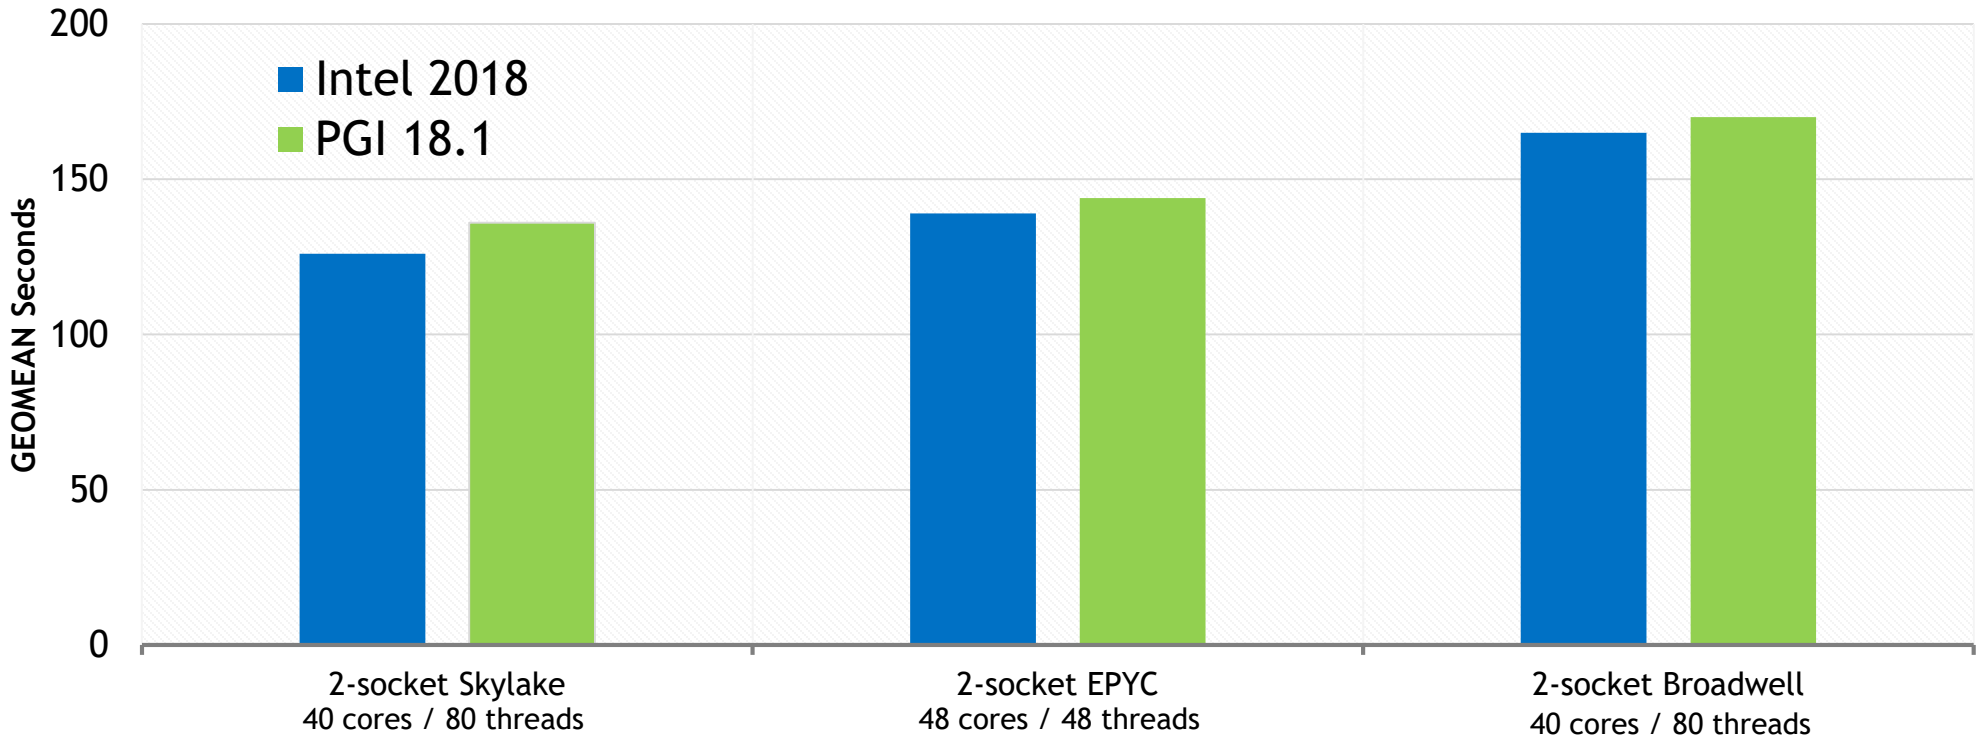

Performance measured February, 2018. Skylake: Two 20 core Intel Xeon Gold 6148 CPUs @ 2.4GHz w/ 376GB memory, hyperthreading enabled. EPYC: Two 24 core AMD EPYC 7451 CPUs @ 2.3GHz w/ 256GB memory. Broadwell: Two 20 core Intel Xeon E5-2698 v4 CPUs @ 3.6GHz w/ 256GB memory, hyperthreading enabled. Volta: NVIDIA DGX1 system with two 20 core Intel Xeon E5-2698 v4 CPUs @ 2.20GHz, 256GB memory, one NVIDIA Tesla V100-SXM2-16GB GPU @ 1.53GHz. SPEC® is a registered trademark of the Standard Performance Evaluation Corporation (www.spec.org).

PGI

SPEC ACCEL 1.2 BENCHMARKS

smaller is better

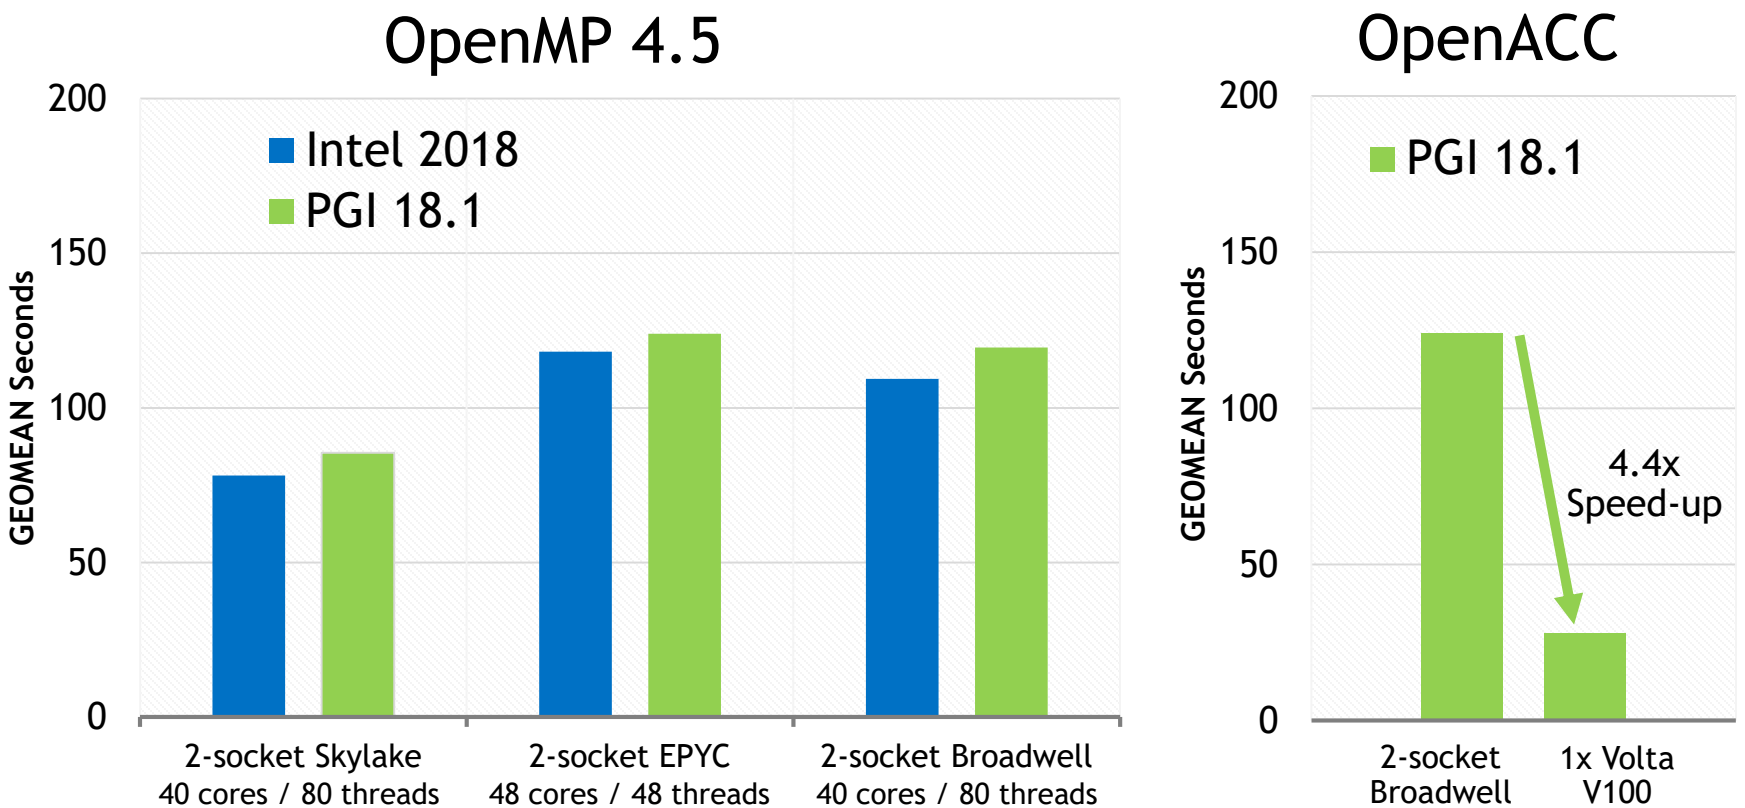

Performance measured February, 2018. Skylake: Two 20 core Intel Xeon Gold 6148 CPUs @ 2.4GHz w/ 376GB memory, hyperthreading enabled. EPYC: Two 24 core AMD EPYC 7451 CPUs @ 2.3GHz w/ 256GB memory. Broadwell: Two 20 core Intel Xeon E5-2698 v4 CPUs @ 3.6GHz w/ 256GB memory, hyperthreading enabled. Volta: NVIDIA DGX1 system with two 20 core Intel Xeon E5-2698 v4 CPUs @ 2.20GHz, 256GB memory, one NVIDIA Tesla V100-SXM2-16GB GPU @ 1.53GHz. SPEC® is a registered trademark of the Standard Performance Evaluation Corporation (www.spec.org).

PGI

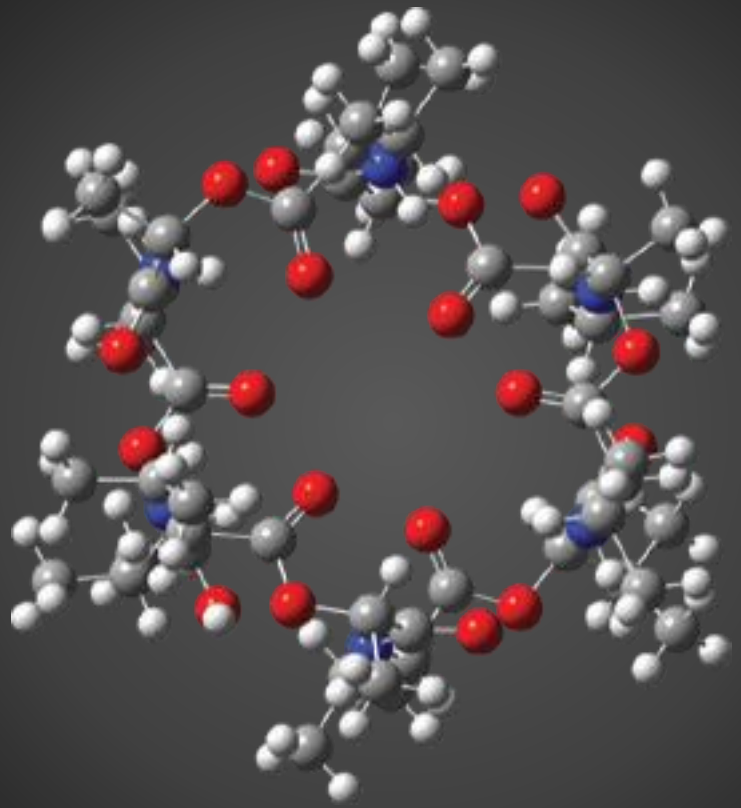

GAUSSIAN 16

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."

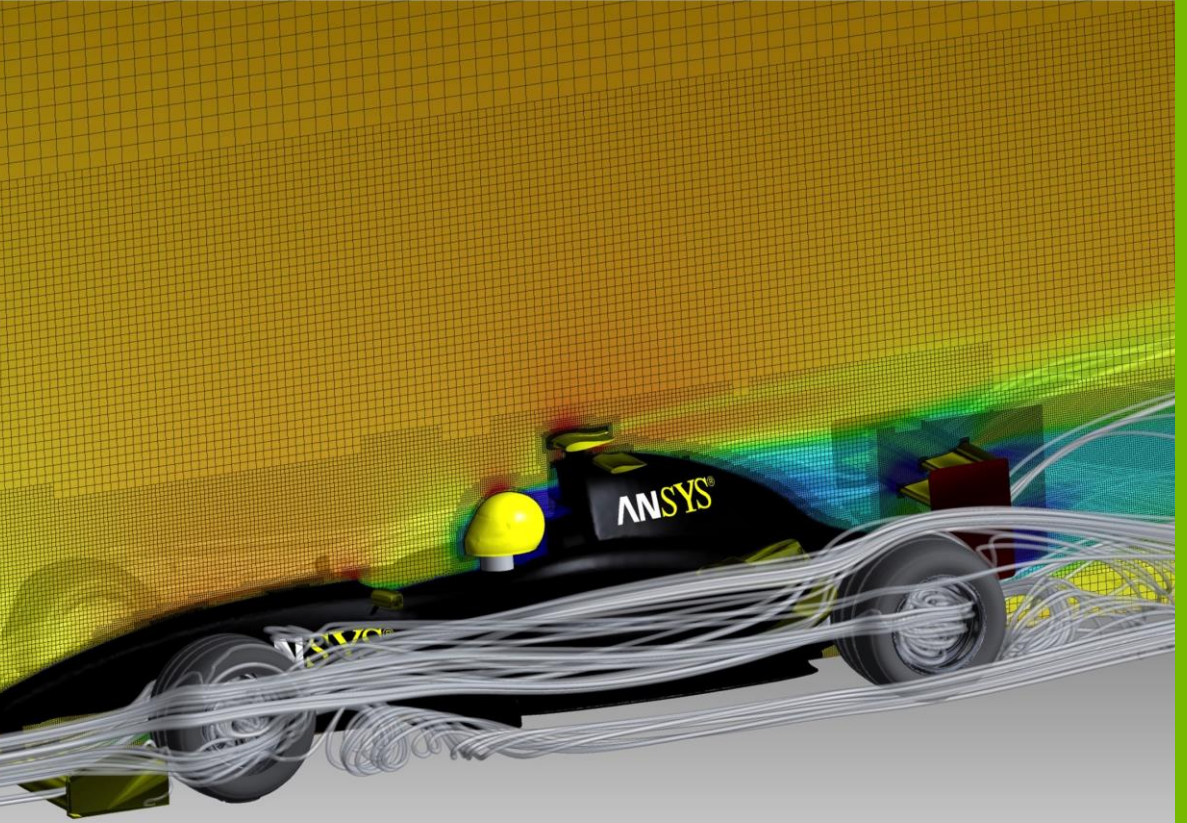

ANSYS FLUENT

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.”

PGI

Image courtesy: NCAR

MPAS-A

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.

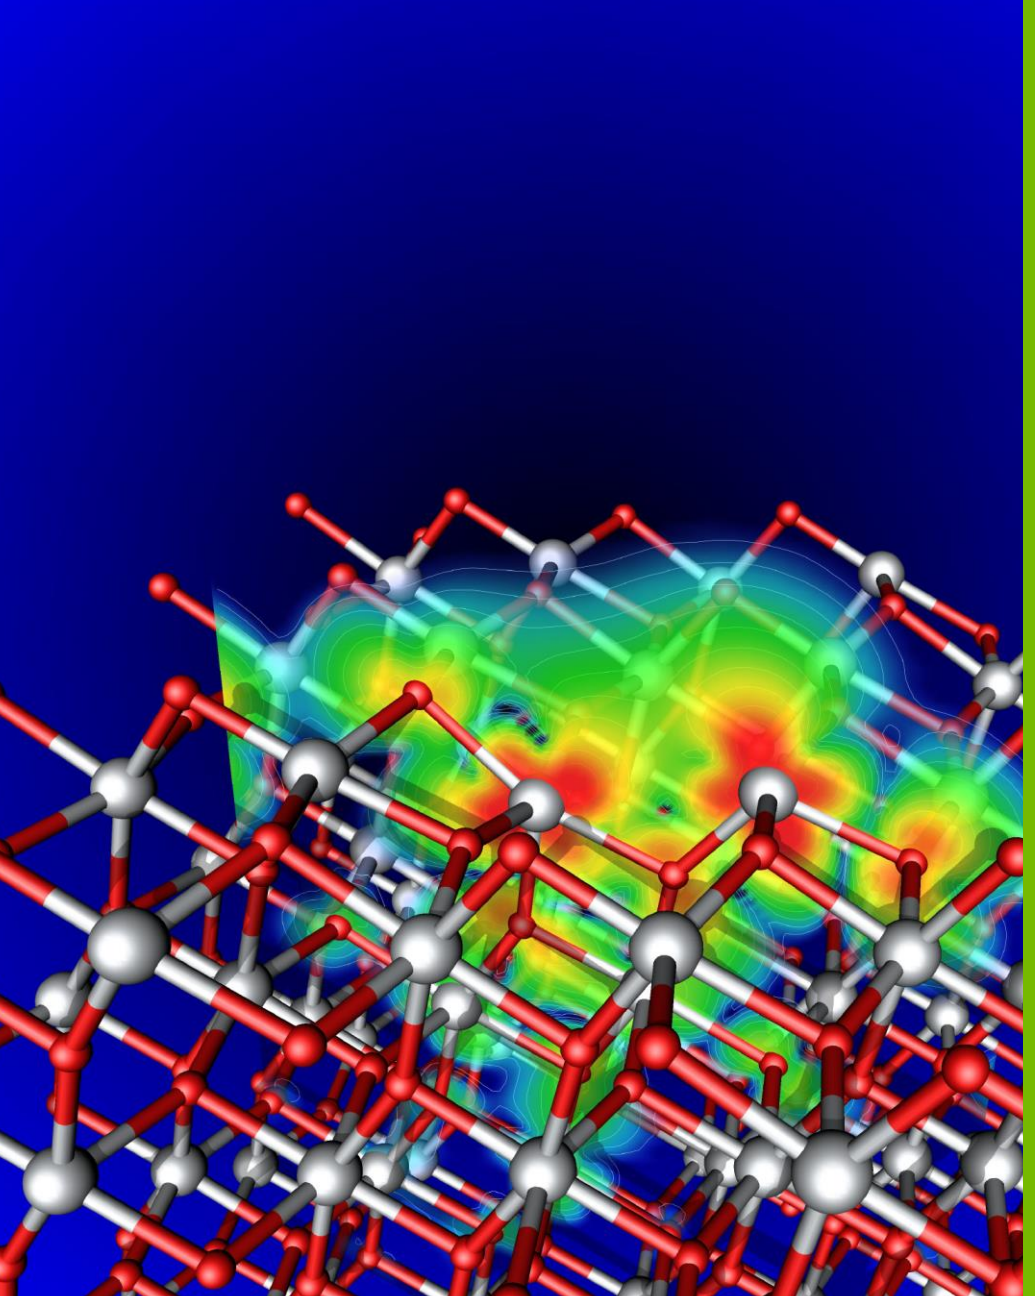

VASP

Prof. Georg Kresse
Computational Materials Physics
University of Vienna

“For VASP, OpenACC is ***the*** way forward for GPU acceleration. Performance is similar and in some cases better than CUDA C, and OpenACC dramatically decreases GPU development and maintenance efforts. We’re excited to collaborate with NVIDIA and PGI as an early adopter of CUDA Unified Memory.”

PGI Presentations @ CUG 2018

OpenACC and CUDA Unified Memory
Sebastien Deldon, PGI Compiler Engineer
4:00pm Wednesday, Technical Session 20B

Strategies to Accelerate VASP with GPUs using OpenACC
Stefan Maintz, NVIDIA Devtech Engineer
4:30PM Wednesday, Technical Session 20B