Adaptive COMPUTING

Intelligent HPC Workload Management

Optimizing Accelerators with Moab® HPC Suite

Harnessing the Power of Intel® Xeon Phi™ and GPGPUs

The mathematical acceleration delivered by many-core General Purpose Graphics Processing Units (GPGPUs) offers significant performance advantages for many classes of numerically intensive applications. Parallel computing tools such as NVIDIA's CUDA and the openCL framework have made harnessing the power of these technologies much more accessible to developers, resulting in the increasing deployment of hybrid GPGPU-based systems and introducing significant challenges for administrators and workload management systems.

The new Intel Xeon Phi coprocessors offer breakthrough performance for highly parallel applications with the benefit of leveraging existing code and flexibility in programming models for faster application development. Moving an application to Intel Xeon Phi coprocessors has the promising benefit of requiring substantially less effort and lower power usage. Such a small investment can reap big performance gains for both data-intensive and numerical or scientific computing applications that can make the most of the Intel Many Integrated Cores (MIC) technology. So it is no surprise that organizations are quickly embracing these new coprocessors in their HPC systems.

The main goal is to create breakthrough discoveries, products and research that improve our world. Harnessing and optimizing the power of these new accelerator technologies is key to doing this as quickly as possible. Moab HPC Suite helps organizations easily integrate accelerator and coprocessor technologies into their HPC systems, optimizing their utilization for the right workloads and problems they are trying to solve.

Moab HPC Suite automatically schedules hybrid systems incorporating Intel Xeon Phi coprocessors and GPGPU accelerators, optimizing their utilization as just another resource type with policies. Organizations can choose the accelerators that best meet their different workload needs.

Managing Hybrid Accelerator Systems

Hybrid accelerator systems add a new dimension of management complexity when allocating workloads to available resources. In addition to the traditional needs of aligning workload placement with software stack dependencies, CPU type, memory, and interconnect requirements, intelligent workload management systems now need to consider:

- Workload's ability to exploit Xeon Phi or GPGPU technology
- Additional software dependenciesand reduce costs, including topology-based allocation
- Current health and usage status of available Xeon Phis or GPGPUs
- Resource configuration for type and number of Xeon Phis or GPGPUs

Moab HPC Suite auto detects which types of accelerators are where in the system to reduce the management effort and costs as these processors are introduced and maintained together in an HPC system. This gives customers maximum choice and performance in selecting the accelerators that work best for each of their workloads, whether Xeon Phi, GPGPU or a hybrid mix of the two.

Optimizing Intel® Xeon Phi™ and GPGPUs Utilization

System and application diversity is one of the first issues workload management software must address in optimizing accelerator utilization. The ability of current codes to use GPGPUs and Xeon Phis effectively, ongoing cluster expansion, and costs usually mean only a portion of a system will be equipped with one or a mix of both types of accelerators. Moab HPC Suite has a two-fold role in reducing the complexity of their administration and optimizing their utilization. First, it must be able to automatically detect new GPGPUs or Intel Xeon Phi coprocessors in the environment and their availability without the cumbersome burden of manual administrator configuration. Second, and most importantly, it must accurately match Xeon Phi-bound or GPGPU-enabled workloads with the appropriately equipped resources in addition to managing contention for those limited resources

Optimizing Accelerators with Moab® HPC Suite

according to organizational policy. Moab's powerful allocation and prioritization policies can ensure the right jobs from users and groups get to the optimal accelerator resources at the right time. This keeps the accelerators at peak utilization for the right priority jobs.

Moab's policies are a powerful ally to administrators in auto determining the optimal Xeon Phi coprocessor or GPGPU to use and ones to avoid when scheduling jobs. These allocation policies can be based on any of the GPGPU or Xeon Phi metrics such as memory (to ensure job needs are met), temperature (to avoid hardware errors or failures), utilization, or other metrics.

Use the following metrics in policies to optimize allocation of Intel Xeon Phi coprocessors:

- Number of Cores
- Number of Hardware Threads
- Physical Memory
- Free Physical Memory
- Swap Memory
- Free Swap Memory
- Max Frequency (in MHz)

Use the following metrics in policies to optimize allocation of GPGPUs and for management diagnostics:

- Error counts; single/double-bit, commands to reset counts
- Temperature
- Fan speed
- Memory; total, used, utilization
- Utilization percent
- Metrics time stamp

Improve GPGPU Job Speed and Success

The GPGPU drivers supplied by vendors today allow multiple jobs to share a GPGPU without corrupting results, but sharing GPGPUs and other key GPGPU factors can significantly impact the speed and success of GPGPU jobs and the final level of service delivered to the system users and the organization.

Moab HPC Suite optimizes accelerator utilization with policies that ensure the right ones get used for the right user and group jobs at the right time.

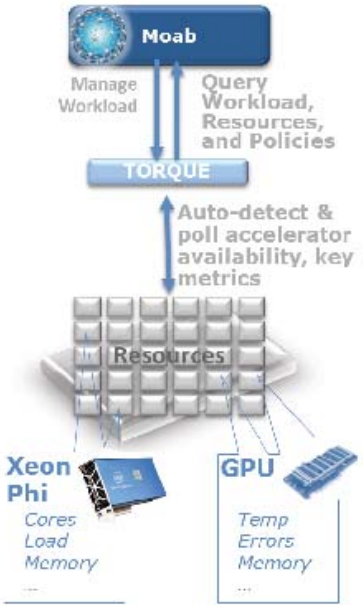

Consider the example of a user estimating the time for a GPGPU job based on exclusive access to a GPGPU, but the workload manager allowing the GPGPU to be shared when the job is scheduled. The job will likely exceed the estimated time and be terminated by the workload manager unsuccessfully, leaving a very unsatisfied user and the organization, group or project they represent. Moab HPC Suite provides the ability to schedule GPGPU jobs to run exclusively and finish in the shortest time possible for certain types, or classes of jobs or users to ensure job speed and success.

Optimize Performance and Innovation

Adaptive Computing's HPC expertise and award winning Moab product help you optimize the utilization and performance from your accelerators so you get from calculations to innovation faster:

- **Simplify management:** auto detection of GPGPUs or Intel Xeon Phi coprocessors
- **Maximize utilization:** policies favor GPGPU or Xeon Phi-oriented jobs over other jobs for accelerator-equipped resources
- **Maximize job success:** policies allocate the optimal accelerator, in the optimal mode, at given time
- **Maximize service level:** policies designate priority scheduling for users, groups, classes using limited accelerator resources

Contact a solutions advisor by phone or email, or visit our Web site today
North America, Latin America +1 (801) 717.3700
Europe, Middle East, Africa +44 (0) 1483 243578
Asia, Pacific, Japan, India +65 6597-7053
Email: solutions@adaptivecomputing.com
www.adaptivecomputing.com

Corporate Headquarters
1712 S. East Bay Blvd. Suite
300 Provo, Utah 84606

Adaptive COMPUTING

441-2012-11-05