

Debugging CUDA-Accelerated Parallel Applications with TotalView

An Integrated Approach to Heterogeneous Multi-tier Parallelism

A White Paper by Rogue Wave Software.

November 2011



Rogue Wave Software
5500 Flatiron Parkway,
Suite 200
Boulder, CO 80301, USA
www.roguewave.com

Debugging CUDA-Accelerated Parallel Applications with TotalView

An Integrated Approach to Heterogeneous Multi-tier Parallelism

by **Rogue Wave Software**

© 2011 by Rogue Wave Software. All Rights Reserved

Printed in the United States of America

Publishing History:

November 2011

Trademark Information

The Rogue Wave Software name and logo, SourcePro, Stingray, HostAccess, TotalView, IMSL and PV-WAVE are registered trademarks of Rogue Wave Software, Inc. or its subsidiaries in the US and other countries. HydraExpress, ThreadSpotter, ReplayEngine, MemoryScape, JMSL, JWAVE, TS-WAVE, and PyIMSL are trademarks of Rogue Wave Software, Inc. or its subsidiaries. All other company, product or brand names are the property of their respective owners.

IMPORTANT NOTICE: The information contained in this document is subject to change without notice. Rogue Wave Software, Inc. makes no warranty of any kind with regards to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Rogue Wave Software, Inc. shall not be liable for errors contained herein or for incidental, consequential, or other indirect damages in connection with the furnishing, performance, or use of this material.

TABLE OF CONTENTS

Abstract.....	4
Introduction	4
HPC Debugging.....	4
CUDA and Heterogeneous Acceleration Architectures	5
CUDA: Powerful but Challenging	6
The TotalView Debugger.....	7
Heterogenous Applications on the IBM Cell.....	8
The NVIDIA GPU Architecture and CUDA	9
The TotalView Model - Extended for CUDA.....	9
Challenges and Features.....	10
Extending the Thread Model	10
Dealing with Different Memory Spaces.....	14
Dealing with Inlined Functions	15
Thread Control and Single Stepping the GPU.....	15
CUDA Memory Exceptions.....	16
Conclusion.....	16

Abstract

CUDA introduces developers to a number of new concepts (such as kernels, streams, warps and explicitly multi-level memory) that are not encountered in serial or other parallel programming paradigms. Visibility into these elements is critical for troubleshooting and tuning applications that make use of CUDA. This paper will highlight CUDA concepts implemented in CUDA 3.0 - 4.0, the impact of those concepts for troubleshooting CUDA, and how TotalView helps users deal with these new CUDA-specific constructs. CUDA is frequently used alongside MPI parallelism and host-side multi-core and multi-thread parallelism. The TotalView parallel debugger provides developers with an integrated view of all three levels of parallelism within a single debugging session.

Introduction

NVIDIA's CUDA language extension provides a very interesting opportunity for scientists and developers to accelerate the runtime performance of computationally intensive applications running on a single workstation or on a cluster of dedicated nodes augmented with GPU cards. However, as developers begin adapting their applications for the GPU, they are finding out that harnessing the available performance requires mastering the concepts of data parallel programming and developing a fairly sophisticated understanding of the GPU architecture. Parallel development and debugging tools are being adapted to help developers tackle both of these objectives. This paper talks about work that has been done to enhance TotalView so that users can seamlessly debug apps that leverage CUDA in the context of multiple tiers of parallelism – device, host and cluster.

HPC Debugging

HPC applications, which typically run as distributed parallel applications on cluster type machines, require significant effort to write. Scientists and developers need to think about breaking tasks down into small units that can be distributed across that cluster. They need to pay special attention to the sequence of operations and the data and resource dependencies that different parts of the application have. They need to be conscious not just about what data is needed where, but also about the bandwidth and latency involved with moving data around the cluster.

Frameworks, libraries and languages like MPI, UPC, Global Arrays, Co-Array FORTRAN, OpenMP, and pthreads provide powerful and flexible mechanisms that can be used to construct parallel applications and address the challenges of instantiating and managing parallel programs, providing synchronization and data



movement between parallel tasks as required. Regardless of whether it is implicit or explicit, there is a tremendous degree of complexity that arises from parceling up the program and data while ensuring data movement and appropriate (but not too much) synchronization. This complexity becomes an issue when it comes to validating, troubleshooting and debugging applications.

Debugging has always been challenging; there is an old adage that debugging a bit of code is at least twice as hard as writing that same bit of code — and that statement was primarily made in the context of serial applications running on a single core. Taking a given algorithm and making it parallel frequently means introducing 1) additional code and 2) more complicated data structures to facilitate distributed execution and data movement. Both cases provide additional opportunities for errors to creep in.

Parallel applications execute many different bits of code at the same time across a large set of distinct nodes. They frequently perform computations that are interconnected enough that an error or oversight anywhere may either cause a local failure or set in motion a series of events that causes incorrect data to be transmitted to another node, leading to a wrong answer or crash somewhere else in this distributed system. A failure to ensure proper synchronization can lead to failures that may happen only rarely — when natural perturbations introduced by the system or other factors external to the program lead to rare and unexpected sequencing of operations. Furthermore, synchronization-dependent errors may become visible only when the program is run at or above some specific scale or is sensitive to details about how parallel tasks are mapped to compute resources.

All of this adds to the complexity of validating, debugging and tuning parallel applications. As if distributed architectures weren't enough, recent trends have introduced heterogeneous computational elements into the mix.

CUDA and Heterogeneous Acceleration Architectures

Recently, as general-purpose processors have grown more complicated and physical limits having to do with thermal dissipation have served to limit raw clock speed increases, there has been a growing interest in exploring heterogeneous architectures. Such architectures attempt to blend the benefits of general-purpose cores that excel at executing a wide range of computation with a set of cores specifically designed to execute more limited types of computations. These special-purpose cores typically focus on performing integer and floating point calculations on vectors of data and may trade off some of the sophisticated predictive control flow capabilities that are featured on the general purpose cores. A few years ago IBM and Toshiba introduced the Cell processor, which blended a PowerPC-based general purpose computing element in a processor package with a set of special vector processors called Synergistic Processing Elements (SPEs). NVIDIA provides a solution that uses their extremely capable, mass-market Graphics Processing Units (GPUs) in the vector processor in conjunction with traditional mass-market Intel or AMD CPUs.



There are several different techniques for taking advantage of GPUs as computational accelerators. In the 1990s curious users who found clever ways to repurpose the OpenGL image processing language to accomplish general-purpose computations. Later NVIDIA introduced the proprietary CUDA for C and C++ language extensions. AMD and Apple sponsored work resulting in OpenCL, and independent companies like PGI and CAPS introduced techniques like the PGI-accelerated Fortran and CAPS HMPP. CUDA has a unique position in the market in terms of widespread adoption and a mature, though still rapidly evolving tool chain and runtime environment. While the jury is still very much out on which approach will be the favored one several years down the road, CUDA is how most ISVs and users are implementing GPU acceleration for NVIDIA hardware today.

CUDA: Powerful but Challenging

Programming for CUDA introduces a series of unique challenges that HPC developers need to overcome to successfully harness the power of the GPU and reduce the time that it takes their applications to run and generate results.

The first challenge is a lack of abstraction; CUDA generally exposes quite a bit about the way that the hardware works to the programmer. In this regard, it is often compared to assembly language for parallel programming. Explicit concepts such as thread arrays, discussed below, map directly to the way the hardware groups computational units and local memory. CUDA compilers and runtimes prior to 3.2 don't support thread stacks, so programmers had to avoid recursion and generally be aware that any functions they wrote in CUDA device code would be inlined for execution. The memory model requires explicit data movement between the host processor and the device and CUDA makes explicit use of a hierarchy of memory address spaces, each of which obeys different sharing rules. This is more low-level detail than the typical application or scientific programmer has to deal with when programming in C, C++ or Fortran and also raises significant concerns with regard to portability and maintainability of code.

The second such challenge is that the programming model for CUDA is one of data parallelism rather than task parallelism. When divvying up work across the nodes of a cluster, HPC programmers are used to looking for and exploiting parallelism at a certain level. The amount of data to assign to each node depends on a number of factors including computational speed of the node, the available memory, the bandwidth and latency of the interconnect and the frequency with which results need to be shared with other processes. Since processors are fast and network bandwidth is relatively scarce, the balance is typically to put quite a bit of data on each compute node and to move data as infrequently as possible. CUDA invites the programmer to think about parallelism of a completely different order, encouraging the developer to break the problem apart into units that are much smaller, often as small as the computation that might be required to assign a single variable into an array. This is often many orders of magnitude more parallelism than was previously expressed in the code and requires reasoning somewhat differently about the



computations themselves.

The third challenge lies in dealing with data movement and multiple address spaces. NVIDIA GPUs are external accelerators attached to the host system via the PCI bus; each GPU has both its own onboard memory and also features smaller bits of memory attached to each one of the compute elements (see the architecture review below for more detail). While there are now mechanisms to address regions of host memory from device kernels, such access is very slow compared to accessing device memory. CUDA programs therefore use a model in which code running on the host processor prepares and explicitly dispatches work to the GPU, pauses for the GPU to complete that work, then reads the resulting data back from the device. Both the units of code representing computational kernels and the associated data on which those computational kernels will execute are dispatched to the device. The data is moved, in whole or part, to the device over the PCI bus for execution and as results are produced they need to be moved, just as explicitly, back from the device to the host computer's main memory.

The device has different kinds of instance memory: either local and reserved for specific sets of computing elements, or globally addressable from any of the computing elements on the CUDA device. CUDA makes this distinction explicit, and requires that developers designate the location of variables at variable declaration time. A variable's location in this memory hierarchy determines how it can be used, forcing the developer to be cognizant of these issues.

The TotalView Debugger

TotalView is a highly interactive graphical source code debugger that provides developers working in C, C++ or Fortran with a way to explore and control their programs. Originally designed to debug one of the first distributed memory architectures, the BBN Butterfly, TotalView has been continually enhanced with a focus on multi-process and multithreaded applications. Today TotalView is used to develop, troubleshoot and maintain applications in a wide range of situations. One group uses it to develop Linux-based commercial embedded computing applications consisting of a single process with a hundred or more separate threads that simultaneously interact with a graphical user interface (GUI), sensors, online databases and network services. Other users troubleshoot sophisticated mathematical models of complex physical systems in astrophysics, geophysics, climate modeling and other areas. These models typically consist of thousands of separate processes that run on large-scale, high performance computing (HPC) clusters on the top 500 list. In both these use cases, TotalView provides the users with a debugging session in which they can examine in detail any one of the many threads or processes that make up the application, controlling each thread or process singly or as a part of various predefined and user-defined groups, and displaying multiple types of data and state information from across many processes and threads. With runtime-performance-conscious developers now exploring CUDA as a way to accelerate computation, Rogue Wave has enhanced TotalView to support users who are developing and debugging CUDA C/C++ code.



Heterogenous Applications on the IBM Cell

The IBM Cell processor technology was a precursor in several ways to GPU technology. Rogue Wave Software worked with Los Alamos National Lab (LANL) to provide support for their leadership-class Roadrunner system. The Cell processor embraces heterogeneity, providing a general purpose POWER core and a set of eight synergistic processing elements that are somewhat simpler and more restricted than the POWER core but each capable of performing computations within a separate memory environment. As with the GPU, user programs run on the general purpose core, which calls on the special purpose processors to accelerate certain types of computation.

The architecture of the Cell presents two general challenges for developers writing new software for or porting existing software to the Cell. The first challenge is in breaking the key components of a program into small chunks that can execute on the eight Synergistic Processing Elements (SPEs). In numerical programs, this often means dividing the data into small independent units. In other cases, individual tasks or pipeline stages may be delegated to specific SPEs. This issue of problem decomposition is similar to that of adapting an application to a distributed memory cluster environment, although the granularity is different due to the limited memory space directly available to each SPE.

The second challenge has to do with moving data between different parts of the processor. Each of the eight SPE cores that is part of the Cell processor has its own independent registers and a small amount of local memory (256KB) used for storing instructions and data. This memory acts similarly to a cache in a general-purpose processor, in that it has a limited size and can be read and written very quickly. Unlike a cache, however, its contents are directly managed by the application. The code units running on the SPE elements are allowed to initiate direct memory access operations that can copy chunks of data from the main memory into the SPE local store or back to main memory from the local store. The same memory is used for the machine instructions and global memory, heap memory and stack. The heap and the stack change size over time and because there is no memory protection, the programmer must take care to avoid collisions between these different memory ranges. Because each processor has a completely separate local store, the structure of a program for the Cell is very different than the structure of a traditional multi-threaded application, where all the threads share the same memory.

In order to support the Cell processor TotalView extended the process and thread model, allowing it to accommodate the idea that a process might contain a heterogeneous mixture of threads, some of them running in the shared address space of the host processor and some of them running off in the separate and more constrained space of the SPE cores. That had implications on the group models that TotalView uses to decide where to plant breakpoints and what processes to hold and what processes to run when performing process control operations. The result was a very natural and



powerful user interface paradigm where the user could quite easily switch between host processor threads and the accelerator threads running on the SPEs.

The NVIDIA GPU Architecture and CUDA

The CUDA hardware model has some significant similarities and differences from the Cell processor. Where the Cell had 8 synergistic processing elements, the NVIDIA GPU has an array of streaming multiprocessors (SMs) each of which provides an execution environment with a bit of local memory and a number of streaming processors (SP). Older hardware may have as few as 8 SPs per SM and newer hardware has up to 48. Devices such as the NVIDIA Tesla C2070 have 14 SMs with 32 SPs each and can execute up to 448 operations in parallel.

Each one of those streaming processors may be working on different data and produce different results but they are not all executing independently. The NVIDIA hardware introduces the concept of a warp, which is a set of 32 threads that execute the same instruction together. Branching within the warp is handled by stalling those threads that are not participating in a branch to be resumed when the shared program counter is again executing instructions that are part of the thread occupying that lane's execution path.

Logically threads are grouped together into arrays of up to 512 threads (16 warps) that execute on an SM and share a bit of memory. The geometry of the array is configurable based on the code and an individual thread is identified by a unique combination of three coordinate indices. Since each thread array executes on a single SM, CUDA programs will typically try to start up many more than one thread array at a time. CUDA provides a way to address sets of thread arrays using a second set of three coordinate indices that are referred to as the grid coordinates. Threads are grouped together into arrays. Arrays are grouped together in a grid. The full combination of six dimensional grid and array coordinates serve to fully specify a unit of work within a discrete CUDA computation.

The TotalView Model - Extended for CUDA

This section of the paper details Rogue Wave Software's adaptations to the TotalView debugger in support of CUDA development and debugging.

TotalView has been extended to provide CUDA support in the same manner that it was extended for Cell. The new version of TotalView builds on and extends the model of processes made up of threads to incorporate CUDA. There are two major differences: first, CUDA threads are heterogeneous, and second, TotalView shows a representative CUDA thread from among all those created as part of running a routine on the GPU.

The GPU thread is shown as part of the process alongside other pthreads or OpenMP threads. Threads created by OpenMP or pthreads in any single process



share an executable image, a single memory address space and other process level resources like file and network handles. This special CUDA thread, on the other hand, executes a different image, on a set of processors with a different instruction set, in a completely distinct memory address space.

Because of the large number of CUDA threads and the fact that only a small number of this large set of logical threads are likely to actually be instantiated on the hardware at any given moment in time, the UI has been adapted to display the thread specific data from a user-selectable representative CUDA device thread.

CUDA is heavily used in conjunction with distributed multi-process programming paradigms such as MPI. As such, the CUDA support in TotalView is complimentary and additive with the MPI support in TotalView. Users debugging an MPI+CUDA job see the same view of all their MPI processes running across the cluster and when they focus on any individual process they can see both what is happening on the host processor and what is happening in the CUDA kernels running on the GPU.

Challenges and Features

This section will discuss some of the challenges that we encountered as tool developers as TotalView was adapted to allow it to debug CUDA accelerated applications and will highlight some of the features that were introduced as a result.

Extending the Thread Model

TotalView already has a very powerful model for multi-process and multithread applications. Each program is modeled as a set of processes, which may run on one or more hosts and which each are comprised of one or more threads. TotalView has features that allow the user to examine and control each thread separately. Most programs are comprised of processes that have anywhere from one to a handful of threads. A few ambitious programs may have a hundred or more threads. The CUDA runtime allows each thread of a UNIX process to dispatch work to be done on the attached GPU device. This work is encapsulated in units of work referred to abstractly as "kernels" and the parallel architecture of the GPU allows it to execute hundreds of instances of these kernels, called device threads, at a time. In order to hide latencies the programmer is encouraged to request the creation of large numbers (many thousands) of such device threads.

The device cannot actually process all of these threads at the same time; instead it "streams" them - scheduling them onto the hardware in batches, running each batch to completion while loading the preconditions for the next batch so that it can immediately run. This streaming technique hides much of the cost of data movement and context creation.



Ideally each CUDA kernel thread is performing self-contained unit of work with clearly defined inputs and with no side effects or dependencies on anything that is meant to happen within any of the other kernel threads that are being executed at the same time. That allows the device and the device runtime environment a large degree of freedom in scheduling the kernels. If resources allowed, the runtime could create all the device threads ahead of time and run them all simultaneously. Alternately it could create one thread at a time, running it to completion and then eliminating it (leaving just its output data). Generally the model allows the runtime to create and run device threads in any order or grouping that make sense based on hardware resources and the data being processed. This device kernel thread concept is very different from host processor threads. The creation of a thread in pthreads or OpenMP is a significant and heavyweight operation; best programming and thread runtime practices involve keeping those threads around for a relatively long period of time and using various techniques to have them operate on data structures that are shared across and between threads. The fact that multiple threads access the same data structures is the key feature and key challenge to multithreaded programming with pthreads or OpenMP.

While CUDA kernel threads are encouraged to be self-contained they are not forced to be — they can access (both reading from and writing to) global and shared memory. This is important in that it avoids having the CUDA program duplicate certain input data structures, but it means that there is the possibility for race conditions and device threads reading incorrect or invalid data.

In order to give visibility into the dynamically changing set of CUDA kernel threads without overwhelming the user, TotalView creates a single thread object for each kernel invocation, called a GPU focus thread, which is displayed next to the host threads in the debugger's display. The user selects this thread to get a view into one of the device threads that are currently executing the kernel on the GPU device. TotalView gives host threads a positive debugger thread ID and CUDA threads a negative thread ID. The initial host thread in process “1” might be labeled “1.1” and the CUDA thread is labeled “1.-1”. See Figure 1 for just such a pair of threads.

At any time the user can control which of the currently active CUDA threads is displayed through the GPU focus thread. This selection is done using a new set of controls, called the GPU thread selector, on the TotalView Process Window. The GPU thread selector reports the device thread currently being displayed, and can be used to change that selection. CUDA threads can be specified using one of two coordinate spaces. One is the logical coordinate space that is in CUDA terms of grid and block indices: <<<(Bx,By,Bz),(Tx,Ty,Tz)>>>. The other is the physical coordinate space that is in hardware terms of device number, streaming multiprocessor (SM) number on the device, warp (WP) number on the SM, and lane (LN) number on the warp. See the GPU thread selector in Figure 1.



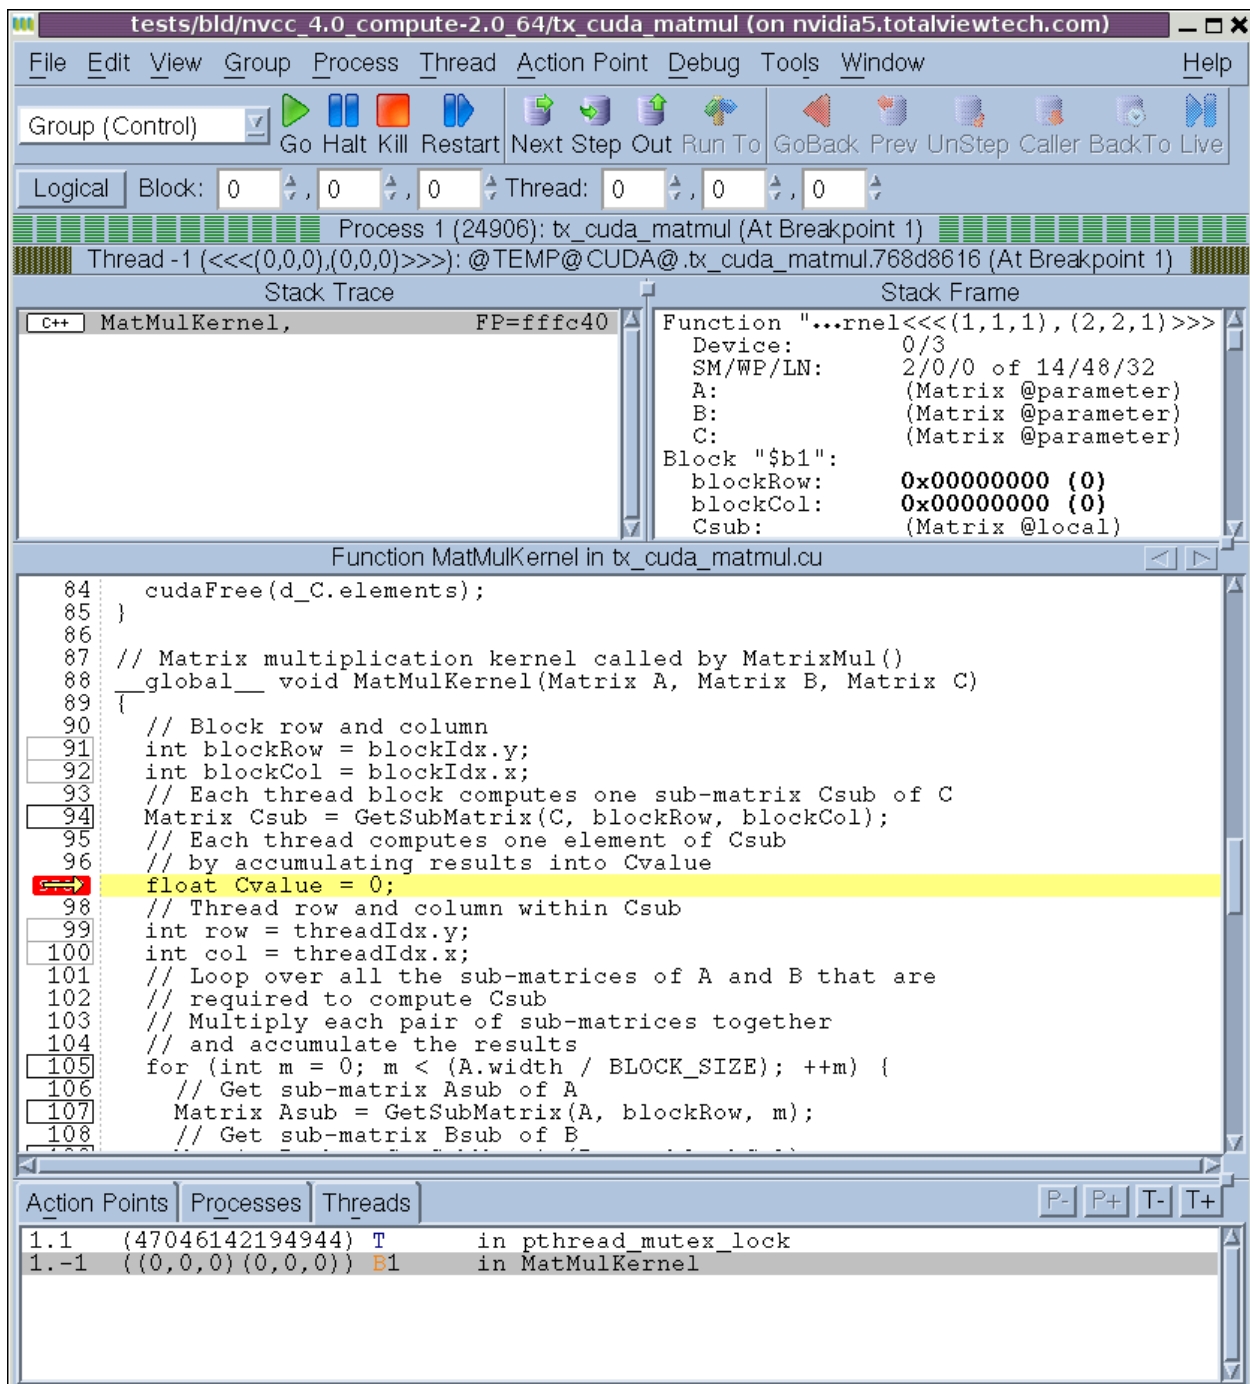


Figure 1. TotalView Process Window focused on a CUDA device thread. Note the CUDA thread selector between the toolbar with the stepping controls and the process status ribbon; the CUDA-specific information in the stack frame, including the qualifiers such as local on the type of variable Csub; and the CUDA device thread 1.-1 and the host thread 1.1 listed in the Threads Pane at the bottom of the window.

Any given thread has both a thread index in this 4D physical coordinate space, and a different thread index in the 6D logical coordinate space. These indices are shown in a series of spin boxes in the Process Window, next to a button that indicates which coordinate system is currently displayed. Pressing this button switches



between the two numbering systems, but does not change the actual thread.

It isn't always obvious how the logical and physical coordinates are mapped to one another, nor is it always clear what logical coordinate range designates the set of threads that are currently mapped to hardware. Therefore TotalView added a CUDA device display which brings up a separate window specifically designed to display the hardware's capabilities (in terms of how many streaming multiprocessors it has and their compute capability rating) and the logical threads currently mapped to the hardware resources.

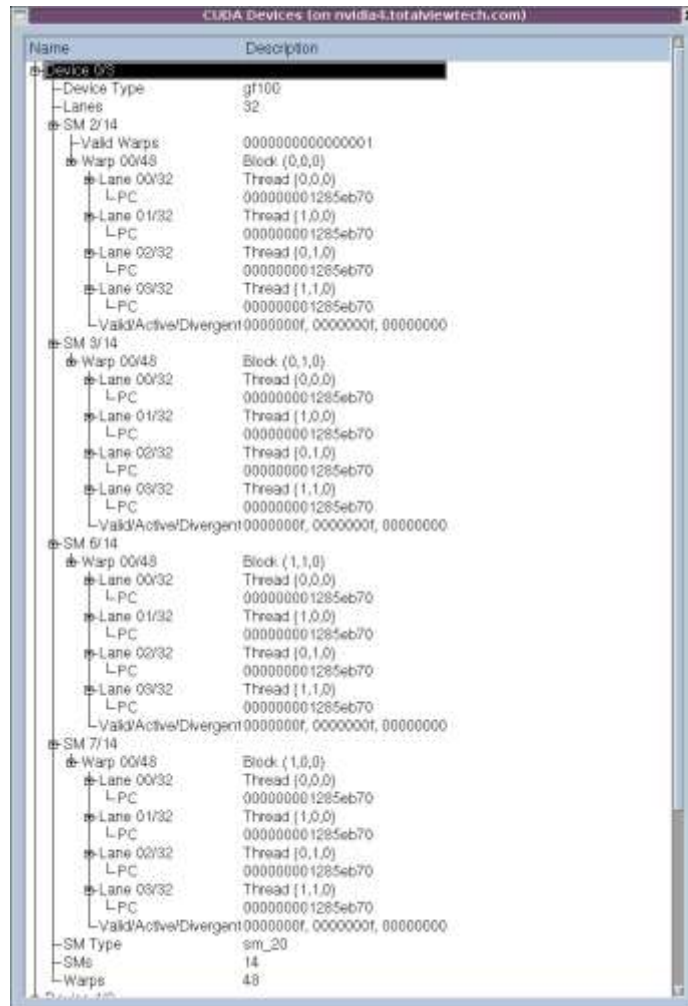


Figure 2. Device Status Window, showing the device hierarchy starting with a numbered set of CUDA devices that are comprised of SMs that include Warps made up of Lanes. This particular display highlights the 1st of three devices on a system, with 4 active SMs and their component Warps and Lanes. The other key information that is presented here is how these hardware resources are mapped to the logical resources of the CUDA program itself. CTA blocks are mapped to specific SMs and threads are mapped to warps and lanes. In the example shows the thread with device coordinates (Device, SM, Warp, Lane) = (0,2,0,3) is mapped to block (0,1,0) thread (1,1,0).

Dealing with Different Memory Spaces

TotalView can display variables and data from a CUDA thread. CUDA variables for a device kernel function are displayed in the stack frame on the Process Window in the same manner that variables are displayed in C, C++ or Fortran code on the host processor. The idioms of hovering over variables with the mouse or diving on variables to open a Variable Window both work just like they do on the host thread, as does the TotalView expression list and expression system. One thing to note is that when you use the CUDA thread selector to change the focus you will see the values of variables in the newly selected thread replace those from the previous thread in all the aforementioned places.

Naturally, all variables have addresses, but depending on the variable's type, those addresses may refer to any one of many different places. That's because CUDA uses an explicitly hierarchical memory. Each thread has its own local memory; so if variable X is a local variable and it has address 0x14 each thread can have its own instance of X, all with the same address but each referring to a private instance of that variable with a different value. Similarly there is a bit of memory associated with each cooperative thread array and the streaming multiprocessor that the CTA is running on. Variables that reference into that memory refer to data that is shared by all the threads in the CTA. Finally there are variables that refer into the global memory that is addressable by and shared by all the threads running on the device.

TotalView expresses these different kinds of memory (register or local, shared, constant or global memory) through the notion of a storage qualifier in the type system. When looking at variables in CUDA these qualifiers appear next to familiar types such as "int" or "float" so that a local variable X expressing an integer would have type of "int @local" in the TotalView Variable Window. Like types, these qualifiers are something that the user can modify. Casting a variable from "int @local" to "int @global" won't change the pointer but it will use that pointer to look in a different place (the same address in a different memory address space) and will likely refer to a different value. These type qualifiers can be seen both in Figure 1 above and in Figure 3 below.

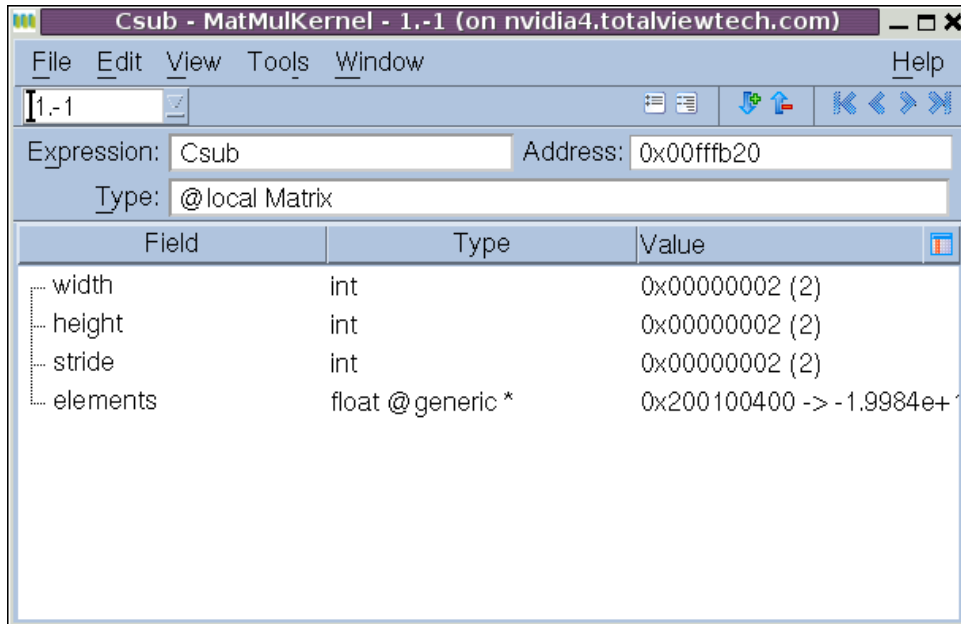


Figure 3. Variable Window displaying a variable called *Csub* that is at 0x00fffb20 in the thread local memory. It has a type of Matrix and it contains a pointer named *elements* that has a type of float* and points to an address in global memory.

Dealing with Inlined Functions

The stack trace pane shows the stack backtrace and may include synthetic stack references for inlined functions. Each stack frame in the stack backtrace represents either the PC location of GPU kernel code, or the expansion of an inlined function. Inlined functions can be nested. The “return PC” of an inlined function is the address of the first instruction following the inline expansion, which is normally within the function containing the inlined-function expansion.

Thread Control and Single Stepping the GPU

TotalView allows you to single-step GPU code just like normal host code, but it’s important to note that a single-step operation steps the entire warp associated with the GPU focus thread. That’s because at the hardware level, all the threads in the warp share a single program counter. It is literally not possible for the GPU to understand the concept of having some of the threads in the warp at location A and the other threads at location B. In some senses this seems like a limitation on the freedom TotalView gives developers to arbitrarily control the execution of host threads. However, the purpose of that control is to allow the developer to explore alternate possible execution scenarios. Since it isn’t possible for the threads in the warp to get out of step with one another there is no need to exercise scenarios other than those with synchronized threads across the warp. Different warps, however, might examine or modify global memory and might run with different timing, so it is both valid and sometimes necessary to explore different sequences of execution



between distinct warps.

CUDA Memory Exceptions

The NVIDIA hardware lacks some of the memory protection logic that developers are likely accustomed to when writing code for the host processor, so mistakes like de-referencing an uninitialized pointer in CUDA code won't result in a segmentation violation, and the program will proceed with undefined consequences. However, NVIDIA provides a mechanism for emulating the memory protection in software; that mechanism can be engaged in TotalView by activating the preference "Enable CUDA Memory Checking". With CUDA memory checking activated you can run the device kernel, and when an invalid pointer operation occurs the debugger will stop the thread with an error message as it does when a host thread encounters a segmentation error. Note that CUDA memory checking is not active when single-stepping the CUDA thread.

Conclusion

This paper has focused on two main topics. First, it briefly reviewed CUDA and some of the challenges that CUDA introduces for HPC developers looking to harness GPU processing power to accelerate their applications. These challenges center on the conceptual transitions related to moving to (or more often layering in) fine-grained data parallelism from (or alongside) medium-grained task parallelism, but also extend to dealing with many low-level details about the NVIDIA GPU device architecture that are exposed to the CUDA programmer as additional language concepts. These features and challenges, while certainly not insurmountable, definitely complicate the picture; the increased complexity potentially hinders troubleshooting, validation, and ongoing software maintenance.

The second section of the paper focuses on changes that have been made to adapt the TotalView parallel debugger for CUDA. The paper shows how the TotalView process and thread model has been extended to allow users to easily switch back and forth between what is happening on the host processors, where MPI-level communication occurs and the program sets the stage for the data parallel work to happen on the GPU, and what is happening over on the GPU itself where lightweight threads are being created and destroyed at a furious pace, each one in existence only long enough to compute one data element. It also details how the TotalView variable display and type system have been modified to provide developers a clear understanding of how their data interacts with the various distinct memory address spaces, each with different characteristics and visibility.

About Rogue Wave Software

Rogue Wave Software, Inc. is the largest independent provider of cross-platform



software development tools and embedded components for the next generation of HPC applications. Rogue Wave marries High Performance Computing with High Productivity Computing to enable developers to harness the power of parallel applications and multicore computing. [Rogue Wave products](#) reduce the complexity of prototyping, developing, debugging, and optimizing multi-processor and data-intensive applications. Rogue Wave customers are industry leaders in the Global 2000, ISVs, OEMs, government laboratories and research institutions that leverage computationally-complex and data-intensive applications to enable innovation and outperform competitors. For more information, visit <http://www.roguewave.com>.

