

HARNESSING GPU COMPUTING IN SYSTEM-LEVEL SOFTWARE

by

Weibin Sun

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Computer Science

School of Computing
The University of Utah

August 2014

Copyright © Weibin Sun 2014

All Rights Reserved

THE UNIVERSITY OF UTAH GRADUATE SCHOOL

STATEMENT OF DISSERTATION APPROVAL

The dissertation of Weibin Sun
has been approved by the following supervisory committee members:

| | | |
|---------------------|---------|-------------------------------------|
| <u>Robert Ricci</u> | , Chair | <u>May 1, 2014</u> Date Approved |
|---------------------|---------|-------------------------------------|

| | | |
|------------------------------|----------|-------------------------------------|
| <u>Jacobus van der Merwe</u> | , Member | <u>May 1, 2014</u> Date Approved |
|------------------------------|----------|-------------------------------------|

| | | |
|--------------------|----------|-------------------------------------|
| <u>John Regehr</u> | , Member | <u>May 1, 2014</u> Date Approved |
|--------------------|----------|-------------------------------------|

| | | |
|------------------|----------|-------------------------------------|
| <u>Mary Hall</u> | , Member | <u>May 1, 2014</u> Date Approved |
|------------------|----------|-------------------------------------|

| | | |
|---------------------|----------|-------------------------------------|
| <u>Eddie Kohler</u> | , Member | <u>May 1, 2014</u> Date Approved |
|---------------------|----------|-------------------------------------|

and by Ross Whitaker, Chair of the
School of Computing

and by David B. Kieda, Dean of The Graduate School.

ABSTRACT

As the base of the software stack, system-level software is expected to provide efficient and scalable storage, communication, security and resource management functionalities. However, there are many computationally expensive functionalities at the system level, such as encryption, packet inspection, and error correction. All of these require substantial computing power.

What's more, today's application workloads have entered gigabyte and terabyte scales, which demand even more computing power. To solve the rapidly increased computing power demand at the system level, this thesis proposes using parallel graphics processing units (GPUs) in system software. GPUs excel at parallel computing, and also have a much faster development trend in parallel performance than central processing units (CPUs). However, system-level software has been originally designed to be latency-oriented. GPUs are designed for long-running computation and large-scale data processing, which are throughput-oriented. Such mismatch makes it difficult to fit the system-level software with the GPUs.

This thesis presents generic principles of system-level GPU computing developed during the process of creating our two general frameworks for integrating GPU computing in storage and network packet processing. The principles are generic design techniques and abstractions to deal with common system-level GPU computing challenges. Those principles have been evaluated in concrete cases including storage and network packet processing applications that have been augmented with GPU computing. The significant performance improvement found in the evaluation shows the effectiveness and efficiency of the proposed techniques and abstractions. This thesis also presents a literature survey of the relatively young system-level GPU computing area, to introduce the state of the art in both applications and techniques, and also their future potentials.

CONTENTS

| | |
|---|------------|
| ABSTRACT | iii |
| LIST OF FIGURES | vii |
| LIST OF TABLES | ix |
| ACKNOWLEDGMENTS | x |
| CHAPTERS | |
| 1. INTRODUCTION | 1 |
| 1.1 Thesis Statement | 3 |
| 1.2 Contributions | 3 |
| 1.3 Findings | 3 |
| 1.4 Thesis Outline | 5 |
| 2. GENERAL PURPOSE GPU COMPUTING | 6 |
| 2.1 Overview | 6 |
| 2.2 Parallel GPGPU | 7 |
| 2.3 GPU Memory Architecture | 8 |
| 2.4 Advanced GPGPU Computing | 9 |
| 2.5 Efficient GPGPU Computing Techniques | 10 |
| 3. SYSTEM-LEVEL GPU COMPUTING PRINCIPLES | 12 |
| 3.1 Throughput-oriented Architecture | 13 |
| 3.1.1 Parallel Workload Abstraction | 13 |
| 3.1.2 Avoid Control Flow Divergence | 14 |
| 3.2 Coalescing Global Memory Accesses | 16 |
| 3.3 Overlapping GPU Computation and Memory Copy | 17 |
| 3.4 Asynchronous GPU Programming | 17 |
| 3.5 Reducing Memory Copy Overhead | 19 |
| 3.5.1 Reducing DMA Memory Copy Overhead | 19 |
| 3.5.2 Avoid Double-Buffering | 19 |
| 3.5.3 Trade Off Different Memory Copy Costs | 20 |
| 3.5.4 Memory Copy Techniques Discussion | 21 |
| 3.6 Miscellaneous Challenges | 21 |
| 3.6.1 OS Kernel Mode and User Mode Communication for GPU Computing | 22 |
| 3.6.2 Implicit Host-Device Synchronization | 22 |
| 3.7 Summary | 22 |

| | |
|---|-----------|
| 4. GPSTORE: HARNESSING GPU COMPUTING FOR OS STORAGE . . | 24 |
| 4.1 Design | 24 |
| 4.1.1 Overview | 25 |
| 4.1.2 Memory Management | 25 |
| 4.1.3 Request Management | 27 |
| 4.1.4 Streams Management | 29 |
| 4.2 Implementation | 29 |
| 4.2.1 Encrypted Storage | 30 |
| 4.2.2 md: Software RAID | 32 |
| 4.3 Evaluation | 33 |
| 4.3.1 Framework Performance | 34 |
| 4.3.2 dm-crypt | 35 |
| 4.3.3 eCryptfs | 37 |
| 4.3.4 md Data Recovery | 40 |
| 4.4 Summary and Future Work | 42 |
| 5. SNAP: FAST AND FLEXIBLE PACKET PROCESSING WITH CLICK AND GPUS | 44 |
| 5.1 Click Background | 45 |
| 5.2 Motivating Experiments | 45 |
| 5.3 Design | 46 |
| 5.3.1 Batched Packet Processing in Wide Click Pipeline | 47 |
| 5.3.2 Handling Packet Processing Divergence | 54 |
| 5.3.3 Asynchronous Processing on GPUs | 56 |
| 5.3.4 Improving Packet I/O | 57 |
| 5.4 Implementation | 58 |
| 5.5 Evaluation | 59 |
| 5.5.1 Packet I/O | 59 |
| 5.5.2 Applications | 60 |
| 5.5.3 Latency and Reordering | 61 |
| 5.5.4 Packet Processing Divergence | 62 |
| 5.5.5 Flexibility and Modularity | 63 |
| 5.6 Summary and Future Work | 63 |
| 6. A LITERATURE SURVEY OF SYSTEM-LEVEL GPU COMPUTING RESEARCH | 70 |
| 6.1 Applications | 71 |
| 6.1.1 Networking | 71 |
| 6.1.2 Storage | 73 |
| 6.1.3 Database Processing | 74 |
| 6.1.4 Security | 75 |
| 6.1.5 Program Analysis | 76 |
| 6.1.6 Data Compression | 77 |
| 6.1.7 Miscellaneous | 77 |
| 6.2 Techniques | 78 |
| 6.2.1 Batched Processing | 78 |
| 6.2.2 Reducing Memory Copy Overhead | 79 |
| 6.2.3 Dealing With Warp Control Flow Divergence | 81 |

| | | |
|-----------|--|-----------|
| 6.2.4 | Improving GPU Cores and DMA Engines Utilization | 82 |
| 6.2.5 | Virtualization and GPGPU Application Migration | 82 |
| 6.2.6 | Resource Management and Scheduling | 83 |
| 6.3 | Future | 84 |
| 6.4 | Summary | 85 |
| 7. | CONCLUSION | 87 |
| | REFERENCES | 89 |

LIST OF FIGURES

| | | |
|------|--|----|
| 2.1 | Architecture of a CUDA GPU | 8 |
| 4.1 | The architecture of GPUstore | 26 |
| 4.2 | The workflow of a GPUstore -based storage stack. | 26 |
| 4.3 | Storage services implemented with GPUstore | 30 |
| 4.4 | Throughput for a GPU kernel that copies data, but performs no computation | 35 |
| 4.5 | GPU AES cipher throughput with different optimizations compared with Linux kernel's CPU implementation. The experiments marked "w/o RB" use the techniques described in Section 4.1.2 to avoid double-buffering. | 35 |
| 4.6 | dm-crypt throughput on an SSD-backed device. | 36 |
| 4.7 | dm-crypt throughput on a RAM-backed device. | 36 |
| 4.8 | eCryptfs throughput on an SSD-backed filesystem. | 37 |
| 4.9 | eCryptfs throughput on a RAM-backed filesystem. | 37 |
| 4.10 | eCryptfs concurrent write throughput on a RAM disk for two different block sizes. | 39 |
| 4.11 | Throughput for the RAID 6 recovery algorithm with and without optimizations to avoid redundant buffers. | 41 |
| 4.12 | RAID 6 read bandwidth in degraded mode | 41 |
| 4.13 | RAID 6 read bandwidth in degraded mode on RAM-backed devices | 42 |
| 5.1 | Simplified Click configurations for motivating experiments | 47 |
| 5.2 | GPU vs. CPU performance on packet processing algorithms. Note that the axes are non-linear. | 48 |
| 5.3 | The <i>PacketBatch</i> structure | 51 |
| 5.4 | A Slicing example | 51 |
| 5.5 | Batching and debatching elements. Serial interfaces are shown as simple arrows, wide interfaces as double arrows. | 53 |
| 5.6 | A batched processing path with one GPU element. | 65 |
| 5.7 | Handling divergent packet paths | 66 |
| 5.8 | Performance of Click and Snap with three different applications. | 67 |
| 5.9 | Forwarding performance when using a GPUSDNClassifier that diverges to two GPUIDSMatcher elements. | 68 |

| | |
|---|----|
| 5.10 Major changes to the standard Click router to implement a GPU-based IDS router. | 68 |
| 5.11 Fully functional IP router + IDS performance | 69 |

LIST OF TABLES

| | | |
|-----|--|----|
| 4.1 | Approximate modified lines of code required to call GPU services using GPUstore. | 30 |
| 4.2 | GPUstore evaluation machine configurations | 34 |
| 4.3 | PostMark results for eCryptfs. Columns marked “M” indicate mixed workloads. CS and GS are the CPU’s and GPU’s results on the SSD; CR and GR are these two on a RAM disk. | 40 |
| 5.1 | Relative throughputs of simple processing pipelines. | 47 |
| 5.2 | Base Forwarding Performance of Snap and Click | 61 |

ACKNOWLEDGMENTS

I want to thank my parents and my grandparents, who always give me the freedom to choose my life and always encourage me to pursue my dream.

Endless thanks to my wife Shuying, who always gives me a backup and cheers me up in any circumstances. She also listened to my practice talks everyday during the week before the defense and gave me valuable comments and suggestions to improve my slides and presentation style. Thank you, for loving me!

I want to thank my daughter, Maggie. She is the main reason that drives me to graduate! She gives us endless laughter and tears.

Many thanks to my advisor, Robert Ricci. I can not use finite number of words to express my thanks to him at all. He is the best advisor and friend you can imagine! He listens to my strange ideas, supports me to study and do research on them, writes funding proposals for my project, modifies and refines every word of my awkward English paper drafts, gives me comments and suggestions on each single slide of my presentations, ... Thank you, for guiding me!

I also want to thank my friend and my office-mate Xing. We discussed a lot on my research project, papers, thesis and presentations, which helps me make this work. Thank you for your help in both study and life as my neighbor!

Special thanks to Mary Hall, who is also in my committee. I learned GPU programming and various techniques from Mary's class. Without her class, I would not be able to do this work at all. Mary also wrote a recommendation letter for me to help me get the NVIDIA fellowship, which funded part of my research. So, also thank NVIDIA!

I also want to thank my other committee members: John, Kobus and Eddie. John may be not aware of the fact that I started learning the low-level PCIe details because of his embedded class, where he showed us how easy the low-level hardware and principle is. Kobus gave me useful suggestions on how to improve my thesis. He also helped me with my Snap paper's publication, which makes me qualified for graduation. Eddie developed the Click modular router, without which Snap can not exist at all!

There are more people that I should thank. Mike, who is always helpful when I did experiments; Gary, who helped editing my GPUstore paper; Eric, the LaTeX guru that helped my papers, and also my presentations; Ren, who helped me with English grammars and words, and many other people who helped me in Flux or in school of computing.

CHAPTER 1

INTRODUCTION

System-level software sits at the bottom of the software stack, including file systems, device drivers, network stacks, and many other operating system (OS) components. They provide basic but essential storage, communication and security functionalities for upper-level applications. It is vital for system software to make those functionalities efficient and scalable, otherwise the entire software stack will be slowed down. Yet, many system-level functionalities require substantial computing power. Examples include: encryption for privacy; deep packet inspection for network protection; error correction code or erasure code for fault tolerance; and lookups in complex data structures (file systems, routing tables, or memory mapping structures). All of these may consume excessive processing power. What's more, today's application workloads are dramatically increasing: gigabytes or even terabytes of multimedia contents, high definition photos and videos, tens of even hundreds gigabits per second network traffic and so on. Thus more and more computing power is needed to process those bulk-data workloads on modern rich functional software stacks.

A very common and important feature of many system-level computational functionalities is that they are inherently parallelizable. Independent processing can be done in parallel at different granularities: data blocks, memory pages, network packets, disk blocks, etc. Highly parallel processors, in the form of graphics processing units (GPUs), are now common in a wide range of systems: from tiny mobile devices up to large scale cloud server clusters [1, 2] and super computers [3, 4]. Modern GPUs provide far more parallel computing power than multicore or many-core central processing units (CPUs): while a CPU may have two to eight cores, a number that is creeping upwards, a modern GPU may have over two thousands [5], and the number of cores is roughly doubling each year [6]. As a result, exploiting parallelism in system-level functionalities to take advantage of today's parallel processor advancement is valuable to satisfy the excessive demand of computing power by modern bulk-data workloads.

GPUs are designed as a *throughput-oriented architecture* [7]: thousands of cores work together to execute very large parallel workloads, attempting to maximize the total throughput, by sacrificing serial performance. Though each single GPU core is slower than a CPU core, when the computing task at hand is highly parallel, GPUs can provide dramatic improvements in throughput. This is especially efficient to process bulk-data workloads, which often bring more parallelism to fully utilize the thousands of GPU cores.

GPUs are leading the way in parallelism: compilers [8, 9, 10], algorithms [11, 12, 13], and computational models [14, 15, 16] for parallel code have made significant advances in recent years. High-level applications such as video processing, computer graphics, artificial intelligence, scientific computing and many other computationally expensive and large scale data processing tasks have benefited with significant performance speedups [17, 18] from the advancement of parallel GPUs. System-level software, however, have been largely left out of this revolution in parallel computing. The major factor in this absence is the lack of techniques for how system-level software should be mapped to and executed on GPUs.

For high-level software, GPU computing is fundamentally based on a computing model derived from the data parallel computation such as graphics processing and high performance computing (HPC). It is designed for long-running computation on large datasets, as seen in graphics and HPC. In contrast, system-level software is built, at the lowest level of the software stack, on sectors, pages, packets, blocks, and other relatively small structures despite modern bulk-data workloads. The scale of the computation required on each small structure is relatively modest. Apart from that, system-level software also has to deal with very low-level computing elements, such as memory management involving pages, caches, and page mappings, block input/output (I/O) scheduling, complex hardwares, device drivers, fine-grained performance tuning, memory optimizations and so on, which are often hidden from the high-level software. As a result, system-level software requires technologies to bridge the gap between the small building structures and the large datasets oriented GPU computing model; also to properly handle those low-level computing elements with careful design trade-offs and optimizations, to take advantage of the parallel throughput-oriented GPU architecture.

This thesis describes the generic principles of system-level GPU computing, which are abstracted and learned from designing, implementing and evaluating general throughput-oriented GPU computing models for two representative categories of system-level software: storage applications and network packet processing applications. Both models are in the form of general frameworks that are designed for seamlessly and efficiently integrating

parallel GPU computing into a large category of system-level software. The principles include unique findings of system-level GPU computing features, and generic design techniques and abstractions to deal with common system-level GPU computing challenges we have identified. The significant performance improvements in storage and network packet processing applications brought by integrating GPU computing with our frameworks shows the effectiveness and efficiency of the proposed techniques and abstractions, and hence supporting the following statement:

1.1 Thesis Statement

The throughput of system software with parallelizable, computationally expensive tasks can be improved by using GPUs and frameworks with memory-efficient and throughput-oriented designs.

1.2 Contributions

The contributions of this thesis work include the followings.

- Two general frameworks, **GPUstore** [19] and **Snap** [20], for integrating parallel GPU computing into storage and network packet processing software, as described in Chapter 4 and Chapter 5.
- Three high throughput GPU-accelerated Linux kernel storage components, including a filesystem and two storage device mappers, which are discussed in Chapter 4.
- A set of Click [21] elements for **Snap** to help build parallel packet processing pipelines with GPUs, which are described in Chapter 5.
- A modified fully functional Click internet protocol (IP) router with fast parallel GPU acceleration built on top of **Snap** as described in Chapter 5.
- A literature survey of system-level GPU computing that covers existing work, potential applications, comparison of the useful techniques applied in surveyed work and the ones proposed in our work as presented in Chapter 6.

1.3 Findings

Besides the above contributions, we also have found and learned valuable facts and lessons from designing and implementing **GPUstore** and **Snap**. We believe they are common and applicable to other system-level GPU computing software, too. We will thoroughly discuss them in the later chapters. For now, they are listed below as an overview.

- **Batching improves GPU utilization.** System code is often latency-oriented. As a result, system software often works with small building blocks. However, GPUs'

architecture is throughput-oriented. Small building blocks lead to constant overhead and low GPU utilization. To bridge the gap between the latency-oriented system code and throughput-oriented GPUs, batching small blocks to process many at once amortizes the overhead and improves GPU utilization. Batching may increase the processing latency of a single block, but for systems dealing with I/O, adding relatively small latency is tolerable.

- **High throughput computing needs truly asynchronous GPU programming.**

With the required synchronization, current GPU programming models are mismatched to asynchronous system components, which widely exist in file systems, block I/O, device drivers, and network stacks at system level. Synchronization stalls the typical GPU computing pipeline which consists of three common stages: (1) host-to-device direct memory access (DMA) copy; (2) GPU kernel execution; (3) device-to-host DMA copy. Event-based or callback-based asynchronous GPU programming is necessary for asynchronous systems code, allowing it to fully utilize the GPU pipeline and achieve high throughput.

- **System code has data usage patterns that are different from traditional GPU code.**

Traditional GPU computing usually reads the data from a file or the network into host memory, copies the entire data buffer into GPU device memory, and does the computation based on all the data. However, the code in a system component often works as a stage of a long data processing pipeline, and may use either the entire dataset passed through the pipeline, or just a few small pieces. In order to improve the system performance, special care must be taken to provide different memory models to system code (both the CPU side code and the GPU side code) according to the code's data usage pattern.

- **The computation needs all the data.** This data usage pattern is similar to traditional GPU computing. In this case, the GPU needs the entire dataset to be copied into GPU memory. To reduce the memory copy overhead, we should focus on reducing the copy within host memory. Different from traditional GPU computing, in the system-level context, processing stages often pass data buffers owned by a third party, such as the page-cache managed memory in the storage stack and network packet buffers in the network stack, through a long pipeline. Remapping the memory to make it GPU DMA-capable avoids redundant memory copy in host memory.
- **The computation uses just a few small pieces of the large trunk of**

data. In contrast to most GPU computing, some system-level computing tasks such as packet processing use only a few small pieces of the entire dataset. In this case, copying the entire dataset into GPU memory wastes a large portion of the bus bandwidth. Considering the much higher host memory bandwidth than the bus bandwidth, it is worth copying the scattered small pieces into a consecutive host memory buffer, allowing one smaller DMA copy to GPU. Such trade-off takes advantage of the faster host memory to reduce the overall memory copy time. Gathering the scattered pieces together also benefits memory access coalescing on the GPU, taking advantage of the wide GPU memory interface.

1.4 Thesis Outline

Chapter 2 gives a simple background introduction of concepts, techniques and features of general purpose GPU computing. This chapter tries to get readers without GPU computing background familiar with it, and also defines the terminology used in later chapters. Readers familiar with traditional GPU computing can skip this.

Chapter 3 describes the system-level GPU computing principles at the high level. It discusses the challenges and special requirements in system-level GPU computing and the proposed techniques and abstractions to deal with them.

Chapter 4 describes the storage framework **GPUstore**, including the design, implementation and the experimental evaluation. The GPU-accelerated file system and block device drivers built on top of **GPUstore** have achieved up to an order of magnitude performance improvements compared with the mature CPU-based implementations.

Chapter 5 describes the network packet processing framework **Snap**, similar to the previous **GPUstore**, including the design, implementation and the experiments. The demonstrated deep packet inspection router built with **Snap** has shown 40 Gbps (gigabits per second) line rate packet processing throughput at very small packet size.

Chapter 6 is the survey of system-level GPU computing, which includes both existing work and the identified possible application areas. The survey discusses techniques used by other system-level GPU systems to compare with what this thesis has proposed and applied.

Chapter 7 reviews the thesis and concludes.

CHAPTER 2

GENERAL PURPOSE GPU COMPUTING

This chapter describes essential general purpose GPU (GPGPU) computing background that is needed to help understanding the rest of the thesis. Currently, there are two most widely used GPGPU frameworks: compute unified device architecture (CUDA) [6] and open computing language (OpenCL) [22]. CUDA is a proprietary framework developed by NVIDIA corporation; OpenCL is a public framework designed by Khronos Group. Despite its proprietary feature, CUDA has several advanced features, such as concurrent streaming and flexible memory management which are helpful to system-level computing and OpenCL missed at the time of my thesis work. As a result, CUDA is used in this thesis to represent the lowest level GPGPU computing framework. And its terminology and concepts are used in this chapter to explain GPGPU computing. For a comprehensive description of CUDA-based modern GPGPU computing, readers may refer to NVIDIA’s CUDA programming guide [6].

2.1 Overview

A GPU works as a coprocessor of the CPU. It has dedicated video memory on the card to save the computing data. The processor cores on GPU can only access the video memory, so any data to be processed by the GPU must be copied into the video memory. To utilize the GPU, a typical workflow includes three steps:

1. CPU code copies the data to be processed from main memory (also called “host memory” in CUDA) to the video memory (also called “device memory”);
2. CPU code starts the GPU kernel, which is the program to be executed on the GPU, to process the copied data and produce the result in the device memory;
3. After the GPU kernel execution, the CPU code copies the result from device memory back to host memory.

Most GPUs sit in the peripheral component interconnect express (PCIe) slots, needing DMA over PCIe for the aforementioned memory copy. The GPU kernel is a program consisting

of GPU binary instructions. Using CUDA, programmers can write C/C++ code, then use `nvcc` to compile them into GPU instructions. CUDA also provides runtime libraries to allow CPU code to use special host memory, device memory, make DMA copy and manage GPU execution.

2.2 Parallel GPGPU

A GPGPU is a special single instruction multiple data (SIMD) processor with hundreds or thousands of cores and a variety of memory types (as showed in Figure 2.1). On recent CUDA GPUs, each 32 cores are grouped into a “warp”. All 32 cores within a warp share the same program counter. A GPU is divided into several “stream multiprocessors”, each of which contains several warps of cores and fast on-chip memory shared by all cores. A GPU kernel is executed by all the physical cores in terms of threads. So for a given GPU kernel, it becomes one thread on each GPU core when executing. High-end GPU models, such as GTX Titan Black [5], can have as many as 2880 cores. In that sense, up to 2880 threads can concurrently run on a GPU to execute a single GPU kernel. Threads on GPU cores may do hardware-assisted context switching, in case of memory operations or synchronization, which is determined and scheduled by the GPU hardware. Such zero-cost hardware context switching makes it possible to run millions of threads on a single GPU without any context switching overhead. Such a “single kernel executed by multiple threads” computing model is called single instruction multiple threads, or SIMT.

SIMT is a simple and efficient parallel processor design. But similar to other SIMD processors, the performance of SIMT may suffer from any control flow divergence because of the shared program counter within a warp. On CUDA GPUs, any control flow divergence within a warp may cause serialized control flow, such as the following code:

```

1  if (condition) {
2      doSomething();
3  } else {
4      doSomethingElse();
5  }
```

The `if-else` structure will be serialized by the SIMT architecture so that each thread executes both `doSomething()` and `doSomethingElse()` branches. When a core executes the control flow branch that its current thread should not jump into, it simply disables memory or register access to disable the execute effects. Different from the powerful sequential cores on CPUs that are armed with large cache, out-of-order execution and accurate branching prediction, a single GPU core is very weak. The performance speedup of many GPGPU programs comes from partitioning the workload into millions of small pieces processed by

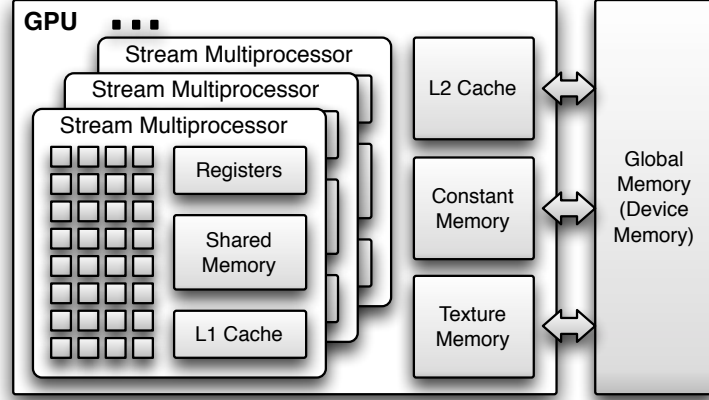


Figure 2.1. Architecture of a CUDA GPU

thousands of lane cores in parallel. As a result, a GPU kernel full of such control flow structures may severely slow down the GPU performance due to the control flow divergence.

SIMT architecture requires programmers to carefully tune their GPU kernels to avoid conditional structures and to reduce loops as many as possible. This can affect the data processing design in system software. For example, when using GPUs to process a batch of network packets through the network stack, packets in the batch may diverge to different paths because of their different protocols. Such path divergence will lead to warp divergence on GPUs because GPU kernel threads must execute different code for packets going to different paths. As we will see in **Snap**, we have proposed techniques and design principles to reduce the overhead caused by this problem.

2.3 GPU Memory Architecture

The original video memory on GPU board has been largely increased to up to 6GB (gigabyte) for a single GPU. Compared with the instruction execution on GPU cores, accessing GPU video memory, which is also called global memory, is much slower and may cost hundreds of cycles. So global memory access is a time consuming operation that should be minimized. Fortunately, the GPU memory bus width is much wider than normal CPUs, e.g. some model [5] has a 384 bits memory bus width which leads to 336 GB/s bandwidth. Though a single thread may not be able to fully utilize the 384 bits bus, a SIMT-oriented GPU kernel may take advantage of the wide bus by issuing memory access operations to consecutive locations from the same warp, which is called “memory access coalescing”.

Besides the global memory, there are some special memory types on GPUs for fast

memory access in certain applications (as showed in Figure 2.1). Constant memory is a small region of memory that supports fast read-only operations. Texture memory is similar to constant memory but has very fast special access patterns. For example, according to the Kargus [23] network intrusion detection system (NIDS), using texture memory for deterministic finite automata (DFA) transition table improves the pattern matching performance by 20%. Shared memory can be viewed as a program-controlled cache for a block of threads. Its access latency can be 100x lower than the uncached global memory. In fact, current CUDA GPUs use a single region of 64KB (kilobyte) on-chip memory per stream multiprocessor to implement both layer one (L1) cache and shared memory, allowing either 48 KB cache and 16 KB shared memory or 16 KB cache and 48 KB shared memory split that is configurable by programmers. Besides program-controlled cache for fast frequent data access, shared memory can also be used to achieve global memory access coalescing, e.g. when using four threads to encrypt a single sixteen-byte advanced encryption standard (AES) data block, each thread can issue a four-byte read operation to be coalesced into a single sixteen-byte transaction.

As a result, to achieve high performance GPU computing, a GPU program should be designed and executed with the following memory related considerations: consecutive threads should issue coalescable memory accesses; memory access latency can be hidden by launching more threads than cores in order to switch out the threads that wait for memory transaction; read-only data should go into constant or texture memory; frequently accessed data should go into shared memory.

2.4 Advanced GPGPU Computing

GPUs have been designed to be more than simple coprocessors. Many advanced techniques have been applied to improve the GPU computing environment. In the hardware perspective: GPUs are capable to have multiple GPU kernels executed concurrently on their cores; some high-end GPUs have more than one DMA engine on a single GPU board; the rapidly development of dynamically random access memory (DRAM) technology has significantly reduced the cost of large volume of DRAM in modern computers, so using large trunks of nonpageable dedicated physical memory for DMA is feasible. To utilize these hardware features to improve the GPU computing performance, CUDA has implemented and exposed software interfaces for programmers to access them.

CUDA provides the “stream” abstraction, which represents a GPU computing context consisting of a DMA engine for memory copy, and GPU cores for GPU kernel execution.

CUDA GPUs support multiple concurrent streams, which essentially enables concurrent execution of multiple GPU kernels on a single GPU, and utilizes multiple DMA engines for concurrently bidirectional PCIe transfers. Each stream is independent to one another: operations (memory copy and GPU kernel execution) in the same stream is sequentially performed, but operations in different streams are totally independent. This provides an asynchronous streaming GPU computing model, which may fully utilize the execution and copy capabilities of GPUs.

DMA memory copy requires nonpageable memory pages that are locked in physical memory. So if a program wants to copy the data in pageable memory with GPU DMA, the GPU driver has to firstly allocate an extra locked memory buffer, then copy the data from pageable memory to locked memory, and finally copy the data from locked memory to the GPU with DMA. This causes a double-buffering problem that causes an extra copy in host memory and also wastes extra memory space. As a result, a program should use locked memory directly to store its data in order to avoid this double-buffering problem in GPU DMA.

2.5 Efficient GPGPU Computing Techniques

To achieve efficient GPGPU computing, the aforementioned GPU features must be considered to design GPU kernels and to design the host side software.

- SIMT architecture requires as few branches as possible in GPU kernel’s control flow to avoid warp divergence. This is a more algorithmic requirement than a system one. However, as the packet processing example shows in Section 2.2, sometime the system-level processing control flow can cause the GPU side divergence, hence it requires the system design to be GPU-aware, to either avoid processing flow divergence, or apply techniques to reduce the divergence overhead like in **Snap**.
- The wide memory bus requires coalescable memory accesses in consecutive threads. This not only requires SIMT GPU kernel to guarantee coalescable memory accesses from consecutive threads, but also needs designing coalescing-friendly data structures or data layout in system-level software.
- The variety of different GPU memory requires GPU kernels to be carefully optimized and tuned to make use of the faster read only memory regions and shared memory. This is mostly a general algorithmic requirement, not a very “system” need. But having said in Section 2.3, system code still needs finely tuned system-level algorithms such as Kargus’ DFA table placement in texture memory that improves 20% pattern

matching performance.

- The CUDA stream technology requires programmers to make use of the concurrency and bidirectional PCIe bus with careful design. This needs system designers to take advantage of concurrent workloads or partitioning large single-threaded workload, properly abstracting the computations on GPUs to utilize the overlapped GPU computation and memory copy pipeline enabled by CUDA stream.
- The page-locked memory requires the programmers to design their memory management mechanism with the consideration of GPU resources at the very beginning. Because memory management in system-level, especially in OS components can be very complex. System components can either have their own memory allocators, or more often, just work as stages in a pipeline that don't own the memory of the processed data buffers at all.

Traditional GPGPU computing considers mostly the GPU side optimization techniques such as the SIMT architecture, wide memory bus and different types of device memory. Their workflows are often very simple, as showed in Section 2.1. For these kind of traditional GPGPU programs, all the complex low-level problems such as memory management, memory copy efficiency, bus bandwidth, device synchronization overhead, interaction with other system components, are hidden and handled efficiently by the system-level code. Now when it comes to the system-level software targeted in this thesis, all of them must be considered in the system design when integrating GPU computing. Special system-level data usage patterns may cause large PCIe bus bandwidth waste when copying memory, and may also cause difficulty to issue coalescable memory accesses from GPU threads. Using page-locked memory may be not as simple as calling the CUDA memory allocator, it may lead to complex memory page remapping in system code. Some inherently asynchronous system code may can't afford the host-device synchronization because that will change their efficient ways of working. All of these system-level issues must be handled in this thesis work, as will be showed in the next chapter.

CHAPTER 3

SYSTEM-LEVEL GPU COMPUTING

PRINCIPLES

A lot of system-level workloads are good candidates of the throughput-oriented GPU accelerations: file systems and storage device layers use encryption, hashing, compression, erasure coding; high bandwidth network traffic processing tasks need routing, forwarding, encryption, error correction coding, intrusion detection, packet classification; a variety of other computing tasks perform virus/malware detection, code verification, program control flow analysis, garbage collection and so on. Not every system-level operation is capable of throughput improvement with GPUs. Those inherently sequential or latency sensitive tasks such as process scheduling, acquiring current date-time, getting process identifier (ID), setting up I/O devices, allocating memory and so on definitely can't afford the relatively high-latency GPU communication that often across external buses.

Some early system-level work has demonstrated the amazing performance speedup brought by GPUs: PacketShader [24], a GPU-accelerated software router, is capable of running IP routing lookup at most 4x faster than the CPU-based mode. SSLShader [25], a GPU-accelerated secure sockets layer (SSL) implementation, runs four times faster than an equivalent CPU version. Gnort [26], as a GPU-accelerated NIDS, has showed 2x speedup over the original CPU-based system. EigenCFA [27], a static program control flow analyzer, has reached as much as 72x speedup when transforming the flow analysis into matrix operations on GPUs.

However, all of them are very specialized systems that have been done in an ad-hoc way to deal with performance obstacles such as memory copy overhead, parallel processing design and those mentioned in previous chapter, without considering any generic systematic design for a wide range of related applications. A general system-level GPU computing model needs to provide generic design principles. It requires a thorough study of system-level behaviors and GPU computing features, to identify the common challenges when integrating GPUs into low-level system components, and then come up with generic techniques and

abstractions to decompose the system, trade off design choices according to contexts, and finally solve the problems.

This chapter discusses the generic challenges faced by system-level GPU computing and our proposed design principles. Similar to traditional general purpose GPU computing, tuning GPU kernel to make parallel algorithms fully utilize the SIMT cores is important to system-level code, too. However, a unique feature of system-level computing tasks is that most of them are much simpler than the application-level tasks in scientific computing and graphics processing. If we take a look at the surveyed computing tasks in system-level in Section 6.1, they are mostly one-round deterministic algorithms without any iterative scheme and convergence requirement such as in partial differential equation (PDE) algorithms and randomized simulations. In that case, other noncomputational costs may become major overhead: for example, as we evaluated for **GPUstore**, the significant AES performance improvements with a variety of memory copy optimization techniques show that the memory copy, rather than the cipher computation, is the bottleneck. As a result, system-level GPU computing is faced with challenges from not only traditional device side GPU kernel optimization, but also the host side GPU computing related elements. The following sections discuss each of the major challenges we’ve identified in the **GPUstore** and **Snap** projects and the generic techniques to deal with them.

3.1 Throughput-oriented Architecture

The SIMT GPU architecture enables simple but highly parallel computing for large throughput processing. But the SIMT architecture is also a big challenge to GPU kernel design and GPU computing workload. As for the GPU kernel design, it means avoiding control flow divergence via tuned GPU code or SIMT-friendly data structures to process. To fit the SIMT architecture, the workload must be able to provide enough parallelism to fully utilize the thousands of parallel SIMT cores. The parallelism may come from the parallelized algorithms in terms of GPU kernels, more importantly, it comes from the amount of data to be processed in one shot, which often requires batched processing in system environments. When trying to provide a generic design for a variety of system-level applications, a right workload abstraction that covers the various processing tasks is needed to guarantee the parallelism, and also to make it easy to design divergenceless GPU kernels.

3.1.1 Parallel Workload Abstraction

Considering the thousands of cores on a modern GPU that processes one basic data unit in one thread, at least the same number of basic units is required to fully utilize all the

cores. This is actually the key to achieve high throughput in GPU computing. Taking the GTX Titan Black [5] as example: even with the parallelized AES algorithm that uses four threads to process a single sixteen-byte block, its 2880 cores still need $2880 \times 4 = 11KB$ data to process in one GPU kernel launch. If we consider the context switching due to global memory access, even more data are required to keep GPU cores busy.

Unfortunately, system-level software is often designed to process data sequentially due to the latency-oriented system design philosophy through the ages. However, the performance of parallel and sequential processors have been significantly improved to be easily achieve very low latency, and also today’s “big data” workloads focus more on high throughput processing, even at the system-level. Batching is a simple yet effective technique to accumulate enough parallelism in one-shot workload. Instead of a single data unit to be processed at one time, a one-shot workload with batching now includes enough units to fully occupy the thousands of GPU cores. “Enough” is ambiguous in the high-level workload abstraction. An intuitive but effective policy for concrete design is to ensure that processing the batched workload on GPUs is at least faster than on CPUs.

An obvious drawback of batching is the increased latency. This may cause serious performance problem in some latency sensitive software, especially the network applications such as video conference, video streaming, etc. As we will see in **Snap**, the packet processing latency introduced by batching is 8 to 9 times larger than the nonbatching latency. So batching is not a pervasive solution to adapt sequential software with the parallel GPUs. This also reveals the fact that not every system-level software can benefit from parallel GPUs, such as those latency sensitive computing tasks whose algorithms can’t be efficiently parallelized.

3.1.2 Avoid Control Flow Divergence

Many research projects [28, 29, 30, 31, 32, 33, 34] have been focused on designing efficient SIMT style parallel implementation of particular algorithms on the GPU. Besides those well-studied techniques, how to efficiently partition the workload into SIMT threads and how to choose SIMT-friendly algorithms and data structures (if possible) are also very important.

The workload partitioning is often very straightforward: each GPU thread processes one single basic data unit. Many system-level computing tasks can use such simple partitioning scheme such as computing erasure code of one disk sector per thread, encrypting a single sixteen-byte AES block per thread, looking up next hop for one IP address per thread, etc. However, exceptions always exist. The AES algorithm is an example, which has been

parallelized by exploiting the intra-block parallelism [35]. **GPUstore** applied this technique in its AES implementation to have four threads to encrypt/decrypt a single AES block.

Designing the data structures used by processing logic to make them SIMT-friendly is another technique to avoid control flow divergence. One example is the pattern matching algorithm in **Snap** that is based on DFA. The DFA matching is a very simple loop: get the current symbol; find the next state indexed by the symbol in current state’s transition table; use the found state as the current state. The algorithm itself can be implemented in SIMT style because all input strings use the same three steps matching. But for the second step, some transition table data structures can not ensure same code execution when finding the next state. For example, a tree-based map may need different lookup steps for different symbol to find the mapped state, which means different GPU threads may have to execute different numbers of lookup loops to find the next state, hence not SIMT-friendly. In the meanwhile, an array-based transition table, which assigns state for every symbol in the DFA alphabet, can guarantee equal steps to find the next state of any given symbol, hence SIMT-friendly. That’s because we can use the symbol as the array index to fetch the mapped state, and the state lookup becomes a single array access. So although the tree-based map is more memory efficient and the array-based map is very memory-consuming because even invalid symbols in the alphabet have transition states, sometimes the memory may be sacrificed to achieve SIMT-friendly GPU code.

3.1.2.1 When It Is Unavoidable

An effective but not efficient way that can always solve the control flow divergence is to partition the computing tasks into multiple GPU kernels, each GPU kernel works for a particular control flow branch of the computing task. This needs some information available at the host side to indicate how many threads a GPU kernel should be executed by, and which data items should be processed by a particular GPU kernel. Such host side information implies a device to host memory copy, the necessary host-device synchronization for its completion and probably a data structure reorganization, which may cost more than admitting divergence at the GPU side. So it is not always unacceptable to introduce divergence into a GPU kernel. The batched packet processing in different GPU elements in **Snap** is a very representative example.

However, allowing divergence in GPU kernel doesn’t mean writing arbitrary code, one still needs to minimize the effects. We have proposed and implemented a predicated execution to achieve the following goals for minimizing the divergence affects:

- **As few branches as possible:** This is obvious: we only allow necessary divergence, and should organize the code to minimize the branches.
- **Creating “empty” branch:** This can be illustrated by the following CUDA kernel code snippet:

```

1 void TheGPUKernel( ... ) {
2     if (shouldRun[threadId]) {
3         ... /* do processing */
4     }
5 }
6

```

The `if` statement has an empty `else` branch, which means even if the `if` statement causes warp divergence, the overhead to execute the other branch is very small because no instructions in that branch at all. So a GPU kernel that can’t avoid divergence should consider organize its code to such style.

More details of the predicated execution are in Section 5.3.2.

3.2 Coalescing Global Memory Accesses

The throughput-oriented GPU architecture is armed with not only high computation throughput, but also high memory access throughput. To fully utilize the 384 bits global memory bus in GTX Titan Black [5], coalescable memory accesses must be issued from neighbor threads in a thread block. Compared with non-coalesced access, coalesced memory access can achieve orders of magnitude faster [28]. So the GPU kernel must be designed to access consecutive memory locations from threads within the same thread block or warp. Besides the GPU kernel code, the data structures to be processed must be carefully organized to make them coalescing-friendly. Some kinds of workloads are quite easy to satisfy the coalescing requirement, for example, data read from a block device are naturally consecutively stored, hence when encrypting them, neighbor GPU threads can always coalesce their memory accesses. However, some workloads are not that coalescing-friendly, needing a coalescing-aware abstraction to achieve a generic solution. A simple but very good example is the IP routing lookup. The lookup GPU kernel needs only the destination IP address of a packet. If we put the entire packet into the global memory for each thread to access, the memory reads for IP addresses issued from neighbor threads will be scattered. One 384 bits memory transaction may only read 32 bits effective data (IPv4 address), which wastes more than 90% memory bandwidth. But if we organize the destination IP addresses into a separate buffer, then up to 12 memory reads for the IP addresses issued from neighbor threads can be coalesced into a single memory transaction, which significantly

reduces the number of global memory transactions. The “region-of-interest”-based (ROI) slicing technique applied in **Snap** is an effective abstraction of the workload data to build such coalescing-friendly data structures for a variety of GPU-accelerated computation (see Chapter 5).

3.3 Overlapping GPU Computation and Memory Copy

The typical GPU computing workflow mentioned in Section 2.1 can be pipelined to improve the utilization of the two major GPU components: SIMT cores and DMA engines. CUDA stream mentioned in Section 2.4 is the pipeline abstraction that enables overlapped GPU computation and memory copy. The pipelined model needs multiple streams, each of which carries proper size workloads for certain computing task. This may require the workload to be split into multiple trunks in order to fill into multiple streams. The host side code will be responsible to do such workload splitting, which needs the task-specific knowledge to ensure the splitting is correct. For a generic GPU computing framework such as **GPUstore**, it is impractical to put those splitting knowledge of each task into the generic framework. The computing tasks need a right abstraction to describe both the computing-specific logic and the task management (such as partitioning) logic. For example, **GPUstore** provides the modular GPU services, which are the abstraction of computing tasks. Each GPU service not only processes computing requests issued from **GPUstore** client code, but also do service-specific request scheduling including workload splitting and merging.

This requirement seems totally contrast to the parallelism one in Section 3.1.1. However, it is not a paradox. An application usually need to find the balance point of these two requirements to decide the optimal size of the workload to be processed on GPUs in one shot, so as to achieve the best performance. This can be done either dynamically or statically. A static approach finds the balance point offline by benchmarking the computing tasks and uses predefined sizes at run time. Both **GPUstore** and **Snap** simply use this approach in their prototypes. The dynamic way would do micro benchmarking at run time to find optimal sizes for the specific execution hardware and software environment. It can adjust workload sizes according to the hardware and system load, hence more accurate.

3.4 Asynchronous GPU Programming

CUDA stream requires asynchronous host side memory copy and GPU kernel launching operations so that a single CPU thread can launch multiple streams. However, current GPU programming needs synchronizations between the host and the device to detect the

completion of the device side operations in each stream, as showed in the following typical CUDA program code:

```

1  for (i=0; i<nrStreams; ++i) {
2      cudaMemcpyAsync(devBuf[i], hostBuf[i], sz, stream[i], ...);
3      gpuKernel<<<..., stream[i], ...>>>(...);
4      cudaMemcpyAsync(hostRet[i], devRet[i], retsz, stream[i], ...);
5  }
6  ...
7  cudaDeviceSynchronize(0);

```

In the code snippet above, all memory copy and GPU kernel launching are asynchronous calls, but the necessary `cudaDeviceSynchronize` call blocks the asynchronous control flow. This may not be a problem at all for traditional GPGPU computing that just focuses on a particular computing task. However, many system-level components are designed to exploit asynchrony in order to achieve high performance.

For example, filesystems work with the Virtual Filesystem (VFS) layer, and often rely on the OS kernel page-cache for reading and writing. By its nature, the page-cache makes all I/O operations asynchronous: read and write requests to the page cache are not synchronous unless an explicit `sync` operation is called or a *sync* flag is set when opening a file.

Other examples include the virtual block device drivers, which work with the OS kernel's block-I/O layer. This layer is an asynchronous request processing system. Once submitted, block-I/O requests are maintained in queues. Device drivers, such as SCSI or SATA drivers, are responsible for processing the queue and invoking callbacks when the I/O is complete.

Some filesystems and block devices, such as NFS, CIFS, and iSCSI, depend on the network stack to provide their functionality. Because of the potential for high, unpredictable, latency on a network, these subsystems are asynchronous by necessity.

When it goes to the network packet processing, it is totally asynchronous workflow. Although there does exist totally synchronous network stacks such as the uIP [36] and lwIP [37], they are designed for memory constraint embedded systems, not for performance. As a result, the synchronization CUDA call is totally unacceptable in those asynchronous environment because it will block the normal workflow and may cause performance slow down.

GPUstore and **Snap** solve this problem by implementing asynchronous CUDA stream callback mechanism to enable completely asynchronous GPU program control flow. Such callback mechanism can be implemented in two different approaches:

- Polling-based, which has a thread keep polling stream state change and invoke callbacks. It is obvious that this approach can get low latency response but will keep a CPU core in busy waiting loop.

- Signal-based, which is implemented with the events invoked by GPU interrupts at the low level. The signal-based response latency is definitely higher than the polling-based one, but its advantage is also obvious: CPU cores can be freed to process other work without busy waiting.

At the time of writing this thesis, latest CUDA release has provided similar signal-based callback for streams after more than two years `GPUstore` did that. This further confirms the effectiveness of the techniques I've proposed.

3.5 Reducing Memory Copy Overhead

Having said at the beginning of this chapter, the memory copy is often the major overhead compared with the computation. Many different aspects are related to the performance of memory copy. Since we are discussing system-level GPU computing, the memory copy is not only the DMA over PCIe bus, we also need to consider that happens in the host main memory as an entire system. We will start from the obvious challenges and problems in GPU related memory copy, then gradually introduce other vital issues and techniques to deal with them.

3.5.1 Reducing DMA Memory Copy Overhead

The overhead of DMA comes from its special requirement: the memory pages in main memory must not be swapped out (paging) during DMA. To satisfy such requirement when copying data in pageable memory to GPUs, CUDA GPU drivers use the aforementioned double-buffering approach as described in Section 2.4. According to the evaluations [38], double-buffering DMA can be two times slower than using page-locked memory directly. Compared with pageable memory, using page-locked memory may lock too many physical pages and reduce the available memory for other applications. However, considering today's cheap DRAM and widely available tens of gigabytes DRAM on a single machine, locking even 6GB memory (for a high end GPU [5]) is totally acceptable.

3.5.2 Avoid Double-Buffering

The aforementioned DMA overhead reducing solution with CUDA page-locked memory requires the system code to use CUDA-allocated memory as its data buffer. This seems trivial: traditional GPGPU computing usually reads the data from a file or the network into host memory, copies the entire data buffer into GPU device memory, and does computation. It is seldom that the host memory used in this scenario may have any other users or complicated dependencies, so it is easy to replace it with CUDA-allocated page-locked

memory. However, the code in a system component often works as a stage of a long data processing pipeline. In that case, it may be impractical to modify the entire system from the beginning of the pipeline to use CUDA’s memory, especially in a large complex system such as the operating system kernel. One way to deal with this is allocating a separate page-locked buffer with CUDA, and copy the data to be processed into this CUDA buffer before DMA (and similar approach for the processing result). This leads to double-buffering, which is similar to the aforementioned early stage CUDA DMA implementation for pageable memory: introducing extra copy in host memory. To avoid the double-buffering problem, I’ve proposed and implemented the page remapping technique, which can remap the non-CUDA page-locked memory pages into CUDA GPU driver’s memory area, and make them DMA-able just like CUDA page-locked memory (refer to Section 4.1.2.) This allows minimum invasive GPU computing integration into an existing complex system: only the component containing the computing tasks needs a few modification.

There does exist some special cases that the component we’d like to put GPU-accelerated computing tasks into may be the beginning of the data processing pipeline, or the one allocates memory for later use. In those cases, replacing previous `malloc`-like memory allocation with CUDA’s memory allocation functions is feasible and efficient. Introducing remapping in those cases doesn’t make any sense due to the added complexity of the extra page mappings.

The techniques in this section try to avoid memory copy in host memory because such copy is not necessary. However, there is an implicit assumption that the data in the memory buffer are all useful to the computation, otherwise it may be unwise to copy the entire buffer to device since it wastes the relatively slow PCIe bandwidth, as we shall discuss in the following section.

3.5.3 Trade Off Different Memory Copy Costs

Some computing tasks in a system component may need only small pieces of a large data unit it receives. Typical examples are the packet processing tasks: IP routing just needs the destination IP address, L2 forwarding just needs the destination MAC address, packet classification just needs five small fields, etc. Copying the entire data unit into GPU memory can waste a large portion of the PCIe bus bandwidth: considering the 4 bytes IP address versus the minimum 64 bytes packet. At the same time, the host side memory bandwidth is much faster than the PCIe bus. Take a look at the machine I used for **Snap** evaluation: the main memory has a 38.4GB/s bandwidth, while the maximum throughput of the PCIe 2.0 16x slot is 8GB/s in each direction, which is a 4.8 times difference. As a

result, for computing tasks with the data usage patterns discussed here, copying the needed data pieces into a much smaller page-locked memory buffer and then launching a PCIe DMA for this small buffer may lead to a much faster total memory copy than copying the entire data unit into GPU memory through PCIe bus. The aforementioned **Snap**'s ROI-based slicing technique is based on this idea. It takes advantage of the much faster host memory to achieve fast host-device memory copy. As we mentioned in Section 3.2, it also makes coalescing-friendly data structures for GPUs (refer to Chapter 5 for details).

3.5.4 Memory Copy Techniques Discussion

Section 3.5.3 advocates a technique that is totally opposite to the one in its previous section (Section 3.5.2). But they are not paradoxical. It is the data usage pattern of the computing task that decides which technique to use. On the one hand, for a computing task that needs the entire data unit received by the enclosing system for its computation, it should avoid extra buffers in host memory for the same data unit content with either the remapping technique or using CUDA page-locked memory directly, depending on the memory management role of the system as we discussed in Section 3.5.2. On the other hand, a computing task needing only small pieces of the entire data unit should trade off the main memory bandwidth and the host-device memory copy bandwidth, and may use techniques similar to **Snap**'s ROI-based slicing for faster total memory copy performance. **Snap**'s ROI-based slicing is a very flexible data abstraction to deal with both kinds of usage patterns: it allows computing tasks to define their own pieces to be processed in a single data unit, for a “use all” usage pattern, the computing task can specify the entire data unit as its interested region.

3.6 Miscellaneous Challenges

There are some challenges caused by the limitations in current GPU computing libraries or GPU hardware. Though we can believe that with the development of GPU technology and the GPU market, these obstacles may disappear in future, it is still worth describing them here for readers who are trying to apply GPU computing into system components at this time. The challenges discussed in this section are only related to my GPU computing experiences learned from **GPUstore** and **Snap**, the survey chapter (Chapter 6) discusses more such challenges and solutions provided by other system-level GPU computing works.

3.6.1 OS Kernel Mode and User Mode

Communication for GPU Computing

Currently, almost all GPU computing libraries, drivers are close-source, not to mention the even more closed GPU hardware. This leads to a big performance problem when applying GPU computing into operating system: OS kernel mode code has to rely on the userspace GPU computing library to use GPUs. So now in OS kernel mode, using GPUs is not as efficient as a function call in userspace, but a cross context communication. Such system needs an efficient kernel-user communication mechanism for invoking GPU computing library functions, and also memory sharing between two modes for computing data. **GPUstore** got this problem on Linux kernel when using CUDA GPU library. Current open source GPU drivers such as nouveau [39] and open source CUDA implementation such as Gdev [40] still can't reach the proprietary software's performance of GPGPU computing. **GPUstore** uses a userspace helper to deal with requests from OS kernel and invoke CUDA calls. The userspace helper is based on polling-based file event mechanism to achieve fast kernel-user communication. The details are in Chapter 4. The Barracuda [41] in-kernel GPU microdriver has evaluated different approaches to implement efficient kernel-user communication for GPU computing.

3.6.2 Implicit Host-Device Synchronization

The host-device synchronization happens not only when the host side explicitly calls `cudaDeviceSynchronize` like functions, but also implicitly when some GPU resource management operations are performed [6]. Such operations include CUDA host or device memory allocation, GPU information query, etc. The host and device side memory allocation is the main trouble maker because it is almost unavoidable. This may stall the aforementioned asynchronous GPU programming in Section 3.4 even when there is no explicit stream synchronization. As a result, as we will see in Chapter 6, a common technique used by many system-level GPU computing works including my **GPUstore** and **Snap** frameworks is to pre-allocate CUDA page-locked memory, and manage the memory allocation on their own. This can easily consume a lot of memory, but due to the current GPU limitations, it is a must to achieve non-stalled asynchronous GPU computing.

3.7 Summary

In this chapter, we discussed generic system-level challenges and also proposed high-level principles and techniques to deal with them. The proposed principles and techniques are designed to make system code efficiently work with the throughput-oriented GPU architecture, take advantage of the wide GPU memory interface, do non-blocking host and device

communication and reduce unnecessary overhead during GPU DMA. These principles are mainly about: batching to provide parallel workloads; truly asynchronously programming GPUs with callbacks or status polling at CPU side; compacting workload data to reduce unnecessary PCIe transfer; and using locked memory directly to avoid double-buffering DMA. In the next two chapters, we will discuss our two concrete frameworks: **GPUstore** and **Snap**, to explain how we apply these generic principles and techniques in practice to deal with their specific problems.

CHAPTER 4

GPUstore: HARNESSING GPU COMPUTING FOR OS STORAGE

This chapter covers the design, implementation and evaluation of GPUstore. GPUstore is a framework for integrating GPU computing into storage systems in Linux kernel. Section 6.1.2 lists typical computationally expensive tasks in storage systems. Those are mostly served for security, confidentiality, reliability, integrity, performance and new functionality in storage systems. Some example storage subsystems containing such tasks include:

- **Security and Confidentiality:** eCryptfs, Cryptfs [42], CFS [43], dm-crypt, etc.
- **Reliability and Integrity:** ZFS [44], software RAID, I³FS [45], etc.
- **Performance:** Venti [46], content-addressing (HydraFS [47], CAPFS [48]), etc.
- **New Functionality:** in-kernel databases [49], steganographic filesystems [50], etc.

These storage subsystems fall into two main categories: filesystems, such as eCryptfs and I³FS, and virtual block device drivers, such as dm-crypt and md.

Different from the works surveyed in Section 6.1.2, GPUstore is a generic framework, not for a particular storage application or subsystem as listed above. It has been designed to collaborate with the storage subsystem in Linux kernel in order to use the GPU as a co-processor. We try to minimize the in-kernel change for a storage component to use GPU computing. So GPUstore follows the OS working style and utilizes existing resource management mechanisms to avoid any fundamental change in the OS kernel. GPUstore has been evaluated with three storage system case studies, showing its efficiency and effectiveness.

4.1 Design

This section will go into the details of how GPUstore has been designed to apply the generic technical principles discussed in Chapter 3 in practice. The following text first takes an overview of the architecture, then discusses specific aspects including memory management, request scheduling and stream management.

4.1.1 Overview

GPUstore has three main functional blocks: memory management, request management and streams management, as showed in Figure 4.1. There are also GPU “services” that don’t belong to the framework, but are managed by **GPUstore** and essential to provide GPU computing tasks. **GPUstore** abstracts the computing tasks into “services”. So storage components request for services the get needed computational functionalities. Services are modular libraries dynamically linked with **GPUstore**, a generic service interface is provided to let **GPUstore** uses and manages them. Due to the close source GPU driver and library, **GPUstore** has to use a userspace helper to interact with the userspace CUDA GPU library. A GPU service also has to be split into two parts: the kernel space part and the userspace part. The kernel part mainly hides the **GPUstore** API calls to turn a function call invoked from storage components into a GPU service request. The userspace part of a service deals with the necessary host-device memory copies and GPU kernel launching to process a request. The workflow of the three example GPU-accelerated storage systems built with **GPUstore** are illustrated in Figure 4.2.

4.1.2 Memory Management

GPUstore manages both host memory and device memory that are used for GPU computing. All the host memory is CUDA page-locked memory to achieve best memory copy performance, and also pre-allocated to avoid stalling the asynchronous GPU computing, as discussed in Section 3.6.2. **GPUstore** simplifies the device memory management by maintaining a one-to-one mapping between the host memory and the device one. The downside of such one-to-one mapping is that it makes sub-optimal use of host memory: buffers must remain allocated even when not in active use for copies.

The Linux kernel storage workflow can be treated as a long pipeline, as illustrated in Figure 4.2, each layer may either allocate memory buffer for new data processing, or accept data buffers from neighbor layers to process. Most computing tasks in the storage systems work on the entire data rather than just small portions of them. That said in Section 3.5.4, such data usage pattern requires avoiding double-buffering when integrating GPU computing. According to our AES cipher evaluation showed in Figure 4.5, avoiding double-buffering adds almost 3x speedup to the cipher performance.

GPUstore provides API to do memory remapping for components processing data buffers received from others, and also in-kernel CUDA memory allocation for components creating its own buffers for data processing. For example, **eCryptfs** in our case studies is a good candidate to use remapping because all the memory pages it uses are allocated and managed

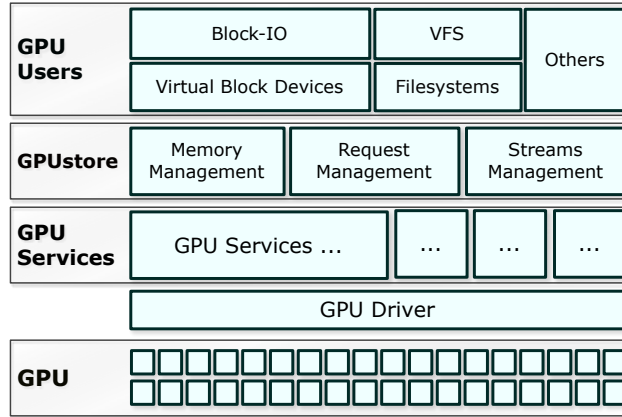


Figure 4.1. The architecture of GPUstore.

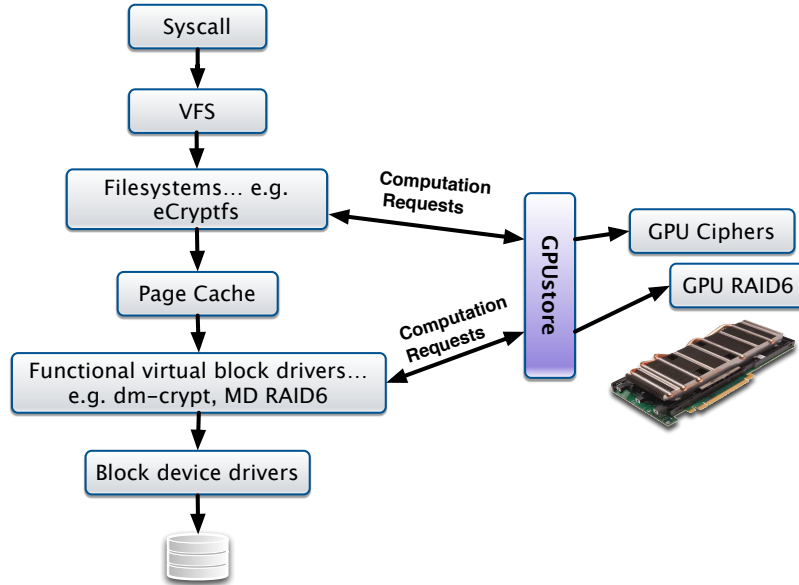


Figure 4.2. The workflow of a GPUstore-based storage stack.

by kernel page cache; while `dm-crypt` is a good fit of allocating its own buffers capable of GPU DMA because its code creates and manages buffers for the encrypted/decrypted data.

“Remapping” means mapping normal memory pages into GPU driver’s memory area to let it treat them same as its page-locked memory that is capable of DMA. Remapping is not a safe way in some cases: some memory pages such as the `mmap`-ed ones may not be allowed to do further remapping in Linux kernel, we “rudely” manipulate the page tables for remapping regardless of the potential problem. Although we haven’t got any errors for our case studies, it may cause bugs in some circumstances.

Allocating CUDA page-locked memory in Linux kernel is achieved with the `GPUstore`

userspace helper. The helper is responsible for all the userspace CUDA calls, and the kernel part takes care of necessary address translation for kernel code use and manages the memory in a memory pool.

4.1.3 Request Management

4.1.3.1 Request Processing

GPU service requests are submitted to **GPUstore** by storage components (via the kernel part GPU service), passed across the kernel-user boundary to the userspace helper, then processed by the requested services, and finally got the response. **GPUstore** processes service requests asynchronously to avoid breaking the normal workflow of storage components as we discussed in Section 3.4. So each request has a callback that is invoked after it has been processed.

GPUstore exposes a Linux character device file for the kernel-user communication. The userspace helper **reads** submitted requests, and after a request has been processed, it **writes** to that file for completion notification. The userspace helper works completely asynchronously to achieve low latency. It does non-blocking **read** on the **GPUstore** device file, and also relies on the non-blocking **write** implemented by the device file. **GPUstore** implements such non-blocking **write** using a separate kernel thread that is blocked on a request completion queue, which is filled by the device file's **write**. The kernel thread is waked up on available elements in the completion queue and responsible for invoking the request callbacks.

GPUstore defines a generic service interface, which includes three common steps to process a request:

- **PrepareProcessing**, this allows a service to do some preparation before processing a request, such as calculating CUDA kernel execution parameters, launching asynchronous host-to-device memory copy, etc.
- **LaunchKernel**, this does the CUDA kernel launching operation, which is an asynchronous operation.
- **PostProcessing**, this performs asynchronous device-to-host memory copy and other necessary operations in specific services.

A service only does asynchronous GPU operations in each step to avoid blocking the userspace helper's workflow. One service doesn't have to follow the typical host-to-device copy, GPU kernel execution and device-to-host copy GPU computing steps. For example, two services can collaborate in such way: the first service produces an intermediate result, without copying it back to host, and then the second service processes the result on GPU

directly, no host-to-device copy needed. This allows **GPUstore** users to do efficient data flow control by saving unnecessary memory copies, which is similar to the pipe data flow model in PTask [51].

4.1.3.2 Request Scheduling

In the userspace helper, request processing is started as first-come-first-serve. However, due to the totally asynchronous CUDA calls, the completion order of requests is not guaranteed. **GPUstore** doesn't try to maintain the order of requests to allow fully asynchronous processing. Users who needs such processing order must maintain it via the request callbacks or special logic in GPU services.

Having said in Section 3.1.1 and Section 3.3: GPUs need enough parallelism in one-shot computing workload which is often accumulated via batching; GPUs also need to pipeline computation and memory copy to improve its component utilization. But most kernel storage code is unaware of GPUs, and hence can hardly issue requests with GPU-friendly workload sizes. Changing exist code to make GPU-friendly workloads may lead to significant amount code refactoring. **GPUstore** tries to reduce the amount of such refactoring by merging or splitting accepted requests.

The *merge* operation is performed on small size requests (for the same service), which is an analogy of batching as said in Section 3.1.1. The *split* operation is on large size requests to utilize the overlapped GPU computation and memory copy technique (see Section 3.3). Both merging and splitting requires the service-specific knowledge to guarantee the correctness of the result requests, so **GPUstore** doesn't perform the actual merging and splitting, instead, it depends on each GPU service to do them. So GPU services can optionally implement merging or splitting logic, which will be utilized by **GPUstore** for request scheduling. In current **GPUstore** system prototype, all GPU services use predefined constant values to decide their optimal request sizes capable of best performance, which may be affected by runtime system load and other factors. Implementing dynamically adjustable merging or splitting parameters is an interesting future work.

merge is not the “ultimate” solution to produce GPU-friendly workloads with enough parallelism. *merge* can only process requests already submitted to **GPUstore**. If a storage component explicitly splits an already large size workload to process one much smaller unit at a time in a loop, it causes unnecessarily sequential constraints on requests: the loop must wait until completion of the previous request in order to do the next round. In that case, *merge* can't find enough requests in request queue to merge at all. Such component side splitting may be done for two main reasons:

- **code simplicity:** This can be illustrated using the `read_page` and `read_pages` example in `eCryptfs` as we shall describe in Section 4.2.1.2: Linux kernel page cache support both single-page and multi-page read/write operations, and by default, if a file system doesn't provide the multi-page implementation, it uses the single-page operation to create a default multi-page one by looping around each single page. Many file systems choose relying on kernel's default implementation and just implement the single-page operation for simplicity. Similar analogies exist in many other kernel components.
- **reducing memory consumption:** some operations need locking the data memory page during processing, hence locking a single page at one time looks more memory-efficient than locking all the pages.

Both cases need refactoring the existing code to eliminate the component side workload splitting. For the later case, one must trade-off the resource-efficiency and performance, although memory is not often the rare resource any more thanks to the rapidly developed DRAM industry.

4.1.4 Streams Management

CUDA streams are assigned to service requests in `GPUstore` to achieve sequentially executed operations on the same request. CUDA stream is also a GPU computing resource abstraction, which is the analogy of a CPU thread or process. So managing the allocation is analogous to process scheduling on CPUs. Due to the functionally limit, GPUs are not capable of preemptive execution, so `GPUstore` uses first-come-first-serve policy to allocate streams to requests, and the request processing is unpreemptable (once started). Even with current GPU functional constraint, it is still possible to enable prioritized request execution with techniques similar to TimeGraph [52] and Gdev [40], though it may need significant engineering work to make use of the immature open source GPU drivers and CUDA libraries they depend on.

4.2 Implementation

`GPUstore` has been prototyped on Linux kernel to accelerate three existing kernel storage components. We enhanced encrypted storage with `dm-crypt` and `eCryptfs`, and the software RAID driver `md`. We chose these three subsystems because they interact with the kernel in different ways: `md` and `dm-crypt` implement the block I/O interface, and `eCryptfs` works with the virtual filesystem (VFS) layer. The architecture of these implementations is shown in Figure 4.3.

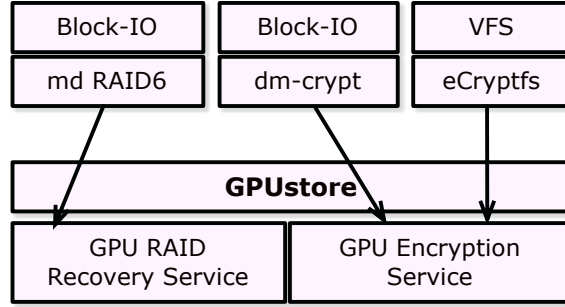


Figure 4.3. Storage services implemented with GPUstore.

Table 4.1. Approximate modified lines of code required to call GPU services using GPUstore.

| Subsystem | Total LOC | Modified LOC | Percent |
|-----------------|-----------|--------------|---------|
| dm-crypt | 1,800 | 50 | 3% |
| eCryptfs | 11,000 | 200 | 2% |
| md | 6,000 | 20 | 0.3% |

The design of **GPUstore** ensures that client subsystems need only minor modifications to call GPU services. Table 4.1 gives the approximate number of lines of code that we had to modify for our example subsystems. The lines of code reported in this table are those in the subsystems that are modified to call **GPUstore**, and do not include the lines of code used to implement the GPU services. Linux storage subsystems typically call out to other re-usable kernel components to perform common operations such as encryption: essentially, we replace these with calls to **GPUstore** and make minor changes to memory management.

4.2.1 Encrypted Storage

Encrypted storage is a good match for the GPU: while encryption and decryption are computationally expensive, most ciphers used by storage systems are parallelizable. This parallelization can happen at different levels, such as files, requests or disk sectors.

We added support for **GPUstore**-accelerated ciphers to the Linux cryptography subsystem, where they can be selected by any client of the Linux cryptography API; **dm-crypt** and **eCryptfs** already make use of this API. To achieve high levels of performance, modest changes are required in the storage systems. These changes address the memory allocation and request size issues discussed in Section 4.1.

Although there exist dedicated encryption accelerators and CPUs with special encryption features (such as AES-NI [53]), GPUs are nonetheless quite competitive. As we will see in Section 4.3, our AES cipher is able to sustain a rate of over 4 GB/s; results reported for

a high-end CPU with AES-NI [25] are less than half of this speed. GPU implementations of ciphers are also much more flexible: since they are “just code,” they allow for any cipher to be implemented, rather than “baking in” support for a particular cipher into hardware as AES-NI does.

4.2.1.1 dm-crypt: Encryption Block Device Layer

dm-crypt is a crypto target in Linux’s device mapper framework. It works as a virtual block device, running on top of other block devices such as disks. **dm-crypt** provides an encrypted view of the underlying devices by performing block-level encryption. It works asynchronously, as do most other block drivers.

dm-crypt interacts with the block I/O layer of the Linux kernel, taking I/O requests and splitting them into 4KB pages, which are asynchronously encrypted or decrypted as required. As we will see in Figure 4.5, for 4KB blocks, the size used by the unmodified **dm-crypt**, GPU overheads dominate and CPU encryption is faster. In our modified version of **dm-crypt**, we leave the task of combining these 4KB pages into larger blocks to the **GPUstore** framework: it relies on the *merge* operation to coalesce several small requests into one large one. **dm-crypt** exposes the unencrypted layer to its upper users, and reads/writes encrypted data to its lower block devices. So it actually allocates new buffers for the cipher operations for most of time (it does support in-place cipher operations, which is rarely used). Having said in Section 4.1.2, **GPUstore**’s remapping doesn’t fit such memory usage, hence we replace **dm-crypt**’s original calls to Linux kernel’s **kmalloc** family functions with **GPUstore**’s in-kernel CUDA page-locked memory allocator. Our changes to **dm-crypt** are minor: our implementation changes only approximately 50 lines of code, less than 3% of the 1,800 total lines in **dm-crypt**.

4.2.1.2 eCryptfs: Encrypted File System

eCryptfs [54], which is derived from **Cryptfs** [42], is a “stacking” filesystem. To use it, one creates another standard Linux filesystem, such as EXT3, and mounts **eCryptfs** “on top” of it. **eCryptfs** uses the underlying filesystem as a backing store and handles encryption and decryption of the files stored to it. It is part of the mainline Linux kernel, and is used for user home directory data security in many Linux distributions such as Ubuntu, Fedora, and Android.

To get good performance from **eCryptfs**, we made two changes. The first was to change the cipher mode from CBC to CTR: CBC encryption cannot be parallelized because ciphertext from each block is used as input to the following block, creating dependencies. CTR,

in contrast, uses a counter to provide non-uniform input, creating no such dependencies and allowing us to evaluate both read and write (decryption and encryption) workloads. In practice, implementations could still choose CBC mode and forgo GPU acceleration for encryption while still taking advantage of it for decryption. This would be a good tradeoff for read-heavy workloads.

The second change was to enable **eCryptfs** to perform cryptography on larger blocks of data. Like **dm-crypt**, **eCryptfs** uses page-sized units for cryptography operations, which means that even when a large I/O buffer is used by higher level applications, **eCryptfs** still splits it into small pages and processes them one by one. This is not a problem when there are many concurrent readers and writers, as **GPUstore**'s *merge* operation can coalesce requests from different I/O streams into large GPU requests. However, with a single reader or writer, large individual I/O requests must be supported in order to attain high performance. We therefore modified **eCryptfs** to support the Linux kernel page cache's **readpages** interface, which allows the kernel's page cache layer to pass multi-page reads directly to the filesystem. Large write support was provided by a replacement **write** operation at the VFS layer instead of **writepages** due to intricacies of the Linux kernel page cache system. Implementing these multi-page read and write interfaces enabled the kernel's existing mechanisms, such as readahead, I/O scheduling, and caching, to also take advantage of **GPUstore**'s *split* request scheduling for GPU performance enhancements.

Despite the fact that we made more changes to **eCryptfs** than **dm-crypt**, our modifications to **eCryptfs** are still relatively small: we modified approximately 200 lines of code, or less than 2%, of **eCryptfs**. Those changes also include remapping **eCryptfs** pages into the GPU driver's address space for GPU service requests; these pages are managed by the kernel's page cache, and it is therefore unfeasible to allocate them from the GPU driver.

4.2.2 md: Software RAID

RAID services may be provided in hardware or in software; the Linux kernel includes a software RAID implementation called **md**. RAID 6, which can continue operating with two failed disks, is widely deployed as a space-efficient, fault-tolerant RAID configuration [55]. RAID 6 employs erasure codes to allow a set number of disks in a system to fail with no data loss. While many proprietary RAID 6 codes exist [56, 57], the Linux kernel includes a RAID 6 implementation based on Reed-Solomon coding [58].

RAID 6 coding, which occurs during write workloads, has efficient SSE implementations for x86 and x86-64 processors. Unfortunately, the available vector instructions are not sufficient for RAID 6 decoding, or data recovery, which must occur during reads when a

disk has failed [59]. The primary issue is that Reed-Solomon coding with an eight-bit word size relies on a lookup table with 256 entries, which is not vectorizable with current CPU instructions, resulting in serialized access to these tables.

We have implemented GPU-accelerated data recovery for RAID 6 and integrated it into the `md` device driver. Our recovery code handles the particularly challenging case where two data strips within a stripe holding data have failed: in this case, all remaining data strips, plus both parity disks, must be used to recover the missing data. Since RAID 6 distributes parity strips across the entire disk array, failure of any two disks will trigger this recovery case.

While RAID stripes can be large enough in size to overcome GPU overheads, the `md` driver, like `dm-crypt`, is fundamentally organized around processing individual 4KB memory pages. `md` is a very complex system that includes working threads, asynchronous parity operations, and multiple internal request queues. We decided not to modify `md` to operate on larger requests, as that would require fundamental changes in its architecture. Instead, we rely on `GPUstore`'s *merge* to deal with small requests. Also similar to `dm-crypt`, `md` allocates new buffers for the recovered data or the parity, which lead us to using `GPUstore`'s in-kernel CUDA memory allocator to replace original Linux kernel's allocator to avoid double-buffering. As a result, our GPU integration for `md` requires only 20 lines of changes, most of which are simply saving buffer pointers for safe sharing of asynchronous calls.

4.3 Evaluation

We benchmarked the `GPUstore` framework itself as well as the three storage subsystems that we adapted to use it. We used two machine configurations, *S1* and *S2* as showed in Table 4.2, for our evaluation. *S1* is used for filesystem and block device tests. *S2* is used for the RAID benchmarks.

All benchmarks were run without use of *hybrid mode* in `GPUstore`; that is, GPU services were not allowed to fall back to the CPU for small requests. This has the effect of clearly illustrating the points where the GPU implementation, by itself, underperforms the CPU, as well as the points where their performance crosses. With *hybrid mode* enabled, `GPUstore` would use the CPU for small requests, and the CPU performance can thus be considered an approximate lower bound for `GPUstore`'s hybrid mode performance.

In many cases, our GPU-accelerated systems are capable of out-performing the physical storage devices in our systems; in those cases, we also evaluate them on RAM-backed storage in order to understand their limits. These RAM-based results suggest that some `GPUstore` accelerated subsystems will be capable of keeping up with multiple fast storage devices

Table 4.2. GPUstore evaluation machine configurations

| Configuration | <i>S1</i> | <i>S2</i> |
|---------------|--|--|
| CPU | Intel Core i7 930 Quad-Core | Intel Core i7 975 Quad-Core |
| DRAM | 6 GB DDR3 | |
| HardDrive | Intel X25-E 32GB SATA SSD, 250 MB/s read, 170 MB/s write | Thirty-two 750 GB 7200 RPM SATA hard disks in two switched JBODs via a 2 port Fibre Channel 4 Gb adapter |
| GPU | NVIDIA GTX 480, 480 cores @ 1.4GHz, 1.5 GB GDDR5 RAM | |
| Software | Ubuntu Linux with 2.6.39.4 kernel, CUDA 4.0 | |

in the same system, or PCI-attached flash storage, which is much faster than the drives available for our benchmarks.

4.3.1 Framework Performance

Our first microbenchmark examines the effect of block sizes on GPUstore’s performance and compares synchronous operation with asynchronous. On *S1*, we called a GPU service which performs no computation: it merely copies data back and forth between host memory and GPU device memory. Note that total data transfer is double the block size, since the data block is first copied to GPU memory and then back to host memory. In Figure 4.4, we can see that at small block sizes, the PCIe bus overheads dominate, limiting throughput. Performance steadily increases along with block size, and reaches approximately 4 GB/s on our system. This benchmark reveals three things. First, it demonstrates the value to be gained from our *merge* operation, which increases block sizes. Second, it shows a performance boost of 30% when using asynchronous, rather than synchronous, requests to the GPU. Finally, it serves as an upper bound for performance of GPU services, since our test service performs no computation.

Our second microbenchmark shows the effects of our optimization to remove redundant buffering and the *split* operation. This benchmark, also run on *S1*, uses the AES cipher service on the GPU, and the results can be seen in Figure 4.5. The baseline GPU result shows a speedup over the CPU cipher, demonstrating the feasibility of GPU acceleration for such computation. Our *split* operation doubles performance at large block sizes, and eliminating redundant buffering triples performance at sizes of 256 KB or larger. Together, these two optimizations give a speedup of approximately four times, and with them, the GPU-accelerated AES cipher achieves a speedup of 36 times over the CPU AES implementation in the Linux kernel. The performance levels approach those seen in Figure 4.4,




Figure 4.4. Throughput for a GPU kernel that copies data, but performs no computation




Figure 4.5. GPU AES cipher throughput with different optimizations compared with Linux kernel’s CPU implementation. The experiments marked “w/o RB” use the techniques described in Section 4.1.2 to avoid double-buffering.

implying that the memory copy, rather than the AES cipher computation, is the bottleneck.

4.3.2 dm-crypt

Next, we use the `dd` tool to measure raw sequential I/O speed in `dm-crypt`. The results shown in Figure 4.6 indicate that with read and write sizes of about 1MB or larger, the GPU-accelerated `dm-crypt` easily reaches our SSD’s maximum throughput (250MB/sq read and 170MB/s write). The CPU version is 60% slower; while it would be fast enough to keep up with a mechanical hard disk, it is unable reach the full potential of the SSD. Substituting




Figure 4.6. `dm-crypt` throughput on an SSD-backed device.




Figure 4.7. `dm-crypt` throughput on a RAM-backed device.

a RAM disk for the SSD (Figure 4.7), we see that the GPU-accelerated `dm-crypt` was limited by the speed of the drive: it is able to achieve a maximum read throughput of 1.4 GB/s, more than six times as fast as the CPU implementation. This is almost exactly the rated read speed for the ioDrive Duo, the third fastest SSD in production [60] at the time of developing `GPUstore` (in 2012). As the throughput of storage systems rises, GPUs present a promising way to place computation into those systems while taking full advantage of the speed of the underlying storage devices.




Figure 4.8. eCryptfs throughput on an SSD-backed filesystem.




Figure 4.9. eCryptfs throughput on a RAM-backed filesystem.

4.3.3 eCryptfs

4.3.3.1 Sequential Throughput

Figure 4.8 and Figure 4.9 compare the sequential performance for the CPU and GPU implementation of eCryptfs. We used the `iozone` tool to do sequential reads and writes using varying block sizes and measured the resulting throughput. Because eCryptfs does not support direct I/O, effects from kernel features such as the page cache and readahead affect our results. To minimize (but not completely eliminate) these effects, we cleared the page cache before running read-only benchmarks, and all writes were done synchronously.

Figure 4.8 shows that on the SSD, the GPU achieves 250 MBps when reading, compared with about 150 MBps for the CPU, a 70% speed increase. Unlike our earlier benchmarks,

read speeds remain nearly constant across all block sizes. This is explained by the Linux page-cache’s readahead behavior: when small reads were performed by `iozone`, the page-cache chose to issue larger reads to the filesystem in anticipation of future reads. The default readahead size of 128 KB is large enough to reach the SSD’s full read speed of 250MB/s. This illustrates an important point: by designing `GPUstore` to fit naturally into existing storage subsystems, we enable it to work smoothly with the rest of the kernel. Thus, by simply implementing the multi-page `readpages` interface for `eCryptfs`, we enabled existing I/O optimizations in the Linux kernel to kick in, maximizing performance even though they are unaware of `GPUstore`.

Another surprising result in Figure 4.8 is that the GPU write speed exceeds the write speed of the SSD, and even its read speed, when block size increases beyond 128 KB. This happens because `eCryptfs` is, by design, “stacked” on top of another filesystem. Even though we take care to `sync` writes to `eCryptfs`, the underlying filesystem still operates asynchronously and caches the writes, returning before the actual disk operation has completed. This demonstrates another important property of `GPUstore`: it does not change the behavior of the storage stack with respect to caching, so client subsystems still get the full effect of these caches without any special effort.

We tested the throughput limits of our GPU `eCryptfs` implementation by repeating the previous experiment on a RAM disk, as shown in Figure 4.9. Our GPU-accelerated `eCryptfs` achieves more than 700 MBps when reading and 420 Mbps when writing. Compared to the CPU, which does not perform much better than it did on the SSD, this is a speed increase of nearly five times for reads and close to three times for writes. It is worth noting that Linux’s readahead mechanism not only “rounds up” read requests to 128 KB, it “rounds down” larger ones as well, preventing `eCryptfs` from reaching even higher levels of performance.

4.3.3.2 Concurrent Throughput

We also used `FileBench` to evaluate `eCryptfs` under concurrent workloads. We varied the number of concurrent writers from one to one hundred, and used the RAM-backed filesystem. Each client writes sequentially to a separate file. The effects of `GPUstore`’s *merge* operation are clearly visible in Figure 4.10: with a single client, performance is low, because we use relatively small block sizes (128 KB and 16 KB) for this test. But with ten clients, `GPUstore` is able to *merge* enough requests to get performance on par with `dm-crypt` at a 1 MB blocksize. This demonstrates that `GPUstore` is useful not only for storage systems with heavy single-threaded workloads, but also for workloads with many




Figure 4.10. eCryptfs concurrent write throughput on a RAM disk for two different block sizes.

simultaneous clients. While block size still has a significant effect on performance, GPUstore is able to amortize overheads across concurrent access streams to achieve high performance even for relatively small I/O sizes.

4.3.3.3 Random Operations

PostMark was designed for measuring operations on a large number of relatively small files, similar to a traditional mail server. We use it to ensure that introduction of GPU computing into eCryptfs has not significantly impacted the latency of such operations. In this benchmark, files are repeatedly chosen at random from a generated pool either for reading or writing. The benchmark then reports how many operations per second it was able to achieve. If the GPU version reports fewer operations per second than the CPU, this is an indication that the latency incurred by the more expensive GPU call has harmed the latency of eCryptfs.

We used 1,000 files with 5,000 transactions and varied file size from 16KB to 1MB. The benchmark was executed on *S1*.

Table 4.3 shows the results as we varied the block sizes. On the SSD (the upper portion of the table), we do see some increased latency (as evidenced by the lower number of operations per second) for small block sizes. For small block sizes, the GPU reduces operations per second by about 30%. The reason for this slowdown is that PostMark is single-threaded, so the GPU cipher service is unable to *merge* small blocks when doing encryption or decryption. This means that when block sizes are small, the overhead required to copy data into the GPU’s memory is greater than the speedup that the GPU processing can provide, and the

Table 4.3. PostMark results for **eCryptfs**. Columns marked “M” indicate mixed workloads. CS and GS are the CPU’s and GPU’s results on the SSD; CR and GR are these two on a RAM disk.

| | Block (KB) | Create (ops/s) | Create ^M (ops/s) | Read (ops/s) | Append (ops/s) | Delete (ops/s) | Delete ^M (ops/s) | Read BW (MB/s) | Write BW (MB/s) |
|----|---------------|-------------------|--------------------------------|-----------------|-------------------|-------------------|--------------------------------|-------------------|--------------------|
| CS | 256 | 250 | 151 | 153 | 150 | 244 | 152 | 91 | 128 |
| | 128 | 213 | 149 | 151 | 148 | 207 | 150 | 89 | 125 |
| | 64 | 225 | 147 | 149 | 146 | 220 | 148 | 87 | 122 |
| | 32 | 229 | 146 | 148 | 146 | 224 | 148 | 87 | 122 |
| | 16 | 225 | 151 | 153 | 150 | 220 | 152 | 90 | 126 |
| GS | 256 | 271 | 151 | 153 | 150 | 264 | 152 | 93 | 131 |
| | 128 | 238 | 156 | 158 | 155 | 232 | 157 | 92 | 129 |
| | 64 | 238 | 158 | 160 | 158 | 232 | 160 | 93 | 131 |
| | 32 | 213 | 151 | 153 | 150 | 207 | 152 | 90 | 126 |
| | 16 | 143 | 107 | 109 | 107 | 139 | 108 | 63 | 89 |
| CR | 256 | 200 | 177 | 180 | 177 | 244 | 179 | 96.15 | 134.34 |
| | 128 | 200 | 177 | 180 | 177 | 244 | 179 | 96.15 | 134.34 |
| | 64 | 200 | 177 | 180 | 177 | 244 | 179 | 96.15 | 134.34 |
| | 32 | 200 | 177 | 180 | 177 | 244 | 179 | 91.34 | 127.62 |
| | 16 | 250 | 165 | 168 | 165 | 244 | 167 | 96.15 | 134.34 |
| GR | 256 | 333 | 248 | 252 | 247 | 244 | 251 | 140.52 | 196.34 |
| | 128 | 333 | 226 | 229 | 225 | 244 | 228 | 130.49 | 182.31 |
| | 64 | 250 | 207 | 210 | 206 | 244 | 209 | 114.17 | 159.53 |
| | 32 | 200 | 177 | 180 | 177 | 244 | 179 | 96.15 | 134.34 |
| | 16 | 125 | 113 | 114 | 112 | 244 | 114 | 60.89 | 85.08 |

result is increased latency. The slowdown could be avoided enabling *hybrid mode* on the GPU service which dynamically switches to the CPU implementation for lower block sizes.

When block sizes are large, bandwidth starts to dominate latency, and we the GPU version run slightly faster. Both CPU and GPU versions, however, are limited by the speed of the SSD. When we move the test to a RAM disk, shown in the lower part of the table, the GPU significantly outperforms the CPU for sufficiently large block sizes.

4.3.4 md Data Recovery

As with encryption, the performance of our GPU-based RAID 6 recovery algorithm increases with larger block sizes, eventually reaching six times the CPU’s performance, as seen in Figure 4.11.

We measured the sequential bandwidth of a degraded RAID 6 array consisting of 32 disks in our *S2* experiment environment. The results are shown in Figure 4.12. We find that GPU accelerated RAID 6 data recovery does not achieve significant speedup unless the array is configured with a large chunk size, or strip size. Interestingly, the speedup is caused by decreasing CPU performance. We believe the decrease is caused by the limit-sized design of *md*’s I/O request memory pool, which does not efficiently handle large numbers of requests on large stripes. Because *GPUstore* merges these requests into larger ones, it




Figure 4.11. Throughput for the RAID 6 recovery algorithm with and without optimizations to avoid redundant buffers.




Figure 4.12. RAID 6 read bandwidth in degraded mode

avoids suffering from the same problem.

We also measured the degraded mode performance of a RAID 6 array in the *S1* system using 6 RAM disks. The results are shown in Figure 4.13. We find that our previous recovery experiment was limited by the speed of the hard disks, and both CPU and GPU implementations would be capable of faster performance given faster disks. With RAM disks, the CPU based recovery reaches the maximum throughput we saw in Figure 4.11, while the GPU version is still far from its own maximum in the figure.

However, with chunk sizes below 16KB, the throughputs on RAM disk arrays are actually much higher than we saw for the raw algorithm in Figure 4.11. This result demonstrates




Figure 4.13. RAID 6 read bandwidth in degraded mode on RAM-backed devices

the effectiveness of the request *merge* operation in **GPUstore**. *merge* was in use for recovery benchmark, but not the raw algorithm test, and the former therefore saw larger effective block sizes.

4.4 Summary and Future Work

GPUstore is a general-purpose framework for using GPU computing power in storage systems within the Linux kernel. By designing **GPUstore** around common storage paradigms, we have made it simple to use from existing code, and have enabled a number of optimizations that are transparent to the calling system. We modified several standard Linux subsystems to use **GPUstore**, and were able to achieve substantial improvements in performance by moving parts of the systems’ computation on to the GPU. Our benchmark results also demonstrate the effectiveness of the optimizations adopted by **GPUstore** for matching the storage subsystems requirements. The source code of **GPUstore** is released under the GPL license at <http://code.google.com/p/kgpu/>.

Because of its in-kernel target applications, **GPUstore** has to suffer from the kernel-user switching overhead and technical workarounds to implement in-kernel CUDA memory management. A promising future work will be migrating **GPUstore** onto open source GPU drivers such as Nouvea [39] and the in-kernel CUDA runtime Gdev [40]. Both of them eliminate the overhead due to kernel-user switching and resource management and provide direct GPU computing primitive access inside Linux kernel. Authors of Gdev [40] have actually done a simple, **eCryptfs** only migration of **GPUstore** to Gdev, and still got pretty large speedup over CPU implementations. We believe that with future open specification

GPUs and their open source drivers and computing libraries, **GPUstore** will be more efficient on top of them.

CHAPTER 5

Snap: FAST AND FLEXIBLE PACKET PROCESSING WITH CLICK AND GPUS

This chapter describes **Snap**, the GPU-accelerate network packet processing framework. As networks advance, the need for high-performance packet processing in the network increases for two reasons: first, networks get faster, and second, we expect more functionality from them [61, 62, 63, 64, 65, 66]. The nature of packet data naturally lends itself to parallel processing [67], and as we shall see in the survey chapter (Chapter 6), a wide variety of network functionalities are capable of parallel GPU accelerations. However, a software router is made up of more than just these heavyweight processing and lookup elements. A range of other elements are needed to build a fully functional router, including “fast path” elements such as TTL decrement, checksum recalculation, and broadcast management, and “slow path” elements such as handling of IP options, ICMP, and ARP. Building new features not present in today’s routers adds even more complexity. To take full advantage of the GPU in a packet processor, what is needed is a flexible, modular framework for building complete processing pipelines by composing GPU programs with each other and with CPU code.

We have designed and implemented **Snap** to address this need. It extends the architecture of the Click modular router [21] to support offloading parts of a packet processor onto the GPU. **Snap** enables individual elements, the building blocks of a Click processing pipeline, to be implemented as GPU code. It extends Click with “wide” ports that pass batches of packets, suitable for processing in parallel, between elements. **Snap** also provides elements that act as adapters between serial portions of the processing pipeline and parallel ones, handling the details of batching up packets, efficiently copying between main memory and the GPU, scheduling GPU execution, and directing the outputs from elements into different paths on the processing pipeline. In addition to these user-visible changes to Click, **Snap** also makes a number of “under the hood” changes to the way that Click manages memory and optimizes its packet I/O mechanisms to support multi-10 Gbit rates.

The following sections start from introducing simple Click background, and several motivating experiments we have done for **Snap** work to convince you the need of **Snap**. Then we shall detailly discuss the design and implementation of **Snap** to describe how it extends Click to efficiently use parallel GPUs to accelerate packet processing. And last we evaluate **Snap** from a variety of aspects to demonstrate its performance and also the flexibility and modularity derived from Click.

5.1 Click Background

Click is a modular software router that provides an efficient pipeline-like abstraction for packet processing on PC hardware. A packet processor is constructed by connecting small software modules called “elements” into a graph called a “configuration.” Click elements have two kinds of “ports”: input ports and output ports, and a single element may have more than one of each. A connection between two elements is made by connecting an output port of one element to an input port of another. Packets move along these connections when they are pushed or pulled: an element at the head of the pipeline can push packets downstream, or packets can be pulled from upstream by elements at the tail of the pipeline. Packets typically enter Click at a **FromDevice** element, which receives them from a physical NIC and pushes them downstream as they arrive. Unless dropped, packets leave through a **ToDevice** element, which pulls them from upstream and transmits them as fast as the outgoing NIC allows. Queues are used to buffer packets between push and pull sections of the configuration.

At the C++ source-code level, elements are written as subclasses of a base **Element** class, ports are instances of a **Port** class, and network packets, which are represented by instances of the **Packet** class, are passed one at a time between **Elements** by calling the elements’ **push()** or **pull()** methods. We run Click at user level—although Click can run directly in the kernel, with the Netmap [68] zero-copy packet I/O engine, user-level Click has a higher forwarding rate than the kernel version [68].

5.2 Motivating Experiments

Our work on **Snap** is motivated by two facts: (1) rich packet processing functionality can represent a major bottleneck in processing pipelines; and (2) by offloading that functionality onto a GPU, large performance improvements are possible, speeding up the entire pipeline. We demonstrate these facts with two motivating experiments. (Our experiment setup and methodologies are described in more detail in Section 5.5.)

Our first experiment starts with the simplest possible Click forwarder, shown at the top of Figure 5.1. It does no processing on packets. It simply forwards them from one interface to another. We then add, one at a time, elements that do IP route lookup, classification based on header bits (as is done in most SDN designs), and string matching (used in many intrusion detection systems and deep-packet-inspection firewalls). The relative throughputs of these four configurations, normalized to the throughput of the Simple Forwarder, are compared in Table 5.1. We can clearly see from this table that the addition of a even a single, common, processing task into the forwarding pipeline can significantly impact performance, cutting throughput by as much as 43%. In short, processing does represent a bottleneck, and if we can speed it up, we can improve router throughput.

Our next set of experiments compares the performance of these three processing element when run on the CPU and on the GPU. These experiments involve no packet I/O—we are simply interested in discovering whether the raw performance of the GPU algorithms offers enough of a speedup to make offloading attractive. We process packets in batches, which is necessary to get parallel speedup on the GPU. The results are shown in Figure 5.2. Two things become clear from these graphs. First, GPUs do indeed offer impressive speedups for these tasks: we see a 16x speedup for IP route lookup (559 Mpps on the GPU vs 34.7 Mpps on the CPU). Second, fairly large batches of packets are needed to achieve this speedup. These results are in line with findings from earlier studies [26, 24].

GPUs are not appropriate for every type of packet processing element. In particular, elements that require a guarantee that they see every packet in a flow in order, or that have heavy state synchronization requirements, are not well-suited to massively parallel processing. Our challenge in **Snap** is to make it possible to take advantage of GPU parallelism in a practical a way that preserves the inherent composability and flexibility of Click, including incorporation of CPU elements into the processing pipeline.

5.3 Design

We designed **Snap** with two goals in mind: enabling *fast* packet processors through GPU offloading while preserving the *flexibility* in Click that allows users to construct complex pipelines from simple elements. **Snap** is designed to offload specific elements to the GPU: parts of the pipeline can continue to be handled by existing elements that run on the CPU, with only those elements that present computational bottlenecks re-implemented on the GPU. From the perspective of a developer or a user, **Snap** appears very similar to regular Click, with the addition of a new type of “batch” element that can be implemented on




Figure 5.1. Simplified Click configurations for motivating experiments


Table 5.1. Relative throughputs of simple processing pipelines.

| Configuration | Throughput | Relative Throughput |
|------------------|------------|---------------------|
| Simple Forwarder | 30.97 Gbps | 100% |
| IP Router | 19.4 Gbps | 62.7% |
| IDS | 17.7 Gbps | 57.3% |
| SDN Forwarder | 18.8 Gbps | 60.7% |


the GPU and a set of adapter elements that move packets to and from batch elements. Internally, **Snap** makes several changes to Click in order to make this pipeline work well at high speeds. Several themes appear in our design choices. In many cases, we find that if we do “extra” work, such as making copies of only the necessary parts of a packet in main memory, or passing along packets that we know will be discarded, we can decrease the need for synchronization and reduce our use of the relatively slow PCIe bus. We also find that scheduling parts of the pipeline asynchronously works well, and fits naturally with Click’s native push/pull scheduling. In this section, we walk through the design and implementation of **Snap**, starting at a high level with the user-visible changes, and progressing through the lower level changes that stem from these high-level decisions.

5.3.1 Batched Packet Processing in Wide Click Pipeline


GPUs need parallel workloads to fully utilize their large number of cores, as discussed in Section 3.1.1. To provide and to process parallel workloads in **Snap**, we designed a batched processing mechanism to batch a large amount of packets and process them on GPUs by



(a) IP lookup algorithm (radix tree)



(b) IDSMatcher algorithm (Modified Aho-Corasick)



(c) SDN forwarder classifier algorithm

Figure 5.2. GPU vs. CPU performance on packet processing algorithms. Note that the axes are non-linear.

one time GPU kernel execution. We extended Click’s single packet processing pipeline to support multiple packets processing in a batch. We also designed special elements and efficient data structures to do the batching and to store the batched packets. The following sections will describe the details of our batched processing design in **Snap**.

5.3.1.1 Widening Click’s Pipeline

In standard Click, the connection between elements is a single packet wide: the `push()` and `pull()` methods that pass packets between elements yield one packet each time they are invoked. To efficiently use a GPU in the pipeline, we added wider versions of the `push()` and `pull()` interfaces, `bpush()` and `bpull()`. These methods exchange a new structure called a *PacketBatch*, which will be described in more detail in the following section. We also made Click’s `Port` class aware of these wider interfaces so that it can correctly pass *PacketBatches* between elements. `bpush()` and `bpull()` belong to a new base class, `BElement`, which derives from Click’s standard `Element` class.

In standard Click, to implement an element, the programmer creates a new class derived from `Element` and overloads the `push()` and `pull()` methods. This is still supported in **Snap**; in fact, most of our pipelines contain many unmodified elements from the standard Click distribution, which we refer to as “serial” elements. To implement a parallel element in **Snap** the programmer simply derives it from `BElement` and overrides the `bpush()` and `bpull()` methods.

A GPU-based parallel element is comprised of two parts: a GPU side, which consists of GPU kernel code, and a CPU side, which receives *PacketBatches* from upstream elements and sends commands to the GPU to invoke the GPU kernel. **Snap** provides a `GPURuntime` object to help Click code interact with the GPU, which is programmed and controlled using NVIDIA CUDA [6]. GPU-based elements interact with `GPURuntime` to request GPU resources such as memory. The GPU kernel is written in CUDA’s variant of C or C++, and is wrapped in an external library that is linked with the element sources when compiling **Snap**. Typically, each packet is processed by its own thread on the GPU.

5.3.1.2 Batching Packets

The `BElement` class leaves us with a design question: how should we collect packets to form *PacketBatches*, and how should we manage copies of *PacketBatches* between host and GPU memory? Our answer to this question takes its cue from the functioning of Click’s `Queue` elements. Parts of a Click configuration operate in a push mode, with packets arriving from a source NIC; other parts of the configuration operate in pull mode, with packets being

pulled along towards output NICs. At some point in the configuration, an adapter must be provided between these two modes of operation. The family of **Queue** elements plays this role. In practice, the way a packet is processed in Click is that it is pushed from the source NIC through a series of elements until it reaches a **Queue**, at which point it is deposited there and Click returns to the input NIC to process the next packet. On the output side of the **Queue**, the packet is dequeued and processed until it reaches the output NIC.

In an analogous manner, we have created a new element, **Batcher**, which collects packets one at a time from a sequential **Element** on one side and pushes them in batches to a **BElement** on the other side. A **Debatcher** element performs the inverse function. **Batcher** can be configured to produce *PacketBatches* with specified *batch-size* packets, or fewer if a specified *timeout* period passed. Implementing this functionality as a new element, rather than changing Click’s infrastructure code, has three advantages. First, it minimizes the changes within Click itself. Second, it makes transitions between CPU and GPU code explicit; since there are overheads associated with packet batching, it is undesirable for it to be completely invisible to the user. Third, and most important, it means that batching and offloading are fully under the control of the creator of the **Snap** configuration—while we provide carefully-tuned implementations of **Batcher** and **Debatcher** elements, it is possible to provide alternate implementations designed for specific uses (such as non-GPU **BElements**) without modifying **Snap**.

5.3.1.3 Slicing-based *PacketBatch*

The *PacketBatch* data structure to represent a batch of packets has been carefully designed specially for offloading computation to GPUs. A *PacketBatch* is associated not only with a collection of packets (represented by Click’s **Packet** objects), but also with allocations of host and GPU device memory. Large consecutive buffers are used in host and GPU memory in order to enable efficient DMA transfers, minimizing the overhead of setting up multiple transfers across the PCIe bus.

The large buffers of a *PacketBatch* are split into small buffers, which contain the slices of packets (such as the headers) that are needed by the **BElement**(s). **Snap** does such slicing and extra host memory buffer because of the special data usage pattern of most packet processing tasks. As we have discussed in Section 3.5.3 and Section 3.5.4, network packet processing is a representative task that only needs small pieces of a given data unit. Such data usage pattern needs a trade-off of the relatively large host memory bandwidth and the relatively small host-device memory bandwidth in order to effectively reduce the total memory copy overhead.




Figure 5.3. The *PacketBatch* structure




Figure 5.4. A Slicing example

We designed Slicing (illustrated in Figure 5.4) for **Snap** to deal with this problem. Slicing allows GPU processing elements to specify the regions of packet data that they operate on, called Regions of Interest (ROIs). An ROI is a consecutive range in the packet data buffer, and a GPU processing element can have multiple ROIs spread throughout the


packet. **Batcher** accepts ROI requests from its downstream **BElements** and takes their union to determine which parts of the packet must be copied to the GPU. It allocates only enough host and device memory to hold these ROIs, and during batching, **Batcher** only copies data in these regions into a packet’s host memory buffer. This reduces both the memory requirements for *PacketBatches* and the overhead associated with copying them to and from the GPU. During debatching, ROIs are selectively copied back into the **Click Packet** structure. Slicing is one reason that we chose not use a zero-copy approach for *PacketBatches*.

Batcher contains optimizations to avoid redundant copies in the case of ROIs that overlap and to combine `memcpy()` calls for consecutive ROIs to reduce function call overhead. For element developers’ convenience, we have provided helper APIs for **BElements** that allow the element to address packet data relative to its ROIs—the true offsets within the *PacketBatch* are computed transparently.

One problem associated with ROIs is that it may be difficult to describe the exact range in numeric values. For example, a **Classifier** element may need TCP port numbers, but the offset of this data within a packet is not constant due to the presence of IP option headers. To support this case, **Batcher** provides some special values to indicate variable offsets, such as the beginning of the TCP header or the end of the IP header. This also enables some special ROIs, such as ROIs that request the entire IP header including all IP options, or ROIs that cover the payload of the packet.

5.3.1.4 Flexible Memory Copy Control

GPU code requires both input data and output results to reside in GPU memory, making it necessary to copy packets back and forth between main memory and the GPU across the PCIe bus. **Snap** factors this task out of the **BElements** that contain processing code: a **HostToDeviceMemcpy** element (provided as part of **Snap**) is placed between the **Batcher** and the first element that runs on the GPU. An analogous **DeviceToHostMemcpy** element is placed before the **Debatcher**. Multiple GPU elements can be placed between a **HostToDeviceMemcpy/DeviceToHostMemcpy** pair, allowing the output ports of one to feed into the input ports of another without incurring a copy back to host memory. This design reduces the host-device packet copy times, reducing the overall memory copy overhead. These memory copy elements, along with the batching and debatching elements, can be seen in Figure 5.5.



(a) Batching packets for copying to the GPU (b) In-order completion management and debatching

Figure 5.5. Batching and debatching elements. Serial interfaces are shown as simple arrows, wide interfaces as double arrows.

5.3.1.5 Avoid Packet Reordering

Previous work on parallel packet processing often causes reordering among packets due to techniques such as load balancing and parallel dispatch across multiple cores [69, 67]. This can hurt TCP or streaming performance [70]. **Snap**, however, does not suffer from this problem: it waits for all threads in a GPU **BElement** to complete before passing the batch to the next element, so **Snap** does not reorder packets within a *PacketBatch*. Because asynchronous scheduling on the GPU (discussed in more detail in Section 5.3.3) may cause reordering of the *PacketBatches* themselves, we add a **GPUCompletionQueue** element between the **DeviceToHostMemcpy** and **Debatcher** elements. **GPUCompletionQueue** keeps a FIFO queue of outstanding GPU operations, and does not release a *PacketBatch* downstream to the **Debatcher** until all previous *PacketBatches* have been released, keeping them in order. Because **GPUCompletionQueue** is simply an **Element** in the configuration graph, a configuration that is not concerned about reordering could simply provide an alternate element that releases batches as soon as they are ready.

5.3.1.6 Putting It Together: A GPU Pipeline

Figure 5.6 shows how all of these elements operate together. Packets enter one at a time from the top. They could come directly from the NIC or from sequential Click **Elements**. A **Batcher** collects them into a *PacketBatch* and sends them along to a **HostToDeviceMemcpy** element, which initiates a copy to the GPU. This, and all other interaction with the GPU, is done through a **GPURuntime** element provided by **Snap**. The GPU element launches the

GPU code kernel and yields to the `DeviceToHostMemcpy` element, which sets up a copy of the GPU element’s results back to host memory; this will occur automatically when the GPU code completes execution. The `GPUCompletionQueue` element waits for the preceding GPU operations to complete, and passes *PacketBatches*, in FIFO order, to a `Debatcher`. The `Debatcher` then passes packets one by one into an output serial `Element`. This picture tells only the part of the story that is visible at the level of a `Snap` configuration. There are several technical challenges remaining that have been dealt with in `Snap` at lower level to make these pieces work efficiently, which will be discussed in the following sections.

5.3.2 Handling Packet Processing Divergence

`Snap` faces a problem not encountered by other GPU processing frameworks [24, 23, 26], namely the fact that packets in Click do not all follow the same path through the `Element` graph. `Elements` may have multiple output `Ports`, reflecting the fact that different packets get routed to different destination NICs, or that different sets of processing elements may be applied depending on decisions made by earlier elements. This means that packets that have been grouped together into a *PacketBatch* may be split up while on the GPU or after being returned to host memory. We encounter two main classes of packet divergence. In *routing or classification divergence*, the number of packets exiting on each port is relatively balanced; with *exception-path divergence* most packets remain on a “fast path” and relatively few are diverted for special processing. Packet divergence may also appear in two places: before the packets are sent to the GPU, or on the GPU, as a result of the decisions made by `BElements`.

Figure 5.7(a) shows an example of exception-path divergence before reaching the GPU: packets may be dropped after the TTL decrement if their TTLs reach zero. Figure 5.7(b) shows an example of routing divergence on the GPU. In this example, different IDS elements (likely applying different sets of rules) are used to process a packet depending on its next hop, as determined by IP routing lookup.

Divergence before reaching the GPU is another reason that we do not attempt to implement zero-copy in the *PacketBatch* structure. The effect of divergence early in the pipeline is memory fragmentation, giving us regions of memory in which only some packets need to be copied to the GPU. This problem is particularly pronounced in the case of routing/classification divergence. Copying all packets, even unnecessary ones, to the GPU would waste time and scarce PCIe bandwidth. Alternately, we could set up a number of small DMA transfers, covering only the necessary packets, but this results in high DMA setup overhead. Instead, we use the relatively plentiful memory bandwidth in host RAM

to copy the necessary packets to one continuous region, which can be sent to the GPU with a single DMA transfer.

For divergence that occurs on the GPU, our experiments show that the overheads associated with splitting up batches and copying them into separate, smaller *PacketBatches* are prohibitive, especially in the case of exception-path divergence. Assembling output *PacketBatches* from selected input packets is also not concurrency-friendly: determining each packet’s place in the output buffer requires knowledge of how many packets precede it, which in turn requires global serialization or synchronization. Having discussed in Section 3.1.2.1, **Snap** should embrace the warp control flow divergence on GPU, but with following predicated execution to minimize its performance impact.

Therefore, **Snap** attaches a set of *predicate bits* to each packet in a *PacketBatch*—these bits are used to indicate which downstream **BElement**(s) should process the packet. Predicates are stored as annotations in Click’s **Packet** structure. The thread processing each packet is responsible for checking and setting its own packet’s predicate; this removes the need for coordination between threads in preparing the output. With predicate bits, a thread simply decides whether to process a given packet or not by checking the packet’s predicate at the very beginning of the GPU kernel, hence it just introduces two branches of the control flow, and one of them is totally “empty”, which follow the guidelines presented in Section 3.1.2.1. Because they are only used to mark divergence that occurs on the GPU, not before it, we save PCIe bus bandwidth by not copying predicate bits *to* GPU memory; we do, however, copy them *from* the GPU once a chain of **BElements** is finished, since they are needed to determine the packets’ next destination **Elements** on the CPU.

Figure 5.7(c) shows how this predicated execution works. The **GPUElement-1** element marks packets with either Predicate 0 or Predicate 1, depending on which downstream element should process them. The packets remain together in a single *PacketBatch* as they move through the element graph, but **GPUElement-2** only processes packets with Predicate 0 set and **GPUElement-3** only processes those with Predicate 1. Eventually, once they have left the GPU, the packets encounter a **Dispatcher** element, which sends them to different downstream destinations depending on their predicate bits. This arrangement can be extended to any number of predicate bits to build arbitrarily complicated paths on the GPU.

We have experimented with two strategies for using predicate bits: scanning all packets’ predicates, and only launching GPU threads for the appropriate packets; and launching threads for all packets, and returning immediately from threads that find that their packet

has the wrong predicate. We found it more efficient to launch threads for all packets: scanning packets for the correct predicates in order to count the number of threads adds to the startup overhead for the **BElement**. The savings in execution time that come from launching fewer threads are typically smaller than the overhead of scanning for the correct threads to launch, and it is faster to simply launch all threads. Because we run many threads per core, the threads that exit early do not necessarily waste a core. Further experiment to evaluate these two strategies on **Snap** applications is discussed in Section 5.5.4.

5.3.3 Asynchronous Processing on GPUs

The host uses two main operations for controlling the GPU: initiating copies between host and device memory; and launching kernels on the GPU. Both can be done asynchronously, with the CPU receiving an interrupt when the copy or code execution completes. Multiple GPU operations can be in flight at once. Based on this, most GPU elements can work asynchronously on the CPU side: when a GPU element is scheduled by **Snap**, it issues appropriate commands to the GPU and schedules its downstream element by passing the *PacketBatch* immediately, without waiting for completion on the GPU. As a result, a push path or a pull path with GPU elements can keep pushing or pulling *PacketBatches* without blocking on the GPU. This allows us to achieve very low latency at the beginning of the path, which is critical when packet rates are high—for example when receiving minimum-sized packets from a 10Gbps interface. The Click **FromDevice** element that receives packets from the NIC must disable interrupts while it pushes packets downstream to first **Queue** or **Batcher** element that they encounter; it is thus critical to have this path run as quickly as possible to avoid lost packets.

Snap uses CUDA stream to achieve overlapped asynchronous kernel execution and asynchronous memory copy. Each stream has a queue of operations which is run in FIFO order. Operations from different streams run concurrently and may complete in any order. We associate each *PacketBatch* with a unique stream. When the *PacketBatch* is first passed to a GPU element (typically, **HostToDeviceMemcpy**), it gets a stream assignment. Each subsequent **BElement** along the path asynchronously queues execution of its operation within the stream and passes control to the next **BElement** immediately, without waiting for the GPU. This sequence of events continues until control reaches a **GPUCompletionQueue**, which is a *push-to-pull* element, much like a **Queue**. When a *PacketBatch* is pushed into the **GPUCompletionQueue**, it simply adds the batch’s stream to its FIFO queue and returns. When the **GPUCompletionQueue**’s `bpull()` method is called, usually by a downstream **Debatcher**, it checks the status of the stream at the head of the FIFO by calling a

non-blocking CUDA stream checking function. If the stream has finished, `bpull()` returns the *PacketBatch*; if not, it indicates to the caller that it has no packets ready.

5.3.4 Improving Packet I/O

Click includes existing support for integration with Netmap [68] for fast, zero-copy packet I/O from userspace. We found, however, that Click’s design for this integration did not perform well enough to handle the packet rates enabled by **Snap**.

Netmap uses the multi-queue (RSS/MQ) support in recent NICs to enable efficient dispatch of packets to multiple threads or CPU cores. The queues maintained by Netmap in RAM are mapped to hardware queues maintained by the NIC; the NICs we use for our prototype fix each receive and transmit queue at 512 packets. When a packet arrives on the NIC, a free slot is found in a queue, and the NIC places the packet in a buffer pointed to by the queue slot. When the packet is passed to Click, the buffer, and thus the queue slot, remains unavailable until Click is either finishes with the packet (by transmitting it on another port or discarding it) or copies it out into another buffer. Since packets may take quite some time to be processed, they tie up these scarce queue slots, which can lead to drops. This problem is exacerbated in **Snap**, which needs to wait for suitably large batches of packets to arrive before sending them to the GPU. Click’s solution is to copy packets out when it notices the Netmap queues getting full. It uses a single global memory pool for all threads, leading to concurrency problems. We found that at the high packet rates supported by **Snap**, these copies occurred for nearly every packet, adding up to high overhead incurred for memory allocation and copying. Note that unlike our *PacketBatch* structure, which copies only regions of interest, the copies discussed here must copy the entire packet.

Unmodified Netmap gives the userspace application a number of packet buffers equal to the number of slots in the hardware queues; while kernel code can request more buffers, userspace code cannot. We added a simple system call that enables applications to request more packet buffers from Netmap. Though the size of the queues themselves remain fixed, **Snap** can now manipulate the queue slots to point to these additional buffers, allowing it to maintain a large number of in-process packets without resorting to copying. **Snap** maintains a pool of available packet buffers—when it receives a packet from the NIC, it changes the queue to point to a free packet buffer, and packet buffers are added back to the free pool when the packets they hold are transmitted or dropped. This eliminates packet buffer copying and overhead from complex memory allocation (`kmalloc()`), and we use per-thread packet buffer pools to avoid overhead from locking.

We also modified Click to pin packet I/O threads to specific cores. This is a well-known

technique that improves cache behavior and interrupt routing when used with multi-queue NICs. Combined, these two optimizations give **Snap** the ability to handle up to 2.4 times as many packets per second as Click’s I/O code—this improvement was critical for small packet sizes, where the unmodified packet I/O path was unable to pull enough packets from the NIC to keep the processing elements busy.

5.4 Implementation

Snap has been implemented on top of the 9200a74 commit in the Click source repository [71]. We used the Netmap release from August 13, 2012 and Linux kernel 3.2.16. **Snap** makes 2,179 lines of changes to Click itself, plus includes 4,636 lines of code for new elements and a 3,815 line library for interacting with the GPU. We modified only 180 lines of code in Netmap. The source for **Snap**, including our modifications to Netmap, can be downloaded from <https://github.com/wbsun/snap>

The new packet processing elements we have implemented are for three kinds of packet processing tasks. Each has a CPU and a GPU version:

- **GPU IPLookup**: this GPU-based IP lookup element implements Click’s `IPRouteTable` class using a radix tree. Its CPU counterpart is Click’s `RadixIPLookup` element. For evaluation, we used a routing table dump from routeview.org [72] that has 167,000 entries.
- **SDNClassifier**, **GPUSDNClassifier**: these two elements classify packets using seven fields from the Ethernet, IP, and UDP or TCP headers. Each entry assigns an action by forwarding the packet out of a specific outbound `Port` on the element. This is roughly analogous to the flow space matching used by many SDN forwarding schemes. **SDNClassifier** is the CPU version. The classification rule set is ClassBench [73]’s ACL1_10K filter set. We randomly assigned an action number to each rule.
- **IDSMatcher**, **GPUIDSMatcher**: these two elements implement the Aho-Corasick [74] string matching on packet payloads. The Aho-Corasick algorithm can match multiple patterns simultaneously by scanning the entire packet payload once. We used Snort’s [75] rules for Mysql, Apache, Webapps, and PHP to simulate a set of rules for a Web application server.

We combine these elements to build three kinds of **Snap** configurations, each of which has both a GPU and a CPU version:

- **SDN Forwarder**: This configuration includes only the **SDNClassifier** or its GPU counterpart. It simulates an SDN-like switch.

- **DPI Router:** This configuration includes an IP lookup element (`RadixIPLookup` or `GPU IPLookup`) and a string matching element (`IDSMatcher` or `GPU IDSMatcher`) as the major processing elements. The intent is to simulate a router with a simple deep packet inspection firewall.
- **IDS Router:** This configurations includes all three elements (IP lookup, IDS matcher, and SDN classifier) to simulate a more sophisticated router with complicated forwarding rules and intrusion detection.

5.5 Evaluation

All experiments were performed on a PC with an Intel Core i7 930 quad-core CPU running at 3.2 GHz. It had 6 x 2 GB of DDR3-1600 triple channel memory having 38.4 GB/s of bandwidth, and an ASUS P6X58D-E motherboard with an Intel X58 chipset having 25.8 GB/s QPI in each direction for PCIe 2.0 devices. The GPU used was an NVIDIA TESLA C2070, which has 448 cores running at 1.15 GHz and 6 GB of GDDR5 RAM. The NICs are two Intel 82599EB dual-port 10 Gbps cards, for a total of four 10 Gbps ports. The GPU is connected via a 16-lane PCIe v2.0 slot and the NICs are each connected to 8-lane slots. All equipment used for evaluation is publicly available as part of the Emulab testbed [76]. Packets were generated at full line rate using a modified version of the packet generator that comes with the Netmap distribution, using a separate set of hosts in Emulab. Packet sizes reported are the Ethernet frame payload. When calculating throughput in Gbps, we add in the Ethernet preamble, SFD, CRC, and interframe gap, so that a reported rate of 10 Gbps represents 100% utilization on a 10 Gbps interface. Forwarding tables were designed such that all packets were forwarded back out the interface they arrived on. This ensured that all outgoing traffic was perfectly balanced so that any drops we observed were due to effects within the **Snap** host, rather than congestion on unbalanced outbound links.

5.5.1 Packet I/O

Our first set of experiments are simple micro-benchmarks that evaluate the packet I/O optimizations described in Section 5.3.4. We measured the forwarding rate for minimum-sized (64 byte) packets using Click’s Netmap packet I/O engine and **Snap**’s improvements to that engine. These experiments use the simplest possible forwarder, which simply passes packets between interfaces with no additional processing. We test both a one-path arrangement, which passes packets from a single input NIC to a single output, and a four-path arrangement that uses all four NICs in our test machine. Click’s existing Netmap support is not thread-safe, allowing only one packet I/O thread to be run. We added

multi-threading support to standard Click’s Netmap code, and also report performance for four threads, one per NIC. **Snap** adds support for multiple threads per NIC, each using a different MQ/RSS queue, so we use sixteen threads for the **Snap** configuration.

The performance numbers are found in Table 5.2. **Snap**’s improvements to the I/O engine introduce a 1.89x speedup for single path forwarding and 2.38x speedup for four-path forwarding. One interesting result is that **Snap**’s four-path performance is not quite four times that of its single-path performance. This suggests that there may be room to improve the forwarding performance of **Snap** using more cores; our test CPU has four physical cores and hyper-threading, meaning that there are two I/O threads mapped to each hyper-threaded core. A recent (at the time of writing this thesis) evaluation of **Snap**’s forwarding performance confirms this guess. We use a recent six-core high-end CPU: Intel Core i7-3930K, and six 10Gb ports to do the basic forwarding, and get 49.62 Gbps forwarding rate, which is a little bit higher than the quad-core machine when considering the per-port rate. That makes sense because the six-core machine still assigns two I/O threads to a single hyper-threaded core. When we use only four 10GB ports on the six-core machine so that in theory each two I/O threads can get one-third more core, the forwarding rate approximately reached 40 Gbps line rate.

5.5.2 Applications

Using the implemented applications, we compared the performance of four configurations: standard Click; **Snap** with only CPU Elements; **Snap** with GPU elements, but with packet slicing disabled; and **Snap** with all optimizations enabled. We experimented with a variety of packet sizes. Each experiment lasted at least one minute, and the numbers reported are the average of three runs. The results are shown in Figure 5.8.

The results show that **Snap** gets significant performance improvements over Click, particularly for small packet sizes. A significant fraction of this speedup comes from our I/O optimizations, which can be seen by comparing the bars for “Click” and “Snap-CPU:” 63% of the 3.1x speedup seen by the SDN forwarder on 64 byte packets comes from this source. Another jump comes from moving the processing-heavy elements to the GPU, with another modest increase with the addition of packet slicing. **Snap** is able to drive all four NICs at full line rate for all but the smallest packets: at and above 128 bytes, it gets full line rate for all configurations (with the exception of the IDS Router, which gets 39.6 GBps at 128 bytes). **Snap** is limited by the availability of PCIe slots in our test machine, which has 32 PCIe lanes: 16 are used by the GPU, and each dual-port NIC uses 8 lanes, meaning that we cannot add any more NICs. The results strongly suggest that **Snap** would be capable

Table 5.2. Base Forwarding Performance of **Snap** and Click

| Configuration | Throughput | |
|---------------------------|------------|-----------|
| Click 1 Path | 4.55 Gbps | 6.5 Mpps |
| Click 4 Paths (1 thread) | 8.28 Gbps | 11.8 Mpps |
| Click 4 Paths (4 threads) | 13.02 Gbps | 18.5 Mpps |
| Snap 1 Path | 8.59 Gbps | 12.2 Mpps |
| Snap 4 Paths | 30.97 Gbps | 44.0 Mpps |

of higher bandwidth on a machine with more or faster PCIe lanes, but such hardware was not available for our tests. We thus leave exploration of **Snap**’s full limits to future work and the availability of suitable hardware, but we are optimistic that it will be capable of exceeding 40 Gbps for large packets. Conversely, this result also means that there is headroom available to do more processing per packet than is performed by our example applications.

For minimum sized packets, **Snap** reaches 29.9 Gbps (75% of the full line rate) for the SDN forwarder; the primary cause of this limitation appears to be due to packet I/O, as the throughput seen in this experiment is very close to the trivial forwarder from Section 5.5.1. The other two GPU applications reach almost the same performance: 29.2 Gbps for the DPI router and 28.0 Gbps for the IDS Router. This matches our intuition, because with all of the complex processing done on the GPU, the CPU only needs to perform simple operations and can spend most of its time on packet I/O. When we use the CPU elements, both packet I/O and the processing algorithms need the CPU, and all three applications slow down significantly: the IDS Router gets 14.73 Gbps, a slowdown of 52% from the trivial forwarder. With the DPI and IDS Router configurations, standard Click is unable to reach much more than 20 Gbps, even for large packet sizes.

The ROI-based slicing mechanism makes a modest improvement in forwarding throughput. For example, the SDN forwarder sees a 13.7% increase in throughput for 64 byte packets. Slicing enables **Snap** to reach nearly the full rate supported by the packet I/O engine for small packets. At larger packet sizes, the improvement disappears because we have reached full line rate on the NICs.

5.5.3 Latency and Reordering

The most obvious drawback of batched processing is an increase in latency, since packets arriving at the beginning of a batch must wait for the batch to fill. To find out how much latency the batching mechanism adds to **Snap**, we measured round-trip time for 64-byte packets using both a CPU and GPU configuration. For the CPU-only configuration, we

saw a mean latency of $57.5\mu s$ (min: $31.4\mu s$, max: $320\mu s$, σ : $25.7\mu s$). For the GPU-based configuration with batched processing (*batch-size* 1024), the mean latency was $508\mu s$ (min: $292\mu s$, max: $606\mu s$, σ : $53.0\mu s$); this represents an increase of less than one quarter of a millisecond in each direction. Reducing the *batch-size* from 1024 to 512, the latency reduces to $380.4\mu s$ on average, but throughput also drops from 28 Gbps to 24 Gbps. The additional latency added by batching for GPU elements is likely to be noticeable, but tolerable, for many LAN applications. On WAN links, this delay will be negligible compared to propagation delay. As part of this experiment, we also checked for packet reordering. We define the RSS dispatching rules of our router NICs to send packets in the same UDP/TCP flow into the same NIC transmit queue, and in the **Snap** configuration, we connected each **FromDevice** element to a **ToDevice** with the same transmit and receive queue IDs. With these settings, we found no reordering in the packet stream.

5.5.4 Packet Processing Divergence

To evaluate whether our design for handling divergent paths is effective, we built an IDS configuration that connects an **GPUSDNClassifier** element with two **GPUIDSMatcher** elements. The classifier marks each packet with a predicate indicating which of the two IDS elements is to process it; this simulates a scenario in which packets are to be handled by different IDSes depending on some property such as source, destination, or port number. In both this configuration and the IDS Router configuration from our earlier experiments, each packet is processed by one IDS element; the difference is that in the diverging configuration, there are two IDS elements, each of which processes half of the packets. Thus, we can expect that, if the overhead of our divergence handling strategy is low, the configuration with two **GPUIDSMatchers** should achieve similar throughput to the configuration with a single one. We evaluated this diverging configuration with different packet sizes and measured the throughput, which is shown in Figure 5.9. The performance under divergence is very similar to the IDS Router result shown in Figure 5.8(c). It is only slightly slower at small packet sizes: the diverging configuration achieves 26.8 Gbps versus the IDS Router’s 28.0 Gbps for 64 byte packets, 39.4 Gbps vs. 39.6 Gbps for 128 byte packets, and 39.9 Gbps vs. 40.0 Gbps for 256 bytes packets. At and above 512 byte packets, both achieve a full 40.0 Gbps. We conclude that the launch of extra GPU threads that have no work to do causes a slight slowdown, but the effects are minimal.

5.5.5 Flexibility and Modularity

Finally, we demonstrate that **Snap** can be used to build not only highly specialized forwarders, but also a complete standards compliant IP router. This task is simple, because such configurations already exist for Click. Specifically, we base our IP router off of the configuration shown in Figure 8 of the Click paper [21], which includes support for error checking of headers, fragmentation, ICMP redirects, TTL expiration, and ARP. We replace the `LookupIPRoute` element with our `GPUIPLookup` element (and the accompanying `Batcher`, etc.), and add an IDS element to both the CPU and GPU configurations.

Due to the complexity of this router, we do not attempt to illustrate the entire **Snap** configuration here. Instead, we illustrate the major changes that we made to the standard Click router configuration in Figure 5.10. The left part of the figure shows our GPU processing path, and the right part is the original CPU route lookup path plus an `IDSMatcher` and its auxiliary alert element. This figure also shows a strategy for handling divergence on the GPU: the `GPUIDSMatcher` sets predicates on packets depending on whether they should raise an alert, then pushes entire the *PacketBatch* downstream. The `GPUIPLookup` is assigned a *CHK_ANNO* argument, which is the predicate controlling processing of each packet. `GPUIPLookup` thus ignores packets flagged for alerts by the IDS, and divergence on the actual element graph is delayed until after the `Debatcher`, using a `Dispatcher` element.

The performance of the CPU-based and GPU-based full router configurations are shown in Figure 5.11. This fully-functional router with built-in IDS is able to achieve 2/3 of the performance of a trivial forwarder for minimum-sized packets, and almost full line rate (38.8 Gbps) for 512-byte and larger packets. This demonstrates the feasibility of composing complex graphs of CPU and GPU code, and shows that existing CPU Click elements can be easily used in **Snap** configurations without modification. The bottleneck in performance appears to be the large number of CPU elements in this configuration—there are 15 types of elements, some of which are duplicated sixteen times, once for each thread. As future work, we believe that the throughput can be significantly improved by moving some of these to the GPU and applying the techniques from Kim et al. [77] to optimize the remaining CPU portions of the configuration.

5.6 Summary and Future Work

Snap expands Click’s composable element structure, adding support for batch processing and offloading of computation. At small packet sizes (128 bytes), **Snap** increases the performance of a combined IP router, SDN forwarder, and IDS on commodity hardware from 10.6 Gbps to 39.6 Gbps. This performance increase comes primarily from two sources:

an improved packet I/O engine for Click that takes advantage of multi-queue NICs, and moving computationally expensive processing tasks to the GPU. A trivial forwarder created with **Snap** can forward at a rate of 44.0 Mpps, while the complex SDN/IDS router reaches 90% of this rate (39.8 Mpps). These results suggest that there is likely potential for elements that are even more computationally complex than the ones we investigated, pointing to future work in complex packet processing. The fact that we are able to saturate all NICs in our test platform with such small packets suggests that it will be possible to reach even higher throughputs when PCIe 3.0 devices are available for testing, allowing us to double the number of NICs on a bus.

The elements and techniques proposed and implemented in **Snap** are not GPU-specific only. While some of the new **Elements** implemented for **Snap**, such as **HostToDeviceMemcpy** and **GPUCompletionQueue**, are GPU-specific, the extensions we made to the Click architecture should be applicable to other parallel offload engines (such as network processors and programmable NICs) as well.

Besides possible throughput improvement with techniques from Kim et al. [77], recent CUDA releases and GPU hardware advancement have provided useful techniques for **Snap** to do fast packet processing. The “dynamic parallelism” technique [78] on latest CUDA GPUs allows a kernel thread on GPU to invoke a new kernel launching, exactly same to what the host side code can do. This indicates an alternative approach to implement predicated execution: packet pointers can be grouped on GPU by their predicates, and then only necessary threads are launched for a following GPU element’s kernel to process the grouped packets only. One severe issue is that such solution needs on GPU synchronization, which is an expensive operation, and hence impractical on GPUs used for evaluating **Snap**. However, dynamic parallelism also makes it possible to only launch a one-thread scheduling kernel at the beginning of the first GPU element in a batched processing path. Such scheduling kernel is responsible to launch the kernels for the following GPU elements according to the configuration and predicates as explained before. In that case, on-GPU synchronization won’t be a problem any more because the scheduling and kernel launching only happen in the single-threaded scheduling kernel.

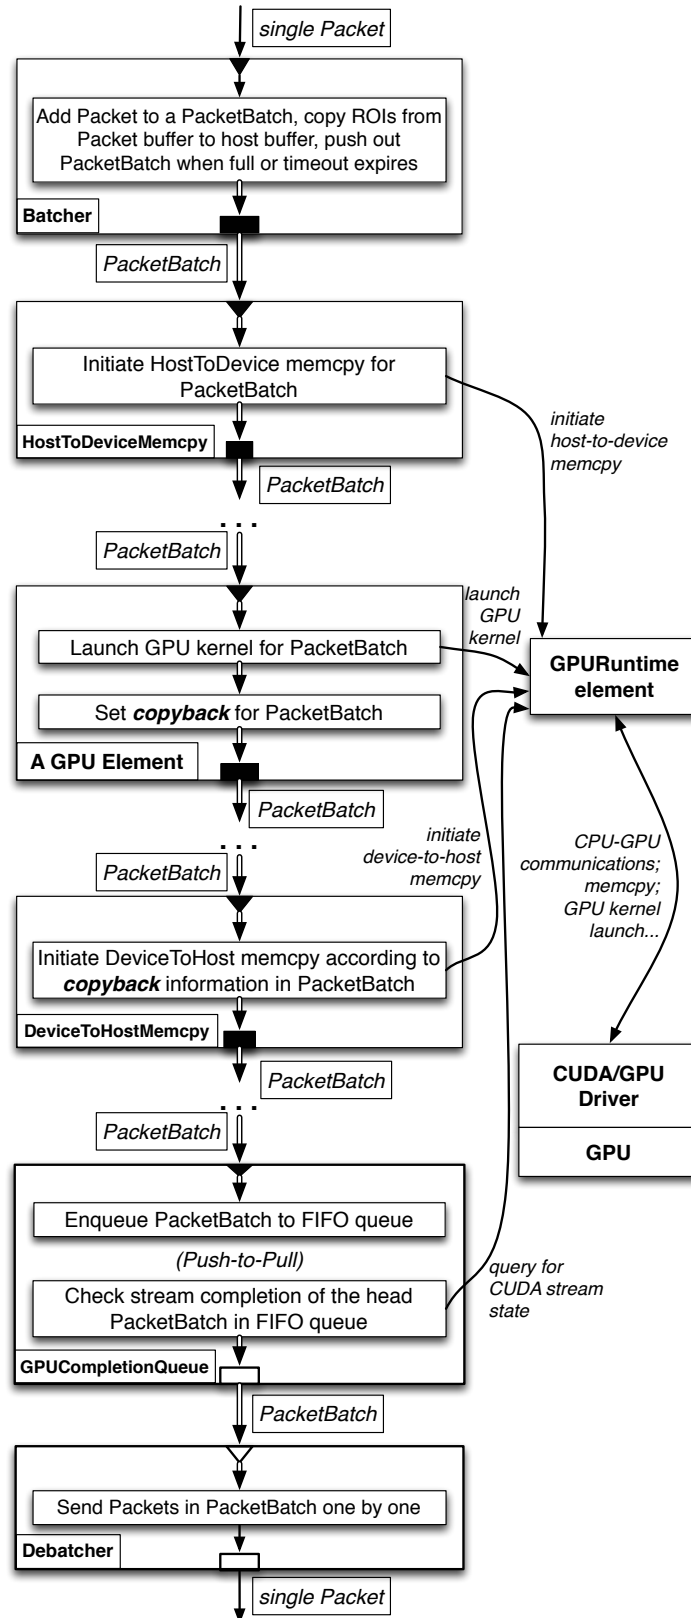




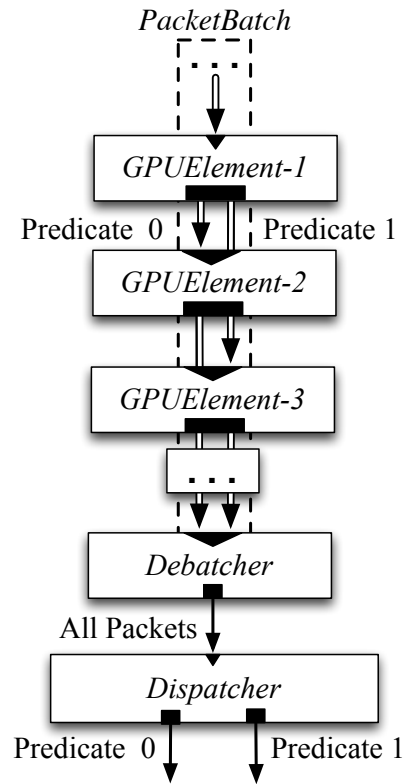
Figure 5.6. A batched processing path with one GPU element.



(a) Exception path divergence before the GPU




(b) Routing divergence on the GPU




(c) Predicated execution on the GPU


Figure 5.7. Handling divergent packet paths



(a) SDN Forwarder



(b) DPI Router



(c) IDS Router

Figure 5.8. Performance of Click and Snap with three different applications.




Figure 5.9. Forwarding performance when using a `GPUSDNClassifier` that diverges to two `GPUIDSMatcher` elements.

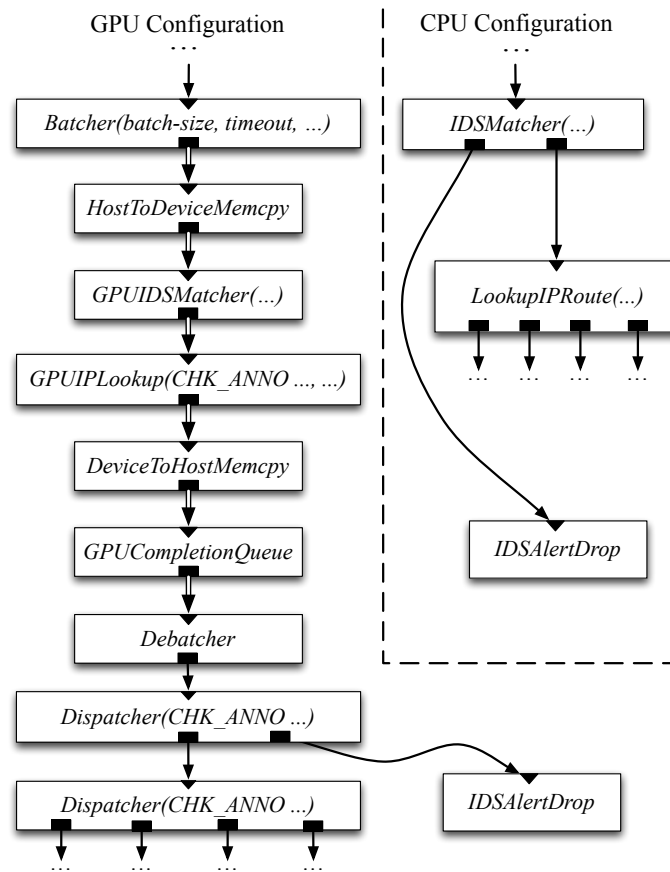


Figure 5.10. Major changes to the standard Click router to implement a GPU-based IDS router.




Figure 5.11. Fully functional IP router + IDS performance

CHAPTER 6

A LITERATURE SURVEY OF SYSTEM-LEVEL GPU COMPUTING RESEARCH

The motivations of this survey are two-fold:

1. To survey the existing work of different system-level GPU computing to demonstrate the wide applications of the GPU in system areas.
2. To analyze techniques used by the related work for efficient GPU integration in system software, comparing them with our techniques proposed in this thesis for pros and cons, as a reference to guide later system-level GPU research.

This survey tries to cover major system-level work that targets the GPU. Most of the papers surveyed are related to various system-level applications with GPU acceleration. However, some work covered in this survey is on system-level support for GPU computing, such as virtualization, resource management, and migration.

This survey has two major sections. The first section is focused on applications of GPUs in system-level software. It discusses the computationally expensive operations or algorithms in different system areas which GPUs are expected to accelerate, with representative work in each area. Similar to my KGPU white paper [79], not only existing work but also potential applications are discussed. For each application area, the attacking targets, which are computationally expensive operations or algorithms, are listed first, then the related work is described. This section serves as a traditional related work survey for this thesis. The second section is focused on techniques proposed and used by applications in the first section to improve the performance of or ease the GPU computing integration into existing system software. Since Chapter 3 has described my own techniques for system-level GPU computing, this section will also compare them with the surveyed techniques. It is not surprising to find a large portion of common techniques during the comparison, which can further confirm the feasibility and effectiveness of techniques proposed in this thesis. Besides those two major sections, there is also a small section at the end of this chapter,

which discusses the potential future GPU development and useful techniques to solve current problems.

6.1 Applications

The system-level work categories listed in the following sections are not orthogonal to each other. Many applications actually belong to two or more categories. For example SSLShader [25] is both networking and security work. Some categories such as networking and storage have much more existing work surveyed due to comprehensive studies done by the research community. Some have just attracted a few efforts, such as program analysis and data compression. However, that doesn't mean they are not important at all. On the contrary, less existing work often means that area has not been thoroughly studied and hence may contain more potential opportunities for future research.

6.1.1 Networking

Currently, the main computing tasks in network systems are the operations and algorithms used to process network packets, providing packet forwarding, security, and data integrity functionalities. The parallelism of packet processing tasks often requires independent processing on a per-packet and per-flow basis. So stateful processing that needs either high-level protocol data or particular packet orders is not a good candidate for parallel acceleration. As we will see in the following text, most computing tasks are packet classification, IP routing, pattern matching, hashing, error correction coding, encryption, name lookup, etc. They are all very packet-independent processing logic.

As for concrete applications, Gnort [26, 80] is a very early attempt at a GPU-accelerated network intrusion detection system (NIDS). It's built on top of Snort [75], which is a widely used open source NIDS. Gnort attacks the pattern matching problem in intrusion detection. It matches multiple patterns in parallel using a DFA built by the Aho-Corasick [74] algorithm, getting 3.2x speedup over the CPU algorithm implementation, and 2x speedup in the macro-benchmark.

Kargus [23] is similar to Gnort. It is also a GPU-accelerated NIDS based on Snort. Different from Gnort, Kargus includes more pattern matching algorithms for different regular expression dialects such as PCRE¹. Kargus also balances the workload between the CPU and the GPU to trade off between the throughput and latency. The load-balanced

¹<http://www.pcre.org/>

processing model in Kargus helps it reach a factor of 1.9 to 4.3 performance improvement over the conventional Snort.

MIDeA [81] is an architecture that utilizes multiple GPUs and CPUs to achieve parallel intrusion detection. MIDeA is more system-oriented rather than algorithm-oriented because it is focused on design and implementation of an efficient intrusion detection pipeline. The way MIDeA utilizes both CPUs and GPUs is different from Kargus. Kargus balances workloads between the CPU and the GPU: the entire packet processing is done completely either on the CPU or on the GPU. MIDeA is similar to **Snap**, it builds a hybrid pipeline that is connected by (interleaved) CPU stages and GPU stages. CPUs and GPUs are assigned to different computing tasks that fit the underlying processors.

PacketShader [24] is a software router with GPU-accelerated IP routing. PacketShader considers the entire network stack from packet I/O on NICs to the routing table lookup, to improve the IP routing throughput. It brings a factor of four performance improvement over the existing software router when doing minimum size IPv4 forwarding. As the pioneer work of GPU-accelerated network systems, PacketShader demonstrates a great opportunity in this area.

SSLShader [25] is an SSL proxy that utilizes both CPU and GPU power. SSLShader uses GPUs to attack the cryptography operations in the SSL protocol, such as SHA1 hashing, RSA [82] asymmetric public key encryption, AES private key encryption and the RC4 encryption algorithm. By combining the CPU and the GPU together to process SSL flows, SSLShader demonstrates a comparable performance to high-end commercial SSL appliances at a much lower price.

G-NetMon [83] is a network traffic monitor with GPU-accelerated packet classification (for packet statistics) for large scale scientific computing networks. G-NetMon does simplified packet classification based on hashing of IP addresses, which are much fewer fields than the classic packet classification fields [84], getting a 5.38x speedup over the CPU-based traffic monitor.

Our **Snap** framework is also focused on network packet processing. The main difference between **Snap** and the above work is **Snap**'s generality as a packet processing platform that is capable of building a variety of system-level network applications, including all of the above ones. Different from **Snap**, which decomposes the abstract GPU-accelerated packet processing task into modular elements for flexible configurable control in different processing tasks, each surveyed system is built as a single box for a particular task only, not suitable for other kinds of packet processing.

There are several algorithm-only efforts: GAMT [85] achieves scalable IP lookup under frequent routing table updates. Wang et.al. [86] did name lookup for the Content-Centric Network (CCN) [87] with GPU-accelerated longest prefix matching on strings. The LDPC (low-density parity-check) coding algorithm that is used in wireless communication for error correction is also attacked by Wang et.al [88] with parallel GPUs. At last, the classic packet classification problem has also benefited from parallel GPU acceleration as done by Nottingham and Irwin [89].

One potential GPU computing application in network systems is the WAN optimization [61, 66] in a wide-area network environment. Candidate computing tasks include computing data signatures for packet deduplication, compressing payload to reduce transfer overhead, computing error correction code for “forward error correction”², etc. The Content-Centric Network is definitely another good target for GPU acceleration. The aforementioned Wang et.al. [86] only did algorithm parallelization with GPUs, the underlying system infrastructure still needs careful optimization to build a high performance CCN. Our **Snap** work provides such complementary infrastructure-level optimization, making it possible to build such a CCN on top of **Snap** with the optimized name lookup algorithm.

6.1.2 Storage

Storage systems also contain many computational tasks. There are content addressable storage systems that require hashing for signatures, secure storage systems that require encryption algorithms, reliable storage systems that require erasure code for error detection and recovery, archival or space-efficient storage systems that compress the data, and high performance storage systems such as Btrfs³ that use complex indexing for fast lookup. Similar to network processing, storage tasks can often be parallelized per data block, per disk sector, per I/O request or per memory page.

Content addressable storage attracts a lot attention due to its expensive hashing operations. Shredder [90], CrystalGPU [91] and GHOST [92] are representative work on incremental storage and deduplication systems with GPU-accelerated content addressing functionality. Shredder works as a content-based chunking framework, which has showed 5x chunking bandwidth speedup over the CPU counterpart, and also large performance improvements in real world applications. CrystalGPU is actually a GPU framework for building distributed content addressable storage systems with two other GPU libraries,

²http://en.wikipedia.org/wiki/forward_error_correction

³<https://btrfs.wiki.kernel.org>

HashGPU [91] and MesaStore [91], which has demonstrated up to 2x speedup. GHOST tries to use data deduplication in primary storage systems with the GPU-accelerated deduplication. It achieves a 1.5 GB/s maximum performance, which is 5x more than the CPU system.

Erasure coding, or error correction coding, is another hot topic: Gibraltar [93] and Barracuda [41] accelerated RAID [94] parity computation with GPUs. Gibraltar is built on top of the Linux kernel’s SCSI target in userspace. It implements a flexible Reed-Solomon coding on GPUs, so it supports broader RAID configurations compared to its counterpart, the `md`’s only $n + 2$ RAID6 scheme in the Linux kernel. Barracuda is similar to GPUstore’s `md` work, they both accelerate the coding tasks in `md`. Similar to Gibraltar, Barracuda implements more schemes than `md`’s default two-failure tolerable one. Different from GPUstore’s `md` acceleration, it didn’t see speedup at the two-failure case, but got a 72x speedup at the eight-failure case.

Besides content addressable storage and erasure coding, secure storage with data encryption can also benefit from GPUs with GPU-accelerated cryptography: GPUstore includes two secure storage systems, the encrypted filesystem and the disk encryption driver.

A potential GPU application in storage is for file system integrity, such as the Featherstitch [95], which exposes the dependencies among writes in a reliable file system. One of the most expensive parts of Featherstitch is analysis of dependencies in its *patch graph*, a task we believe could be done efficiently on the GPU with parallel graph algorithms.

There exists other work that tries to ease the storage access in GPU programming. GPUfs [96] is a POSIX file API available to CUDA GPU kernels. With GPUfs, GPU kernel programmers can access file content directly from within the kernel, no need to go back to the host side any more, which leads to complex interleaved host code and device code.

6.1.3 Database Processing

For databases, the major computing tasks are in data query, which may require large scale linear scanning, sorting and hashing. Generic large scale data processing, which often involves the MapReduce [97] model, requires not only sorting during key shuffling, but also a GPU computing friendly platform for potential GPU-accelerated mappers or reducers. Note that we won’t survey the specific GPU-accelerated MapReduce applications for particular tasks such as stock analysis, because they are not system-oriented. Instead, we focus on the system support by means of a framework or runtime to build GPU-accelerated MapReduce applications.

The main algorithmic operations in relational database queries: scan, join and sorting have been studied to take advantage of parallel GPUs by a lot of existing work [98, 99, 100, 101, 102, 103], however, only a few have built practical database systems with GPU accelerations. The most practical work is the PG-Strom [104] engine, which accelerates the linear scan of massive records in the industrial-level PostgreSQL [105] database management system with GPUs. The in-GPU-memory column-oriented database [106] is designed for processing analytical workloads. The main idea, or motivation of its in-memory design is to address the low bandwidth memory copy between host and device. By keeping data in GPU device memory, memory copy through PCIe bus is unnecessary. Key-value store databases also benefit from the GPU-accelerated lookup algorithms, such as the GPU-accelerated Memcached [107] which has got up to 7.5x performance increase with integrated CPU+GPU key-value lookup.

For large scale data processing, GPUTeraSort [108] is a pipelined disk-based external sorting system to process large scale data, and shows a good price-performance in practice. As for the GPU-accelerated MapReduce frameworks, Mars [109] is the pioneer work that eases the GPU-accelerated MapReduce programming. Mars works on a single machine, so its performance is limited compared with the other CPU-based distributed MapReduce system such as Hadoop⁴. Mars also uses a simple workload mapping mechanism: it maps a single data item to one GPU thread. Such single mapping prevents programmers from exploiting the inter-thread collaboration to improve performance. Recent GPU-accelerated MapReduce frameworks such as GPMR [110] and MGMR [111] relax such mapping constraint, allowing flexible one-to-many or many-to-one mappings. More over, they have been extended to multiple GPU clusters to embrace more computing power.

6.1.4 Security

Cryptography algorithms are a main class of computing tasks for security. Another computing task is pattern matching. We have gone through the NIDS that heavily relies on pattern matching for network security, in this section, we mainly focus on general security systems.

The cryptography algorithms fit quite well on GPUs because of their simple control flow logic: symmetric cryptography [112], asymmetric cryptography [113, 114] and hashing [25], all have been exploited to use parallel GPUs to accelerate. The pattern matching based security systems are not only network intrusion detection systems as we have discussed

⁴<http://hadoop.apache.org/>

in previous sections, but also antivirus software that requires heavy signature matching. GrAVity [115] is a parallel antivirus engine with GPU-accelerated signature matching via the Boyer-Moore algorithm [116] and the Aho-Corasick algorithm [74]. It is built on top of ClamAV⁵, a popular open source antivirus software, to achieve 20 Gb/s, 100 times faster than the CPU-only performance. Kaspersky, which is a very famous antivirus software, also tried using GPUs to match virus signatures and got two orders of magnitude speedup [117]. This is the most practical GPU-accelerated antivirus solution right now, though as far as we know, Kaspersky hasn't released any product with such GPU acceleration available to end users.

A very useful potential GPU application in security is for the trusted computing in virtualized environment based on Trusted Platform Module (TPM). A TPM is traditionally hardware, but recent software implementations of the TPM specification, such as vTPM [118], are developed for hypervisors to provide trusted computing in virtualized environments where virtual machines cannot access the host TPM directly. Because TPM operations are cryptography-heavy, they can also potentially be accelerated with GPUs.

6.1.5 Program Analysis

A very interesting work is the EigenCFA [27] static program analyzer, which maps flow analysis onto matrix multiplication accelerated by GPUs to get up to 72x speedup. Though EigenCFA just does the basic OCFA [119] and can't fundamentally solve the exponential complexity control flow analysis problem, it makes the flow analysis further practical with much faster GPUs. One of my side projects, the GodelHash [120], also reveals another potential computation target for GPUs to attack: the time-consuming flow entry subsumption operations in kCFA [119]. GodelHash does perfect hashing that maps a flow entry to a prime number, then uses the product of the prime number to represent a flow set. Hence the subset and membership, which are actually the subsumption operation in control flow analysis, are done by the divisibility. It requires multiple precision arithmetic to deal with the huge integers produced by multiplying a set of prime numbers, which has demonstrated great speedup with GPU acceleration [25], hence GodelHash-based flow analysis is a good candidate for parallel acceleration.

⁵<http://www.clamav.net>

6.1.6 Data Compression

Some data compression algorithms can also be parallelized. For general compression scheme, the LZSS [121] has been exploited by Ozsoy et.al [122] to enable pipelined parallel compression. They accelerate the two main stages in LZSS, substring matching and encoding, to get a 34x speedup over a serial CPU implementation, and a 2.21x speedup compared with a parallel CPU implementation. The aforementioned GPU-accelerated database compression [123] is a specialized compression scheme for column-oriented databases. It implements and evaluates several compression schemes including Null Suppression with Fixed Length [124], Null Suppression with Variable Length [124], Dictionary, Bitmap and Run-Length-Encoding.

A potential opportunity is the recent migratory compression [125] work. Migratory compression does three steps to achieve better compression than traditional methods like **gzip**: it firstly deduplicates the data, then relocates the similar data chunks to put them together, and last do the traditional compression on the relocated data. The chunk relocation improves the data similarity, and helps migratory compression achieve 44-157% compression ratio improvement on archival data. The relocation step needs to find similar data chunks, which relies on the computation intensive “super-feature” generating and matching. The feature generating is totally chunk-independent, hence a good parallel target for GPUs. The matching needs hashing for hash table lookup and large scale sorting for similarity comparison, which are all parallelizable.

6.1.7 Miscellaneous

Prefetcher COMPASS [126] is a programmable data prefetcher using GPU cores. COMPASS assumes a CPU-GPU shared cache architecture that both CPU and GPU can access the cache, and hence the same physical memory. This is more advanced than the current integrated GPU architecture, which only allows shared memory between CPU and GPU. Compared with hardware based prefetcher, COMPASS is much more flexible because of its programmability. The authors have developed three prefetching policies from simple stride prefetching to complex Markov [127], delta prefetching [128] and region prefetching [129]. The simulated result shows a 68% performance improvement of memory-intensive applications on average.

Garbage Collection An early patent issued by AMD shows their research on using GPUs to accelerate garbage collection in language runtime. Garbage collection needs to traverse the object reference graph, which demonstrates great speedup in [130]. Martin Mass et.al. [131] tried to use GPUs in a research Java virtual machine, but failed to see any

speedup due to the inefficient memory copy in OpenCL. The graph algorithms have been demonstrated with their GPU accelerations [30, 132, 130], so we can believe that garbage collection, which heavily relies on them, is definitely a very promising GPU acceleration target.

6.2 Techniques

In this section, we will discuss the techniques used by a variety of existing system-level GPU work. As said at the beginning of the chapter, we focus on the system-level techniques rather than algorithm parallelization. The general goal of the techniques is to improve the system performance, however, there are some technologies that have been exploited to make GPU computing available in more environments, such as Gdev for in-kernel GPU computing, and the virtualized GPU solutions discussed at the end of this section that makes GPGPU available to virtualization environments and cloud computing. The techniques are categorized by the challenges or problems they solved, which we will see a large portion of common categories shared with the ones discussed in Chapter 3. For those categories, readers should refer to Chapter 3 for the description or origins of the challenges and problems.

6.2.1 Batched Processing

All the surveyed work needs some kind of batching to ensure enough parallelism in the workload to be processed on GPUs. However, there are two different kinds of batching approaches: static batch size and dynamic batch size.

The static batch size approach is easier, like mentioned in most surveyed papers [24, 93, 41, 91, 108] and my own **GPUstore** and **Snap** work, the number of data units in a single batch is constant. Such constant value is decided by performing algorithmic benchmarking, finding the optimal size of the batch, and then using it forever. The optimal batch size can be the size where GPU performance goes down, or the first size that the GPU outperforms CPU. The drawback of such an approach is obvious: it can be easily affected by the system load and any other runtime overheads. Because of its simplicity, this has been applied by most system-level work.

The dynamic batch size is complex: it still needs some pre-benchmarking, but instead of finding a single size, it builds a relationship between the batch size and the performance. During the run time, it dynamically adjusts the batch size, or disables GPU offloading according to the built relationship, system load and other factors. **SSLShader** [25] and **Kargus** [23] use this approach to implement flexible GPU offloading control to balance the

latency and throughput. They also implement dynamic load balancing between the CPU and the GPU with the help of the built relationship: small packets, low packet rates and high GPU load all lead to CPU processing.

The dynamic approach is for sure the better way to finely tune the overall system performance. But it requires a comprehensive study and evaluation of the GPU-accelerated algorithms under a variety of different environments, and also possibly complex runtime workload balancer. So the static approach is still useful to demonstrate GPU accelerations without considering the combined CPU+GPU power and external environment conditions.

6.2.2 Reducing Memory Copy Overhead

Memory copy is the main bottleneck for most system-level GPU computing. Taking the NVIDIA GTX Titan as example, compared with the 336 GB/s GPU device memory bandwidth and more than tera-flops computing power, even the maximum PCIe 3.0 x16 can only run at 16 GB/s, which is 31 times slower! Besides that, Section 3.5.2 also discusses the double-buffering problem caused by inefficient GPU computing integration into system software in practice. Several techniques have been applied by previous work to reduce the memory copy overhead. Here they are categorized into the following approaches.

6.2.2.1 Using Page-Locked Memory

Although there are some systems, such as the GPUTeraSort [108], Barracuda [41], database compression on MonetDB [123], EigenCFA [27] and Mars [109], that haven't used page-locked memory to improve PCIe DMA performance, most of the aforementioned work applies this simple but effective memory copy optimization. However, not all of them have efficiently used page-locked memory in practice. Many systems just use page-locked memory as a separate memory area dedicated for GPU computing, rather than treating it as normal host memory that is used for any host side computation, which leads to the inefficient double-buffering problem.

6.2.2.2 Avoid Double-Buffering

Many systems surveyed perform a "use all" data usage pattern as discussed in Section 3.5.4. However, even when using page-locked memory to improve memory copy, most of them still failed to efficiently integrate the memory into the host side computation as normal host memory. Some are totally unaware of this problem such as the CrystalGPU-based content addressable storage [91] and GHOST deduplication storage [92]. Some of them are due to the technical obstacle when getting data from the operating system

kernel [93, 25, 90, 41, 107]. Such kernel-user memory sharing obstacles can be solved by allocating the GPU computing host side memory directly in the OS kernel as the kernel-space GPU driver Gdev [40] does, or supporting such in-kernel allocation with pre-allocated page-locked memory passed into the kernel from userspace, like **GPUstore** does. Only a few of the surveyed papers have carefully designed their systems to deal with the double-buffering problem. For example, the GrAVity [115] anti-virus engine has modified the memory buffers used in the mature anti-virus software, ClamAV, to efficiently use CUDA page-locked memory for virus pattern detection. Another example of such engineering work is the PG-Strom GPU query engine for PostgreSQL [104], which has efficiently integrated both page-locked memory and asynchronous overlapped GPU computation and memory copy into a large complex existing industrial-level database system.

6.2.2.3 Trade Off Host Memory and PCIe Bandwidths

Having discussed in Section 3.5.3, packet processing is the typical task that follows a “use part” of the entire data unit usage pattern. But not all the network systems in Section 6.1.1 have considered the potential performance improvement by taking advantage of the much larger host memory bandwidth. SSLShader [25] is such an example that is totally unaware of such potential speedup. Although the rest, such as those NIDS systems [26, 81, 23] that only copy packet payloads into device memory and the PacketShader [24] that copies only an IP address and a packet identifier to GPU, still do not use the more generic ROI-based slicing technique as **Snap** has proposed in Section 5.3.1.3. Actually, such copy in most of those papers is not described as a memory copy optimization, but instead, an overhead that must be suffered due to the technical difficulty of sharing DMA-able memory between the NIC driver and the GPU driver.

Interestingly, some papers other than network systems are aware of this issue, and have carefully traded off not only the host memory bandwidth and PCIe’s, but also considering the CPU performance in some cases. Such as in EigenCFA [27], which does not use page-locked memory though, authors try to significantly reduce the size of the flow data representation to be copied into device memory with pre-converted bit-packed matrix at the CPU side. Another example is the GPUTeraSort [108] system, which is similar to EigenCFA, using pageable memory only, but still, authors do CPU side pre-processing on the data entries read from the disks to produce a much smaller (**key**, **pointer**) tuple for each entry. Although the GPUTeraSort case is very straightforward, it still represent the application of an effective memory copy optimization technique, and actually, has inspired

my generic ROI-based slicing design in **Snap**.

6.2.2.4 Exploiting The Blackbox GPU Driver

There are some papers that try to improve the underlying blackbox GPU driver in order to get better PCIe DMA performance. Although we can optimistically expect that future GPUs will be completely open, it is still unclear how long the “future” can be. Hence those papers studying the close source GPU drivers make sense to the relatively long “present”. The work done by Y.Fuji et.al [133, 134] has exploited the micro-controllers on GPUs that manages the CPU-GPU communication and GPU command queues. By modifying the firmware for the GPU micro-controller using an open source GPU driver [39], they are able to achieve up to 10x speedup when copying small size data (ranging from 16 bytes to 256 bytes) controlled by the micro-controller rather than the dedicated DMA engines. Gdev [40] implements a CUDA runtime inside the Linux kernel based on the same open source GPU driver, and has implemented across-driver DMA-able memory pages sharing for arbitrary kernel components, which is supposed to be provided by GPUDirect [135], but is rarely available in practice. Gdev is complementary to all the in-kernel GPU computing work such as **GPUstore** and Barracuda [41] because of its efficient in-kernel GPU driver access without the need of a userspace helper.

6.2.3 Dealing With Warp Control Flow Divergence

Besides the algorithmic design techniques that are used to eliminate as many control flow branches as possible, most system work doesn’t encounter the on-GPU data divergence problem as **Snap** has solved in Section 5.3.2. Even in packet processing work, systems such as those NIDSes [26, 23, 81] do pre-processing on CPUs to classify packets into different “port groups” according to their layer 3 protocol ports. This pre-processing is acceptable in those systems because they just have a one-stage (the pattern matching only) processing control flow, rather than a multi-stage one with multiple layers of branching as in **Snap**.

Some non-system work such as the GPU accelerated packet classification algorithms [89] and non-GPU work such as the batched processing in Click [77] have identified and dealt with this problem. The classification paper uses a similar approach to **Snap**’s predicated execution: it uses a boolean value to notify the following classification rule kernel whether to process a particular packet or not. The batched processing Click paper encountered this problem when doing the CPU side batching, they simply split the batch into multiple small batches, which has not got any evaluation result due to being in the early stage of their work.

Another similar work, PTask [51], which tries to build a data flow programming model for GPU tasks by providing a UNIX pipe-like abstraction, simply does not support such flow-divergence, or conditional data flow, in the flow graph, though it supports the opposite one: merging of multiple outputs into a single input. Besides that, PTask is a more application-level abstraction that is designed to deal with inter-process data flow. It is a heavy weighted abstraction compared with data sharing mechanisms in system code, such as memory page sharing in an OS kernel.

6.2.4 Improving GPU Cores and DMA Engines Utilization

Most of the recent papers use CUDA’s overlapped GPU computation and memory copy, and asynchronous GPU kernel execution features to improve their performance by fully utilizing the GPU cores and DMA engines. Those failed to benefit from them are mostly very old papers before NVIDIA released its CUDA 4.0, which first introduce the techniques. However, there are exceptions such as PacketShader, which claims, though without any further explanation, a performance slow down when using such techniques [24]. Despite the fact that many systems use such overlapped scheme and asynchronous execution, none of them have studied the synchronization problem in GPU programming even with CUDA stream like technique, as discussed in Section 3.4.

6.2.5 Virtualization and GPGPU Application Migration

The virtualization and application migration techniques are mainly for enabling GPU computing in more environments, especially in the cloud computing. Cloud providers such as Amazon EC2⁶ have provided GPGPU devices in their cloud instances [1]. However, those GPUs are completely exposed to users rather than through virtualization. This not only brings security issues, but also prevents resource consolidation and instance migration, which are the most important features of a cloud platform.

Several GPU virtualization solutions have been proposed by both industry and academia. The VMWare’s GPU virtualization [136] aims at the graphics processing functionality. It is based on VMWare’s hosted I/O architecture to enable GPU command queue isolation, GPU memory isolation and I/O space isolation so as to achieve independent virtualized GPUs. GViM [137] and gVirtuS [138] are very simple virtualization layers just for CUDA computing. They both provide an intercept layer for CUDA library calls to forward them

⁶<http://aws.amazon.com/ec2/>

to the actual CUDA runtime in the host. Gdev [40] supports both graphics processing virtualization and the CUDA computing virtualization. Gdev can achieve the fully functional virtualization because it is built on top of an open source GPU driver, and hence it can access and control very fine-grained and low-level GPU functionality.

Migration of GPGPU computing programs is in a very early stage at this time. The currently available two solutions, CheCL [139] and CheCUDA [140] both use an interception layer for CUDA calls, which is similar to the aforementioned simple virtualization solutions. The proprietary GPU driver and closed GPU specifications prevent efficient GPU virtualization and migration. As a result, the obstacle is not any technical issue, but some marketing policies of GPU vendors. Recent good news is that AMD has gradually released its GPU specification [141], and NVIDIA has released some open specifications of its mobile Tegra GPUs⁷. We should believe that in the future, open GPUs will be available. Both GPU virtualization and GPGPU application migration are going to be improved in functionality and qualities.

6.2.6 Resource Management and Scheduling

These are actually not techniques to improve system-level GPU computing performance, but system-level work that improves all GPGPU computing applications.

TimeGraph [52] is a GPU resource scheduling system based on an open source GPU driver. TimeGraph schedules GPU commands issued from different tasks based on their priorities to achieve resource isolation and reservation for real-time environments. A similar work is PTask [51], which also deals with the scheduling problem of GPU tasks. Different from TimeGraph, PTask considers the priorities of both GPU tasks and CPU tasks to achieve fair scheduling by counting GPU execution time into a task's total execution time. As said in Section 6.2.3, PTask also provides a UNIX pipe-like model to synthesize an efficient data flow of multiple GPU tasks with dependencies. The aforementioned Gdev [40] is also an OS-level GPU resource management and scheduling solution. Gdev uses the same open source GPU driver as TimeGraph, so it can provide similar scheduling functionality, though not for real-time environments.

⁷http://www.phoronix.com/scan.php?page=news_item&px=MTU40Tg

6.3 Future

GPUs have been developed much faster than CPUs: the number of cores is doubled almost in only one year [6] instead of CPU's 18 months following Moore's Law⁸. Thus we can confidently believe that future GPUs will come with tens or hundreds of thousands of cores. Such amazingly large number of cores are ideal to deal with current "big data" trend in both application level and the system level.

Another technical improvement should be the CPU-GPU communication. Today's fastest PCIe 3.0 bus only reaches 16 GB/s peak bandwidth, compared with the more than 300 GB/s GPU device memory bandwidth, it easily becomes the bottleneck in many GPU computing as we have seen in previous sections and chapters. Besides dramatically increasing PCIe bandwidth, another potentially useful technique is the integrated GPUs that sit on the same chip of the CPU. Integrated GPUs are expected to eliminate the PCIe bus overhead, however, most current integrated GPUs are limited to only access their own dedicated main memory region. This can still cause possible data copy in main memory due to the dedicated GPU region. What's more, the host side memory bus interface is much narrower than the dedicated GPU's, e.g., DDR3 RAM's 64 bits width v.s. NVIDIA GTX Titan Black's 384 bits. As a result, integrated GPUs suffer more memory copy issues. Recent AMD's Kaveri APU⁹ has been designed to provide a better technique: heterogeneous Uniform Memory Access (hUMA) [142]. hUMA allows integrated GPUs to directly access any memory locations, and even further, the CPU cache and GPU cache benefit from hardware-based cache coherence. Though hUMA still can't provide the wide device memory interface, at least it eliminates any possible copy in main memory. As a result, the main memory techniques such as DDR4 or DDR5 should try to provide wider and faster memory for integrated GPUs, too.

Besides technical improvements, as we said in Section 6.2.5, the openness of GPU hardware specification and software drivers is also vital to applications, especially system-level applications, to get rid of current overheads introduced by many technical workarounds due to the blackbox GPU driver and close GPU specification. Though the open process is very slow according to recent years' experience, some good news mentioned in Section 6.2.5 gives us confidence that it is possible for future GPUs to become completely open.

Another system-level GPU computing area that is lacking attention is with mobile

⁸http://en.wikipedia.org/wiki/Moore's_law

⁹<http://www.amd.com/en-us/products/processors/desktop/a-series-apu>

devices. Traditional GPGPU computing has already explored mobile GPUs for mainly visual computing [143, 144, 145], but none of the system-level software has tried such Small but ubiquitous GPUs. The main problem in current mobile GPU computing is their relatively low performance compared with the PC GPUs. According to a mobile GPGPU computing evaluation done by Arian Maghazeh et al. [146], mobile GPU performance is still not on par their CPU competitors. Fortunately, mobile GPUs are getting more powerful. Recent NVIDIA Tegra K1 mobile processor has significantly increased its GPU cores to up to 192 CUDA cores¹⁰. This makes it possible to get much faster GPU computing than the CPU on mobile devices like on PCs. Mobile GPUs are more closely combined with CPUs than the PC ones. It may be possible for mobile GPUs to share and control more integrated components such as memory, DMA, and I/O on the highly integrated system-on-chip (SoC) for mobile devices. But mobile devices also suffer from a common energy consumption problem. So future system-level GPU computing on mobile devices will definitely need new techniques to achieve high performance yet power efficiency.

6.4 Summary

This literature survey covers recent applications of and techniques for improving the performance and functionality of system-level GPU computing. It has been done with strong system flavor, from the applications to the techniques, which are very different from traditional GPU computing surveys such as those done by John D. Owens et.al [17]. Besides existing work, it tries to identify several potential system-level tasks that can be improved with GPU computing. The system techniques surveyed in this chapter share a large portion with the ones applied in my work, which confirms the effectiveness of each other.

As showed in this survey, some techniques such as batched processing, overlapped GPU computation and memory copy, and using locked memory directly to avoid double buffering GPU DMA, are commonly used by the surveyed system work. However, among the generic principles and techniques described in Chapter 3, the truly asynchronous GPU programming, SIMT-friendly data structures and using locked memory according to different data usage patterns are newly proposed in this study. Previous system designs either did not study the related problems, or used very problem specific approaches. For example, as shown in Section 6.2.2.3, PacketShader avoids wasting PCIe bandwidth by copying only an IP address and packet identifier, but it is a very IP routing specific design compared with Snap’s generic ROI-based packet slicing that is available to all the packet processing tasks.

¹⁰<http://www.nvidia.com/object/tegra-k1-processor.html>

The surveyed papers are actually origins of most techniques proposed in this thesis work, after careful trade-off, abstracting, generalization and many rounds of design considerations. Those techniques have been categorized by the problems and challenges they've solved in order to make it convenient for readers to refer to in practice.

CHAPTER 7

CONCLUSION

Because of its low-level position in the software stack, efficient system-level software is vital to high performance computer systems. However, computationally expensive tasks in system-level software require more and more computing power due to the rapidly increased bulk-data workloads. They may consume excessive processing power, slowing down the entire software stack.

We deal with this problem using parallel GPUs. The highly parallel GPU processors provide a throughput-oriented architecture, which is designed for parallel bulk-data processing. Integrating GPU computing into originally latency-oriented low-level system software has faced many technical challenges that prevent high performance GPU computing in system level. We have proposed generic principles for designing and implementing GPU-accelerated system-level applications. The two specific instances, **GPUstore** and **Snap** frameworks for two representative kinds of system software, have successfully enabled efficient GPU computing in typical system-level tasks, and both have got significant performance improvement with their efficient memory and throughput-oriented design guided by our proposed principles. The success of these two not only validates the feasibility of system-level GPU computing, but also provides a variety of effective generic techniques and abstractions that are designed to deal with common challenges encountered during the GPU computing integration, achieving better performance speedup. We then surveyed the current system-level GPU computing area, identifying potential applications, discussing useful techniques applied in existing work, comparing them with our techniques, and giving interesting readers and researchers a literature to get a picture of the state of the art and more importantly, a technical reference for their own system-level GPU computing work.

Although the techniques and principles described in this work are mainly for parallel GPUs, they may also be applied to other similar computing platforms. For example,

the wide range of co-processors includes cryptography co-processors ¹, network processors, DSP co-processors, computing co-processors such as Intel Xeon Phi ² and FPGAs. These co-processors share a lot of common architecture features with GPUs: SIMD style throughput-oriented computing, PCIe-based high latency communication, indirect or limited host memory access capability, and dedicated on-device memory. All of these common features indicate that similar challenges also exist when using those co-processors in system code. As a result, our generic system-level GPU computing principles and techniques that are designed to deal with those challenges may also apply to those platforms. Systems working with the co-processors may also benefit from the study in this work. Besides the co-processors, today's multicore CPU platform may benefit from this study as well. A recent study [77] that did batched processing in Click on CPUs shows that the CPU-based packet processing can also get speedup from batching packets. The similar techniques we implemented in **Snap**, such as batching elements and packet batch data structures, are the candidate techniques applicable to multicore CPUs. This shows the much broader impact of our work, which is not limited to GPUs only, but also valuable to a wide range of computing platforms.

¹<http://www-03.ibm.com/security/cryptocards>

²<https://software.intel.com/en-us/mic-developer>

REFERENCES

- [1] Amazon Web Service LLC., “Amazon EC2 Instance Types.” <http://aws.amazon.com/ec2/instance-types/>.
- [2] W. Wade, “How GRID Brings Amazing Graphics to the Enterprise Cloud,” in *GPU Technology Conference*, GTC’13, 2013.
- [3] O. R. L. C. Facility, “Introducing Titan.” <http://www.olcf.ornl.gov/titan>.
- [4] Nvidia, “High Performance Computing (HPC) and Supercomputing.” <http://www.nvidia.com/object/tesla-supercomputing-solutions.html>.
- [5] NVIDIA, “GeForce GTX Titan Black GPU.” <http://www.nvidia.com/gtx-700-graphics-cards/gtx-titan-black>.
- [6] NVIDIA, “CUDA C Programming Guide.” <http://docs.nvidia.com/cuda/cuda-c-programming-guide/>.
- [7] M. Garland and D. B. Kirk, “Understanding throughput-oriented architectures,” *Comm. ACM*, vol. 53, pp. 58–66, 2010.
- [8] K. Yang, B. He, Q. Luo, P. V. Sander, and J. Shi, “Stack-based parallel recursion on graphics processors,” in *Proceedings of the 14th ACM SIGPLAN symposium on Principles and practice of parallel programming*, PPOPP ’09, pp. 299–300, ACM, 2009.
- [9] Y. Yang, P. Xiang, J. Kong, and H. Zhou, “A GPGPU compiler for memory optimization and parallelism management,” in *Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation*, PLDI ’10, pp. 86–97, ACM, 2010.
- [10] T. B. Jablin, P. Prabhu, J. A. Jablin, N. P. Johnson, S. R. Beard, and D. I. August, “Automatic CPU-GPU communication management and optimization,” in *Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation*, PLDI ’11, pp. 142–151, ACM, 2011.
- [11] J. Soman, K. Kishore, and P. Narayanan, “A fast GPU algorithm for graph connectivity,” in *Parallel Distributed Processing, Workshops and Phd Forum (IPDPSW), 2010 IEEE International Symposium on*, pp. 1–8, april 2010.
- [12] R. R. Amossen and R. Pagh, “A New Data Layout for Set Intersection on GPUs,” in *Proceedings of the 2011 IEEE International Parallel & Distributed Processing Symposium*, IPDPS ’11, pp. 698–708, IEEE Computer Society, 2011.
- [13] M. S. Rehman, K. Kothapalli, and P. J. Narayanan, “Fast and scalable list ranking on the GPU,” in *Proceedings of the 23rd international conference on Supercomputing*, ICS ’09, pp. 235–243, ACM, 2009.

- [14] G. Capannini, F. Silvestri, and R. Baraglia, “K-Model: A New Computational Model for Stream Processors,” in *High Performance Computing and Communications (HPCC), 2010 12th IEEE International Conference on*, pp. 239–246, sept. 2010.
- [15] B. Domonkos and G. Jakab, “A programming model for GPU-based parallel computing with scalability and abstraction,” in *Proceedings of the 2009 Spring Conference on Computer Graphics, SCCG '09*, pp. 103–111, ACM, 2009.
- [16] S. Hong and H. Kim, “An analytical model for a GPU architecture with memory-level and thread-level parallelism awareness,” in *Proceedings of the 36th annual international symposium on Computer architecture, ISCA '09*, pp. 152–163, ACM, 2009.
- [17] J. D. Owens, D. Luebke, N. Govindaraju, M. Harris, J. Krger, A. Lefohn, and T. J. Purcell, “A Survey of General-Purpose Computation on Graphics Hardware,” *Computer Graphics Forum*, vol. 26, no. 1, pp. 80–113, 2007.
- [18] J. D. Owens, M. Houston, D. Luebke, S. Green, J. Stone, and J. Phillips, “GPU computing,” *Proceedings of the IEEE*, vol. 96, no. 5, pp. 879–899, 2008.
- [19] W. Sun, R. Ricci, and M. J. Curry, “GPUstore: Harnessing GPU computing for storage systems in the OS kernel,” in *Proceedings of the Fifth ACM International Systems and Storage Conference (SYSTOR)*, June 2012.
- [20] W. Sun and R. Ricci, “Fast and flexible: Parallel packet processing with GPUs and Click,” in *Proceedings of the Ninth ACM/IEEE Symposium on Architecture for Networking and Communications Systems (ANCS)*, Oct. 2013.
- [21] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The Click Modular Router,” *ACM Transactions on Computer Systems*, vol. 18, pp. 263–297, Aug. 2000.
- [22] Khronos Group, “OpenCL Specification 2.0.” <https://www.khronos.org/opencl/>.
- [23] M. Jamshed, J. Lee, S. Moon, D. Kim, S. Lee, and K. Park, “Kargus : A Highly-scalable Software-based Intrusion Detection System Categories and Subject Descriptors,” in *Proceedings of the 2012 ACM Conference on Computer and Communications Security, CCS'12*, pp. 317–328, 2012.
- [24] S. Han, K. Jang, K. Park, and S. Moon, “PacketShader : A GPU-Accelerated Software Router,” in *Proceedings of the ACM SIGCOMM 2010 Conference, SIGCOMM'10*, pp. 195–206, ACM, Oct. 2010.
- [25] K. Jang, S. Han, S. Han, S. Moon, and K. Park, “SSLShader : Cheap SSL Acceleration with Commodity Processors,” in *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation, NSDI'11*, 2011.
- [26] G. Vasiliadis, S. Antonatos, M. Polychronakis, E. P. Markatos, and S. Ioannidis, “Gnort : High Performance Network Intrusion Detection Using Graphics Processors,” in *Proceedings of the 11th International Symposium On Recent Advances In Intrusion Detection, RAID'08*, Sept. 2008.
- [27] T. Prabhu, S. Ramalingam, M. Might, and M. Hall, “EigenCFA: Accelerating Flow Analysis with GPUs,” in *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '11*, pp. 511–522, ACM, 2011.

- [28] NVIDIA, “CUDA C Best Practice Guide.” <http://docs.nvidia.com/cuda/cuda-c-best-practices-guide/>.
- [29] P. Kipfer and R. Westermann, “Improved GPU Sorting,” *GPU Gems*, Apr. 2005.
- [30] P. Harish and P. J. Narayanan, “Accelerating Large Graph Algorithms on the GPU Using CUDA,” in *Proceedings of the 14th International Conference on High Performance Computing*, HiPC’07, pp. 197–208, Springer-Verlag, 2007.
- [31] NVIDIA, “CUDA Basic Linear Algebra Subroutines (cuBLAS).” <https://developer.nvidia.com/cublas>.
- [32] M. Harris, S. Sengupta, and J. D. Owens, “Parallel Prefix Sum (Scan) with CUDA,” *GPU Gems*, 2007.
- [33] D. Roger, U. Assarsson, and N. Holzschuch, “Efficient Stream Reduction on the GPU,” in *Workshop on General Purpose Processing on Graphics Processing Units*, oct 2007.
- [34] M. Billeter, O. Olsson, and U. Assarsson, “Efficient Stream Compaction on Wide SIMD Many-core Architectures,” in *Proceedings of the Conference on High Performance Graphics 2009*, HPG ’09, pp. 159–166, 2009.
- [35] D. Le, J. Chang, X. Gou, A. Zhang, and C. Lu, “Parallel AES Algorithm for Fast Data Encryption on GPU,” in *Proceedings of the 2nd International Conference on Computer Engineering and Technology*, ICCET’10, pp. 1–6, IEEE, Apr. 2010.
- [36] A. Dunkels, “uIP (micro IP) TCP/IP stack.” [http://en.wikipedia.org/wiki/UIP_\(micro_IP\)](http://en.wikipedia.org/wiki/UIP_(micro_IP)).
- [37] A. Dunkels, “lwIP - A Lightweight TCP/IP stack.” <http://savannah.nongnu.org/projects/lwip>.
- [38] M. Harris, “How to Optimize Data Transfers in CUDA C/C++,” Dec. 2012. <http://devblogs.nvidia.com/paralleforall/how-optimize-data-transfers-cuda-cc>.
- [39] freedesktop.org, “Nouveau: Accelerated Open Source driver for nVidia Cards.” <http://nouveau.freedesktop.org/wiki/>.
- [40] S. Kato, M. Mchrow, C. Maltzahn, and S. Brandt, “Gdev : First-Class GPU Resource Management in the Operating System,” in *Proceedings of the 2012 USENIX conference on Annual Technical Conference*, ATC’12, 2012.
- [41] A. Brinkmann and D. Eschweiler, “A Microdriver Architecture for Error Correcting Codes inside the Linux Kernel,” in *Proceedings of the ACM/IEEE Conference on High Performance Computing*, SC’09, (Portland, OR), Nov. 2009.
- [42] E. Zadok, I. Badulescu, and A. Shender, “Cryptfs: A Stackable Vnode Level Encryption File System,” Tech. Rep. CUCS-021-98, Computer Science, Columbia University, 1998.
- [43] M. Blaze, “A cryptographic file system for UNIX,” in *Proceedings of the 1st ACM Conference on Computer and Communications Security*, 1993.

- [44] S. Watanabe, *Solaris 10 ZFS Essentials*. Prentice Hall, 2009.
- [45] S. Patil, G. Sivathanu, and E. Zadok, “T³fs: An in-kernel integrity checker and intrusion detection file system,” in *Proceedings of the 18th Annual Large Installation System Administration Conference (LISA)*, pp. 67–78, 2004.
- [46] S. Quinlan and S. Dorward, “Venti: A New Approach to Archival Data Storage,” in *Proceedings of the 1st USENIX Conference on File and Storage Technologies (FAST)*, 2002.
- [47] C. Ungureanu, B. Atkin, A. Aranya, S. Gokhale, S. Rago, G. Calkowski, C. Dubnicki, and A. Bohra, “HydraFS: a high-throughput file system for the HYDRASstor content-addressable storage system,” in *Proceedings of the 8th USENIX conference on File and Storage Technologies (FAST)*, 2010.
- [48] P. Nath, B. Urgaonkar, and A. Sivasubramaniam, “Evaluating the usefulness of content addressable storage for high-performance data intensive applications,” in *Proceedings of the 17th International Symposium on High Performance Distributed Computing (HPDC)*, 2008.
- [49] A. Kashyap and A. Kashyap, “File System Extensibility and Reliability Using an in-Kernel Database,” tech. rep., Stony Brook University, 2004.
- [50] H. Pang, K.-L. Tan, and X. Zhou, “StegFS: A Steganographic File System,” *International Conference on Data Engineering*, 2003.
- [51] C. J. Rossbach, J. Currey, and B. Ray, “PTask : Operating System Abstractions To Manage GPUs as Compute Devices,” in *Proceedings of the 23rd ACM Symposium on Operating Systems Principles, SOSP’11*, 2011.
- [52] S. Kato, K. Lakshmanan, R. R. Rajkumar, and Y. Ishikawa, “TimeGraph : GPU Scheduling for Real-Time Multi-Tasking Environments,” in *Proceedings of the 2011 USENIX Annual Technical Conference, ATC’11*, 2011.
- [53] L. Xu, *Securing the Enterprise with Intel AES-NI*. Intel Corporation, 2010.
- [54] “eCryptfs - Enterprise Cryptographic Filesystem.” <https://launchpad.net/ecryptfs>.
- [55] P. M. Chen, E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson, “RAID: High-Performance, Reliable Secondary Storage,” *ACM Computing Surveys*, vol. 26, no. 2, pp. 145–185, 1994.
- [56] M. Blaum, J. Brady, J. Bruck, and J. Menon, “EVENODD: An optimal scheme for tolerating double disk failures in RAID architectures,” in *Proceedings of the 21st Annual International Symposium on Computer Architecture (ISCA)*, 1994.
- [57] P. Corbett, B. English, A. Goel, T. Grcanac, S. Kleiman, J. Leong, and S. Sankar, “Row-Diagonal Parity for Double Disk Failure Correction,” in *Proceedings of the 3rd USENIX Symposium on File and Storage Technologies (FAST)*, pp. 1–14, 2004.
- [58] I. S. Reed and G. Solomon, “Polynomial Codes Over Certain Finite Fields,” *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960.

- [59] R. Bhaskar, P. K. Dubey, V. Kumar, and A. Rudra, "Efficient Galois field arithmetic on SIMD architectures," in *Proceedings of the Symposium on Parallel Algorithms and Architectures*, 2003.
- [60] FastestSSD.com, "SSD Ranking: The Fastest Solid State Drives," Apr. 2012. <http://www.fastestssd.com/featured/ssd-rankings-the-fastest-solid-state-drives/#pcie>; accessed April 27, 2012.
- [61] Wikipedia, "WAN optimization." http://en.wikipedia.org/wiki/WAN_optimization.
- [62] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, 2008.
- [63] R. Heady, G. Luger, A. Maccabe, and M. Servilla, "The Architecture of a Network Level Intrusion Detection System," tech. rep., Department of Computer Science, University of New Mexico, 1990.
- [64] T. Anderson, T. Roscoe, and D. Wetherall, "Preventing Internet Denial-of-Service with Capabilities," in *Proceedings of the HotNets Workshop*, 2003.
- [65] D. Han, A. Anand, F. Dogar, B. Li, H. Lim, M. Machado, A. Mukundan, W. Wu, A. Akella, D. G. Andersen, J. W. Byers, S. Seshan, and P. Steenkiste, "XIA: Efficient Support for Evolvable Internetworking," in *Proceedings of NSDI*, 2012.
- [66] A. Anand, A. Gupta, A. Akella, S. Seshan, and S. Shenker, "Packet Caches on Routers: the Implications of Universal Redundant Traffic Elimination," in *ACM SIGCOMM Computer Communication Review*, vol. 38, pp. 219–230, ACM, 2008.
- [67] M. Dobrescu, N. Egi, K. Argyraki, B.-G. Chun, K. Fall, G. Iannaccone, A. Knies, M. Manesh, and S. Ratnasamy, "RouteBricks: Exploiting Parallelism to Scale Software Routers," in *Proceedings of the 22nd SOSP*, 2009.
- [68] L. Rizzo, "Netmap: A Novel Framework For Fast Packet I/O," in *Proceedings of USENIX ATC*, 2012.
- [69] J. C. R. Bennett, C. Partridge, and N. Shectman, "Packet reordering is not pathological network behavior," *IEEE/ACM Trans. Networking*, vol. 7, pp. 789–798, Dec. 1999.
- [70] M. Laor and L. Gendel, "The Effect of Packet Reordering in a Backbone Link on Application Throughput," *Network, IEEE*, vol. 16, pp. 28 – 36, Sep/Oct 2002.
- [71] E. Kohler, "Click Source Code Repository." <https://github.com/kohler/click>.
- [72] "RouteView Routing Table Dump." <http://www.read.cs.ucla.edu/click/routetabletest-167k.click.gz>.
- [73] D. E. Taylor and J. S. Turner, "ClassBench: a packet classification benchmark," *IEEE/ACM Trans. Netw.*, vol. 15, pp. 499–511, June 2007.
- [74] A. V. Aho and M. J. Corasick, "Efficient String Matching: An Aid to Bibliographic Search," *Commun. ACM*, vol. 18, pp. 333–340, June 1975.

- [75] M. Roesch, “Snort - Lightweight Intrusion Detection for Networks,” in *Proceedings of the 13th USENIX LISA*, 1999.
- [76] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar, “An Integrated Experimental Environment for Distributed Systems and Networks,” in *Proceedings of OSDI*, 2002.
- [77] J. Kim, S. Huh, K. Jang, K. Park, and S. Moon, “The Power of Batching In The Click Modular Router,” in *Proceedings of the ACM Asia-Pacific Workshop on Systems*, 2012.
- [78] NVIDIA, “CUDA Dynamic Parallelism Programming Guide.” <http://docs.nvidia.com/cuda/cuda-dynamic-parallelism>.
- [79] W. Sun and R. Ricci, “Augmenting Operating Systems With the GPU,” *CoRR*, vol. abs/1305.3345, 2013.
- [80] G. Vasiliadis, M. Polychronakis, S. Antonatos, E. P. Markatos, and S. Ioannidis, “Regular Expression Matching on Graphics Hardware for Intrusion Detection,” in *Proceedings of the 12th International Symposium On Recent Advances In Intrusion Detection*, RAID’09, Sept. 2009.
- [81] G. Vasiliadis, M. Polychronakis, and S. Ioannidis, “MIDeA : A Multi-Parallel Intrusion Detection Architecture,” in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, CCS’11, Oct. 2011.
- [82] R. L. Rivest, A. Shamir, and L. Adleman, “A Method for Obtaining Digital Signatures and Public-key Cryptosystems,” *Commun. ACM*, vol. 21, pp. 120–126, Feb. 1978.
- [83] W. Wu, P. DeMar, D. Holmgren, A. Singh, and R. Pordes, “G-NetMon: A GPU-accelerated network performance monitoring system for large scale scientific collaborations,” in *Proceedings of the 36th IEEE Conference on Local Computer Networks*, LCN’11, pp. 195–198, IEEE, Oct. 2011.
- [84] D. E. Taylor, “Survey and Taxonomy of Packet Classification Techniques,” *ACM Computing Surveys*, vol. 37, pp. 238–275, Sept. 2005.
- [85] Y. Li, D. Zhang, A. X. Liu, and J. Zheng, “GAMT: A fast and scalable IP lookup engine for GPU-based software routers,” in *Proceedings of the 9th Architectures for Networking and Communications Systems*, ANCS’13, pp. 1–12, IEEE, Oct. 2013.
- [86] Y. Wang, Y. Zu, T. Zhang, K. Peng, Q. Dong, and B. Liu, “Wire Speed Name Lookup : A GPU-based Approach,” in *Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation*, NSDI’13, pp. 199–212, 2013.
- [87] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, “Networking Named Content,” in *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT ’09, pp. 1–12, 2009.
- [88] G. Wang, M. Wu, B. Yin, and J. R. Cavallaro, “High Throughput Low Latency LDPC Decoding on GPU for SDR Systems,” in *Proceedings of the 1st IEEE Global Conference on Signal and Information Processing*, no. December in GlobalSIP’13, (Austin, Texas USA), pp. 3–6, IEEE, Dec. 2013.

- [89] A. Nottingham and B. Irwin, "Parallel packet classification using GPU co-processors," in *Proceedings of the 2010 Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists*, SAICSIT '10, 2010.
- [90] P. Bhatotia and R. Rodrigues, "Shredder : GPU-Accelerated Incremental Storage and Computation," in *Proceedings of the 10th USENIX Conference on File and Storage Technologies*, FAST'12, 2012.
- [91] A. Gharaibeh, S. Al-Kiswany, S. Gopalakrishnan, and M. Ripeanu, "A GPU accelerated storage system," in *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, HPDC '10, p. 167, ACM Press, 2010.
- [92] C. Kim, K.-w. Park, and K. H. Park, "GHOST : GPGPU-Offloaded High Performance Storage I/O Deduplication for Primary Storage System Categories and Subject Descriptors," in *Proceedings of the 2012 International Workshop on Programming Models and Applications for Multicores and Manycores*, PMAM'12, pp. 17–26, ACM, 2012.
- [93] M. L. Curry, H. L. Ward, A. Skjellum, and R. Brightwell, "A Lightweight, GPU-Based Software RAID System," in *Proceedings of the 39th International Conference on Parallel Processing*, ICPP'10, (San Diego, CA), pp. 565–572, IEEE, Sept. 2010.
- [94] D. A. Patterson, G. Gibson, and R. H. Katz, "A case for redundant arrays of inexpensive disks (RAID)," in *SIGMOD '88: Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data*, 1988.
- [95] C. Frost, M. Mammarella, E. Kohler, A. de los Reyes, S. Hovsepian, A. Matsuoka, and L. Zhang, "Generalized file system dependencies," SOSP 2007, ACM.
- [96] M. Silberstein, B. Ford, and E. Witchel, "GPUfs : Integrating a File System with GPUs," in *Proceedings of the 18th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS'13, pp. 485–497, 2013.
- [97] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," in *Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation*, OSDI'04, 2004.
- [98] M. D. Lieberman, J. Sankaranarayanan, and H. Samet, "A Fast Similarity Join Algorithm Using Graphics Processing Units," in *Proceedings of the 24th IEEE International Conference on Data Engineering*, no. April in ICDE'08, pp. 1111–1120, Apr. 2008.
- [99] B. He, K. Yang, R. Fang, M. Lu, N. K. Govindaraju, Q. Luo, and P. V. Sander, "Relational Joins on Graphics Processors," in *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD'08, (New York, USA), ACM Press, 2008.
- [100] N. K. Govindaraju, B. Lloyd, W. Wang, M. Lin, and D. Manocha, "Fast Computation of Database Operations using Graphics Processors," in *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, SIGMOD'04, June 2004.
- [101] P. B. Volk, D. Habich, and W. Lehner, "GPU-Based Speculative Query Processing for Database Operations," in *Proceedings of the 1st International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures*, ADMS'10, Sept. 2010.

- [102] G. Diamos, A. Lele, J. Wang, and S. Yalamanchili, "Relational Algorithms for Multi-Bulk-Synchronous Processors," tech. rep., NVIDIA, June 2012.
- [103] C.-F. Lin and S.-M. Yuan, "The Design and Evaluation of GPU Based Memory Database," in *Proceedings of the 2011 Fifth International Conference on Genetic and Evolutionary Computing, ICGEC '11*, pp. 224–231, 2011.
- [104] K. Kohei, "PG-Strom: A FDW Module of PostgreSQL using GPU for Asynchronous Super-Parallel Query Execution." https://github.com/kaigai/pg_strom.
- [105] "PostgreSQL Database." <http://www.postgresql.org/>.
- [106] P. Ghodsnia, "An In-GPU-Memory Column-Oriented Database for Processing Analytical Workloads," in *The VLDB 2012 PhD Workshop*, (Istanbul, Turkey), pp. 54–59, Aug. 2012.
- [107] T. H. Hetherington, T. G. Rogers, L. Hsu, M. O. Connor, and T. M. Aamodt, "Characterizing and Evaluating a Key-value Store Application on Heterogeneous CPU-GPU Systems," in *Proceedings of the 2012 IEEE International Symposium on Performance Analysis of Systems and Software*, ISPASS'12, pp. 88–98, Apr. 2012.
- [108] N. K. Govindaraju, J. Gray, R. Kumar, and D. Manocha, "GPUTeraSort : High Performance Graphics Co-processor Sorting for Large Database Management," in *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, SIGMOD'06, pp. 325–336, June 2006.
- [109] B. He, W. Fang, Q. Luo, N. K. Govindaraju, and T. Wang, "Mars : A MapReduce Framework on Graphics Processors," in *Proceedings of the 7th International Conference on Parallel Architectures and Compilation Techniques*, PACT'08, (Toronto, Canada), Oct. 2008.
- [110] J. A. Stuart and J. D. Owens, "Multi-GPU MapReduce on GPU Clusters," in *Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium*, IPDPS'11, IEEE, May 2011.
- [111] Y. Chen, Z. Qiao, H. Jiang, K.-C. Li, and W. W. Ro, "MGMR: Multi-GPU Based MapReduce," in *Proceedings of the 8th International Conference on Grid and Pervasive Computing*, GPC '13, pp. 433–442, Springer-Verlag, May 2013.
- [112] O. Harrison and J. Waldron, "Practical Symmetric Key Cryptography on Modern Graphics Hardware," in *Proceedings of the 17th USENIX Security Symposium*, pp. 195–209, 2008.
- [113] O. Harrison and J. Waldron, "Efficient Acceleration of Asymmetric Cryptography on Graphics Hardware," in *Proceedings of the 2nd International Conference on Cryptology in Africa: Progress in Cryptology*, AfricaCrypt'09, pp. 350–367, Springer-Verlag, June 2009.
- [114] A. Moss, D. Page, and N. P. Smart, "Toward Acceleration of RSA Using 3D Graphics Hardware," in *Proceedings of the 11th IMA International Conference on Cryptography and Coding*, pp. 369–388, Springer-Verlag, Dec. 2007.
- [115] G. Vasiliadis and S. Ioannidis, "GrAVity : A Massively Parallel Antivirus Engine," in *Proceedings of the 13th International Symposium On Recent Advances In Intrusion Detection*, RAID'10, Sept. 2010.

- [116] R. S. Boyer and J. S. Moore, “A Fast String Searching Algorithm,” *Commun. ACM*, vol. 20, pp. 762–772, Oct. 1977.
- [117] Kaspersky Lab, “Kaspersky Lab utilizes NVIDIA technologies to enhance protection.” <http://www.kaspersky.com/news?id=207575979>.
- [118] S. Berger, R. Cáceres, K. A. Goldman, R. Perez, R. Sailer, and L. van Doorn, “vTPM: virtualizing the trusted platform module,” *USENIX Security* 2006.
- [119] O. Shivers, *Control-Flow Analysis of Higher-Order Languages*. PhD thesis, Carnegie Mellon University, May 1991.
- [120] S. Liang, M. Might, and W. Sun, “Godel Hashes: Fast, Compact, Monotonic, Dynamic, Incremental and Perfect,” in *Submission*, 2014.
- [121] J. A. Storer and T. G. Szymanski, “Data Compression via Textual Substitution,” *Journal of ACM*, vol. 29, pp. 928–951, Oct. 1982.
- [122] A. Ozsoy, M. Swamy, and A. Chauhan, “Pipelined Parallel LZSS for Streaming Data Compression on GPGPUs,” in *Proceedings of the IEEE 18th International Conference on Parallel and Distributed Systems*, pp. 37–44, IEEE, Dec. 2012.
- [123] W. Fang, B. He, and Q. Luo, “Database Compression on Graphics Processors,” *Proceedings of the VLDB Endowment*, vol. 3, pp. 670–680, Sept. 2010.
- [124] P. Alsberg, “Space and time savings through large data base compression and dynamic restructuring,” *Proceedings of the IEEE*, vol. 63, pp. 1114–1122, Aug 1975.
- [125] X. Lin, G. Lu, F. Douglass, P. Shilane, and G. Wallace, “Migratory Compression: Coarse-grained Data Reordering to Improve Compressibility,” in *Proceedings of the 12th USENIX Conference on File and Storage Technologies (FAST)*, Feb. 2014.
- [126] D. H. Woo and H.-h. S. Lee, “COMPASS : A Programmable Data Prefetcher Using Idle GPU Shaders,” in *Proceedings of the 15th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS’10*, pp. 297–310, Mar. 2010.
- [127] D. Joseph and D. Grunwald, “Prefetching Using Markov Predictors,” in *Proceedings of the 24th Annual International Symposium on Computer Architecture, ISCA ’97*, pp. 252–263, 1997.
- [128] K. J. Nesbit and J. E. Smith, “Data Cache Prefetching Using a Global History Buffer,” in *Proceedings of the 10th International Symposium on High Performance Computer Architecture, HPCA ’04*, pp. 96–, IEEE Computer Society, 2004.
- [129] W.-f. Lin, “Reducing DRAM Latencies with an Integrated Memory Hierarchy Design,” in *Proceedings of the 7th International Symposium on High-Performance Computer Architecture, HPCA ’01*, pp. 301–, IEEE Computer Society, 2001.
- [130] D. Merrill, M. Garland, and A. Grimshaw, “Scalable GPU Graph Traversal,” in *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, vol. 47 of *PPoPP’12*, pp. 117–128, Sept. 2012.
- [131] M. Maas, P. Reames, J. Morlan, K. Asanović, A. D. Joseph, and J. Kubiawicz, “GPUs As an Opportunity for Offloading Garbage Collection,” in *Proceedings of the 2012 International Symposium on Memory Management, ISMM ’12*, pp. 25–36, 2012.

- [132] S. Hong, S. K. Kim, T. Oguntebi, and K. Olukotun, "Accelerating CUDA graph algorithms at maximum warp," *ACM SIGPLAN Notices*, vol. 46, no. 8, pp. 267–276, 2011.
- [133] Y. Fujii, T. Azumi, N. Nishio, S. Kato, and M. Edahiro, "Data Transfer Matters for GPU Computing," in *Proceedings of the 19th IEEE International Conference on Parallel and Distributed System*, 2013.
- [134] Y. Fujii, T. Azumi, N. Nishio, and S. Kato, "Exploring microcontrollers in GPUs," in *Proceedings of the 4th Asia-Pacific Workshop on Systems*, APSys '13, pp. 1–6, ACM Press, 2013.
- [135] NVIDIA, "RDMA for GPUDirect." <http://docs.nvidia.com/cuda/gpudirect-rdma>.
- [136] M. Dowty and J. Sugerman, "GPU Virtualization on VMware s Hosted I/O Architecture," *ACM SIGOPS Operating System Review*, vol. 43, pp. 73–82, July 2009.
- [137] V. Gupta, A. Gavrilovska, and P. Ranganathan, "GViM : GPU-accelerated Virtual Machines," in *Proceedings of the 3rd ACM Workshop on System-level Virtualization for High Performance Computing*, HPCVirt'09, (Nuremberg, Germany), ACM, Mar. 2009.
- [138] G. Giunta, R. Montella, G. Agrillo, and G. Coviello, "A GPGPU Transparent Virtualization Component for High Performance Computing Clouds," in *Proceedings of the 16th International Euro-Par Conference on Parallel Processing*, EuroPar'10, pp. 379–391, Springer-Verlag, 2010.
- [139] H. Takizawa, K. Koyama, K. Sato, K. Komatsu, and H. Kobayashi, "CheCL: Transparent Checkpointing and Process Migration of OpenCL Applications," in *Proceedings of the 2011 IEEE International Parallel & Distributed Processing Symposium*, IPDPS'11, pp. 864–876, IEEE, May 2011.
- [140] H. Takizawa, K. Sato, K. Komatsu, and H. Kobayashi, "CheCUDA: A Checkpoint/Restart Tool for CUDA Applications," in *Proceedings of the 10th International Conference on Parallel and Distributed Computing, Applications and Technologies*, PDCAT'09, pp. 408–413, IEEE, Dec. 2009.
- [141] AMD, "AMD Open Source." <http://developer.amd.com/tools-and-sdks/open-source/>.
- [142] P. Bright, "AMDs "heterogeneous uniform memory access" coming this year in Kaveri," Apr. 2013. <http://arstechnica.com/information-technology/2013/04/amds-heterogeneous-uniform-memory-access-coming-this-year-in-kaveri>.
- [143] K.-T. Cheng and Y.-C. Wang, "Using mobile GPU for general-purpose computing, a case study of face recognition on smartphones," in *Proceedings of the 2011 International Symposium on VLSI Design, Automation and Test*, pp. 1–4, Apr 2011.
- [144] G. Wang, Y. Xiong, J. Yun, and J. Cavallaro, "Accelerating computer vision algorithms using OpenCL framework on the mobile GPU - A case study," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2629–2633, May 2013.

- [145] M. Huang and C. Lai, “Accelerating Applications using GPUs in Embedded Systems and Mobile Devices,” in *Proceedings of the 15th IEEE International Conference on High Performance Computing and Communication (HPCC)*, pp. 1031–1038, Nov. 2013.
- [146] A. Maghazeh, U. D. Bordoloi, P. Eles, and Z. Peng, “General purpose computing on low-power embedded GPUs: Has it come of age?,” in *Proceedings of the International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIII)*, pp. 1–10, July 2013.