

Programming on Parallel Machines

Norman Matloff
University of California, Davis ¹

¹ **Licensing:** This work is licensed under a Creative Commons Attribution-No Derivative Works 3.0 United States License. Copyright is retained by N. Matloff in all non-U.S. jurisdictions, but permission to use these materials in teaching is still granted, provided the authorship and licensing information here is displayed in each unit. I would appreciate being notified if you use this book for teaching, just so that I know the materials are being put to use, but this is not required.

Author's Biographical Sketch

Dr. Norm Matloff is a professor of computer science at the University of California at Davis, and was formerly a professor of mathematics and statistics at that university. He is a former database software developer in Silicon Valley, and has been a statistical consultant for firms such as the Kaiser Permanente Health Plan.

Dr. Matloff was born in Los Angeles, and grew up in East Los Angeles and the San Gabriel Valley. He has a PhD in pure mathematics from UCLA, specializing in probability theory. He has published numerous papers in computer science and statistics, including the *ACM Trans. on Database Systems*, the *ACM Trans. on Modeling and Computer Simulation*, the *Annals of Probability*, *Biometrika*, the *Communications of the ACM*, the *IEEE/ACM Trans. on Networking*, the *IEEE Trans. on Data Engineering*, the *IEEE Trans. on Communications*, and the *IEEE Trans. on Reliability*, and the University of Michigan J. of Law Reform, as well as highly selective conferences such as the ACM International Conference on Supercomputing, INFOCOM, the International Conference on Data Engineering, the SIAM Conference on Data Mining and so on.

Prof. Matloff is a former appointed member of IFIP Working Group 11.3, an international committee concerned with database software security, established under UNESCO. He was a founding member of the UC Davis Department of Statistics, and participated in the formation of the UCD Computer Science Department as well. He is a recipient of the campuswide Distinguished Teaching Award at UC Davis, as well as the departmental teaching award.

Dr. Matloff is the author of two published textbooks, and of a number of widely-used Web tutorials on computer topics, such as the Linux operating system and the Python programming language. He and Dr. Peter Salzman are authors of *The Art of Debugging with GDB, DDD, and Eclipse*. Prof. Matloff's book on the R programming language, *The Art of R Programming*, is due to be published in 2011. He is also the author of several open-source textbooks, including *From Algorithms to Z-Scores: Probabilistic and Statistical Modeling in Computer Science* (<http://heather.cs.ucdavis.edu/probstatbook>), and *Programming on Parallel Machines* (<http://heather.cs.ucdavis.edu/~matloff/ParProcBook.pdf>).

Contents

1	Introduction to Parallel Processing	1
1.1	Overview: Why Use Parallel Systems?	1
1.1.1	Execution Speed	1
1.1.2	Memory	2
1.2	Parallel Processing Hardware	2
1.2.1	Shared-Memory Systems	2
1.2.1.1	Basic Architecture	2
1.2.1.2	Example: SMP Systems	3
1.2.2	Message-Passing Systems	4
1.2.2.1	Basic Architecture	4
1.2.2.2	Example: Networks of Workstations (NOWs)	4
1.2.3	SIMD	4
1.3	Programmer World Views	5
1.3.1	Shared-Memory	5
1.3.1.1	Programmer View	5
1.3.1.2	Example	5
1.3.2	Message Passing	10
1.3.2.1	Programmer View	10
1.3.3	Example	11

1.4	Relative Merits: Shared-Memory Vs. Message-Passing	14
1.5	Issues in Parallelizing Applications	15
1.5.1	Communication Bottlenecks	15
1.5.2	Load Balancing	15
1.5.3	“Embarrassingly Parallel” Applications	15
2	Shared Memory Parallelism	17
2.1	What Is Shared?	17
2.2	Memory Modules	18
2.2.1	Interleaving	18
2.2.2	Bank Conflicts and Solutions	19
2.3	Interconnection Topologies	20
2.3.1	SMP Systems	20
2.3.2	NUMA Systems	21
2.3.3	NUMA Interconnect Topologies	22
2.3.3.1	Crossbar Interconnects	22
2.3.3.2	Omega (or Delta) Interconnects	24
2.3.4	Comparative Analysis	25
2.3.5	Why Have Memory in Modules?	26
2.4	Test-and-Set Type Instructions	27
2.5	Cache Issues	28
2.5.1	Cache Coherency	28
2.5.2	Example: the MESI Cache Coherency Protocol	31
2.5.3	The Problem of “False Sharing”	33
2.6	Memory-Access Consistency Policies	33
2.7	Fetch-and-Add and Packet-Combining Operations	35
2.8	Multicore Chips	37

2.9	Illusion of Shared-Memory through Software	37
2.9.0.1	Software Distributed Shared Memory	37
2.9.0.2	Case Study: JIAJIA	39
2.10	Barrier Implementation	43
2.10.1	A Use-Once Version	43
2.10.2	An Attempt to Write a Reusable Version	44
2.10.3	A Correct Version	44
2.10.4	Refinements	45
2.10.4.1	Use of Wait Operations	45
2.10.4.2	Parallelizing the Barrier Operation	47
2.10.4.2.1	Tree Barriers	47
2.10.4.2.2	Butterfly Barriers	47
3	The Python Threads and Multiprocessing Modules	49
3.1	Python Threads Modules	49
3.1.1	The <code>thread</code> Module	49
3.1.2	The <code>threading</code> Module	58
3.2	Condition Variables	62
3.2.1	General Ideas	62
3.2.2	Event Example	63
3.2.3	Other <code>threading</code> Classes	65
3.3	Threads Internals	65
3.3.1	Kernel-Level Thread Managers	66
3.3.2	User-Level Thread Managers	66
3.3.3	Comparison	66
3.3.4	The Python Thread Manager	66
3.3.4.1	The GIL	67

3.3.4.2	Implications for Randomness and Need for Locks	68
3.4	The multiprocessing Module	68
3.5	The Queue Module for Threads and Multiprocessing	71
3.6	Debugging Threaded and Multiprocessing Python Programs	74
3.6.1	Using PDB to Debug Threaded Programs	75
3.6.2	RPDB2 and Winpdb	76
4	Introduction to OpenMP	77
4.1	Overview	77
4.2	Running Example	77
4.2.1	The Algorithm	80
4.2.2	The OpenMP parallel Pragma	80
4.2.3	Scope Issues	81
4.2.4	The OpenMP single Pragma	82
4.2.5	The OpenMP barrier Pragma	82
4.2.6	Implicit Barriers	82
4.2.7	The OpenMP critical Pragma	83
4.3	The OpenMP for Pragma	83
4.3.1	Basic Example	83
4.3.2	Nested Loops	86
4.3.3	Controlling the Partitioning of Work to Threads	86
4.3.4	The OpenMP reduction Clause	87
4.4	The Task Directive	88
4.5	Other OpenMP Synchronization Issues	90
4.5.1	The OpenMP atomic Clause	90
4.5.2	Memory Consistency and the flush Pragma	91
4.6	Compiling, Running and Debugging OpenMP Code	92

4.6.1	Compiling	92
4.6.2	Running	92
4.6.3	Debugging	92
4.7	Combining Work-Sharing Constructs	93
4.8	Performance	93
4.8.1	The Effect of Problem Size	93
4.8.2	Some Fine Tuning	94
4.8.3	OpenMP Internals	98
4.9	The Rest of OpenMP	98
4.10	Further Examples	99
5	Introduction to GPU Programming with CUDA	101
5.1	Overview	101
5.2	Sample Program	102
5.3	Understanding the Hardware Structure	105
5.3.1	Processing Units	105
5.3.2	Thread Operation	106
5.3.2.1	SIMT Architecture	106
5.3.2.2	The Problem of Thread Divergence	106
5.3.2.3	“OS in Hardware”	107
5.3.3	Memory Structure	107
5.3.3.1	Shared and Global Memory	107
5.3.3.2	Global-Memory Performance Issues	111
5.3.3.3	Shared-Memory Performance Issues	111
5.3.3.4	Host/Device Memory Transfer Performance Issues	112
5.3.3.5	Other Types of Memory	112
5.3.4	Threads Hierarchy	113

5.3.5	What's NOT There	115
5.4	Synchronization	115
5.5	Hardware Requirements, Installation, Compilation, Debugging	116
5.6	Improving the Sample Program	117
5.7	More Examples	118
5.7.1	Finding the Mean Number of Mutual Outlinks	118
5.7.2	Finding Prime Numbers	120
5.8	CUBLAS	123
5.9	Error Checking	125
5.10	The New Generation	126
5.11	Further Examples	126
6	Message Passing Systems	127
6.1	Overview	127
6.2	A Historical Example: Hypercubes	128
6.2.0.0.1	Definitions	128
6.3	Networks of Workstations (NOWs)	130
6.3.1	The Network Is Literally the Weakest Link	130
6.3.2	Other Issues	131
6.4	Systems Using Nonexplicit Message-Passing	131
6.4.1	MapReduce	131
7	Introduction to MPI	135
7.1	Overview	135
7.1.1	History	135
7.1.2	Structure and Execution	136
7.1.3	Implementations	136

7.1.4	Performance Issues	136
7.2	Running Example	137
7.2.1	The Algorithm	137
7.2.2	The Code	137
7.2.3	Introduction to MPI APIs	141
7.2.3.1	MPI_Init() and MPI_Finalize()	141
7.2.3.2	MPI_Comm_size() and MPI_Comm_rank()	141
7.2.3.3	MPI_Send()	141
7.2.3.4	MPI_Recv()	142
7.3	Collective Communications	143
7.3.1	Example	143
7.3.2	MPI_Bcast()	145
7.3.2.1	MPI_Reduce()/MPI_Allreduce()	146
7.3.2.2	MPI_Gather()/MPI_Allgather()	147
7.3.2.3	The MPI_Scatter()	148
7.3.2.4	The MPI_Barrier()	148
7.3.3	Creating Communicators	148
7.4	Buffering, Synchrony and Related Issues	149
7.4.1	Buffering, Etc.	149
7.4.2	Safety	151
7.4.3	Living Dangerously	151
7.4.4	Safe Exchange Operations	151
7.5	Use of MPI from Other Languages	152
7.5.1	Python: pyMPI	152
7.5.2	R	154
8	Introduction to Parallel Matrix Operations	155

8.1	It's Not Just Physics Anymore	155
8.2	CUBLAS	156
8.3	Partitioned Matrices	156
8.4	Matrix Multiplication	157
8.4.1	Message-Passing Case	158
8.4.1.1	Fox's Algorithm	158
8.4.1.2	Performance Issues	159
8.4.2	Shared-Memory Case	160
8.4.2.1	OpenMP	160
8.4.2.2	CUDA	160
8.4.3	Finding Powers of Matrices	163
8.5	Solving Systems of Linear Equations	164
8.5.1	Gaussian Elimination	164
8.5.2	The Jacobi Algorithm	165
8.6	OpenMP Implementation of the Jacobi Algorithm	166
8.7	Matrix Inversion	167
8.7.1	Using the Methods for Solving Systems of Linear Equations	167
8.7.2	Power Series Method	167
9	Parallel Combinatorial Algorithms	169
9.1	Overview	169
9.2	The 8 Queens Problem	169
9.3	The 8-Square Puzzle Problem	170
9.4	Itemset Analysis in Data Mining	172
9.4.1	What Is It?	172
9.4.2	The Market Basket Problem	173
9.4.3	Serial Algorithms	173

9.4.4	Parallelizing the Apriori Algorithm	174
10	Introduction to Parallel Sorting	175
10.1	Quicksort	175
10.1.1	Shared-Memory Quicksort	176
10.1.2	Hyperquicksort	177
10.2	Mergesorts	178
10.2.1	Sequential Form	178
10.2.2	Shared-Memory Mergesort	178
10.2.3	Message Passing Mergesort on a Tree Topology	178
10.2.4	Compare-Exchange Operations	179
10.2.5	Bitonic Mergesort	179
10.3	The Bubble Sort and Its Cousins	181
10.3.1	The Much-Maligned Bubble Sort	181
10.3.2	A Popular Variant: Odd-Even Transposition	182
10.3.3	CUDA Implementation of Odd/Even Transposition Sort	182
10.4	Shearsort	184
10.5	Bucket Sort with Sampling	184
10.6	Enumeration Sort	185
11	Parallel Computation of Fourier Series, with an Introduction to Parallel Imaging	187
11.1	General Principles	187
11.1.1	One-Dimensional Fourier Series	187
11.1.2	Two-Dimensional Fourier Series	191
11.2	Discrete Fourier Transforms	191
11.2.1	One-Dimensional Data	192
11.2.2	Two-Dimensional Data	193

11.3	Parallel Computation of Discrete Fourier Transforms	193
11.3.1	CUFFT	193
11.3.2	The Fast Fourier Transform	193
11.3.3	A Matrix Approach	194
11.3.4	Parallelizing Computation of the Inverse Transform	194
11.3.5	Parallelizing Computation of the Two-Dimensional Transform	195
11.4	Applications to Image Processing	195
11.4.1	Smoothing	195
11.4.2	Edge Detection	196
11.5	The Cosine Transform	197
11.6	Keeping the Pixel Intensities in the Proper Range	198
11.7	Does the Function $g()$ Really Have to Be Repeating?	198
11.8	Vector Space Issues (optional section)	199
11.9	Bandwidth: How to Read the <i>San Francisco Chronicle</i> Business Page (optional section)	200
12	Applications to Statistics/Data Mining	203
12.1	Itemset Analysis	203
12.1.1	What Is It?	203
12.1.2	The Market Basket Problem	204
12.1.3	Serial Algorithms	205
12.1.4	Parallelizing the Apriori Algorithm	205
12.2	Probability Density Estimation	206
12.2.1	Kernel-Based Density Estimation	206
12.2.2	Histogram Computation for Images	209
12.3	Clustering	210
12.4	Principal Component Analysis (PCA)	212
12.5	Parallel Processing in R	213

12.5.1 Rmpi	213
12.5.2 The R snow Package	214
12.5.3 Rdsm	217
12.5.4 R with GPUs	219
12.5.4.1 The gputools Package	219
12.5.4.2 The rgpu Package	220
12.5.4.3 Debugging R Applications	221
A Review of Matrix Algebra	223
A.1 Terminology and Notation	223
A.1.1 Matrix Addition and Multiplication	224
A.2 Matrix Transpose	225
A.3 Linear Independence	225
A.4 Determinants	226
A.5 Matrix Inverse	226
A.6 Eigenvalues and Eigenvectors	226

Chapter 1

Introduction to Parallel Processing

Parallel machines provide a wonderful opportunity for applications with large computational requirements. Effective use of these machines, though, requires a keen understanding of how they work. This chapter provides an overview.

1.1 Overview: Why Use Parallel Systems?

1.1.1 Execution Speed

There is an ever-increasing appetite among some types of computer users for faster and faster machines. This was epitomized in a statement by Steve Jobs, founder/CEO of Apple and Pixar. He noted that when he was at Apple in the 1980s, he was always worried that some other company would come out with a faster machine than his. But now at Pixar, whose graphics work requires extremely fast computers, he is always hoping someone produces faster machines, so that he can use them!

A major source of speedup is the parallelizing of operations. Parallel operations can be either within-processor, such as with pipelining or having several ALUs within a processor, or between-processor, in which many processor work on different parts of a problem in parallel. Our focus here is on between-processor operations.

For example, the Registrar's Office at UC Davis uses shared-memory multiprocessors for processing its on-line registration work. Online registration involves an enormous amount of database computation. In order to handle this computation reasonably quickly, the program partitions the work to be done, assigning different portions of the database to different processors. The database field has contributed greatly to the commercial success of large shared-memory machines.

As the Pixar example shows, highly computation-intensive applications like computer graphics also have a

need for these fast parallel computers. No one wants to wait hours just to generate a single image, and the use of parallel processing machines can speed things up considerably. For example, consider **ray tracing** operations. Here our code follows the path of a ray of light in a scene, accounting for reflection and absorption of the light by various objects. Suppose the image is to consist of 1,000 rows of pixels, with 1,000 pixels per row. In order to attack this problem in a parallel processing manner with, say, 25 processors, we could divide the image into 25 squares of size 200x200, and have each processor do the computations for its square.

Note, though, that it may be much more challenging than this implies. First of all, the computation will need some communication between the processors, which hinders performance if it is not done carefully. Second, if one really wants good speedup, one may need to take into account the fact that some squares require more computation work than others. More on this below.

1.1.2 Memory

Yes, execution speed is the reason that comes to most people's minds when the subject of parallel processing comes up. But in many applications, an equally important consideration is memory capacity. Parallel processing application often tend to use huge amounts of memory, and in many cases the amount of memory needed is more than can fit on one machine. If we have many machines working together, especially in the message-passing settings described below, we can accommodate the large memory needs.

1.2 Parallel Processing Hardware

This is not a hardware course, but since the goal of using parallel hardware is speed, the efficiency of our code is a major issue. That in turn means that we need a good understanding of the underlying hardware that we are programming. In this section, we give an overview of parallel hardware.

1.2.1 Shared-Memory Systems

1.2.1.1 Basic Architecture

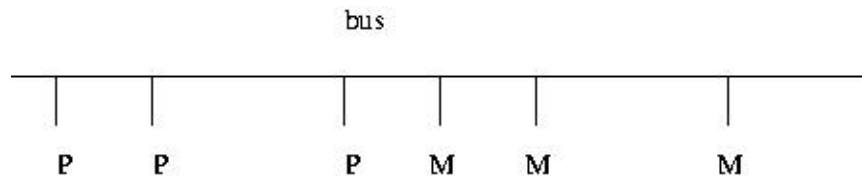
Here many CPUs share the same physical memory. This kind of architecture is sometimes called MIMD, standing for Multiple Instruction (different CPUs are working independently, and thus typically are executing different instructions at any given instant), Multiple Data (different CPUs are generally accessing different memory locations at any given time).

Until recently, shared-memory systems cost hundreds of thousands of dollars and were affordable only by large companies, such as in the insurance and banking industries. The high-end machines are indeed still

quite expensive, but now **dual-core** machines, in which two CPUs share a common memory, are commonplace in the home.

1.2.1.2 Example: SMP Systems

A Symmetric Multiprocessor (SMP) system has the following structure:



Here and below:

- The Ps are processors, e.g. off-the-shelf chips such as Pentiums.
- The Ms are **memory modules**. These are physically separate objects, e.g. separate boards of memory chips. It is typical that there will be the same number of memory modules as processors. In the shared-memory case, the memory modules collectively form the entire shared address space, but with the addresses being assigned to the memory modules in one of two ways:
 - (a)
 - High-order interleaving. Here consecutive addresses are in the same M (except at boundaries). For example, suppose for simplicity that our memory consists of addresses 0 through 1023, and that there are four Ms. Then M0 would contain addresses 0-255, M1 would have 256-511, M2 would have 512-767, and M3 would have 768-1023.
 - We need 10 bits for addresses (since $1024 = 2^{10}$). The two most-significant bits would be used to select the module number (since $4 = 2^2$); hence the term *high-order* in the name of this design. The remaining eight bits are used to select the word within a module.
 - (b)
 - Low-order interleaving. Here consecutive addresses are in consecutive memory modules (except when we get to the right end). In the example above, if we used low-order interleaving, then address 0 would be in M0, 1 would be in M1, 2 would be in M2, 3 would be in M3, 4 would be back in M0, 5 in M1, and so on.
 - Here the two least-significant bits are used to determine the module number.
- To make sure only one processor uses the bus at a time, standard bus arbitration signals and/or arbitration devices are used.
- There may also be **coherent caches**, which we will discuss later.

1.2.2 Message-Passing Systems

1.2.2.1 Basic Architecture

Here we have a number of independent CPUs, each with its own independent memory. The various processors communicate with each other via networks of some kind.

1.2.2.2 Example: Networks of Workstations (NOWs)

Large shared-memory multiprocessor systems are still very expensive. A major alternative today is networks of workstations (NOWs). Here one purchases a set of commodity PCs and networks them for use as parallel processing systems. The PCs are of course individual machines, capable of the usual uniprocessor (or now multiprocessor) applications, but by networking them together and using parallel-processing software environments, we can form very powerful parallel systems.

The networking does result in a significant loss of performance. This will be discussed in Chapter 6. But even without these techniques, the price/performance ratio in NOW is much superior in many applications to that of shared-memory hardware.

One factor which can be key to the success of a NOW is the use of a fast network, fast both in terms of hardware and network protocol. Ordinary Ethernet and TCP/IP are fine for the applications envisioned by the original designers of the Internet, e.g. e-mail and file transfer, but is slow in the NOW context. A good network for a NOW is, for instance, Infiniband.

NOWs have become so popular that there are now “recipes” on how to build them for the specific purpose of parallel processing. The term **Beowulf** came to mean a cluster of PCs, usually with a fast network connecting them, used for parallel processing. Software packages such as ROCKS (<http://www.rocksclusters.org/wordpress/>) have been developed to make it easy to set up and administer such systems.

1.2.3 SIMD

In contrast to MIMD systems, processors in SIMD—Single Instruction, Multiple Data—systems execute in lockstep. At any given time, all processors are executing the same machine instruction on different data.

Some famous SIMD systems in computer history include the ILLIAC and Thinking Machines Corporation’s CM-1 and CM-2. Also, DSP (“digital signal processing”) chips tend to have an SIMD architecture.

But today the most prominent example of SIMD is that of GPUs—graphics processing units. In addition to powering your PC’s video cards, GPUs can now be used for general-purpose computation. The architecture is fundamentally shared-memory, but the individual processors do execute in lockstep, SIMD-fashion.

1.3 Programmer World Views

To explain the two paradigms, we will use the term **nodes**, where roughly speaking one node corresponds to one processor, and use the following example:

Suppose we wish to multiply an $n \times 1$ vector X by an $n \times n$ matrix A , putting the product in an $n \times 1$ vector Y , and we have p processors to share the work.

1.3.1 Shared-Memory

1.3.1.1 Programmer View

In the shared-memory paradigm, the arrays for A , X and Y would be held in common by all nodes. If for instance node 2 were to execute

```
Y[3] = 12;
```

and then node 15 were to subsequently execute

```
print ("%d\n", Y[3]);
```

then the outputted value from the latter would be 12.

1.3.1.2 Example

Today, programming on shared-memory multiprocessors is typically done via **threading**. (Or, as we will see in other chapters, by higher-level code that runs threads underneath.) A **thread** is similar to a **process** in an operating system (OS), but with much less overhead. Threaded applications have become quite popular in even uniprocessor systems, and Unix,¹ Windows, Python, Java and Perl all support threaded programming.

In the typical implementation, a thread is a special case of an OS process. One important difference is that the various threads of a program share memory. (One can arrange for processes to share memory too in some OSs, but they don't do so by default.)

On a uniprocessor system, the threads of a program take turns executing, so that there is only an illusion of parallelism. But on a multiprocessor system, one can genuinely have threads running in parallel.

¹Here and below, the term *Unix* includes Linux.

One of the most popular threads systems is Pthreads, whose name is short for POSIX threads. POSIX is a Unix standard, and the Pthreads system was designed to standardize threads programming on Unix. It has since been ported to other platforms.

Following is an example of Pthreads programming, in which we determine the number of prime numbers in a certain range. Read the comments at the top of the file for details; the threads operations will be explained presently.

```

1 // PrimesThreads.c
2
3 // threads-based program to find the number of primes between 2 and n;
4 // uses the Sieve of Eratosthenes, deleting all multiples of 2, all
5 // multiples of 3, all multiples of 5, etc.
6
7 // for illustration purposes only; NOT claimed to be efficient
8
9 // Unix compilation: gcc -g -o primesthreads PrimesThreads.c -lpthread -lm
10
11 // usage: primesthreads n num_threads
12
13 #include <stdio.h>
14 #include <math.h>
15 #include <pthread.h> // required for threads usage
16
17 #define MAX_N 100000000
18 #define MAX_THREADS 25
19
20 // shared variables
21 int nthreads, // number of threads (not counting main())
22     n, // range to check for primeness
23     prime[MAX_N+1], // in the end, prime[i] = 1 if i prime, else 0
24     nextbase; // next sieve multiplier to be used
25 // lock for the shared variable nextbase
26 pthread_mutex_t nextbaselock = PTHREAD_MUTEX_INITIALIZER;
27 // ID structs for the threads
28 pthread_t id[MAX_THREADS];
29
30 // "crosses out" all odd multiples of k
31 void crossout(int k)
32 { int i;
33   for (i = 3; i*k <= n; i += 2) {
34     prime[i*k] = 0;
35   }
36 }
37
38 // each thread runs this routine
39 void *worker(int tn) // tn is the thread number (0,1,...)
40 { int lim,base,
41     work = 0; // amount of work done by this thread
42   // no need to check multipliers bigger than sqrt(n)
43   lim = sqrt(n);
44   do {
45     // get next sieve multiplier, avoiding duplication across threads
46     // lock the lock
47     pthread_mutex_lock(&nextbaselock);

```

```

48     base = nextbase;
49     nextbase += 2;
50     // unlock
51     pthread_mutex_unlock(&nextbaselock);
52     if (base <= lim) {
53         // don't bother crossing out if base known composite
54         if (prime[base]) {
55             crossout(base);
56             work++; // log work done by this thread
57         }
58     }
59     else return work;
60 } while (1);
61 }
62
63 main(int argc, char **argv)
64 { int nprimes, // number of primes found
65   i,work;
66   n = atoi(argv[1]);
67   nthreads = atoi(argv[2]);
68   // mark all even numbers nonprime, and the rest "prime until
69   // shown otherwise"
70   for (i = 3; i <= n; i++) {
71       if (i%2 == 0) prime[i] = 0;
72       else prime[i] = 1;
73   }
74   nextbase = 3;
75   // get threads started
76   for (i = 0; i < nthreads; i++) {
77       // this call says to create a thread, record its ID in the array
78       // id, and get the thread started executing the function worker(),
79       // passing the argument i to that function
80       pthread_create(&id[i],NULL,worker,i);
81   }
82
83   // wait for all done
84   for (i = 0; i < nthreads; i++) {
85       // this call said to wait until thread number id[i] finishes
86       // execution, and to assign the return value of that thread to our
87       // local variable work here
88       pthread_join(id[i],&work);
89       printf("%d values of base done\n",work);
90   }
91
92   // report results
93   nprimes = 1;
94   for (i = 3; i <= n; i++)
95       if (prime[i]) {
96           nprimes++;
97       }
98   printf("the number of primes found was %d\n",nprimes);
99
100 }

```

To make our discussion concrete, suppose we are running this program with two threads. Suppose also the both threads are running simultaneously most of the time. This will occur if they aren't competing for turns

with other big threads, say if there are no other big threads, or more generally if the number of other big threads is less than or equal to the number of processors minus two.

Note the global variables:

```
int nthreads, // number of threads (not counting main())
    n, // range to check for primeness
    prime[MAX_N+1], // in the end, prime[i] = 1 if i prime, else 0
    nextbase; // next sieve multiplier to be used
pthread_mutex_t nextbaselock = PTHREAD_MUTEX_INITIALIZER;
pthread_t id[MAX_THREADS];
```

This will require some adjustment for those who've been taught that global variables are “evil.” All communication between threads is via global variables, so if they are evil, they are a necessary evil. Personally I think the stern admonitions against global variables is overblown anyway. See <http://heather.cs.ucdavis.edu/~matloff/globals.html>.

As mentioned earlier, the globals are shared by all processors.² If one processor, for instance, assigns the value 0 to `prime[35]` in the function `crossout()`, then that variable will have the value 0 when accessed by any of the other processors as well. On the other hand, local variables have different values at each processor; for instance, the variable `i` in that function has a different value at each processor.

Note that in the statement

```
pthread_mutex_t nextbaselock = PTHREAD_MUTEX_INITIALIZER;
```

the right-hand side is not a constant. It is a macro call, and is thus something which is executed.

In the code

```
pthread_mutex_lock(&nextbaselock);
base = nextbase
nextbase += 2
pthread_mutex_unlock(&nextbaselock);
```

we see a **critical section** operation which is typical in shared-memory programming. In this context here, it means that we cannot allow more than one thread to execute

```
base = nextbase;
nextbase += 2;
```

²Technically, we should say “shared by all threads” here, as a given thread does not always execute on the same processor, but at any instant in time each executing thread is at some processor, so the statement is all right.

at the same time. The calls to `pthread_mutex_lock()` and `pthread_mutex_unlock()` ensure this. If thread A is currently executing inside the critical section and thread B tries to lock the lock by calling `pthread_mutex_lock()`, the call will block until thread B executes `pthread_mutex_unlock()`.

Here is why this is so important: Say currently `nextbase` has the value 11. What we want to happen is that the next thread to read `nextbase` will “cross out” all multiples of 11. But if we allow two threads to execute the critical section at the same time, the following may occur:

- thread A reads `nextbase`, setting its value of `base` to 11
- thread B reads `nextbase`, setting its value of `base` to 11
- thread A adds 2 to `nextbase`, so that `nextbase` becomes 13
- thread B adds 2 to `nextbase`, so that `nextbase` becomes 15

Two problems would then occur:

- Both threads would do “crossing out” of multiples of 11, duplicating work and thus slowing down execution speed.
- We will never “cross out” multiples of 13.

Thus the lock is crucial to the correct (and speedy) execution of the program.

Note that these problems could occur either on a uniprocessor or multiprocessor system. In the uniprocessor case, thread A’s turn might end right after it reads `nextbase`, followed by a turn by B which executes that same instruction. In the multiprocessor case, A and B could literally be running simultaneously, but still with the action by B coming an instant after A.

This problem frequently arises in parallel database systems. For instance, consider an airline reservation system. If a flight has only one seat left, we want to avoid giving it to two different customers who might be talking to two agents at the same time. The lines of code in which the seat is finally assigned (the **commit** phase, in database terminology) is then a critical section.

A critical section is always a potential bottleneck in a parallel program, because its code is serial instead of parallel. In our program here, we may get better performance by having each thread work on, say, five values of `nextbase` at a time. Our line

```
nextbase += 2;
```

would become

```
nextbase += 10;
```

That would mean that any given thread would need to go through the critical section only one-fifth as often, thus greatly reducing overhead. On the other hand, near the end of the run, this may result in some threads being idle while other threads still have a lot of work to do.

Note this code.

```
for (i = 0; i < nthreads; i++) {
    pthread_join(id[i], &work);
    printf("%d values of base done\n", work);
}
```

This is a special case of of **barrier**.

A **barrier** is a point in the code that all threads must reach before continuing. In this case, a barrier is needed in order to prevent premature execution of the later code

```
for (i = 3; i <= n; i++)
    if (prime[i]) {
        nprimes++;
    }
```

which would result in possibly wrong output if we start counting primes before some threads are done.

The **pthread_join()** function actually causes the given thread to exit, so that we then “join” the thread that created it, i.e. **main()**. Thus some may argue that this is not really a true barrier.

Barriers are very common in shared-memory programming, and will be discussed in more detail in Chapter 2.

1.3.2 Message Passing

1.3.2.1 Programmer View

By contrast, in the message-passing paradigm, all nodes would have separate copies of A, X and Y. In this case, in our example above, in order for node 2 to send this new value of Y[3] to node 15, it would have to execute some special function, which would be something like

```
send(15, 12, "Y[3]");
```

and node 15 would have to execute some kind of **receive()** function.

1.3.3 Example

Here we use the MPI system, with our hardware being a NOW.

MPI is a popular public-domain set of interface functions, callable from C/C++, to do message passing. We are again counting primes, though in this case using a **pipelining** method. It is similar to hardware pipelines, but in this case it is done in software, and each “stage” in the pipe is a different computer.

The program is self-documenting, via the comments.

```

1
2  /* MPI sample program; NOT INTENDED TO BE EFFICIENT as a prime
3     finder, either in algorithm or implementation
4
5     MPI (Message Passing Interface) is a popular package using
6     the "message passing" paradigm for communicating between
7     processors in parallel applications; as the name implies,
8     processors communicate by passing messages using "send" and
9     "receive" functions
10
11     finds and reports the number of primes less than or equal to N
12
13     uses a pipeline approach:  node 0 looks at all the odd numbers
14     (i.e. has already done filtering out of multiples of 2) and
15     filters out those that are multiples of 3, passing the rest
16     to node 1; node 1 filters out the multiples of 5, passing
17     the rest to node 2; in this simple example, we just have node
18     2 filter out all the rest and then report the number of primes
19
20     note that we should NOT have a node run through all numbers
21     before passing them on to the next node, since we would then
22     have no parallelism at all; on the other hand, passing on just
23     one number at a time isn't efficient either, due to the high
24     overhead of sending a message if it is a network (tens of
25     microseconds until the first bit reaches the wire, due to
26     software delay); thus efficiency would be greatly improved if
27     each node saved up a chunk of numbers before passing them to
28     the next node */
29
30 // this include file is mandatory
31 #include <mpi.h>
32
33 #define MAX_N 100000
34 #define PIPE_MSG 0 // type of message containing a number to
35                   // be checked
36 #define END_MSG 1 // type of message indicating no more data will
37                   // be coming
38
39 int NNodes, /* number of nodes in computation*/
40     N, /* find all primes from 2 to N */
41     Me, /* my node number */
42     ToCheck; /* current number to check for passing on to next node;
43              // stylistically this might be nicer as a local in
44              // Node*(), but I have placed it here to dramatize
45              // the fact that the globals are NOT shared among

```

```

46         the nodes */
47
48 double T1,T2; /* start and finish times */
49
50 Init (Argc,Argv)
51     int Argc;  char **Argv;
52
53 { int DebugWait;
54
55     N = atoi(Argv[1]);
56     DebugWait = atoi(Argv[2]);
57
58     /* this loop is here to synchronize all nodes for debugging;
59        if DebugWait is specified as 1 on the command line, all nodes
60        wait here until the debugging programmer starts GDB at all
61        nodes and within GDB sets DebugWait to 0 to then proceed */
62     while (DebugWait) ;
63
64     /* mandatory to begin any MPI program */
65     MPI_Init(&Argc,&Argv);
66
67     /* puts the number of nodes in NNodes */
68     MPI_Comm_size(MPI_COMM_WORLD,&NNodes);
69     /* puts the node number of this node in Me */
70     MPI_Comm_rank(MPI_COMM_WORLD,&Me);
71
72     /* OK, get started; first record current time in T1 */
73     if (Me == 2) T1 = MPI_Wtime();
74 }
75
76 Node0()
77
78 { int I,Dummy,
79     Error; /* not checked in this example */
80     for (I = 1; I <= N/2; I++) {
81         ToCheck = 2 * I + 1;
82         if (ToCheck > N) break;
83         /* MPI_Send -- send a message
84            parameters:
85            pointer to place where message is to be drawn from
86            number of items in message
87            item type
88            destination node
89            message type ("tag") programmer-defined
90            node group number (in this case all nodes) */
91         if (ToCheck % 3 > 0)
92             Error = MPI_Send(&ToCheck,1,MPI_INT,1,PIPE_MSG,MPI_COMM_WORLD);
93     }
94     Error = MPI_Send(&Dummy,1,MPI_INT,1,END_MSG,MPI_COMM_WORLD);
95 }
96
97 Node1()
98
99 { int Error, /* not checked in this example */
100     Dummy;
101     MPI_Status Status; /* see below */
102
103     while (1) {

```

```

104     /* MPI_Recv -- receive a message
105         parameters:
106         pointer to place to store message
107         number of items in message (see notes on
108             this at the end of this file)
109         item type
110         accept message from which node(s)
111         message type ("tag"), programmer-defined (in this
112             case any type)
113         node group number (in this case all nodes)
114         status (see notes on this at the end of this file) */
115     Error = MPI_Recv(&ToCheck,1,MPI_INT,0,MPI_ANY_TAG,
116         MPI_COMM_WORLD,&Status);
117     if (Status.MPI_TAG == END_MSG) break;
118     if (ToCheck % 5 > 0)
119         Error = MPI_Send(&ToCheck,1,MPI_INT,2,PIPE_MSG,MPI_COMM_WORLD);
120 }
121 /* now send our end-of-data signal, which is conveyed in the
122     message type, not the message (we have a dummy message just
123     as a placeholder */
124 Error = MPI_Send(&Dummy,1,MPI_INT,2,END_MSG,MPI_COMM_WORLD);
125 }
126
127 Node2()
128
129 { int ToCheck, /* current number to check from Node 0 */
130     Error, /* not checked in this example */
131     PrimeCount,I,IsComposite;
132 MPI_Status Status; /* see below */
133
134 PrimeCount = 3; /* must account for the primes 2, 3 and 5, which
135     won't be detected below */
136 while (1) {
137     Error = MPI_Recv(&ToCheck,1,MPI_INT,1,MPI_ANY_TAG,
138         MPI_COMM_WORLD,&Status);
139     if (Status.MPI_TAG == END_MSG) break;
140     IsComposite = 0;
141     for (I = 7; I*I <= ToCheck; I += 2)
142         if (ToCheck % I == 0) {
143             IsComposite = 1;
144             break;
145         }
146     if (!IsComposite) PrimeCount++;
147 }
148 /* check the time again, and subtract to find run time */
149 T2 = MPI_Wtime();
150 printf("elapsed time = %f\n", (float)(T2-T1));
151 /* print results */
152 printf("number of primes = %d\n",PrimeCount);
153 }
154
155 main(argc,argv)
156     int argc; char **argv;
157
158 { Init(argc,argv);
159     /* note: instead of having a switch statement, we could write
160         three different programs, each running on a different node */
161     switch (Me) {

```

```

162     case 0: Node0();
163         break;
164     case 1: Node1();
165         break;
166     case 2: Node2();
167 };
168 /* mandatory for all MPI programs */
169 MPI_Finalize();
170 }
171
172 /* explanation of "number of items" and "status" arguments at the end
173    of MPI_Recv():
174
175    when receiving a message you must anticipate the longest possible
176    message, but the actual received message may be much shorter than
177    this; you can call the MPI_Get_count() function on the status
178    argument to find out how many items were actually received
179
180    the status argument will be a pointer to a struct, containing the
181    node number, message type and error status of the received
182    message
183
184    say our last parameter is Status; then Status.MPI_SOURCE
185    will contain the number of the sending node, and
186    Status.MPI_TAG will contain the message type; these are
187    important if used MPI_ANY_SOURCE or MPI_ANY_TAG in our
188    node or tag fields but still have to know who sent the
189    message or what kind it is */

```

The set of machines can be heterogeneous, but MPI “translates” for you automatically. If say one node has a big-endian CPU and another has a little-endian CPU, MPI will do the proper conversion.

1.4 Relative Merits: Shared-Memory Vs. Message-Passing

It is generally believed in the parallel processing community that the shared-memory paradigm produces code that is easier to write, debug and maintain than message-passing.

On the other hand, in some cases message-passing can produce faster code. Consider the Odd/Even Transposition Sort algorithm, for instance. Here pairs of processes repeatedly swap sorted arrays with each other. In a shared-memory setting, this might produce a bottleneck at the shared memory, slowing down the code. Of course, the obvious solution is that if you are using a shared-memory machine, you should just choose some other sorting algorithm, one tailored to the shared-memory setting.

There used to be a belief that message-passing was more **scalable**, i.e. amenable to very large systems. However, GPU has demonstrated that one can achieve extremely good scalability with shared-memory.

My own preference, obviously, is shared-memory.

1.5 Issues in Parallelizing Applications

The available parallel hardware systems sound wonderful at first. But many people have had the experience of enthusiastically writing their first parallel program, anticipating great speedups, only to find that their parallel code actually runs more slowly than their original nonparallel program. In this section, we highlight some major issues that will pop up throughout the book.

1.5.1 Communication Bottlenecks

Whether you are on a shared-memory, message-passing or other platform, communication is always a potential bottleneck. On a shared-memory system, the threads must contend with each other for memory access, and memory access itself can be slow, e.g. due to cache coherency transactions. On a NOW, even a very fast network is very slow compared to CPU speeds.

1.5.2 Load Balancing

Another major issue is **load balancing**, i.e. keeping all the processors busy as much as possible. In the graphics ray-tracing application, for instance, suppose a ray is coming from the “northeast” section of the image, and is reflected by a solid object. Then the ray won’t reach some of the “southwest” portions of the image, which then means that the processors assigned to those portions will not have any work to do which is associated with this ray. What we need to do is then try to give these processors some other work to do; the more they are idle, the slower our system will be.

1.5.3 “Embarrassingly Parallel” Applications

The term **embarrassingly parallel** is heard often in talk about parallel programming.

Consider a matrix multiplication application, for instance, in which we compute AX for a matrix A and a vector X . One way to parallelize this problem would be for have each processor handle a group of rows of A , multiplying each by X in parallel with the other processors, which are handling other groups of rows. We call the problem embarrassingly parallel, with the word “embarrassing” meaning that the problems are so easy to parallelize that there is no intellectual challenge involved. It is pretty obvious that the computation $Y = AX$ can be parallelized very easily by splitting the rows of A into groups.

By contrast, most parallel sorting algorithms require a great deal of interaction. For instance, consider Mergesort. It breaks the vector to be sorted into two (or more) independent parts, say the left half and right half, which are then sorted in parallel by two processes. So far, this is embarrassingly parallel, at least after the vector is broken in half. But then the two sorted halves must be merged to produce the sorted version

of the original vector, and that process is *not* embarrassingly parallel; it can be parallelized, but in a more complex manner.

Of course, it's no shame to have an embarrassingly parallel problem! On the contrary, except for showoff academics, having an embarrassingly parallel application is a cause for celebration, as it is easy to program. More importantly, embarrassingly parallel problems tend to have low communication needs, crucial to performance as discussed earlier.

Chapter 2

Shared Memory Parallelism

Shared-memory programming is considered by many in the parallel processing community as being the clearest of the various parallel paradigms available.

2.1 What Is Shared?

The term **shared memory** means that the processors all share a common address space. Say this is occurring at the hardware level, and we are using Intel Pentium CPUs. Suppose processor P3 issues the instruction

```
movl 200, %eax
```

which reads memory location 200 and places the result in the EAX register in the CPU. If processor P4 does the same, they both will be referring to the same physical memory cell. In non-shared-memory machines, each processor has its own private memory, and each one will then have its own location 200, completely independent of the locations 200 at the other processors' memories.

Say a program contains a global variable **X** and a local variable **Y** on share-memory hardware (and we use shared-memory software). If for example the compiler assigns location 200 to the variable **X**, i.e. $\&X = 200$, then the point is that all of the processors will have that variable in common, because any processor which issues a memory operation on location 200 will access the same physical memory cell.

On the other hand, each processor will have its own separate run-time stack. All of the stacks are in shared memory, but they will be accessed separately, since each CPU has a different value in its SP (Stack Pointer) register. Thus each processor will have its own independent copy of the local variable **Y**.

To make the meaning of “shared memory” more concrete, suppose we have a bus-based system, with all the processors and memory attached to the bus. Let us compare the above variables **X** and **Y** here. Suppose

again that the compiler assigns X to memory location 200. Then in the machine language code for the program, every reference to X will be there as 200. Every time an instruction that writes to X is executed by a CPU, that CPU will put 200 into its Memory Address Register (MAR), from which the 200 flows out on the address lines in the bus, and goes to memory. This will happen in the same way no matter which CPU it is. Thus the same physical memory location will end up being accessed, no matter which CPU generated the reference.

By contrast, say the compiler assigns a local variable Y to something like $ESP+8$, the third item on the stack (on a 32-bit machine), 8 bytes past the word pointed to by the stack pointer, ESP . The OS will assign a different ESP value to each thread, so the stacks of the various threads will be separate. Each CPU has its own ESP register, containing the location of the stack for whatever thread that CPU is currently running. So, the value of Y will be different for each thread.

2.2 Memory Modules

Parallel execution of a program requires, to a large extent, parallel accessing of memory. To some degree this is handled by having a cache at each CPU, but it is also facilitated by dividing the memory into separate **modules** or **banks**. This way several memory accesses can be done simultaneously.

In this section, assume for simplicity that our machine has 32-bit words. This is still true for many GPUs, in spite of the widespread use of 64-bit general-purpose machines today, and in any case, the numbers here can easily be converted to the 64-bit case.

Note that this means that consecutive words differ in address by 4. Let's thus define the word-address of a word to be its ordinary address divided by 4. Note that this is also its address with the lowest two bits deleted.

2.2.1 Interleaving

There is a question of how to divide up the memory into modules. There are two main ways to do this:

- (a) **High-order interleaving:** Here consecutive words are in the same module (except at boundaries). For example, suppose for simplicity that our memory consists of word-addresses 0 through 1023, and that there are four modules, M_0 through M_3 . Then M_0 would contain word-addresses 0-255, M_1 would have 256-511, M_2 would have 512-767, and M_3 would have 768-1023.
- (b) **Low-order interleaving:** Here consecutive addresses are in consecutive modules (except when we get to the right end). In the example above, if we used low-order interleaving, then word-address 0 would be in M_0 , 1 would be in M_1 , 2 would be in M_2 , 3 would be in M_3 , 4 would be back in M_0 , 5 in M_1 , and so on.

Say we will have eight modules. Then under high-order interleaving, the first two bits of a word-address would be taken to be the module number, with the remaining bits being address within module. Under low-order interleaving, the two least significant bits would be used.

Low-order interleaving has often been used for **vector processors**. On such a machine, we might have both a regular add instruction, `ADD`, and a vector version, `VADD`. The latter would add two vectors together, so it would need to read two vectors from memory. If low-order interleaving is used, the elements of these vectors are spread across the various modules, so fast access is possible.

A more modern use of low-order interleaving, but with the same motivation as with the vector processors, is in GPUs. See Chapter 5.

High-order interleaving might work well in matrix applications, for instance, where we can partition the matrix into blocks, and have different processors work on different blocks. In image processing applications, we can have different processors work on different parts of the image. Such partitioning almost never works perfectly—e.g. computation for one part of an image may need information from another part—but if we are careful we can get good results.

2.2.2 Bank Conflicts and Solutions

C/C++ stores two-dimensional arrays in **row-major** order, which we'll assume here.¹ In this scheme, all of row 0 is stored contiguously, then all of row 1, and so on.

Consider a matrix `m` with four columns. The its storage in memory would first have `m[0][0]`, then `m[0][1]`, then `m[0][2]`, then `m[0][3]`. After that would come `m[1][0]`, `m[1][1]`, and so on. Now consider two pieces of code to find the sum of the entire array, assuming say 10 rows. Here's the first:

```
sum = 0;
for (i = 0; i < 10; i++)
  for (j = 0; j < 4; j++)
    sum += m[i][j];
```

And here's the alternative:

```
sum = 0;
for (j = 0; j < 4; j++)
  for (i = 0; i < 10; i++)
    sum += m[i][j];
```

Suppose, again for simplicity, that we have four memory modules.

¹FORTRAN uses column-major order, as does the R statistical language.

Remember, whichever version of the code is used, it might be executed in parallel, by different threads. Which will be better?

The answer is that since we are using row-major order, the first marches through consecutive addresses in memory, thus consecutive modules (with wraparound). Thus we may be able to keep all the modules busy, thus achieve good parallelism.

By contrast, the second version of the code will have repeated bank conflicts, because each consecutive access will be four words apart in memory—and thus will be in the *same* module!

In the first version of the program, consecutive memory accesses were one word apart—which we term a **stride** of 1. In the second version, the stride is 4. This motivates the following:

Suppose we have q memory modules, and consider the word-addresses of q consecutive memory requests addresses issued from our program, having a stride of s . We hope those q requests are to q different modules. Let's call this "full usage" of the memory.

Theorem: We will achieve full usage if and only if $\gcd(q,s) = 1$, i.e. q and s are relatively prime.

Proof: This is comes from a theorem in group theory, not presented here.


What if q and s are not relatively prime? Here are a couple of fairly easy solutions:

- We can rewrite our algorithm, e.g. use the first version of the above code instead of the second.
- We can add padding to the array. If for instance we add a fifth column to our matrix, then in the second version of the code, consecutive memory requests will be five words apart instead of four—and thus in different modules.

2.3 Interconnection Topologies

2.3.1 SMP Systems

A Symmetric Multiprocessor (SMP) system has the following structure:




Here and below:

- The Ps are processors, e.g. off-the-shelf chips such as Pentiums.
- The Ms are **memory modules**. These are physically separate objects, e.g. separate boards of memory chips. It is typical that there will be the same number of Ms as Ps.
- To make sure only one P uses the bus at a time, standard bus arbitration signals and/or arbitration devices are used.
- There may also be **coherent caches**, which we will discuss later.

2.3.2 NUMA Systems

In a **Nonuniform Memory Access** (NUMA) architecture, each CPU has a memory module physically next to it, and these processor/memory (P/M) pairs are connected by some kind of network.

Here is a simple version:



Each P/M/R set here is called a **processing element** (PE). Note that each PE has its own local bus, and is also connected to the global bus via R, the router.

Suppose for example that P3 needs to access location 200, and suppose that high-order interleaving is used. If location 200 is in M3, then P3's request is satisfied by the local bus.² On the other hand, suppose location 200 is in M8. Then the R3 will notice this, and put the request on the global bus, where it will be seen by R8, which will then copy the request to the local bus at PE8, where the request will be satisfied. (E.g. if it was a read request, then the response will go back from M8 to R8 to the global bus to R3 to P3.)

It should be obvious now where NUMA gets its name. P8 will have much faster access to M8 than P3 will to M8, if none of the buses is currently in use—and if say the global bus is currently in use, P3 will have to wait a long time to get what it wants from M8.

²This sounds similar to the concept of a cache. However, it is very different. A cache contains a local copy of some data stored elsewhere. Here it is the data itself, not a copy, which is being stored locally.

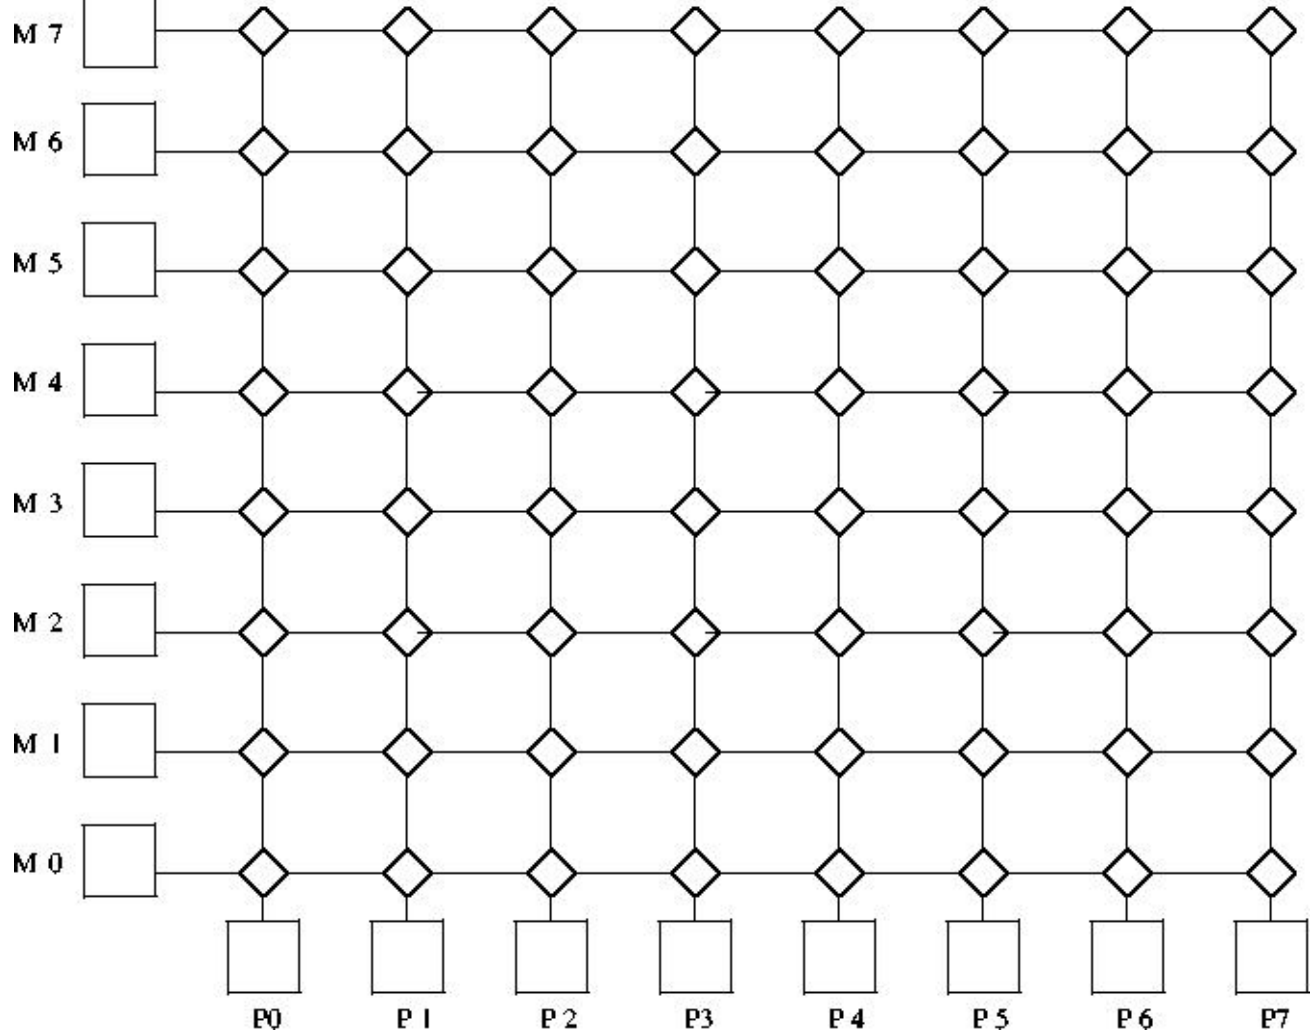
Today almost all high-end MIMD systems are NUMAs. One of the attractive features of NUMA is that by good programming we can exploit the nonuniformity. In matrix problems, for example, we can write our program so that, for example, P8 usually works on those rows of the matrix which are stored in M8, P3 usually works on those rows of the matrix which are stored in M3, etc. In order to do this, we need to make use of the C language's & address operator, and have some knowledge of the memory hardware structure, i.e. the interleaving.

2.3.3 NUMA Interconnect Topologies

The problem with a bus connection, of course, is that there is only one pathway for communication, and thus only one processor can access memory at the same time. If one has more than, say, two dozen processors on the bus, the bus becomes saturated, even if traffic-reducing methods such as adding caches are used. Thus multipathway topologies are used for all but the smallest systems. In this section we look at two alternatives to a bus topology.

2.3.3.1 Crossbar Interconnects

Consider a shared-memory system with n processors and n memory modules. Then a crossbar connection would provide n^2 pathways. E.g. for $n = 8$:



Generally serial communication is used from node to node, with a packet containing information on both source and destination address. E.g. if P2 wants to read from M5, the source and destination will be 3-bit strings in the packet, coded as 010 and 101, respectively. The packet will also contain bits which specify which word within the module we wish to access, and bits which specify whether we wish to do a read or a write. In the latter case, additional bits are used to specify the value to be written.

Each diamond-shaped node has two inputs (bottom and right) and two outputs (left and top), with buffers at the two inputs. If a buffer fills, there are two design options: (a) Have the node from which the input comes block at that output. (b) Have the node from which the input comes discard the packet, and retry later, possibly outputting some other packet for now. If the packets at the heads of the two buffers both need to go out the same output, the one (say) from the bottom input will be given priority.

There could also be a return network of the same type, with this one being memory → processor, to return

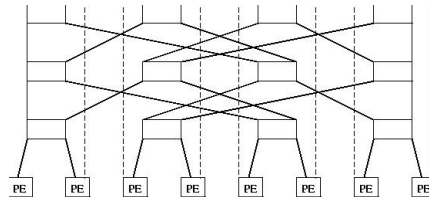
the result of the read requests.³

Another version of this is also possible. It is not shown here, but the difference would be that at the bottom edge we would have the PE_i and at the left edge the memory modules M_i would be replaced by lines which wrap back around to PE_i, similar to the Omega network shown below.

Crossbar switches are too expensive for large-scale systems, but are useful in some small systems. The 16-CPU Sun Microsystems Enterprise 10000 system includes a 16x16 crossbar.

2.3.3.2 Omega (or Delta) Interconnects

These are multistage networks similar to crossbars, but with fewer paths. Here is an example of a NUMA 8x8 system:



Recall that each PE is a processor/memory pair. PE₃, for instance, consists of P₃ and M₃.

Note the fact that at the third stage of the network (top of picture), the outputs are routed back to the PEs, each of which consists of a processor and a memory module.⁴

At each network node (the nodes are the three rows of rectangles), the output routing is done by destination bit. Let's number the stages here 0, 1 and 2, starting from the bottom stage, number the nodes within a stage 0, 1, 2 and 3 from left to right, number the PEs from 0 to 7, left to right, and number the bit positions in a destination address 0, 1 and 2, starting from the most significant bit. Then at stage *i*, bit *i* of the destination address is used to determine routing, with a 0 meaning routing out the left output, and 1 meaning the right one.

Say P₂ wishes to read from M₅. It sends a read-request packet, including 5 = 101 as its destination address, to the switch in stage 0, node 1. Since the first bit of 101 is 1, that means that this switch will route the packet out its right-hand output, sending it to the switch in stage 1, node 3. The latter switch will look at the next bit in 101, a 0, and thus route the packet out its left output, to the switch in stage 2, node 2. Finally, that switch will look at the last bit, a 1, and output out its right-hand output, sending it to PE₅, as desired. M₅ will process the read request, and send a packet back to PE₂, along the same

Again, if two packets at a node want to go out the same output, one must get priority (let's say it is the one

³For safety's sake, i.e. fault tolerance, even writes are typically acknowledged in multiprocessor systems.

⁴The picture may be cut off somewhat at the top and left edges. The upper-right output of the rectangle in the top row, leftmost position should connect to the dashed line which leads down to the second PE from the left. Similarly, the upper-left output of that same rectangle is a dashed lined, possibly invisible in your picture, leading down to the leftmost PE.

from the left input).

Here is how the more general case of $N = 2^n$ PEs works. Again number the rows of switches, and switches within a row, as above. So, S_{ij} will denote the switch in the i -th row from the bottom and j -th column from the left (starting our numbering with 0 in both cases). Row i will have a total of N input ports I_{ik} and N output ports O_{ik} , where $k = 0$ corresponds to the leftmost of the N in each case. Then if row i is not the last row ($i < n - 1$), O_{ik} will be connected to I_{jm} , where $j = i+1$ and

$$m = (2k + \lfloor (2k)/N \rfloor) \bmod N \quad (2.1)$$

If row i is the last row, then O_{ik} will be connected to, PE k .

2.3.4 Comparative Analysis

In the world of parallel architectures, a key criterion for a proposed feature is **scalability**, meaning how well the feature performs as we go to larger and larger systems. Let n be the system size, either the number of processors and memory modules, or the number of PEs. Then we are interested in how fast the latency, bandwidth and cost grow with n :

critierion	bus	Omega	crossbar
latency	$O(1)$	$O(\log_2 n)$	$O(n)$
bandwidth	$O(1)$	$O(n)$	$O(n)$
cost	$O(1)$	$O(n \log_2 n)$	$O(n^2)$

Let us see where these expressions come from, beginning with a bus: No matter how large n is, the time to get from, say, a processor to a memory module will be the same, thus $O(1)$. Similarly, no matter how large n is, only one communication can occur at a time, thus again $O(1)$.⁵

Again, we are interested only in “ $O(\)$ ” measures, because we are only interested in growth rates as the system size n grows. For instance, if the system size doubles, the cost of a crossbar will quadruple; the $O(n^2)$ cost measure tells us this, with any multiplicative constant being irrelevant.

For Omega networks, it is clear that $\log_2 n$ network rows are needed, hence the latency value given. Also, each row will have $n/2$ switches, so the number of network nodes will be $O(n \log_2 n)$. This figure then gives the cost (in terms of switches, the main expense here). It also gives the bandwidth, since the maximum number of simultaneous transmissions will occur when all switches are sending at once.

Similar considerations hold for the crossbar case.

⁵ Note that the ‘1’ in “ $O(1)$ ” does not refer to the fact that only one communication can occur at a time. If we had, for example, a two-bus system, the bandwidth would still be $O(1)$, since multiplicative constants do not matter. What $O(1)$ means, again, is that as n grows, the bandwidth stays at a multiple of 1, i.e. stays constant.

The crossbar's big advantage is that it is guaranteed that n packets can be sent simultaneously, providing they are to distinct destinations.

That is not true for Omega-networks. If for example, PE0 wants to send to PE3, and at the same time PE4 wishes to send to PE2, the two packets will clash at the leftmost node of stage 1, where the packet from PE0 will get priority.

On the other hand, a crossbar is very expensive, and thus is dismissed out of hand in most modern systems. Note, though, that an equally troublesome aspect of crossbars is their high latency value; this is a big drawback when the system is not heavily loaded.

The bottom line is that Omega-networks amount to a compromise between buses and crossbars, and for this reason have become popular.

2.3.5 Why Have Memory in Modules?

In the shared-memory case, the Ms collectively form the entire shared address space, but with the addresses being assigned to the Ms in one of two ways:

- (a)

High-order interleaving. Here consecutive addresses are in the same M (except at boundaries). For example, suppose for simplicity that our memory consists of addresses 0 through 1023, and that there are four Ms. Then M0 would contain addresses 0-255, M1 would have 256-511, M2 would have 512-767, and M3 would have 768-1023.

- (b)

Low-order interleaving. Here consecutive addresses are in consecutive M's (except when we get to the right end). In the example above, if we used low-order interleaving, then address 0 would be in M0, 1 would be in M1, 2 would be in M2, 3 would be in M3, 4 would be back in M0, 5 in M1, and so on.

The idea is to have several modules busy at once, say in conjunction with a **split-transaction bus**. Here, after a processor makes a memory request, it relinquishes the bus, allowing others to use it while the memory does the requested work. Without splitting the memory into modules, this wouldn't achieve parallelism. The bus does need extra lines to identify which processor made the request.

2.4 Test-and-Set Type Instructions

Consider a bus-based system. In addition to whatever memory read and memory write instructions the processor included, there would also be a TAS instruction.⁶ This instruction would control a TAS pin on the processor chip, and the pin in turn would be connected to a TAS line on the bus.

Applied to a location *L* in memory and a register *R*, say, TAS does the following:

```
copy L to R
if R is 0 then write 1 to L
```

And most importantly, these operations are done in an **atomic** manner; no bus transactions by other processors may occur between the two steps.

The TAS operation is applied to variables used as **locks**. Let's say that 1 means locked and 0 unlocked. Then the guarding of a critical section *C* by a lock variable *L* would be done by having the following code in the program being run:

```
TRY:  TAS R,L
      JNZ TRY
C:    ... ; start of critical section
      ...
      ... ; end of critical section
      MOV L,0 ; unlock
```

where of course JNZ is a jump-if-nonzero instruction, and we are assuming that the copying from the Memory Data Register to *R* results in the processor *N* and *Z* flags (condition codes) being affected.

On Pentium machines, the LOCK prefix can be used to get atomicity for certain instructions.⁷ For example,

```
lock add $2, x
```

would add the constant 2 to the memory location labeled *x* in an atomic manner.

The LOCK prefix locks the bus for the entire duration of the instruction. Note that the ADD instruction here involves two memory transactions—one to read the old value of *x*, and the second the write the new, incremented value back to *x*. So, we are locking for a rather long time, but the benefits can be huge.

A good example of this kind of thing would be our program **PrimesThreads.c** in Chapter 1, where our critical section consists of adding 2 to **nextbase**. There we surrounded the add-2 code by Pthreads lock

⁶This discussion is for a mythical machine, but any real system works in this manner.

⁷The instructions ADD, ADC, AND, BTC, BTR, BTS, CMPXCHG, DEC, INC, NEG, NOT, OR, SBB, SUB, XOR, XADD. Also, XCHG asserts the LOCK# bus signal even if the LOCK prefix is specified. Locking only applies to these instructions in forms in which there is an operand in memory.

and unlock operations. These involve system calls, which are very time consuming, involving hundreds of machine instructions. Compare that to the one-instruction solution above! The very heavy overhead of pthreads would be thus avoided.

In crossbar or Ω -network systems, some 2-bit field in the packet must be devoted to transaction type, say 00 for Read, 01 for Write and 10 for TAS. In a system with 16 CPUs and 16 memory modules, say, the packet might consist of 4 bits for the CPU number, 4 bits for the memory module number, 2 bits for the transaction type, and 32 bits for the data (for a write, this is the data to be written, while for a read, it would be the requested value, on the trip back from the memory to the CPU).

But note that the atomicity here is best done at the memory, i.e. some hardware should be added at the memory so that TAS can be done; otherwise, an entire processor-to-memory path (e.g. the bus in a bus-based system) would have to be locked up for a fairly long time, obstructing even the packets which go to other memory modules.

There are many variations of test-and-set, so don't expect that all processors will have an instruction with this name, but they all will have some kind of synchronization instruction like it.

Note carefully that in many settings it may not be crucial to get the most up-to-date value of a variable. For example, a program may have a data structure showing work to be done. Some processors occasionally add work to the queue, and others take work from the queue. Suppose the queue is currently empty, and a processor adds a task to the queue, just as another processor is checking the queue for work. As will be seen later, it is possible that even though the first processor has written to the queue, the new value won't be visible to other processors for some time. But the point is that if the second processor does not see work in the queue (even though the first processor has put it there), the program will still work correctly, albeit with some performance loss.

2.5 Cache Issues

2.5.1 Cache Coherency

Consider, for example, a bus-based system. Relying purely on TAS for interprocessor synchronization would be unthinkable: As each processor contending for a lock variable spins in the loop shown above, it is adding tremendously to bus traffic.

An answer is to have caches at each processor.⁸ These will store copies of the values of lock variables. (Of course, non-lock variables are stored too. However, the discussion here will focus on effects on lock variables.) The point is this: Why keep looking at a lock variable L again and again, using up the bus bandwidth? L may not change value for a while, so why not keep a copy in the cache, avoiding use of the

⁸The reader may wish to review the basics of caches. See for example <http://heather.cs.ucdavis.edu/~matloff/50/PLN/CompOrganization.pdf>.

bus?

The answer of course is that eventually *L* will change value, and this causes some delicate problems. Say for example that processor P5 wishes to enter a critical section guarded by *L*, and that processor P2 is already in there. During the time P2 is in the critical section, P5 will spin around, always getting the same value for *L* (1) from C5, P5's cache. When P2 leaves the critical section, P2 will set *L* to 0—and now C5's copy of *L* will be incorrect. This is the **cache coherency problem**, inconsistency between caches.

A number of solutions have been devised for this problem. For bus-based systems, **snoopy** protocols of various kinds are used, with the word “snoopy” referring to the fact that all the caches monitor (“snoop on”) the bus, watching for transactions made by other caches.

The most common protocols are the **invalidate** and **update** types. This relation between these two is somewhat analogous to the relation between **write-back** and **write-through** protocols for caches in uniprocessor systems:

- Under an invalidate protocol, when a processor writes to a variable in a cache, it first (i.e. before actually doing the write) tells each other cache to mark as invalid its cache line (if any) which contains a copy of the variable.⁹ Those caches will be updated only later, the next time their processors need to access this cache line.
- For an update protocol, the processor which writes to the variable tells all other caches to immediately update their cache lines containing copies of that variable with the new value.

Let's look at an outline of how one implementation (many variations exist) of an invalidate protocol would operate:

In the scenario outlined above, when P2 leaves the critical section, it will write the new value 0 to *L*. Under the invalidate protocol, P2 will post an invalidation message on the bus. All the other caches will notice, as they have been monitoring the bus. They then mark their cached copies of the line containing *L* as invalid.

Now, the next time P5 executes the TAS instruction—which will be very soon, since it is in the loop shown above—P5 will find that the copy of *L* in C5 is invalid. It will respond to this cache miss by going to the bus, and requesting P2 to supply the “real” (and valid) copy of the line containing *L*.

But there's more. Suppose that all this time P6 had also been executing the loop shown above, along with P5. Then P5 and P6 may have to contend with each other. Say P6 manages to grab possession of the bus first.¹⁰ P6 then executes the TAS again, which finds *L* = 0 and changes *L* back to 1. P6 then relinquishes the bus, and enters the critical section. Note that in changing *L* to 1, P6 also sends an invalidate signal to all the

⁹We will follow commonly-used terminology here, distinguishing between a *cache line* and a *memory block*. Memory is divided in blocks, some of which have copies in the cache. The cells in the cache are called *cache lines*. So, at any given time, a given cache line is either empty or contains a copy (valid or not) of some memory block.

¹⁰Again, remember that ordinary bus arbitration methods would be used.

other caches. So, when P5 tries its execution of the TAS again, it will have to ask P6 to send a valid copy of the block. P6 does so, but L will be 1, so P5 must resume executing the loop. P5 will then continue to use its valid local copy of L each time it does the TAS, until P6 leaves the critical section, writes 0 to L, and causes another cache miss at P5, etc.

At first the update approach seems obviously superior, and actually, if our shared, cacheable¹¹ variables were only lock variables, this might be true.

But consider a shared, cacheable vector. Suppose the vector fits into one block, and that we write to each vector element sequentially. Under an update policy, we would have to send a new message on the bus/network for each component, while under an invalidate policy, only one message (for the first component) would be needed. If during this time the other processors do not need to access this vector, all those update messages, and the bus/network bandwidth they use, would be wasted.

Or suppose for example we have code like

```
Sum += X[I];
```

in the middle of a **for** loop. Under an update protocol, we would have to write the value of Sum back many times, even though the other processors may only be interested in the final value when the loop ends. (This would be true, for instance, if the code above were part of a critical section.)

Thus the invalidate protocol works well for some kinds of code, while update works better for others. The CPU designers must try to anticipate which protocol will work well across a broad mix of applications.¹²

Now, how is cache coherency handled in non-bus shared-memory systems, say crossbars? Here the problem is more complex. Think back to the bus case for a minute: The very feature which was the biggest negative feature of bus systems—the fact that there was only one path between components made bandwidth very limited—is a very positive feature in terms of cache coherency, because it makes broadcast very easy: Since everyone is attached to that single pathway, sending a message to all of them costs no more than sending it to just one—we get the others for free. That’s no longer the case for multipath systems. In such systems, extra copies of the message must be created for each path, adding to overall traffic.

A solution is to send messages only to “interested parties.” In **directory-based** protocols, a list is kept of all caches which currently have valid copies of all blocks. In one common implementation, for example, while P2 is in the critical section above, it would be the **owner** of the block containing L. (Whoever is the latest node to write to L would be considered its current owner.) It would maintain a directory of all caches having valid copies of that block, say C5 and C6 in our story here. As soon as P2 wrote to L, it would then send either invalidate or update packets (depending on which type was being used) to C5 and C6 (and not to other caches which didn’t have valid copies).

¹¹ Many modern processors, including Pentium and MIPS, allow the programmer to mark some blocks as being noncacheable.

¹²Some protocols change between the two modes dynamically.

There would also be a directory at the memory, listing the current owners of all blocks. Say for example P0 now wishes to “join the club,” i.e. tries to access L, but does not have a copy of that block in its cache C0. C0 will thus not be listed in the directory for this block. So, now when it tries to access L and it will get a cache miss. P0 must now consult the **home** of L, say P14. The home might be determined by L’s location in main memory according to high-order interleaving; it is the place where the main-memory version of L resides. A table at P14 will inform P0 that P2 is the current owner of that block. P0 will then send a message to P2 to add C0 to the list of caches having valid copies of that block. Similarly, a cache might “resign” from the club, due to that cache line being replaced, e.g. in a LRU setting, when some other cache miss occurs.

2.5.2 Example: the MESI Cache Coherency Protocol

Many types of cache coherency protocols have been proposed and used, some of them quite complex. A relatively simple one for snoopy bus systems which is widely used is MESI, which for example is the protocol used in the Pentium series.

MESI is an invalidate protocol for bus-based systems. Its name stands for the four states a given cache line can be in for a given CPU:

- Modified
- Exclusive
- Shared
- Invalid

Note that *each memory block* has such a state at *each cache*. For instance, block 88 may be in state S at P5’s and P12’s caches but in state I at P1’s cache.

Here is a summary of the meanings of the states:

state	meaning
M	written to more than once; no other copy valid
E	valid; no other cache copy valid; memory copy valid
S	valid; at least one other cache copy valid
I	invalid (block either not in the cache or present but incorrect)

Following is a summary of MESI state changes.¹³ When reading it, keep in mind again that there is a separate state for each cache/memory block combination.

¹³See *Pentium Processor System Architecture*, by D. Anderson and T. Shanley, Addison-Wesley, 1995. We have simplified the presentation here, by eliminating certain programmable options.

In addition to the terms **read hit**, **read miss**, **write hit**, **write miss**, which you are already familiar with, there are also **read snoop** and **write snoop**. These refer to the case in which our CPU observes on the bus a block request by another CPU that has attempted a read or write action but encountered a miss in its own cache; if our cache has a valid copy of that block, we must provide it to the requesting CPU (and in some cases to memory).

So, here are various events and their corresponding state changes:

If our CPU does a read:

present state	event	new state
M	read hit	M
E	read hit	E
S	read hit	S
I	read miss; no valid cache copy at any other CPU	E
I	read miss; at least one valid cache copy in some other CPU	S

If our CPU does a memory write:

present state	event	new state
M	write hit; do not put invalidate signal on bus; do not update memory	M
E	same as M above	M
S	write hit; put invalidate signal on bus; update memory	E
I	write miss; update memory but do nothing else	I

If our CPU does a read or write snoop:

present state	event	newstate
M	read snoop; write line back to memory, picked up by other CPU	S
M	write snoop; write line back to memory, signal other CPU now OK to do its write	I
E	read snoop; put shared signal on bus; no memory action	S
E	write snoop; no memory action	I
S	read snoop	S
S	write snoop	I
I	any snoop	I

Note that a write miss does NOT result in the associated block being brought in from memory.

Example: Suppose a given memory block has state M at processor A but has state I at processor B, and B attempts to write to the block. B will see that its copy of the block is invalid, so it notifies the other CPUs via the bus that it intends to do this write. CPU A sees this announcement, tells B to wait, writes its own copy of the block back to memory, and then tells B to go ahead with its write. The latter action means that A's copy of the block is not correct anymore, so the block now has state I at A. B's action does not cause

loading of that block from memory to its cache, so the block still has state I at B.

2.5.3 The Problem of “False Sharing”

Consider the C declaration

```
int W,Z;
```

Since **W** and **Z** are declared adjacently, most compilers will assign them contiguous memory addresses. Thus, unless one of them is at a memory block boundary, when they are cached they will be stored in the same cache line. Suppose the program writes to **Z**, and our system uses an invalidate protocol. Then **W** will be considered invalid at the other processors, even though its values at those processors’ caches are correct. This is the **false sharing** problem, alluding to the fact that the two variables are sharing a cache line even though they are not related.

This can have very adverse impacts on performance. If for instance our variable **W** is now written to, then **Z** will suffer unfairly, as its copy in the cache will be considered invalid even though it is perfectly valid. This can lead to a “ping-pong” effect, in which alternate writing to two variables leads to a cyclic pattern of coherency transactions.

One possible solution is to add padding, e.g. declaring **W** and **Z** like this:

```
int Q,U[1000],Z;
```

to separate **Q** and **Z** so that they won’t be in the same cache block. Of course, we must take block size into account, and check whether the compiler really has placed the two variables in widely separated locations. To do this, we could for instance run the code

```
printf("%x %x\n",&Q,&Z);
```

2.6 Memory-Access Consistency Policies

Though the word *consistency* in the title of this section may seem to simply be a synonym for *coherency* from the last section, and though there actually is some relation, the issues here are quite different. In this case, it is a timing issue: After one processor changes the value of a shared variable, when will that value be visible to the other processors?

There are various reasons why this is an issue. For example, many processors, especially in multiprocessor systems, have **write buffers**, which save up writes for some time before actually sending them to memory.

(For the time being, let's suppose there are no caches.) The goal is to reduce memory access costs. Sending data to memory in groups is generally faster than sending one at a time, as the overhead of, for instance, acquiring the bus is amortized over many accesses. Reads following a write may proceed, without waiting for the write to get to memory, except for reads to the same address. So in a multiprocessor system in which the processors use write buffers, there will often be some delay before a write actually shows up in memory.

A related issue is that operations may occur, or appear to occur, out of order. As noted above, a read which follows a write in the program may execute before the write is sent to memory. Also, in a multiprocessor system with multiple paths between processors and memory modules, two writes might take different paths, one longer than the other, and arrive "out of order." In order to simplify the presentation here, we will focus on the case in which the problem is due to write buffers, though.

The designer of a multiprocessor system must adopt some **consistency model** regarding situations like this. The above discussion shows that the programmer must be made aware of the model, or risk getting incorrect results. Note also that different consistency models will give different levels of performance. The "weaker" consistency models make for faster machines but require the programmer to do more work.

The strongest consistency model is Sequential Consistency. It essentially requires that memory operations done by one processor are observed by the other processors to occur in the same order as executed on the first processor. Enforcement of this requirement makes a system slow, and it has been replaced on most systems by weaker models.

One such model is **release consistency**. Here the processors' instruction sets include instructions ACQUIRE and RELEASE. Execution of an ACQUIRE instruction at one processor involves telling all other processors to flush their write buffers. However, the ACQUIRE won't execute until pending RELEASEs are done. Execution of a RELEASE basically means that you are saying, "I'm done writing for the moment, and wish to allow other processors to see what I've written." An ACQUIRE waits for all pending RELEASEs to complete before it executes.¹⁴

A related model is **scope consistency**. Say a variable, say **Sum**, is written to within a critical section guarded by LOCK and UNLOCK instructions. Then under scope consistency any changes made by one processor to **Sum** within this critical section would then be visible to another processor when the latter next enters this critical section. The point is that memory update is postponed until it is actually needed. Also, a barrier operation (again, executed at the hardware level) forces all pending memory writes to complete.

All modern processors include instructions which implement consistency operations. For example, Sun Microsystems' SPARC has a MEMBAR instruction. If used with a STORE operand, then all pending writes at this processor will be sent to memory. If used with the LOAD operand, all writes will be made visible to this processor.

Now, how does cache coherency fit into all this? There are many different setups, but for example let's consider a design in which there is a write buffer between each processor and its cache. As the processor

¹⁴There are many variants of all of this, especially in the software distributed shared memory realm, to be discussed later.

does more and more writes, the processor saves them up in the write buffer. Eventually, some programmer-induced event, e.g. a MEMBAR instruction,¹⁵ will cause the buffer to be flushed. Then the writes will be sent to “memory”—actually meaning that they go to the cache, and then possibly to memory.

The point is that (in this type of setup) before that flush of the write buffer occurs, the cache coherency system is quite unaware of these writes. Thus the cache coherency operations, e.g. the various actions in the MESI protocol, won’t occur until the flush happens.

To make this notion concrete, again consider the example with **Sum** above, and assume release or scope consistency. The CPU currently executing that code (say CPU 5) writes to **Sum**, which is a memory operation—it affects the cache and thus eventually the main memory—but that operation will be invisible to the cache coherency protocol for now, as it will only be reflected in this processor’s write buffer. But when the unlock is finally done (or a barrier is reached), the write buffer is flushed and the writes are sent to this CPU’s cache. That then triggers the cache coherency operation (depending on the state). The point is that the cache coherency operation would occur only now, not before.

What about reads? Suppose another processor, say CPU 8, does a read of **Sum**, and that page is marked invalid at that processor. A cache coherency operation will then occur. Again, it will depend on the type of coherency policy and the current state, but in typical systems this would result in **Sum**’s cache block being shipped to CPU 8 from whichever processor the cache coherency system thinks has a valid copy of the block. That processor may or may not be CPU 5, but even if it is, that block won’t show the recent change made by CPU 5 to **Sum**.

The analysis above assumed that there is a write buffer between each processor and its cache. There would be a similar analysis if there were a write buffer between each cache and memory.

Note once again the performance issues. Instructions such as ACQUIRE or MEMBAR will use a substantial amount of interprocessor communication bandwidth. A consistency model must be chosen carefully by the system designer, and the programmer must keep the communication costs in mind in developing the software.

The recent Pentium models use Sequential Consistency, with any write done by a processor being immediately sent to its cache as well.

2.7 Fetch-and-Add and Packet-Combining Operations

Another form of interprocessor synchronization is a **fetch-and-add** (FA) instruction. The idea of FA is as follows. For the sake of simplicity, consider code like

¹⁵We call this “programmer-induced,” since the programmer will include some special operation in her C/C++ code which will be translated to MEMBAR.

```
LOCK (K) ;
Y = X++;
UNLOCK (K) ;
```

Suppose our architecture's instruction set included an F&A instruction. It would add 1 to the specified location in memory, and return the old value (to **Y**) that had been in that location before being incremented. And all this would be an atomic operation.

We would then replace the code above by a library call, say,

```
FETCH_AND_ADD (X, 1) ;
```

The C code above would compile to, say,

```
F&A X, R, 1
```

where **R** is the register into which the old (pre-incrementing) value of **X** would be returned.

There would be hardware adders placed at each memory module. That means that the whole operation could be done in one round trip to memory. Without F&A, we would need two round trips to memory just for the

```
X++;
```

(we would load **X** into a register in the CPU, increment the register, and then write it back to **X** in memory), and then the `LOCK()` and `UNLOCK()` would need trips to memory too. This could be a huge time savings, especially for long-latency interconnects.

In addition to read and write operations being specifiable in a network packet, an F&A operation could be specified as well (a 2-bit field in the packet would code which operation was desired). Again, there would be adders included at the memory modules, i.e. the addition would be done at the memory end, not at the processors. When the F&A packet arrived at a memory module, our variable **X** would have 1 added to it, while the old value would be sent back in the return packet (and put into **R**).

Another possibility for speedup occurs if our system uses a multistage interconnection network such as a crossbar. In that situation, we can design some intelligence into the network nodes to do **packet combining**: Say more than one CPU is executing an F&A operation at about the same time for the same variable **X**. Then more than one of the corresponding packets may arrive at the same network node at about the same time. If each one requested an incrementing of **X** by 1, the node can replace the two packets by one, with an increment of 2. Of course, this is a delicate operation, and we must make sure that different CPUs get different return values, etc.

2.8 Multicore Chips

A recent trend has been to put several CPUs on one chip, termed a **multicore** chip. As of March 2008, dual-core chips are common in personal computers, and quad-core machines are within reach of the budgets of many people. Just as the invention of the integrated circuit revolutionized the computer industry by making computers affordable for the average person, multicore chips will undoubtedly revolutionize the world of parallel programming.

A typical dual-core setup might have the two CPUs sharing a common L2 cache, with each CPU having its own L3 cache. The chip may interface to the bus or interconnect network of via an L1 cache.

Multicore is extremely important these days. However, they are just SMPs, for the most part, and thus should not be treated differently.

2.9 Illusion of Shared-Memory through Software

2.9.0.1 Software Distributed Shared Memory

There are also various shared-memory software packages that run on message-passing hardware such as NOWs, called **software distributed shared memory** (SDSM) systems. Since the platforms do not have any physically shared memory, the shared-memory view which the programmer has is just an illusion. But that illusion is very useful, since the shared-memory paradigm is believed to be the easier one to program in. Thus SDSM allows us to have “the best of both worlds”—the convenience of the shared-memory world view with the inexpensive cost of some of the message-passing hardware systems, particularly networks of workstations (NOWs).

SDSM itself is divided into two main approaches, the **page-based** and **object-based** varieties. The page-based approach is generally considered clearer and easier to program in, and provides the programmer the “look and feel” of shared-memory programming better than does the object-based type.¹⁶ We will discuss only the page-based approach here. The most popular SDSM system today is the page-based Treadmarks (Rice University). Another excellent page-based system is JIAJIA (Academy of Sciences, China).

To illustrate how page-paged SDSMs work, consider the line of JIAJIA code

```
Prime = (int *) jia_alloc(N*sizeof(int));
```

The function **jia_alloc()** is part of the JIAJIA library, **libjia.a**, which is linked to one’s application program during compilation.

¹⁶The term *object-based* is not related to the term *object-oriented programming*.

At first this looks a little like a call to the standard `malloc()` function, setting up an array `Prime` of size `N`. In fact, it does indeed allocate some memory. Note that each node in our JIAJIA group is executing this statement, so each node allocates some memory at that node. Behind the scenes, not visible to the programmer, each node will then have its own copy of `Prime`.

However, JIAJIA sets things up so that when one node later accesses this memory, for instance in the statement

```
Prime[I] = 1;
```

this action will eventually trigger a network transaction (not visible to the programmer) to the other JIAJIA nodes.¹⁷ This transaction will then update the copies of `Prime` at the other nodes.¹⁸

How is all of this accomplished? It turns out that it relies on a clever usage of the nodes' virtual memory (VM) systems. To understand this, let's review how VM systems work.

Suppose a variable `X` has the virtual address 1200, i.e. `&X = 1200`. The actual physical address may be, say, 5000. When the CPU executes a machine instruction that specifies access to 1200, the CPU will do a lookup on the **page table**, and find that the true location is 5000, and then access 5000. On the other hand, `X` may not be **resident** in memory at all, in which case the page table will say so. If the CPU finds that `X` is nonresident, it will cause an internal interrupt, which in turn will cause a jump to the operating system (OS). The OS will then read `X` in from disk,¹⁹ place it somewhere in memory, and then update the page table to show that `X` is now someplace in memory. The OS will then execute a return from interrupt instruction,²⁰ and the CPU will restart the instruction which triggered the page fault.

Here is how this is exploited to develop SDSMs on Unix systems. The SDSM will call a system function such as `mprotect()`. This allows the SDSM to deliberately mark a page as nonresident (even if the page *is* resident). Basically, anytime the SDSM knows that a node's local copy of a variable is invalid, it will mark the page containing that variable as nonresident. Then, the next time the program at this node tries to access that variable, a page fault will occur.

As mentioned in the review above, normally a page fault causes a jump to the OS. However, technically any page fault in Unix is handled as a signal, specifically SIGSEGV. Recall that Unix allows the programmer to write his/her own signal handler for any signal type. In this case, that means that the programmer—meaning the people who developed JIAJIA or any other page-based SDSM—writes his/her own page fault handler, which will do the necessary network transactions to obtain the latest valid value for `X`.

Note that although SDSMs are able to create an illusion of almost all aspects of shared memory, it really is not possible to create the illusion of shared pointer variables. For example on shared memory hardware we might have a variable like `P`:

¹⁷There are a number of important issues involved with this word *eventually*, as we will see later.

¹⁸The update may not occur immediately. More on this later.

¹⁹Actually, it will read the entire page containing `X` from disk, but to simplify language we will just say `X` here.

²⁰E.g. `iret` on Pentium chips.

```
int Y, *P;
...
...
P = &Y;
...
```

There is no simple way to have a variable like **P** in an SDSM. This is because a pointer is an address, and each node in an SDSM has its own memory separate address space. The problem is that even though the underlying SDSM system will keep the various copies of **Y** at the different nodes consistent with each other, **Y** will be at a potentially different address on each node.

All SDSM systems must deal with a software analog of the cache coherency problem. Whenever one node modifies the value of a shared variable, that node must notify the other nodes that a change has been made. The designer of the system must choose between update or invalidate protocols, just as in the hardware case.²¹ Recall that in non-bus-based shared-memory multiprocessors, one needs to maintain a directory which indicates at which processor a valid copy of a shared variable exists. Again, SDSMs must take an approach similar to this.

Similarly, each SDSM system must decide between sequential consistency, release consistency etc. More on this later.

Note that in the NOW context the internode communication at the SDSM level is typically done by TCP/IP network actions. Treadmarks uses UDP, which is faster than TCP, but still part of the slow TCP/IP protocol suite. TCP/IP was simply not designed for this kind of work. Accordingly, there have been many efforts to use more efficient network hardware and software. The most popular of these is the Virtual Interface Architecture (VIA).

Not only are coherency actions more expensive in the NOW SDSM case than in the shared-memory hardware case due to network slowness, there is also expense due to granularity. In the hardware case we are dealing with cache blocks, with a typical size being 512 bytes. In the SDSM case, we are dealing with pages, with a typical size being 4096 bytes. The overhead for a cache coherency transaction can thus be large.

2.9.0.2 Case Study: JIAJIA

Programmer Interface

We will not go into detail on JIAJIA programming here. There is a short tutorial on JIAJIA at <http://heather.cs.ucdavis.edu/~matloff/jiajia.html>, but here is an overview:

- One writes in C/C++ (or FORTRAN), making calls to the JIAJIA library, which is linked in upon

²¹Note, though, that we are not actually dealing with a cache here. Each node in the SDSM system will have a cache, of course, but a node's cache simply stores parts of that node's set of pages. The coherency across nodes is across pages, not caches. We must insure that a change made to a given page is eventually propagated to pages on other nodes which correspond to this one.

compilation.

- The library calls include standard shared-memory operations for lock, unlock, barrier, processor number, etc., plus some calls aimed at improving performance.

Following is a JIAJIA example program, performing Odd/Even Transposition Sort. This is a variant on Bubble Sort, sometimes useful in parallel processing contexts.²² The algorithm consists of n phases, in which each processor alternates between trading with its left and right neighbors.

```

1 // JIAJIA example program: Odd-Even Tranposition Sort
2
3 // array is of size n, and we use n processors; this would be more
4 // efficient in a "chunked" versions, of course (and more suited for a
5 // message-passing context anyway)
6
7 #include <stdio.h>
8 #include <stdlib.h>
9 #include <jia.h> // required include; also must link via -ljia
10
11 // pointer to shared variable
12 int *x; // array to be sorted
13
14 int n, // range to check for primeness
15     debug; // 1 for debugging, 0 else
16
17 // if first arg is bigger, then replace it by the second
18 void cpsmaller(int *p1,int *p2)
19 { int tmp;
20     if (*p1 > *p2) *p1 = *p2;
21 }
22
23 // if first arg is smaller, then replace it by the second
24 void cpbigger(int *p1,int *p2)
25 { int tmp;
26     if (*p1 < *p2) *p1 = *p2;
27 }
28
29 // does sort of m-element array y
30 void oddeven(int *y, int m)
31 { int i,left=jiapid-1,right=jiapid+1,newval;
32     for (i=0; i < m; i++) {
33         if ((i+jiapid)%2 == 0) {
34             if (right < m)
35                 if (y[jiapid] > y[right]) newval = y[right];
36         }
37         else {
38             if (left >= 0)
39                 if (y[jiapid] < y[left]) newval = y[left];
40         }
41         jia_barrier();
42         if ((i+jiapid)%2 == 0 && right < m || (i+jiapid)%2 == 1 && left >= 0)
43             y[jiapid] = newval;

```

²²Though, as mentioned in the comments, it is aimed more at message-passing contexts.

```

44     jia_barrier();
45 }
46 }
47
48 main(int argc, char **argv)
49 { int i,mywait=0;
50   jia_init(argc,argv); // required init call
51   // get command-line arguments (shifted for nodes > 0)
52   if (jiapid == 0) {
53     n = atoi(argv[1]);
54     debug = atoi(argv[2]);
55   }
56   else {
57     n = atoi(argv[2]);
58     debug = atoi(argv[3]);
59   }
60   jia_barrier();
61   // create a shared array x of length n
62   x = (int *) jia_alloc(n*sizeof(int));
63   // barrier recommended after allocation
64   jia_barrier();
65   // node 0 gets simple test array from command-line
66   if (jiapid == 0) {
67     for (i = 0; i < n; i++)
68       x[i] = atoi(argv[i+3]);
69   }
70   jia_barrier();
71   if (debug && jiapid == 0)
72     while (mywait == 0) { ; }
73   jia_barrier();
74   oddeven(x,n);
75   if (jiapid == 0) {
76     printf("\nfinal array\n");
77     for (i = 0; i < n; i++)
78       printf("%d\n",x[i]);
79   }
80   jia_exit();
81 }

```

System Workings

JIAJIA's main characteristics as an SDSM are:

- page-based
- scope consistency
- home-based
- multiple writers

Let's take a look at these.

As mentioned earlier, one first calls `jia_alloc()` to set up one's shared variables. Note that this will occur at each node, so there are multiple copies of each variable; the JIAJIA system ensures that these copies are consistent with each other, though of course subject to the laxity afforded by scope consistency.

Recall that under scope consistency, a change made to a shared variable at one processor is guaranteed to be made visible to another processor if the first processor made the change between lock/unlock operations and the second processor accesses that variable between lock/unlock operations on that same lock.²³

Each page—and thus each shared variable—has a **home** processor. If another processor writes to a page, then later when it reaches the unlock operation it must send all changes it made to the page back to the home node. In other words, the second processor calls `jia_unlock()`, which sends the changes to its sister invocation of `jia_unlock()` at the home processor.²⁴ Say later a third processor calls `jia_lock()` on that same lock, and then attempts to read a variable in that page. A page fault will occur at that processor, resulting in the JIAJIA system running, which will then obtain that page from the first processor.

Note that all this means the JIAJIA system at each processor must maintain a page table, listing where each home page resides.²⁵ At each processor, each page has one of three states: Invalid, Read-Only, Read-Write. State changes, though, are reported when lock/unlock operations occur. For example, if CPU 5 writes to a given page which had been in Read-Write state at CPU 8, the latter will not hear about CPU 5's action until some CPU does a lock. This CPU need not be CPU 8. When one CPU does a lock, it must coordinate with all other nodes, at which time state-change messages will be piggybacked onto lock-coordination messages.

Note also that JIAJIA allows the programmer to specify which node should serve as the home of a variable, via one of several forms of the `jia_alloc()` call. The programmer can then tailor his/her code accordingly. For example, in a matrix problem, the programmer may arrange for certain rows to be stored at a given node, and then write the code so that most writes to those rows are done by that processor.

The general principle here is that writes performed at one node can be made visible at other nodes on a “need to know” basis. If for instance in the above example with CPUs 5 and 8, CPU 2 does not access this page, it would be wasteful to send the writes to CPU 2, or for that matter to even inform CPU 2 that the page had been written to. This is basically the idea of all non-Sequential consistency protocols, even though they differ in approach and in performance for a given application.

JIAJIA allows multiple writers of a page. Suppose CPU 4 and CPU 15 are simultaneously writing to a particular page, and the programmer has relied on a subsequent barrier to make those writes visible to other

²³Writes will also be propagated at barrier operations, but two successive arrivals by a processor to a barrier can be considered to be a lock/unlock pair, by considering a departure from a barrier to be a “lock,” and considering reaching a barrier to be an “unlock.” So, we'll usually not mention barriers separately from locks in the remainder of this subsection.

²⁴The set of changes is called a **diff**, reminiscent of the Unix file-compare command. A copy, called a **twin**, had been made of the original page, which now will be used to produce the diff. This has substantial overhead. The Treadmarks people found that it took 167 microseconds to make a twin, and as much as 686 microseconds to make a diff.

²⁵In JIAJIA, that location is normally fixed, but JIAJIA does include advanced programmer options which allow the location to migrate.

processors.²⁶ When the barrier is reached, each will be informed of the writes of the other.²⁷ Allowing multiple writers helps to reduce the performance penalty due to false sharing.

2.10 Barrier Implementation

Recall that a **barrier** is program code²⁸ which has a processor do a wait-loop action until all processors have reached that point in the program.²⁹

A function **Barrier()** is often supplied as a library function; here we will see how to implement such a library function in a correct and efficient manner. Note that since a barrier is a serialization point for the program, efficiency is crucial to performance.

Implementing a barrier in a fully correct manner is actually a bit tricky. We'll see here what can go wrong, and how to make sure it doesn't.

In this section, we will approach things from a shared-memory point of view. But the methods apply in the obvious way to message-passing systems as well, as will be discussed later.

2.10.1 A Use-Once Version

```

1 struct BarrStruct {
2     int NNodes, // number of threads participating in the barrier
3     Count, // number of threads that have hit the barrier so far
4     pthread_mutex_t Lock = PTHREAD_MUTEX_INITIALIZER;
5 } ;
6
7 Barrier(struct BarrStruct *PB)
8 { pthread_mutex_lock(&PB->Lock);
9   PB->Count++;
10  pthread_mutex_unlock(&PB->Lock);
11  while (PB->Count < PB->NNodes) ;
12 }
```

This is very simple, actually overly so. This implementation will work once, so if a program using it doesn't make two calls to **Barrier()** it would be fine. But not otherwise. If, say, there is a call to **Barrier()** in a loop, we'd be in trouble.

²⁶The only other option would be to use lock/unlock, but then their writing would not be simultaneous.

²⁷If they are writing to the same variable, not just the same page, the programmer would use locks instead of a barrier, and the situation would not arise.

²⁸Some hardware barriers have been proposed.

²⁹I use the word *processor* here, but it could be just a thread on the one hand, or on the other hand a processing element in a message-passing context.

What is the problem? Clearly, something must be done to reset **Count** to 0 at the end of the call, but doing this safely is not so easy, as seen in the next section.

2.10.2 An Attempt to Write a Reusable Version

Consider the following attempt at fixing the code for **Barrier()**:

```

1 Barrier(struct BarrStruct *PB)
2 { int OldCount;
3   pthread_mutex_lock(&PB->Lock);
4   OldCount = PB->Count++;
5   pthread_mutex_unlock(&PB->Lock);
6   if (OldCount == PB->NNodes-1) PB->Count = 0;
7   while (PB->Count < PB->NNodes) ;
8 }
```

Unfortunately, this doesn't work either. To see why, consider a loop with a barrier call at the end:

```

1 struct BarrStruct B; // global variable
2 .....
3 while (.....) {
4   .....
5   Barrier(&B);
6   .....
7 }
```

At the end of the first iteration of the loop, all the processors will wait at the barrier until everyone catches up. After this happens, one processor, say 12, will reset **B.Count** to 0, as desired. But if we are unlucky, some other processor, say processor 3, will then race ahead, perform the second iteration of the loop in an extremely short period of time, and then reach the barrier and increment the **Count** variable before processor 12 resets it to 0. This would result in disaster, since processor 3's increment would be canceled, leaving us one short when we try to finish the barrier the second time.

Another disaster scenario which might occur is that one processor might reset **B.Count** to 0 before another processor had a chance to notice that **B.Count** had reached **B.NNodes**.

2.10.3 A Correct Version

One way to avoid this would be to have *two* **Count** variables, and have the processors alternate using one then the other. In the scenario described above, processor 3 would increment the *other* **Count** variable, and

thus would not conflict with processor 12's resetting. Here is a safe barrier function based on this idea:

```
1 struct BarrStruct {
2     int NNodes, // number of threads participating in the barrier
3     Count[2], // number of threads that have hit the barrier so far
4     pthread_mutex_t Lock = PTHREAD_MUTEX_INITIALIZER;
5 } ;
6
7 Barrier(struct BarrStruct *PB)
8 { int Par, OldCount;
9   Par = PB->EvenOdd;
10  pthread_mutex_lock(&PB->Lock);
11  OldCount = PB->Count[Par]++;
12  pthread_mutex_unlock(&PB->Lock);
13  if (OldCount == PB->NNodes-1) {
14      PB->Count[Par] = 0;
15      PB->EvenOdd = 1 - Par;
16  }
17  else while (PB->Count[Par] > 0) ;
18 }
```

2.10.4 Refinements

2.10.4.1 Use of Wait Operations

The code

```
else while (PB->Count[Par] > 0) ;
```

is harming performance, since it has the processor spinning around doing no useful work. In the Pthreads context, we can use a condition variable:

```

1  struct BarrStruct {
2      int NNodes, // number of threads participating in the barrier
3          Count[2], // number of threads that have hit the barrier so far
4          pthread_mutex_t Lock = PTHREAD_MUTEX_INITIALIZER;
5          pthread_cond_t CV = PTHREAD_COND_INITIALIZER;
6  };
7
8  Barrier(struct BarrStruct *PB)
9  { int Par, I;
10     Par = PB->EvenOdd;
11     pthread_mutex_lock(&PB->Lock);
12     PB->Count[Par]++;
13     if (PB->Count < PB->NNodes)
14         pthread_cond_wait(&PB->CV, &PB->Lock);
15     else {
16         PB->Count[Par] = 0;
17         PB->EvenOdd = 1 - Par;
18         for (I = 0; I < PB->NNodes-1; I++)
19             pthread_cond_signal(&PB->CV);
20     }
21     pthread_mutex_unlock(&PB->Lock);
22 }
```

Here, if a thread finds that not everyone has reached the barrier yet, it still waits for the rest, but does so passively, via the wait for the condition variable **CV**. This way the thread is not wasting valuable time on that processor, which can run other useful work.

Note that the call to **pthread_cond_wait()** requires use of the lock. Your code must lock the lock before making the call. The call itself immediately unlocks that lock after it registers the wait with the threads manager. But the call blocks until awakened when another thread calls **pthread_cond_signal()** or **pthread_cond_broadcast()**.

It is required that your code lock the lock before calling **pthread_cond_signal()**, and that it unlock the lock after the call.

By using **pthread_cond_wait()** and placing the unlock operation later in the code, as seen above, we actually could get by with just a single **Count** variable, as before.

Even better, the **for** loop could be replaced by a single call

```
pthread_cond_broadcast(&PB->PB->CV);
```

This still wakes up the waiting threads one by one, but in a much more efficient way, and it makes for clearer code.

2.10.4.2 Parallelizing the Barrier Operation

2.10.4.2.1 Tree Barriers It is clear from the code above that barriers can be costly to performance, since they rely so heavily on critical sections, i.e. serial parts of a program. Thus in many settings it is worthwhile to parallelize not only the general computation, but also the barrier operations themselves.

Consider for instance a barrier in which 16 threads are participating. We could speed things up by breaking this barrier down into two sub-barriers, with eight threads each. We would then set up three barrier operations: one of the first group of eight threads, another for the other group of eight threads, and a third consisting of a “competition” between the two groups. The variable **NNodes** above would have the value 8 for the first two barriers, and would be equal to 2 for the third barrier.

Here thread 0 could be the representative for the first group, with thread 4 representing the second group. After both groups’s barriers were hit by all of their members, threads 0 and 4 would participated in the third barrier.

Note that then the notification phase would the be done in reverse: When the third barrier was complete, threads 0 and 4 would notify the members of their groups.

This would parallelize things somewhat, as critical-section operations could be executing simultaneously for the first two barriers. There would still be quite a bit of serial action, though, so we may wish to do further splitting, by partitioning each group of four threads into two subgroups of two threads each.

In general, for n threads (with n , say, equal to a power of 2) we would have a tree structure, with $\log_2 n$ levels in the tree. The i^{th} level (starting with the root as level 0) with consist of 2^i parallel barriers, each one representing $n/2^i$ threads.

2.10.4.2.2 Butterfly Barriers Another method basically consists of each node “shaking hands” with every other node. In the shared-memory case, handshaking could be done by having a global array **ReachedBarrier**. When thread 3 and thread 7 shake hands, for instance, would reach the barrier, thread 3 would set **ReachedBarrier[3]** to 1, and would then wait for **ReachedBarrier[7]** to become 1. The wait, as before, could either be a **while** loop or a call to **pthread_cond_wait()**. Thread 7 would do the opposite.

If we have n nodes, again with n being a power of 2, then the barrier process would consist of $\log_2 n$ phases, which we’ll call phase 0, phase 1, etc. Then the process works as follows.

For any node i , let $i(k)$ be the number obtained by inverting bit k in the binary representation of i , with bit 0 being the least significant bit. Then in the k^{th} phase, node i would shake hands with node $i(k)$.

For example, say $n = 8$. In phase 0, node $5 = 101_2$, say, would shake hands with node $4 = 100_2$.

Actually, a butterfly exchange amounts to a number of simultaneously tree operations.

Chapter 3

The Python Threads and Multiprocessing Modules

Python's thread system builds on the underlying OS threads. They are thus pre-emptible. Note, though, that Python adds its own threads manager on top of the OS thread system; see Section 3.3.

3.1 Python Threads Modules

Python threads are accessible via two modules, **thread.py** and **threading.py**. The former is more primitive, thus easier to learn from, and we will start with it.

3.1.1 The `thread` Module

The example here involves a client/server pair.¹ As you'll see from reading the comments at the start of the files, the program does nothing useful, but is a simple illustration of the principles. We set up two invocations of the client; they keep sending letters to the server; the server concatenates all the letters it receives.

Only the server needs to be threaded. It will have one thread for each client.

Here is the client code, **clnt.py**:

```
1 # simple illustration of thread module
2
```

¹It is preferable here that the reader be familiar with basic network programming. See my tutorial at <http://heather.cs.ucdavis.edu/~matloff/Python/PyNet.pdf>. However, the comments preceding the various network calls would probably be enough for a reader without background in networks to follow what is going on.

```

3 # two clients connect to server; each client repeatedly sends a letter,
4 # stored in the variable k, which the server appends to a global string
5 # v, and reports v to the client; k = '' means the client is dropping
6 # out; when all clients are gone, server prints the final string v
7
8 # this is the client; usage is
9
10 # python clnt.py server_address port_number
11
12 import socket # networking module
13 import sys
14
15 # create Internet TCP socket
16 s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
17
18 host = sys.argv[1] # server address
19 port = int(sys.argv[2]) # server port
20
21 # connect to server
22 s.connect((host, port))
23
24 while(1):
25     # get letter
26     k = raw_input('enter a letter:')
27     s.send(k) # send k to server
28     # if stop signal, then leave loop
29     if k == '': break
30     v = s.recv(1024) # receive v from server (up to 1024 bytes)
31     print v
32
33 s.close() # close socket

```

And here is the server, `svr.py`:

```

1 # simple illustration of thread module
2
3 # multiple clients connect to server; each client repeatedly sends a
4 # letter k, which the server adds to a global string v and echos back
5 # to the client; k = '' means the client is dropping out; when all
6 # clients are gone, server prints final value of v
7
8 # this is the server
9
10 import socket # networking module
11 import sys
12
13 import thread
14
15 # note the globals v and nclnt, and their supporting locks, which are
16 # also global; the standard method of communication between threads is
17 # via globals
18
19 # function for thread to serve a particular client, c
20 def serveclient(c):
21     global v, nclnt, vlock, nclntlock
22     while 1:

```



```

23     # receive letter from c, if it is still connected
24     k = c.recv(1)
25     if k == '': break
26     # concatenate v with k in an atomic manner, i.e. with protection
27     #     by locks
28     vlock.acquire()
29     v += k
30     vlock.release()
31     # send new v back to client
32     c.send(v)
33     c.close()
34     nclntlock.acquire()
35     nclnt -= 1
36     nclntlock.release()
37
38 # set up Internet TCP socket
39 lstn = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
40
41 port = int(sys.argv[1]) # server port number
42 # bind lstn socket to this port
43 lstn.bind(('', port))
44 # start listening for contacts from clients (at most 2 at a time)
45 lstn.listen(5)
46
47 # initialize concatenated string, v
48 v = ''
49 # set up a lock to guard v
50 vlock = thread.allocate_lock()
51
52 # nclnt will be the number of clients still connected
53 nclnt = 2
54 # set up a lock to guard nclnt
55 nclntlock = thread.allocate_lock()
56
57 # accept calls from the clients
58 for i in range(nclnt):
59     # wait for call, then get a new socket to use for this client,
60     #     and get the client's address/port tuple (though not used)
61     (clnt,ap) = lstn.accept()
62     # start thread for this client, with serveclient() as the thread's
63     #     function, with parameter clnt; note that parameter set must be
64     #     a tuple; in this case, the tuple is of length 1, so a comma is
65     #     needed
66     thread.start_new_thread(serveclient, (clnt,))
67
68 # shut down the server socket, since it's not needed anymore
69 lstn.close()
70
71 # wait for both threads to finish
72 while nclnt > 0: pass
73
74 print 'the final value of v is', v

```

Make absolutely sure to run the programs before proceeding further.² Here is how to do this:

²You can get them from the `.tex` source file for this tutorial, located wherever you picked up the `.pdf` version.

I'll refer to the machine on which you run the server as **a.b.c**, and the two client machines as **u.v.w** and **x.y.z**.³ First, on the server machine, type

```
python srvr.py 2000
```

and then on each of the client machines type

```
python clnt.py a.b.c 2000
```

(You may need to try another port than 2000, anything above 1023.)

Input letters into both clients, in a rather random pattern, typing some on one client, then on the other, then on the first, etc. Then finally hit Enter without typing a letter to one of the clients to end the session for that client, type a few more characters in the other client, and then end that session too.

The reason for threading the server is that the inputs from the clients will come in at unpredictable times. At any given time, the server doesn't know which client will send input next, and thus doesn't know on which client to call **recv()**. One way to solve this problem is by having threads, which run "simultaneously" and thus give the server the ability to read from whichever client has sent data.⁴

So, let's see the technical details. We start with the "main" program.⁵

```
vlock = thread.allocate_lock()
```

Here we set up a **lock variable** which guards **v**. We will explain later why this is needed. Note that in order to use this function and others we needed to import the **thread** module.

```
ncnt = 2
ncntlock = thread.allocate_lock()
```

We will need a mechanism to insure that the "main" program, which also counts as a thread, will be passive until both application threads have finished. The variable **ncnt** will serve this purpose. It will be a count of how many clients are still connected. The "main" program will monitor this, and wrap things up later when the count reaches 0.

```
thread.start_new_thread(serveclient, (clnt,))
```

³You could in fact run all of them on the same machine, with address name **localhost** or something like that, but it would be better on separate machines.

⁴Another solution is to use nonblocking I/O. See this example in that context in <http://heather.cs.ucdavis.edu/~matloff/Python/PyNet.pdf>

⁵Just as you should write the main program first, you should read it first too, for the same reasons.

Having accepted a client connection, the server sets up a thread for serving it, via `thread.start_new_thread()`. The first argument is the name of the application function which the thread will run, in this case `serveclient()`. The second argument is a tuple consisting of the set of arguments for that application function. As noted in the comment, this set is expressed as a tuple, and since in this case our tuple has only one component, we use a comma to signal the Python interpreter that this is a tuple.

So, here we are telling Python's threads system to call our function `serveclient()`, supplying that function with the argument `clnt`. The thread becomes "active" immediately, but this does not mean that it starts executing right away. All that happens is that the threads manager adds this new thread to its list of threads, and marks its current state as Run, as opposed to being in a Sleep state, waiting for some event.

By the way, this gives us a chance to show how clean and elegant Python's threads interface is compared to what one would need in C/C++. For example, in `pthread`, the function analogous to `thread.start_new_thread()` has the signature

```
pthread_create (pthread_t *thread_id, const pthread_attr_t *attributes,
               void *(*thread_function)(void *), void *arguments);
```

What a mess! For instance, look at the types in that third argument: A pointer to a function whose argument is pointer to `void` and whose value is a pointer to `void` (all of which would have to be `cast` when called). It's such a pleasure to work in Python, where we don't have to be bothered by low-level things like that.

Now consider our statement

```
while nclnt > 0: pass
```

The statement says that as long as at least one client is still active, do nothing. Sounds simple, and it is, but you should consider what is really happening here.

Remember, the three threads—the two client threads, and the "main" one—will take turns executing, with each turn lasting a brief period of time. Each time "main" gets a turn, it will loop repeatedly on this line. But all that empty looping in "main" is wasted time. What we would really like is a way to prevent the "main" function from getting a turn at all until the two clients are gone. There are ways to do this which you will see later, but we have chosen to remain simple for now.

Now consider the function `serveclient()`. Any thread executing this function will deal with only one particular client, the one corresponding to the connection `c` (an argument to the function). So this `while` loop does nothing but read from that particular client. If the client has not sent anything, the thread will block on the line

```
k = c.recv(1)
```

This thread will then be marked as being in Sleep state by the thread manager, thus allowing the other client thread a chance to run. If neither client thread can run, then the “main” thread keeps getting turns. When a user at one of the clients finally types a letter, the corresponding thread unblocks, i.e. the threads manager changes its state to Run, so that it will soon resume execution.

Next comes the most important code for the purpose of this tutorial:

```
vlock.acquire()
v += k
vlock.release()
```

Here we are worried about a **race condition**. Suppose for example `v` is currently `'abx'`, and Client 0 sends `k` equal to `'g'`. The concern is that this thread's turn might end in the middle of that addition to `v`, say right after the Python interpreter had formed `'abxg'` but before that value was written back to `v`. This could be a big problem. The next thread might get to the same statement, take `v`, still equal to `'abx'`, and append, say, `'w'`, making `v` equal to `'abxw'`. Then when the first thread gets its next turn, it would finish its interrupted action, and set `v` to `'abxg'`—which would mean that the `'w'` from the other thread would be lost.

All of this hinges on whether the operation

```
v += k
```

is interruptible. Could a thread's turn end somewhere in the midst of the execution of this statement? If not, we say that the operation is **atomic**. If the operation were atomic, we would not need the lock/unlock operations surrounding the above statement. I did this, using the methods described in Section 3.3.4.1, and it appears to me that the above statement is *not* atomic.

Moreover, it's safer not to take a chance, especially since Python compilers could vary or the virtual machine could change; after all, we would like our Python source code to work even if the machine changes.

So, we need the lock/unlock operations:

```
vlock.acquire()
v += k
vlock.release()
```

The lock, **vlock** here, can only be held by one thread at a time. When a thread executes this statement, the Python interpreter will check to see whether the lock is locked or unlocked right now. In the latter case, the interpreter will lock the lock and the thread will continue, and will execute the statement which updates `v`. It will then release the lock, i.e. the lock will go back to unlocked state.

If on the other hand, when a thread executes **acquire()** on this lock when it is locked, i.e. held by some other thread, its turn will end and the interpreter will mark this thread as being in Sleep state, waiting for the lock

to be unlocked. When whichever thread currently holds the lock unlocks it, the interpreter will change the blocked thread from Sleep state to Run state.

Note that if our threads were non-preemptive, we would not need these locks.

Note also the crucial role being played by the global nature of `v`. Global variables are used to communicate between threads. In fact, recall that this is one of the reasons that threads are so popular—easy access to global variables. Thus the dogma so often taught in beginning programming courses that global variables must be avoided is wrong; on the contrary, there are many situations in which globals are necessary and natural.⁶

The same race-condition issues apply to the code

```
nc1ntlock.acquire()
nc1nt -= 1
nc1ntlock.release()
```

Following is a Python program that finds prime numbers using threads. Note carefully that it is not claimed to be efficient at all (it may well run more slowly than a serial version); it is merely an illustration of the concepts. Note too that we are using the simple **thread** module, rather than **threading**.

```
1  #!/usr/bin/env python
2
3  import sys
4  import math
5  import thread
6
7  def dowork(tn): # thread number tn
8      global n,prime,nexti,nextilock,nstarted,nstartedlock,donelock
9      donelock[tn].acquire()
10     nstartedlock.acquire()
11     nstarted += 1
12     nstartedlock.release()
13     lim = math.sqrt(n)
14     nk = 0
15     while 1:
16         nextilock.acquire()
17         k = nexti
18         nexti += 1
19         nextilock.release()
20         if k > lim: break
21         nk += 1
22         if prime[k]:
23             r = n / k
24             for i in range(2,r+1):
25                 prime[i*k] = 0
26     print 'thread', tn, 'exiting; processed', nk, 'values of k'
27     donelock[tn].release()
```

⁶I think that dogma is presented in a far too extreme manner anyway. See <http://heather.cs.ucdavis.edu/~matloff/globals.html>.

```

28
29 def main():
30     global n, prime, nexti, nextilock, nstarted, nstartedlock, donelock
31     n = int(sys.argv[1])
32     prime = (n+1) * [1]
33     nthreads = int(sys.argv[2])
34     nstarted = 0
35     nexti = 2
36     nextilock = thread.allocate_lock()
37     nstartedlock = thread.allocate_lock()
38     donelock = []
39     for i in range(nthreads):
40         d = thread.allocate_lock()
41         donelock.append(d)
42         thread.start_new_thread(dowork, (i,))
43     while nstarted < nthreads: pass
44     for i in range(nthreads):
45         donelock[i].acquire()
46     print 'there are', reduce(lambda x,y: x+y, prime) - 2, 'primes'
47
48 if __name__ == '__main__':
49     main()

```

So, let's see how the code works.

The algorithm is the famous Sieve of Erathosthenes: We list all the numbers from 2 to **n**, then cross out all multiples of 2 (except 2), then cross out all multiples of 3 (except 3), and so on. The numbers which get crossed out are composite, so the ones which remain at the end are prime.

Line 32: We set up an array **prime**, which is what we will be “crossing out.” The value 1 means “not crossed out,” so we start everything at 1. (Note how Python makes this easy to do, using list “multiplication.”)

Line 33: Here we get the number of desired threads from the command line.

Line 34: The variable **nstarted** will show how many threads have already started. This will be used later, in Lines 43-45, in determining when the **main()** thread exits. Since the various threads will be writing this variable, we need to protect it with a lock, on Line 37.

Lines 35-36: The variable **nexti** will say which value we should do “crossing out” by next. If this is, say, 17, then it means our next task is to cross out all multiples of 17 (except 17). Again we need to protect it with a lock.

Lines 39-42: We create the threads here. The function executed by the threads is named **dowork()**. We also create locks in an array **donelock**, which again will be used later on as a mechanism for determining when **main()** exits (Line 44-45).

Lines 43-45: There is a lot to discuss here.

To start, recall that in **svr.py**, our example in Section 3.1.1, we didn't want the main thread to exit until the

child threads were done.⁷ So, Line 50 was a **busy wait**, repeatedly doing nothing (**pass**). That’s a waste of time—each time the main thread gets a turn to run, it repeatedly executes **pass** until its turn is over.

Here in our primes program, a premature exit by **main()** result in printing out wrong answers. On the other hand, we don’t want **main()** to engage in a wasteful busy wait. We could use **join()** from **threading.Thread** for this purpose, to be discussed later, but here we take a different tack: We set up a list of locks, one for each thread, in a list **donelock**. Each thread initially acquires its lock (Line 9), and releases it when the thread finishes its work (Lin 27). Meanwhile, **main()** has been waiting to acquire those locks (Line 45). So, when the threads finish, **main()** will move on to Line 46 and print out the program’s results.

But there is a subtle problem (threaded programming is notorious for subtle problems), in that there is no guarantee that a thread will execute Line 9 before **main()** executes Line 45. That’s why we have a busy wait in Line 43, to make sure all the threads acquire their locks before **main()** does. Of course, we’re trying to avoid busy waits, but this one is quick.

Line 13: We need not check any “crosser-outers” that are larger than \sqrt{n} .

Lines 15-25: We keep trying “crosser-outers” until we reach that limit (Line 20). Note the need to use the lock in Lines 16-19. In Line 22, we check the potential “crosser-outer” for primeness; if we have previously crossed it out, we would just be doing duplicate work if we used this **k** as a “crosser-outer.”

Here’s one more example, a type of Web crawler. This one continually monitors the access time of the Web, by repeatedly accessing a list of “representative” Web sites, say the top 100. What’s really different about this program, though, is that we’ve reserved one thread for human interaction. The person can, whenever he/she desires, find for instance the mean of recent access times.

```

1 import sys
2 import time
3 import os
4 import thread
5
6 class gblbs:
7     acctimes = [] # access times
8     acclock = thread.allocate_lock() # lock to guard access time data
9     nextprobe = 0 # index of next site to probe
10    nextprobelock = thread.allocate_lock() # lock to guard access time data
11    sites = open(sys.argv[1]).readlines() # the sites to monitor
12    ww = int(sys.argv[2]) # window width
13
14 def probe(me):
15     if me > 0:
16         while 1:
17             # determine what site to probe next
18             gblbs.nextprobelock.acquire()
19             i = gblbs.nextprobe
20             i1 = i + 1
21             if i1 >= len(gblbs.sites): i1 = 0

```

⁷The effect of the main thread ending earlier would depend on the underlying OS. On some platforms, exit of the parent may terminate the child threads, but on other platforms the children continue on their own.

```

22     glbls.nextprobe = i1
23     glbls.nextprobelock.release()
24     # do probe
25     t1 = time.time()
26     os.system('wget --spider -q '+glbls.sites[i1])
27     t2 = time.time()
28     accesstime = t2 - t1
29     glbls.accllock.acquire()
30     # list full yet?
31     if len(glbls.acctimes) < glbls.ww:
32         glbls.acctimes.append(accesstime)
33     else:
34         glbls.acctimes = glbls.acctimes[1:] + [accesstime]
35     glbls.accllock.release()
36 else:
37     while 1:
38         rsp = raw_input('monitor: ')
39         if rsp == 'mean': print mean(glbls.acctimes)
40         elif rsp == 'median': print median(glbls.acctimes)
41         elif rsp == 'all': print all(glbls.acctimes)
42
43 def mean(x):
44     return sum(x)/len(x)
45
46 def median(x):
47     y = x
48     y.sort()
49     return y[len(y)/2] # a little sloppy
50
51 def all(x):
52     return x
53
54 def main():
55     nthr = int(sys.argv[3]) # number of threads
56     for thr in range(nthr):
57         thread.start_new_thread(probe, (thr,))
58     while 1: continue
59
60 if __name__ == '__main__':
61     main()
62

```

3.1.2 The threading Module

The program below treats the same network client/server application considered in Section 3.1.1, but with the more sophisticated **threading** module. The client program stays the same, since it didn't involve threads in the first place. Here is the new server code:

```

1 # simple illustration of threading module
2
3 # multiple clients connect to server; each client repeatedly sends a
4 # value k, which the server adds to a global string v and echos back
5 # to the client; k = '' means the client is dropping out; when all

```



```

6 # clients are gone, server prints final value of v
7
8 # this is the server
9
10 import socket # networking module
11 import sys
12 import threading
13
14 # class for threads, subclassed from threading.Thread class
15 class srvr(threading.Thread):
16     # v and vlock are now class variables
17     v = ''
18     vlock = threading.Lock()
19     id = 0 # I want to give an ID number to each thread, starting at 0
20     def __init__(self,clntsock):
21         # invoke constructor of parent class
22         threading.Thread.__init__(self)
23         # add instance variables
24         self.myid = srvr.id
25         srvr.id += 1
26         self.myclntsock = clntsock
27     # this function is what the thread actually runs; the required name
28     # is run(); threading.Thread.start() calls threading.Thread.run(),
29     # which is always overridden, as we are doing here
30     def run(self):
31         while 1:
32             # receive letter from client, if it is still connected
33             k = self.myclntsock.recv(1)
34             if k == '': break
35             # update v in an atomic manner
36             srvr.vlock.acquire()
37             srvr.v += k
38             srvr.vlock.release()
39             # send new v back to client
40             self.myclntsock.send(srvr.v)
41             self.myclntsock.close()
42
43 # set up Internet TCP socket
44 lstn = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
45 port = int(sys.argv[1]) # server port number
46 # bind lstn socket to this port
47 lstn.bind(('', port))
48 # start listening for contacts from clients (at most 2 at a time)
49 lstn.listen(5)
50
51 nclnt = int(sys.argv[2]) # number of clients
52
53 mythreads = [] # list of all the threads
54 # accept calls from the clients
55 for i in range(nclnt):
56     # wait for call, then get a new socket to use for this client,
57     # and get the client's address/port tuple (though not used)
58     (clnt,ap) = lstn.accept()
59     # make a new instance of the class srvr
60     s = srvr(clnt)
61     # keep a list all threads
62     mythreads.append(s)
63     # threading.Thread.start calls threading.Thread.run(), which we

```

```

64     # overrode in our definition of the class srvr
65     s.start()
66
67 # shut down the server socket, since it's not needed anymore
68 lstn.close()
69
70 # wait for all threads to finish
71 for s in mythreads:
72     s.join()
73
74 print 'the final value of v is', srvr.v

```

Again, let's look at the main data structure first:

```
class srvr(threading.Thread):
```

The **threading** module contains a class **Thread**, any instance of which represents one thread. A typical application will subclass this class, for two reasons. First, we will probably have some application-specific variables or methods to be used. Second, the class **Thread** has a member method **run()** which is meant to be overridden, as you will see below.

Consistent with OOP philosophy, we might as well put the old globals in as class variables:

```
v = ''
vlock = threading.Lock()
```

Note that class variable code is executed immediately upon execution of the program, as opposed to when the first object of this class is created. So, the lock is created right away.

```
id = 0
```

This is to set up ID numbers for each of the threads. We don't use them here, but they might be useful in debugging or in future enhancement of the code.

```
def __init__(self, clntsock):
    ...
    self.myclntsock = clntsock

# ``main`` program
...
    (clnt, ap) = lstn.accept()
    s = srvr(clnt)

```

The "main" program, in creating an object of this class for the client, will pass as an argument the socket for that client. We then store it as a member variable for the object.

```
def run(self):
    ...
```

As noted earlier, the **Thread** class contains a member method **run()**. This is a dummy, to be overridden with the application-specific function to be run by the thread. It is invoked by the method **Thread.start()**, called here in the main program. As you can see above, it is pretty much the same as the previous code in Section 3.1.1 which used the **thread** module, adapted to the class environment.

One thing that is quite different in this program is the way we end it:

```
for s in mythreads:
    s.join()
```

The **join()** method in the class **Thread** blocks until the given thread exits. (The threads manager puts the main thread in Sleep state, and when the given thread exits, the manager changes that state to Run.) The overall effect of this loop, then, is that the main program will wait at that point until all the threads are done. They “join” the main program. This is a much cleaner approach than what we used earlier, and it is also more efficient, since the main program will not be given any turns in which it wastes time looping around doing nothing, as in the program in Section 3.1.1 in the line

```
while nclnt > 0: pass
```

Here we maintained our own list of threads. However, we could also get one via the call **threading.enumerate()**. If placed after the **for** loop in our server code above, for instance as

```
print threading.enumerate()
```

we would get output like

```
[<_MainThread(MainThread, started)>, <srvr(Thread-1, started)>,
<srvr(Thread-2, started)>]
```

Here’s another example, which finds and counts prime numbers, again not assumed to be efficient:

```
1 #!/usr/bin/env python
2
3 # prime number counter, based on Python threading class
4
5 # usage: python PrimeThreading.py n nthreads
6 #   where we wish the count of the number of primes from 2 to n, and to
7 #   use nthreads to do the work
8
```

```

9 # uses Sieve of Erathosthenes: write out all numbers from 2 to n, then
10 # cross out all the multiples of 2, then of 3, then of 5, etc., up to
11 # sqrt(n); what's left at the end are the primes
12
13 import sys
14 import math
15 import threading
16
17 class prmfinder(threading.Thread):
18     n = int(sys.argv[1])
19     nthreads = int(sys.argv[2])
20     thrdlist = [] # list of all instances of this class
21     prime = (n+1) * [1] # 1 means assumed prime, until find otherwise
22     nextk = 2 # next value to try crossing out with
23     nextklock = threading.Lock()
24     def __init__(self, id):
25         threading.Thread.__init__(self)
26         self.myid = id
27     def run(self):
28         lim = math.sqrt(prmfinder.n)
29         nk = 0 # count of k's done by this thread, to assess load balance
30         while 1:
31             # find next value to cross out with
32             prmfinder.nextklock.acquire()
33             k = prmfinder.nextk
34             prmfinder.nextk += 1
35             prmfinder.nextklock.release()
36             if k > lim: break
37             nk += 1 # increment workload data
38             if prmfinder.prime[k]: # now cross out
39                 r = prmfinder.n / k
40                 for i in range(2, r+1):
41                     prmfinder.prime[i*k] = 0
42             print 'thread', self.myid, 'exiting; processed', nk, 'values of k'
43
44 def main():
45     for i in range(prmfinder.nthreads):
46         pf = prmfinder(i) # create thread i
47         prmfinder.thrdlist.append(pf)
48         pf.start()
49     for thrd in prmfinder.thrdlist: thrd.join()
50     print 'there are', reduce(lambda x,y: x+y, prmfinder.prime) - 2, 'primes'
51
52 if __name__ == '__main__':
53     main()

```

3.2 Condition Variables

3.2.1 General Ideas

We saw in the last section that **threading.Thread.join()** avoids the need for wasteful looping in **main()**, while the latter is waiting for the other threads to finish. In fact, it is very common in threaded programs to

have situations in which one thread needs to wait for something to occur in another thread. Again, in such situations we would not want the waiting thread to engage in wasteful looping.

The solution to this problem is **condition variables**. As the name implies, these are variables used by code to wait for a certain condition to occur. Most threads systems include provisions for these, and Python's **threading** package is no exception.

The **threads** package, for instance, has a type **pthread_cond** for such variables, and has functions such as **pthread_cond_wait()**, which a thread calls to wait for an event to occur, and **pthread_cond_signal()**, which another thread calls to announce that the event now has occurred.

But as is typical with Python in so many things, it is easier for us to use condition variables in Python than in C. At the first level, there is the class **threading.Condition**, which corresponds well to the condition variables available in most threads systems. However, at this level condition variables are rather cumbersome to use, as not only do we need to set up condition variables but we also need to set up extra locks to guard them. This is necessary in any threading system, but it is a nuisance to deal with.

So, Python offers a higher-level class, **threading.Event**, which is just a wrapper for **threading.Condition**, but which does all the condition lock operations behind the scenes, alleviating the programmer of having to do this work.

3.2.2 Event Example

Following is an example of the use of **threading.Event**. It searches a given network host for servers at various ports on that host. (This is called a **port scanner**.) As noted in a comment, the threaded operation used here would make more sense if many hosts were to be scanned, rather than just one, as each **connect()** operation does take some time. But even on the same machine, if a server is active but busy enough that we never get to connect to it, it may take a long for the attempt to timeout. It is common to set up Web operations to be threaded for that reason. We could also have each thread check a block of ports on a host, not just one, for better efficiency.

The use of threads is aimed at checking many ports in parallel, one per thread. The program has a self-imposed limit on the number of threads. If **main()** is ready to start checking another port but we are at the thread limit, the code in **main()** waits for the number of threads to drop below the limit. This is accomplished by a condition wait, implemented through the **threading.Event** class.

```

1 # portscanner.py: checks for active ports on a given machine; would be
2 # more realistic if checked several hosts at once; different threads
3 # check different ports; there is a self-imposed limit on the number of
4 # threads, and the event mechanism is used to wait if that limit is
5 # reached
6
7 # usage: python portscanner.py host maxthreads
8
```

```
9 import sys, threading, socket
10
11 class scanner(threading.Thread):
12     tlist = [] # list of all current scanner threads
13     maxthreads = int(sys.argv[2]) # max number of threads we're allowing
14     evnt = threading.Event() # event to signal OK to create more threads
15     lck = threading.Lock() # lock to guard tlist
16     def __init__(self,tn,host):
17         threading.Thread.__init__(self)
18         self.threadnum = tn # thread ID/port number
19         self.host = host # checking ports on this host
20     def run(self):
21         s = socket.socket(socket.AF_INET,socket.SOCK_STREAM)
22         try:
23             s.connect((self.host, self.threadnum))
24             print "%d: successfully connected" % self.threadnum
25             s.close()
26         except:
27             print "%d: connection failed" % self.threadnum
28             # thread is about to exit; remove from list, and signal OK if we
29             # had been up against the limit
30             scanner.lck.acquire()
31             scanner.tlist.remove(self)
32             print "%d: now active --" % self.threadnum, scanner.tlist
33             if len(scanner.tlist) == scanner.maxthreads-1:
34                 scanner.evnt.set()
35                 scanner.evnt.clear()
36             scanner.lck.release()
37     def newthread(pn,hst):
38         scanner.lck.acquire()
39         sc = scanner(pn,hst)
40         scanner.tlist.append(sc)
41         scanner.lck.release()
42         sc.start()
43         print "%d: starting check" % pn
44         print "%d: now active --" % pn, scanner.tlist
45     newthread = staticmethod(newthread)
46
47 def main():
48     host = sys.argv[1]
49     for i in range(1,100):
50         scanner.lck.acquire()
51         print "%d: attempting check" % i
52         # check to see if we're at the limit before starting a new thread
53         if len(scanner.tlist) >= scanner.maxthreads:
54             # too bad, need to wait until not at thread limit
55             print "%d: need to wait" % i
56             scanner.lck.release()
57             scanner.evnt.wait()
58         else:
59             scanner.lck.release()
60             scanner.newthread(i,host)
61     for sc in scanner.tlist:
62         sc.join()
63
64 if __name__ == '__main__':
65     main()
```

As you can see, when **main()** discovers that we are at our self-imposed limit of number of active threads, we back off by calling **threading.Event.wait()**. At that point **main()**—which, recall, is also a thread—blocks. It will not be given any more timeslices for the time being. When some active thread exits, we have it call **threading.Event.set()** and **threading.Event.clear()**. The threads manager reacts to the former by moving all threads which had been waiting for this event—in our case here, only **main()**—from Sleep state to Run state; **main()** will eventually get another timeslice.

The call to **threading.Event.clear()** is crucial. The word *clear* here means that **threading.Event.clear()** is clearing the occurrence of the event. Without this, any subsequent call to **threading.Event.wait()** would immediately return, even though the condition has not been met yet.

Note carefully the use of locks. The **main()** thread adds items to **tlist**, while the other threads delete items (delete themselves, actually) from it. These operations must be atomic, and thus must be guarded by locks.

I've put in a lot of extra **print** statements so that you can get an idea as to how the threads' execution is interleaved. Try running the program.⁸ But remember, the program may appear to hang for a long time if a server is active but so busy that the attempt to connect times out.

3.2.3 Other threading Classes

The function **Event.set()** “wakes” all threads that are waiting for the given event. That didn't matter in our example above, since only one thread (**main()**) would ever be waiting at a time in that example. But in more general applications, we sometimes want to wake only one thread instead of all of them. For this, we can revert to working at the level of **threading.Condition** instead of **threading.Event**. There we have a choice between using **notify()** or **notifyAll()**.

The latter is actually what is called internally by **Event.set()**. But **notify()** instructs the threads manager to wake just one of the waiting threads (we don't know which one).

The class **threading.Semaphore** offers semaphore operations. Other classes of advanced interest are **threading.RLock** and **threading.Timer**.

3.3 Threads Internals

The thread manager acts like a “mini-operating system.” Just like a real OS maintains a table of processes, a thread system's thread manager maintains a table of threads. When one thread gives up the CPU, or has its turn pre-empted (see below), the thread manager looks in the table for another thread to activate. Whichever thread is activated will then resume execution where it had left off, i.e. where its last turn ended.

⁸Disclaimer: Not guaranteed to be bug-free.

Just as a process is either in Run state or Sleep state, the same is true for a thread. A thread is either ready to be given a turn to run, or is waiting for some event. The thread manager will keep track of these states, decide which thread to run when another has lost its turn, etc.

3.3.1 Kernel-Level Thread Managers

Here each thread really is a process, and for example will show up on Unix systems when one runs the appropriate **ps** process-list command, say **ps axH**. The threads manager is then the OS.

The different threads set up by a given application program take turns running, among all the other processes.

This kind of thread system is used in the Unix **pthreads** system, as well as in Windows threads.

3.3.2 User-Level Thread Managers

User-level thread systems are “private” to the application. Running the **ps** command on a Unix system will show only the original application running, not all the threads it creates. Here the threads are not pre-empted; on the contrary, a given thread will continue to run until it voluntarily gives up control of the CPU, either by calling some “yield” function or by calling a function by which it requests a wait for some event to occur.⁹

A typical example of a user-level thread system is **pth**.

3.3.3 Comparison

Kernel-level threads have the advantage that they can be used on multiprocessor systems, thus achieving true parallelism between threads. This is a major advantage.

On the other hand, in my opinion user-level threads also have a major advantage in that they allow one to produce code which is much easier to write, is easier to debug, and is cleaner and clearer. This in turn stems from the non-preemptive nature of user-level threads; application programs written in this manner typically are not cluttered up with lots of lock/unlock calls (details on these below), which are needed in the pre-emptive case.

3.3.4 The Python Thread Manager

Python “piggybacks” on top of the OS’ underlying threads system. A Python thread is a real OS thread. If a Python program has three threads, for instance, there will be three entries in the **ps** output.

⁹In typical user-level thread systems, an external event, such as an I/O operation or a signal, will also cause the current thread to relinquish the CPU.

However, Python’s thread manager imposes further structure on top of the OS threads. It keeps track of how long a thread has been executing, in terms of the number of Python **byte code** instructions that have executed.¹⁰ When that reaches a certain number, by default 100, the thread’s turn ends. In other words, the turn can be pre-empted either by the hardware timer and the OS, or when the interpreter sees that the thread has executed 100 byte code instructions.¹¹

3.3.4.1 The GIL

In the case of CPython (but not Jython or Iron Python) Most importantly, there is a global interpreter lock, the famous (or infamous) GIL. It is set up to ensure that only one thread runs at a time, in order to facilitate easy garbage collection.

Suppose we have a C program with three threads, which I’ll call X, Y and Z. Say currently Y is running. After 30 milliseconds (or whatever the quantum size has been set to by the OS), Y will be interrupted by the timer, and the OS will start some other process. Say the latter, which I’ll call Q, is a different, unrelated program. Eventually Q’s turn will end too, and let’s say that the OS then gives X a turn. From the point of view of our X/Y/Z program, i.e. ignoring Q, control has passed from Y to X. The key point is that the point within Y at which that event occurs is random (with respect to where Y is at the time), based on the time of the hardware interrupt.

By contrast, say my Python program has three threads, U, V and W. Say V is running. The hardware timer will go off at a random time, and again Q might be given a turn, *but* definitely neither U nor W will be given a turn, because the Python interpreter had earlier made a call to the OS which makes U and W wait for the GIL to become unlocked.

Let’s look at this a little closer. The key point to note is that the Python interpreter itself is threaded, say using **pthreads**. For instance, in our X/Y/Z example above, when you ran **ps axH**, you would see three Python processes/threads. I just tried that on my program **thsvr.py**, which creates two threads, with a command-line argument of 2000 for that program. Here is the relevant portion of the output of **ps axH**:

```
28145 pts/5    Rl      0:09 python thsvr.py 2000
28145 pts/5    Sl      0:00 python thsvr.py 2000
28145 pts/5    Sl      0:00 python thsvr.py 2000
```

What has happened is the Python interpreter has spawned two child threads, one for each of my threads in **thsvr.py**, in addition to the interpreter’s original thread, which runs my **main()**. Let’s call those threads UP, VP and WP. Again, these are the threads that the OS sees, while U, V and W are the threads that I see—or think I see, since they are just virtual.

¹⁰This is the “machine language” for the Python virtual machine.

¹¹The author thanks Alex Martelli for a helpful clarification.

The GIL is a **pthread**s lock. Say V is now running. Again, what that actually means on my real machine is that VP is running. VP keeps track of how long V has been executing, in terms of the number of Python **byte code** instructions that have executed. When that reaches a certain number, by default 100, UP will release the GIL by calling **pthread_mutex_unlock()** or something similar.

The OS then says, “Oh, were any threads waiting for that lock?” It then basically gives a turn to UP or WP (we can’t predict which), which then means that from my point of view U or W starts, say U. Then VP and WP are still in Sleep state, and thus so are my V and W.

So you can see that it is the Python interpreter, not the hardware timer, that is determining how long a thread’s turn runs, relative to the other threads in my program. Again, Q might run too, but within this Python program there will be no control passing from V to U or W simply because the timer went off; such a control change will only occur when the Python interpreter wants it to. This will be either after the 100 byte code instructions or when U reaches an I/O operation or other wait-event operation.

So, the bottom line is that while Python uses the underlying OS threads system as its base, it superimposes further structure in terms of transfer of control between threads.

3.3.4.2 Implications for Randomness and Need for Locks

I mentioned in Section 3.3.2 that non-pre-emptive threading is nice because one can avoid the code clutter of locking and unlocking (details of lock/unlock below). Since, barring I/O issues, a thread working on the same data would seem to always yield control at exactly the same point (i.e. at 100 byte code instruction boundaries), Python would seem to be deterministic and non-pre-emptive. However, it will not quite be so simple.

First of all, there is the issue of I/O, which adds randomness. There may also be randomness in how the OS chooses the first thread to be run, which could affect computation order and so on.

Finally, there is the question of atomicity in Python operations: The interpreter will treat any Python virtual machine instruction as indivisible, thus not needing locks in that case. But the bottom line will be that unless you know the virtual machine well, you should use locks at all times.

3.4 The multiprocessing Module

CPython’s GIL is the subject of much controversy. As we saw in Section 3.3.4.1, it prevents running true parallel applications when using the **thread** or **threading** modules.

That might not seem to be too severe a restriction—after all if you really need the speed, you probably won’t use a scripting language in the first place. But a number of people took the point of view that, given that they have decided to use Python no matter what, they would like to get the best speed subject to that restriction.

So, there was much grumbling about the GIL.

Thus, later the **multiprocessing** module was developed, which enables true parallel processing with Python on a multiprocessor machine, with an interface very close to that of the **threading** module.

Moreover, one can run a program across machines! In other words, the **multiprocessing** module allows to run several threads not only on the different cores of one machine, but on many machines at once, in cooperation in the same manner that threads cooperate on one machine. By the way, this idea is similar to something I did for Perl some years ago (PerlDSM: A Distributed Shared Memory System for Perl. *Proceedings of PDPTA 2002*, 63-68). We will not cover the cross-machine case here.

So, let's go to our first example, a simulation application that will find the probability of getting a total of exactly k dots when we roll n dice:

```

1 # dice probability finder, based on Python multiprocessing class
2
3 # usage: python DiceProb.py n k nreps nthreads
4 #   where we wish to find the probability of getting a total of k dots
5 #   when we roll n dice; we'll use nreps total repetitions of the
6 #   simulation, dividing those repetitions among nthreads threads
7
8 import sys
9 import random
10 from multiprocessing import Process, Lock, Value
11
12 class gbls: # globals, other than shared
13     n = int(sys.argv[1])
14     k = int(sys.argv[2])
15     nreps = int(sys.argv[3])
16     nthreads = int(sys.argv[4])
17     thrdlist = [] # list of all instances of this class
18
19 def worker(id,tot,totlock):
20     mynreps = gbls.nreps/gbls.nthreads
21     r = random.Random() # set up random number generator
22     count = 0 # number of times get total of k
23     for i in range(mynreps):
24         if rolldice(r) == gbls.k:
25             count += 1
26     totlock.acquire()
27     tot.value += count
28     totlock.release()
29     # check for load balance
30     print 'thread', id, 'exiting; total was', count
31
32 def rolldice(r):
33     ndots = 0
34     for roll in range(gbls.n):
35         dots = r.randint(1,6)
36         ndots += dots
37     return ndots
38
39 def main():
40     tot = Value('i',0)

```

```

41     totlock = Lock()
42     for i in range(glbls.nthreads):
43         pr = Process(target=worker, args=(i,tot,totlock))
44         glbls.thrdlist.append(pr)
45         pr.start()
46     for thrd in glbls.thrdlist: thrd.join()
47     # adjust for truncation, in case nthreads doesn't divide nreps evenly
48     actualnreps = glbls.nreps/ghbls.nthreads * glbls.nthreads
49     print 'the probability is',float(tot.value)/actualnreps
50
51 if __name__ == '__main__':
52     main()

```

As in any simulation, the longer one runs it, the better the accuracy is likely to be. Here we run the simulation **nreps** times, but divide those repetitions among the threads. This is an example of an “embarrassingly parallel” application, so we should get a good speedup (not shown here).

So, how does it work? The general structure looks similar to that of the Python **threading** module, using **Process()** to create a thread, **start()** to get it running, **Lock()** to create a lock, **acquire()** and **release()** to lock and unlock a lock, and so on.

The main difference, though, is that globals are not automatically shared. Instead, shared variables must be created using **Value** for a scalar and **Array** for an array. Unlike Python in general, here one must specify a data type, ‘i’ for integer and ‘d’ for double (floating-point). (One can use **Namespace** to create more complex types, at some cost in performance.) One also specifies the initial value of the variable. One must pass these variables explicitly to the functions to be run by the threads, in our case above the function **worker()**. Note carefully that the shared variables are still accessed syntactically as if they were globals.

Here’s the prime number-finding program from before, now using **multiprocessing**:

```

1  #!/usr/bin/env python
2
3  # prime number counter, based on Python multiprocessing class
4
5  # usage:  python PrimeThreading.py n nthreads
6  #   where we wish the count of the number of primes from 2 to n, and to
7  #   use nthreads to do the work
8
9  # uses Sieve of Erathosthenes:  write out all numbers from 2 to n, then
10 # cross out all the multiples of 2, then of 3, then of 5, etc., up to
11 # sqrt(n); what's left at the end are the primes
12
13 import sys
14 import math
15 from multiprocessing import Process, Lock, Array, Value
16
17 class glbls: # globals, other than shared
18     n = int(sys.argv[1])
19     nthreads = int(sys.argv[2])
20     thrdlist = [] # list of all instances of this class
21

```

```

22 def prmfinder(id,prm,nxtk,nxtklock):
23     lim = math.sqrt(glbls.n)
24     nk = 0 # count of k's done by this thread, to assess load balance
25     while 1:
26         # find next value to cross out with
27         nxtklock.acquire()
28         k = nxtk.value
29         nxtk.value = nxtk.value + 1
30         nxtklock.release()
31         if k > lim: break
32         nk += 1 # increment workload data
33         if prm[k]: # now cross out
34             r = glbls.n / k
35             for i in range(2,r+1):
36                 prm[i*k] = 0
37     print 'thread', id, 'exiting; processed', nk, 'values of k'
38
39 def main():
40     prime = Array('i', (glbls.n+1) * [1]) # 1 means prime, until find otherwise
41     nextk = Value('i',2) # next value to try crossing out with
42     nextklock = Lock()
43     for i in range(glbls.nthreads):
44         pf = Process(target=prmfinder, args=(i,prime,nextk,nextklock))
45         glbls.thrdlist.append(pf)
46         pf.start()
47     for thrd in glbls.thrdlist: thrd.join()
48     print 'there are', reduce(lambda x,y: x+y, prime) - 2, 'primes'
49
50 if __name__ == '__main__':
51     main()

```

The main new item in this example is use of **Array()**.

One can use the **Pool** class to create a set of threads, rather than doing so “by hand” in a loop as above. You can start them with various initial values for the threads using **Pool.map()**, which works similarly to Python’s ordinary **map()** function.

The **multiprocessing** documentation warns that shared items may be costly, and suggests using **Queue** and **Pipe** where possible. We will cover the former in the next section. Note, though, that in general it’s difficult to get much speedup (or difficult even to avoid slowdown!) with non-“embarrassingly parallel” applications.

3.5 The Queue Module for Threads and Multiprocessing

Threaded applications often have some sort of work queue data structure. When a thread becomes free, it will pick up work to do from the queue. When a thread creates a task, it will add that task to the queue.

Clearly one needs to guard the queue with locks. But Python provides the **Queue** module to take care of all the lock creation, locking and unlocking, and so on. This means we don’t have to bother with it, and the code will probably be faster.

Queue is implemented for both **threading** and **multiprocessing**, in almost identical forms. This is good, because the documentation for **multiprocessing** is rather sketchy, so you can turn to the docs for `threading` for more details.

The function `put()` in `Queue` adds an element to the end of the queue, while `get()` will remove the head of the queue, again without the programmer having to worry about race conditions.

Note that `get()` will block if the queue is currently empty. An alternative is to call it with `block=False`, within a `try/except` construct. One can also set timeout periods.

Here once again is the prime number example, this time done with **Queue**:

```

1  #!/usr/bin/env python
2
3  # prime number counter, based on Python multiprocessing class with
4  # Queue
5
6  # usage: python PrimeThreading.py n nthreads
7  #   where we wish the count of the number of primes from 2 to n, and to
8  #   use nthreads to do the work
9
10 # uses Sieve of Erathosthenes: write out all numbers from 2 to n, then
11 # cross out all the multiples of 2, then of 3, then of 5, etc., up to
12 # sqrt(n); what's left at the end are the primes
13
14 import sys
15 import math
16 from multiprocessing import Process, Array, Queue
17
18 class gblbs: # globals, other than shared
19     n = int(sys.argv[1])
20     nthreads = int(sys.argv[2])
21     thrdlist = [] # list of all instances of this class
22
23 def prmfinder(id,prm,nxtk):
24     nk = 0 # count of k's done by this thread, to assess load balance
25     while 1:
26         # find next value to cross out with
27         try: k = nxtk.get(False)
28         except: break
29         nk += 1 # increment workload data
30         if prm[k]: # now cross out
31             r = gblbs.n / k
32             for i in range(2,r+1):
33                 prm[i*k] = 0
34     print 'thread', id, 'exiting; processed', nk, 'values of k'
35
36 def main():
37     prime = Array('i', (gblbs.n+1) * [1]) # 1 means prime, until find otherwise
38     nextk = Queue() # next value to try crossing out with
39     lim = int(math.sqrt(gblbs.n)) + 1 # fill the queue with 2...sqrt(n)
40     for i in range(2,lim): nextk.put(i)
41     for i in range(gblbs.nthreads):
42         pf = Process(target=prmfinder, args=(i,prime,nextk))
43         gblbs.thrdlist.append(pf)

```

```

44     pfs.append(pf)
45     pf.start()
46     for thrd in gbls.thrdlist: thrd.join()
47     print 'there are', reduce(lambda x,y: x+y, prime) - 2, 'primes'
48
49 if __name__ == '__main__':
50     main()

```

The way **Queue** is used here is to put all the possible “crosser-outers,” obtained in the variable **nextk** in the previous versions of this code, into a queue at the outset. One then uses `get()` to pick up work from the queue. Look Ma, no locks!

Below is an example of queues in an in-place quicksort. (Again, the reader is warned that this is just an example, not claimed to be efficient.)

The work items in the queue are a bit more involved here. They have the form **(i,j,k)**, with the first two elements of this tuple meaning that the given array chunk corresponds to indices **i** through **j** of **x**, the original array to be sorted. In other words, whichever thread picks up this chunk of work will have the responsibility of handling that particular section of **x**.

Quicksort, of course, works by repeatedly splitting the original array into smaller and more numerous chunks. Here a thread will split its chunk, taking the lower half for itself to sort, but placing the upper half into the queue, to be available for other chunks that have not been assigned any work yet. I’ve written the algorithm so that as soon as all threads have gotten some work to do, no more splitting will occur. That’s where the value of **k** comes in. It tells us the split number of this chunk. If it’s equal to **nthreads-1**, this thread won’t split the chunk.

```

1  # Quicksort and test code, based on Python multiprocessing class and
2  # Queue
3
4  # code is incomplete, as some special cases such as empty subarrays
5  # need to be accounted for
6
7  # usage:  python QSort.py n nthreads
8  #   where we wish to test the sort on a random list of n items,
9  #   using nthreads to do the work
10
11 import sys
12 import random
13 from multiprocessing import Process, Array, Queue
14
15 class gbls: # globals, other than shared
16     nthreads = int(sys.argv[2])
17     thrdlist = [] # list of all instances of this class
18     r = random.Random(9876543)
19
20 def sortworker(id,x,q):
21     chunkinfo = q.get()
22     i = chunkinfo[0]
23     j = chunkinfo[1]

```

```

24     k = chunkinfo[2]
25     if k < gblbs.nthreads - 1: # need more splitting?
26         splitpt = separate(x,i,j)
27         q.put((splitpt+1,j,k+1))
28         # now, what do I sort?
29         rightend = splitpt + 1
30     else: rightend = j
31     tmp = x[i:(rightend+1)] # need copy, as Array type has no sort() method
32     tmp.sort()
33     x[i:(rightend+1)] = tmp
34
35 def separate(xc, low, high): # common algorithm; see Wikipedia
36     pivot = xc[low] # would be better to take, e.g., median of 1st 3 elts
37     (xc[low],xc[high]) = (xc[high],xc[low])
38     last = low
39     for i in range(low,high):
40         if xc[i] <= pivot:
41             (xc[last],xc[i]) = (xc[i],xc[last])
42             last += 1
43     (xc[last],xc[high]) = (xc[high],xc[last])
44     return last
45
46 def main():
47     tmp = []
48     n = int(sys.argv[1])
49     for i in range(n): tmp.append(gblbs.r.uniform(0,1))
50     x = Array('d',tmp)
51     # work items have form (i,j,k), meaning that the given array chunk
52     # corresponds to indices i through j of x, and that this is the kth
53     # chunk that has been created, x being the 0th
54     q = Queue() # work queue
55     q.put((0,n-1,0))
56     for i in range(gblbs.nthreads):
57         p = Process(target=sortworker, args=(i,x,q))
58         gblbs.thrdlist.append(p)
59         p.start()
60     for thrd in gblbs.thrdlist: thrd.join()
61     if n < 25: print x[:]
62
63 if __name__ == '__main__':
64     main()

```

3.6 Debugging Threaded and Multiprocessing Python Programs

Debugging is always tough with parallel programs, including threads programs. It's especially difficult with pre-emptive threads; those accustomed to debugging non-threads programs find it rather jarring to see sudden changes of context while single-stepping through code. Tracking down the cause of deadlocks can be very hard. (Often just getting a threads program to end properly is a challenge.)

Another problem which sometimes occurs is that if you issue a “next” command in your debugging tool, you may end up inside the internal threads code. In such cases, use a “continue” command or something

like that to extricate yourself.

Unfortunately, as of April 2010, I know of no debugging tool that works with **multiprocessing**. However, one can do well with **thread** and **threading**.

3.6.1 Using PDB to Debug Threaded Programs

Using PDB is a bit more complex when threads are involved. One cannot, for instance, simply do something like this:

```
pdb.py buggyprog.py
```

because the child threads will not inherit the PDB process from the main thread. You can still run PDB in the latter, but will not be able to set breakpoints in threads.

What you can do, though, is invoke PDB from *within* the function which is run by the thread, by calling **pdb.set_trace()** at one or more points within the code:

```
import pdb
pdb.set_trace()
```

In essence, those become breakpoints.

For example, in our program **srvr.py** in Section 3.1.1, we could add a PDB call at the beginning of the loop in **serveclient()**:

```
while 1:
    import pdb
    pdb.set_trace()
    # receive letter from client, if it is still connected
    k = c.recv(1)
    if k == '': break
```

You then run the program directly through the Python interpreter as usual, NOT through PDB, but then the program suddenly moves into debugging mode on its own. At that point, one can then step through the code using the **n** or **s** commands, query the values of variables, etc.

PDB's **c** (“continue”) command still works. Can one still use the **b** command to set additional breakpoints? Yes, but it might be only on a one-time basis, depending on the context. A breakpoint might work only once, due to a scope problem. Leaving the scope where we invoked PDB causes removal of the trace object. Thus I suggested setting up the trace inside the loop above.

Of course, you can get fancier, e.g. setting up “conditional breakpoints,” something like:

```
debugflag = int(sys.argv[1])
...
if debugflag == 1:
    import pdb
    pdb.set_trace()
```

Then, the debugger would run only if you asked for it on the command line. Or, you could have multiple **debugflag** variables, for activating/deactivating breakpoints at various places in the code.

Moreover, once you get the (Pdb) prompt, you could set/reset those flags, thus also activating/deactivating breakpoints.

Note that local variables which were set before invoking PDB, including parameters, are not accessible to PDB.

Make sure to insert code to maintain an ID number for each thread. This really helps when debugging.

3.6.2 RPDB2 and Winpdb

The Winpdb debugger (www.digitalpeers.com/pythondebugger/),¹² is very good. Among other things, it can be used to debug threaded code, curses-based code and so on, which many debuggers can't. Winpdb is a GUI front end to the text-based RPDB2, which is in the same package. I have a tutorial on both at <http://heather.cs.ucdavis.edu/~matloff/winpdb.html>.

Another very promising debugger that handles threads is PYDB, by Rocky Bernstein (not to be confused with an earlier debugger of the same name). You can obtain it from <http://code.google.com/p/pydbgr/> or the older version at <http://bashdb.sourceforge.net/pydb/>. Invoke it on your code **x.py** by typing

```
$ pydb --threading x.py your_command_line_args_for_x
```

¹²No, it's not just for Microsoft Windows machines, in spite of the name.

Chapter 4

Introduction to OpenMP

OpenMP has become the *de facto* standard for shared-memory programming.

4.1 Overview

OpenMP has become the environment of choice for many, if not most, practitioners of shared-memory parallel programming. It consists of a set of directives which are added to one's C/C++/FORTRAN code that manipulate threads, without the programmer him/herself having to deal with the threads directly. This way we get “the best of both worlds”—the true parallelism of (nonpreemptive) threads and the pleasure of avoiding the annoyances of threads programming.

Most OpenMP constructs are expressed via **pragmas**, i.e. directives. The syntax is

```
#pragma omp .....
```

The number sign must be the first nonblank character in the line.

4.2 Running Example

The following example, implementing Dijkstra's shortest-path graph algorithm, will be used throughout this tutorial, with various OpenMP constructs being illustrated later by modifying this code:

```
1 // Dijkstra.c
2
```

```

3 // OpenMP example program: Dijkstra shortest-path finder in a
4 // bidirectional graph; finds the shortest path from vertex 0 to all
5 // others
6
7 // usage: dijkstra nv print
8
9 // where nv is the size of the graph, and print is 1 if graph and min
10 // distances are to be printed out, 0 otherwise
11
12 #include <omp.h>
13
14 // global variables, shared by all threads by default
15
16 int nv, // number of vertices
17     *notdone, // vertices not checked yet
18     nth, // number of threads
19     chunk, // number of vertices handled by each thread
20     md, // current min over all threads
21     mv, // vertex which achieves that min
22     largeint = -1; // max possible unsigned int
23
24 unsigned *ohd, // 1-hop distances between vertices; "ohd[i][j]" is
25 // ohd[i*nv+j]
26     *mind; // min distances found so far
27
28 void init(int ac, char **av)
29 { int i,j,tmp;
30   nv = atoi(av[1]);
31   ohd = malloc(nv*nv*sizeof(int));
32   mind = malloc(nv*sizeof(int));
33   notdone = malloc(nv*sizeof(int));
34   // random graph
35   for (i = 0; i < nv; i++)
36     for (j = i; j < nv; j++) {
37       if (j == i) ohd[i*nv+i] = 0;
38       else {
39         ohd[nv*i+j] = rand() % 20;
40         ohd[nv*j+i] = ohd[nv*i+j];
41       }
42     }
43   for (i = 1; i < nv; i++) {
44     notdone[i] = 1;
45     mind[i] = ohd[i];
46   }
47 }
48
49 // finds closest to 0 among notdone, among s through e
50 void findmymin(int s, int e, unsigned *d, int *v)
51 { int i;
52   *d = largeint;
53   for (i = s; i <= e; i++)
54     if (notdone[i] && mind[i] < *d) {
55       *d = ohd[i];
56       *v = i;
57     }
58 }
59
60 // for each i in [s,e], ask whether a shorter path to i exists, through

```

```

61 // mv
62 void updatemind(int s, int e)
63 { int i;
64   for (i = s; i <= e; i++)
65     if (mind[mv] + ohd[mv+nv+i] < mind[i])
66       mind[i] = mind[mv] + ohd[mv+nv+i];
67 }
68
69 void dowork()
70 {
71   #pragma omp parallel
72   { int startv, endv, // start, end vertices for my thread
73     step, // whole procedure goes nv steps
74     mymv, // vertex which attains the min value in my chunk
75     me = omp_get_thread_num();
76     unsigned mymd; // min value found by this thread
77     #pragma omp single
78     { nth = omp_get_num_threads();
79       if (nv % nth != 0) {
80         printf("nv must be divisible by nth\n");
81         exit(1);
82       }
83       chunk = nv/nth;
84       printf("there are %d threads\n", nth);
85     }
86     startv = me * chunk;
87     endv = startv + chunk - 1;
88     for (step = 0; step < nv; step++) {
89       // find closest vertex to 0 among notdone; each thread finds
90       // closest in its group, then we find overall closest
91       #pragma omp single
92       { md = largeint; mv = 0; }
93       findmymin(startv, endv, &mymd, &mymv);
94       // update overall min if mine is smaller
95       #pragma omp critical
96       { if (mymd < md)
97         { md = mymd; mv = mymv; }
98       }
99       #pragma omp barrier
100      // mark new vertex as done
101      #pragma omp single
102      { notdone[mv] = 0; }
103      // now update my section of mind
104      updatemind(startv, endv);
105      #pragma omp barrier
106    }
107  }
108 }
109
110 int main(int argc, char **argv)
111 { int i, j, print;
112   double starttime, endtime;
113   init(argc, argv);
114   starttime = omp_get_wtime();
115   // parallel
116   dowork();
117   // back to single thread
118   endtime = omp_get_wtime();

```

```

119     printf("elapsed time:  %f\n",endtime-starttime);
120     print = atoi(argv[2]);
121     if (print) {
122         printf("graph weights:\n");
123         for (i = 0; i < nv; i++) {
124             for (j = 0; j < nv; j++)
125                 printf("%u  ",ohd[nv*i+j]);
126             printf("\n");
127         }
128         printf("minimum distances:\n");
129         for (i = 1; i < nv; i++)
130             printf("%u\n",mind[i]);
131     }
132 }

```

The constructs will be presented in the following sections, but first the algorithm will be explained.

4.2.1 The Algorithm

The code implements the Dijkstra algorithm for finding the shortest paths from vertex 0 to the other vertices in an N-vertex undirected graph. Pseudocode for the algorithm is shown below, with the array G assumed to contain the one-hop distances from 0 to the other vertices.

```

1 Done = {0} # vertices checked so far
2 NewDone = None # currently checked vertex
3 NonDone = {1,2,...,N-1} # vertices not checked yet
4 for J = 0 to N-1 Dist[J] = G(0,J) # initialize shortest-path lengths
5
6 for Step = 1 to N-1
7     find J such that Dist[J] is min among all J in NonDone
8     transfer J from NonDone to Done
9     NewDone = J
10    for K = 1 to N-1
11        if K is in NonDone
12            # check if there is a shorter path from 0 to K through NewDone
13            # than our best so far
14            Dist[K] = min(Dist[K],Dist[NewDone]+G[NewDone,K])

```

At each iteration, the algorithm finds the closest vertex J to 0 among all those not yet processed, and then updates the list of minimum distances to each vertex from 0 by considering paths that go through J. Two obvious potential candidate part of the algorithm for parallelization are the “find J” and “for K” lines, and the above OpenMP code takes this approach.

4.2.2 The OpenMP parallel Pragma

As can be seen in the comments in the lines

```
// parallel
dowork();
// back to single thread
```

the function **main()** is run by a **master thread**, which will then branch off into many threads running **dowork()** in parallel. The latter feat is accomplished by the directive in the lines

```
void dowork()
{
    #pragma omp parallel
    { int startv,endv, // start, end vertices for this thread
      step, // whole procedure goes nv steps
      mymv, // vertex which attains that value
      me = omp_get_thread_num();
```

That directive sets up a team of threads (which includes the master), all of which execute the block following the directive in parallel.¹ Note that, unlike the **for** directive which will be discussed below, the **parallel** directive leaves it up to the programmer as to how to partition the work. In our case here, we do that by setting the range of vertices which this thread will process:

```
startv = me * chunk;
endv = startv + chunk - 1;
```

Again, keep in mind that *all* of the threads execute this code, but we've set things up with the variable **me** so that different threads will work on different vertices. This is due to the OpenMP call

```
me = omp_get_thread_num();
```

which sets **me** to the thread number for this thread.

4.2.3 Scope Issues

Note carefully that in

```
#pragma omp parallel
{ int startv,endv, // start, end vertices for this thread
  step, // whole procedure goes nv steps
  mymv, // vertex which attains that value
  me = omp_get_thread_num();
```

¹There is an issue here of thread startup time. The OMPi compiler sets up threads at the outset, so that that startup time is incurred only once. When a **parallel** construct is encountered, they are awakened. At the end of the construct, they are suspended again, until the next **parallel** construct is reached.

the pragma comes *before* the declaration of the local variables. That means that all of them are “local” to each thread, i.e. not shared by them. But if a work sharing directive comes within a function but *after* declaration of local variables, those variables are actually “global” to the code in the directive, i.e. they *are* shared in common among the threads.

This is the default, but you can change these properties, e.g. using the **shared** keyword. For instance,

```
#pragma omp parallel private(x,y)
```

would make **x** and **y** nonshared even if they were declared above the directive line.

It is crucial to keep in mind that variables which are global to the program (in the C/C++ sense) are automatically global to all threads. This is the primary means by which the threads communicate with each other.

4.2.4 The OpenMP `single` Pragma

In some cases we want just one thread to execute some code, even though that code is part of a **parallel** or other **work sharing** block.² We use the **single** directive to do this, e.g.:

```
#pragma omp single
{
  nth = omp_get_num_threads();
  if (nv % nth != 0) {
    printf("nv must be divisible by nth\n");
    exit(1);
  }
  chunk = nv/nth;
  printf("there are %d threads\n",nth);
}
```

Since the variables **nth** and **chunk** are global and thus shared, we need not have all threads set them, hence our use of **single**.

4.2.5 The OpenMP `barrier` Pragma

As seen in the example above, the **barrier** implements a standard barrier, applying to all threads.

4.2.6 Implicit Barriers

Note that there is an implicit barrier at the end of each **single** block, which is also the case for **parallel**, **for**, and **sections** blocks. This can be overridden via the **nowait** clause, e.g.

²This is an OpenMP term. The **for** directive is another example of it. More on this below.


```
#pragma omp for nowait
```

Needless to say, the latter should be used with care, and in most cases will not be usable. On the other hand, putting in a barrier where it is not needed would severely reduce performance.

4.2.7 The OpenMP `critical` Pragma

The last construct used in this example is **critical**, for critical sections.

```
#pragma omp critical
{ if (mymd < md)
  { md = mymd; mv = mymv; }
}
```

It means what it says, allowing entry of only one thread at a time while others wait. Here we are updating global variables **md** and **mv**, which has to be done atomically, and **critical** takes care of that for us. This is much more convenient than setting up lock variables, etc., which we would do if we were programming threads code directly.

4.3 The OpenMP for Pragma

This one breaks up a C/C++ **for** loop, assigning various iterations to various threads. This way the iterations are done in parallel. Of course, that means that they need to be independent iterations, i.e. one iteration cannot depend on the result of another.

4.3.1 Basic Example

Here's how we could use this construct in the Dijkstra program :

```
1 // Dijkstra.c
2
3 // OpenMP example program (OMP version): Dijkstra shortest-path finder
4 // in a bidirectional graph; finds the shortest path from vertex 0 to
5 // all others
6
7 // usage: dijkstra nv print
8
9 // where nv is the size of the graph, and print is 1 if graph and min
10 // distances are to be printed out, 0 otherwise
11
12 #include <omp.h>
```

```

13
14 // global variables, shared by all threads by default
15
16 int nv, // number of vertices
17     *notdone, // vertices not checked yet
18     nth, // number of threads
19     chunk, // number of vertices handled by each thread
20     md, // current min over all threads
21     mv, // vertex which achieves that min
22     largeint = -1; // max possible unsigned int
23
24 unsigned *ohd, // 1-hop distances between vertices; "ohd[i][j]" is
25               // ohd[i*nv+j]
26             *mind; // min distances found so far
27
28 void init(int ac, char **av)
29 { int i,j,tmp;
30   nv = atoi(av[1]);
31   ohd = malloc(nv*nv*sizeof(int));
32   mind = malloc(nv*sizeof(int));
33   notdone = malloc(nv*sizeof(int));
34   // random graph
35   for (i = 0; i < nv; i++)
36     for (j = i; j < nv; j++) {
37       if (j == i) ohd[i*nv+i] = 0;
38       else {
39         ohd[nv*i+j] = rand() % 20;
40         ohd[nv*j+i] = ohd[nv*i+j];
41       }
42     }
43   for (i = 1; i < nv; i++) {
44     notdone[i] = 1;
45     mind[i] = ohd[i];
46   }
47 }
48
49 void dowork()
50 {
51   #pragma omp parallel
52   { int step, // whole procedure goes nv steps
53     mymv, // vertex which attains that value
54     me = omp_get_thread_num(),
55     i;
56     unsigned mymd; // min value found by this thread
57     #pragma omp single
58     { nth = omp_get_num_threads();
59       printf("there are %d threads\n",nth); }
60     for (step = 0; step < nv; step++) {
61       // find closest vertex to 0 among notdone; each thread finds
62       // closest in its group, then we find overall closest
63       #pragma omp single
64       { md = largeint; mv = 0; }
65       mymd = largeint;
66       #pragma omp for nowait
67       for (i = 1; i < nv; i++) {
68         if (notdone[i] && mind[i] < mymd) {
69           mymd = ohd[i];
70           mymv = i;

```

```

71         }
72     }
73     // update overall min if mine is smaller
74     #pragma omp critical
75     { if (mynd < md)
76         { md = mynd; mv = mymv; }
77     }
78     // mark new vertex as done
79     #pragma omp single
80     { notdone[mv] = 0; }
81     // now update ohd
82     #pragma omp for
83     for (i = 1; i < nv; i++)
84         if (mind[mv] + ohd[mv*nv+i] < mind[i])
85             mind[i] = mind[mv] + ohd[mv*nv+i];
86     }
87 }
88 }
89
90 int main(int argc, char **argv)
91 { int i,j,print;
92   init(argc,argv);
93   // parallel
94   dowork();
95   // back to single thread
96   print = atoi(argv[2]);
97   if (print) {
98       printf("graph weights:\n");
99       for (i = 0; i < nv; i++) {
100           for (j = 0; j < nv; j++)
101               printf("%u ",ohd[nv*i+j]);
102           printf("\n");
103       }
104       printf("minimum distances:\n");
105       for (i = 1; i < nv; i++)
106           printf("%u\n",mind[i]);
107   }
108 }
109

```

The work which used to be done in the function **findmymin()** is now done here:

```

#pragma omp for
for (i = 1; i < nv; i++) {
    if (notdone[i] && mind[i] < mynd) {
        mynd = ohd[i];
        mymv = i;
    }
}

```

Each thread executes one or more of the iterations, i.e. takes responsibility for one or more values of *i*. This occurs in parallel, so as mentioned earlier, the programmer must make sure that the iterations are independent; there is no predicting which threads will do which values of *i*, in which order. By the way, for obvious reasons OpenMP treats the loop index, *i* here, as private even if by context it would be shared.

4.3.2 Nested Loops

If we use the **for** pragma to nested loops, by default the pragma applies only to the outer loop. We can of course insert another **for** pragma inside, to parallelize the inner loop.

Or, starting with OpenMP version 3.0, one can use the **collapse** clause, e.g.

```
#pragma omp parallel for collapse(2)
```

to specify two levels of nesting in the assignment of threads to tasks.

4.3.3 Controlling the Partitioning of Work to Threads

In this default version of the **for** construct, iterations are executed by threads *in unpredictable order*, with each thread taking on *one iteration's worth of work at a time*. Both of these can be changed by the programmer, using the **schedule** clause.

For instance, our original version of our program in Section 4.2 broke the work into chunks, with chunk size being the number vertices divided by the number of threads.

For the Dijkstra algorithm, for instance, we could have this:

```
...
    #pragma omp for schedule(static,chunk)
    for (i = 1; i < nv; i++) {
        if (notdone[i] && mind[i] < mymd) {
            mymd = ohd[i];
            mymv = i;
        }
    }
...
    #pragma omp for schedule(static,chunk)
    for (i = 1; i < nv; i++)
        if (mind[mv] + ohd[mv*nv+i] < mind[i])
            mind[i] = mind[mv] + ohd[mv*nv+i];
...
```

But one can enhance performance by considering other chunk sizes (in which case a thread would be responsible for more than one chunk). On the one hand, large chunks are good, due to there being less overhead—every time a thread finishes a chunk, it must go through the critical section, which serializes our parallel program and thus slows things down. On the other hand, if chunk sizes are large, then toward the end of the work, some threads may be working on their last chunks while others have finished and are now idle, thus foregoing potential speed enhancement. So it would be nice to have large chunks at the beginning of the run, to reduce the overhead, but smaller chunks at the end. This can be done using the **guided** clause.

For the Dijkstra algorithm, for instance, we could have this:

```

...
    #pragma omp for schedule(guided)
    for (i = 1; i < nv; i++) {
        if (notdone[i] && mind[i] < mymd) {
            mymd = ohd[i];
            mymv = i;
        }
    }
...
    #pragma omp for schedule(guided)
    for (i = 1; i < nv; i++)
        if (mind[mv] + ohd[mv*nv+i] < mind[i])
            mind[i] = mind[mv] + ohd[mv*nv+i];
...

```

4.3.4 The OpenMP reduction Clause

The name of this OpenMP clause alludes to the term **reduction** in functional programming. Many parallel programming languages include such operations, to enable the programmer to more conveniently (and often more efficiently) have threads/processors cooperate in computing sums, products, etc. OpenMP does this via the **reduction** clause.

For example, consider

```

1 int z;
2 ...
3 #pragma omp for reduction(+:z)
4 for (i = 0; i < n; i++) z += x[i];

```

The pragma says that the threads will share the work as in our previous discussion of the **for** pragma. In addition, though, there will be independent copies of **z** maintained for each thread, each initialized to 0 before the loop begins. When the loop is entirely done, the values of **z** from the various threads will be summed, of course in an atomic manner.

Note that the **+** operator not only indicates that the values of **z** are to be summed, but also that their initial values are to be 0. If the operator were *****, say, then the product of the values would be computed, and their initial values would be 1.

One can specify several reduction variables to the right of the colon, separated by commas.

Our use of the **reduction** clause here makes our programming much easier. Indeed, if had old serial code that we wanted to parallelize, we would have to make no change to it! OpenMP is taking care of both the work splitting across values of **i**, and the atomic operations. Moreover—note this carefully—it is efficient, because by maintaining separate copies of **z** until the loop is done, we are reducing the number of serializing atomic actions, and are avoiding time-costly cache coherency transactions and the like.

Without this construct, we would have to do

```

int z,myz=0;
...
#pragma omp for private(myz)
for (i = 0; i < n; i++) myz += x[i];
#pragma omp critical
{ z += myz; }

```

Here are the eligible operators and the corresponding initial values:

In C/C++, you can use **reduction** with +, -, *, &, |, && and || (and the exclusive-or operator).

operator	initial value
+	0
-	0
*	1
&	bit string of 1s
	bit string of 0s
^	0
&&	1
	0

The lack of other operations typically found in other parallel programming languages, such as min and max, is due to the lack of these operators in C/C++. The FORTRAN version of OpenMP does have min and max.³

Note that the reduction variables must be shared by the threads, so they must be declared before the **parallel**, or with a **shared** clause.

4.4 The Task Directive

This is new to OpenMP 3.0. The basic idea is this: When a thread encounters a **task** directive, it arranges for some thread to execute the associated block. The first thread can continue. Note that the task might not execute right away; it may have to wait for some thread to become free after finishing another task.

Here's a Quicksort example:

```

1 // OpenMP example program: quicksort; not necessarily efficient
2
3 void swap(int *yi, int *yj)
4 { int tmp = *yi;
5   *yi = *yj;
6   *yj = tmp;
7 }

```

³Note, though, that plain min and max would not help in our Dijkstra example above, as we not only need to find the minimum value, but also need the vertex which attains that value.

```

8
9 int *separate(int *x, int low, int high)
10 { int i,pivot,last;
11     pivot = x[low]; // would be better to take, e.g., median of 1st 3 elts
12     swap(x+low,x+high);
13     last = low;
14     for (i = low; i < high; i++) {
15         if (x[i] <= pivot) {
16             swap(x+last,x+i);
17             last += 1;
18         }
19     }
20     swap(x+last,x+high);
21     return last;
22 }
23
24 // quicksort of the array z, elements zstart through zend; set the
25 // latter to 0 and m-1 in first call, where m is the length of z;
26 // firstcall is 1 or 0, according to whether this is the first of the
27 // recursive calls
28 void qs(int *z, int zstart, int zend, int firstcall)
29 {
30     #pragma omp parallel
31     { int part;
32         if (firstcall == 1) {
33             #pragma omp single nowait
34             qs(z,0,zend,0);
35         } else {
36             if (zstart < zend) {
37                 part = separate(z,zstart,zend);
38                 #pragma omp task
39                 qs(z,zstart,part-1,0);
40                 #pragma omp task
41                 qs(z,part+1,zend,0);
42             }
43         }
44     }
45 }
46 }
47
48 main(int argc, char**argv)
49 { int i,n,*w;
50     n = atoi(argv[1]);
51     w = malloc(n*sizeof(int));
52     for (i = 0; i < n; i++) w[i] = rand();
53     qs(w,0,n-1,1);
54     if (n < 25)
55         for (i = 0; i < n; i++) printf("%d\n",w[i]);
56 }

```

The code

```

if (firstcall == 1) {
    #pragma omp single nowait
    qs(z,0,zend,0);
}

```

gets things going. We want only one thread to execute the root of the recursion tree, hence the need for the **single** clause. After that, the code

```
part = separate(z, zstart, zend);
#pragma omp task
qs(z, zstart, part-1, 0);
```

sets up a call to a subtree, with the **task** directive stating, “OMP system, please make sure that this subtree is handled by some thread.”

This really simplifies the programming. Compare this to the Python **multiprocessing** version in Section 3.5, where the programmer needed to write code to handle the work queue.

There are various refinements, such as the barrier-like **taskwait** clause.

4.5 Other OpenMP Synchronization Issues

Earlier we saw the **critical** and **barrier** constructs. There is more to discuss, which we do here.

4.5.1 The OpenMP `atomic` Clause

The **critical** construct not only serializes your program, but also it adds a lot of overhead. If your critical section involves just a one-statement update to a shared variable, e.g.

```
x += y;
```

etc., then the OpenMP compiler can take advantage of an atomic hardware instruction, e.g. the **LOCK** prefix on Intel, to set up an extremely efficient critical section, e.g.

```
#pragma omp atomic
x += y;
```

Since it is a single statement rather than a block, there are no braces.

The eligible operators are:

```
++, --, +=, *=, <<=, &=, |=
```


4.5.2 Memory Consistency and the `flush` Pragma

Consider a shared-memory multiprocessor system with coherent caches, and a shared, i.e. global, variable `x`. If one thread writes to `x`, you might think that the cache coherency system will ensure that the new value is visible to other threads. But it is not quite so simple as this.

For example, the compiler may store `x` in a register, and update `x` itself at certain points. In between such updates, since the memory location for `x` is not written to, the cache will be unaware of the new value, which thus will not be visible to other threads. If the processors have write buffers etc., the same problem occurs.

In other words, we must account for the fact that our program could be run on different kinds of hardware with different memory consistency models. Thus OpenMP must have its own memory consistency model, which is then translated by the compiler to mesh with the hardware.

OpenMP takes a **relaxed consistency** approach, meaning that it forces updates to memory (“flushes”) at all synchronization points, i.e. at:

- **barrier**
- entry/exit to/from **critical**
- entry/exit to/from **ordered**
- entry/exit to/from **parallel**
- exit from **parallel for**
- exit from **parallel sections**
- exit from **single**

In between synchronization points, one can force an update to `x` via the **flush** pragma:

```
#pragma omp flush (x)
```

The flush operation is obviously architecture-dependent. OpenMP compilers will typically have the proper machine instructions available for some common architectures. For the rest, it can force a flush at the hardware level by doing lock/unlock operations, though this may be costly in terms of time.

4.6 Compiling, Running and Debugging OpenMP Code

4.6.1 Compiling

There are a number of open source compilers available for OpenMP, including:

- Omni: This is available at (<http://phase.hpcc.jp/Omni/>). To compile an OpenMP program in `x.c` and create an executable file `x`, run

```
omcc -g -o x x.c
```

- Ompi: You can download this at <http://www.cs.uoi.gr/~ompi/index.html>. Compile `x.c` by

```
ompicc -g -o x x.c
```

- GCC, version 4.2 or later:⁴ Compile `x.c` via

```
gcc -fopenmp -g -o x x.c
```

4.6.2 Running

Just run the executable as usual.

The number of threads will be the number of processors, by default. To change that value, set the `OMP_NUM_THREADS` environment variable. For example, to get four threads in the C shell, type

```
setenv OMP_NUM_THREADS 4
```

4.6.3 Debugging

OpenMP's use of pragmas makes it difficult for the compilers to maintain your original source code line numbers, and your function and variable names. But with a little care, a symbolic debugger such as GDB can still be used. Here are some tips for the compilers mentioned above, using GDB as our example debugging tool:

- Omni: The function `main()` in your executable is actually in the OpenMP library, and your function `main()` is renamed `_ompc_main()`. So, when you enter GDB, first set a breakpoint at your own code:

⁴You may find certain subversions of GCC 4.1 can be used too.

```
(gdb) b _ompc_main
```

Then run your program to this breakpoint, and set whatever other breakpoints you want.

You should find that your other variable and function names are unchanged.

- **Ompi:** During preprocessing of your file **x.c**, the compiler produces a file **x_ompi.c**, and the latter is what is actually compiled. Your function **main** is renamed to **_ompi_originalMain()**. Your other functions and variables are renamed. For example in our Dijkstra code, the function **dowork()** is renamed to **dowork_parallel_0**. And by the way, all indenting is lost! Keep these points in mind as you navigate through your code in GDB.
- **GCC:** GCC maintains line numbers and names well. In earlier versions, it had a problem in that it did not retain names of local variables within blocks controlled by **omp parallel** at all. That problem is now fixed (e.g. in version 4.4 of the GCC suite).

4.7 Combining Work-Sharing Constructs

In our examples of the **for** pragma above, that pragma would come within a block headed by a **parallel** pragma. The latter specifies that a team of threads is to be created, with each one executing the given block, while the former specifies that the various iterations of the loop are to be distributed among the threads. As a shortcut, we can combine the two pragmas:

```
#pragma omp parallel for
```

This also works with the **sections** pragma.

4.8 Performance

As is usually the case with parallel programming, merely parallelizing a program won't necessarily make it faster, even on shared-memory hardware. Operations such as critical sections, barriers and so on serialize an otherwise-parallel program, sapping much of its speed. In addition, there are issues of cache coherency transactions, false sharing etc.

4.8.1 The Effect of Problem Size

To illustrate this, I ran our original Dijkstra example (Section 4.2 on various graph sizes, on a quad core machine. Here are the timings:

nv	nth	time
1000	1	0.005472
1000	2	0.011143
1000	4	0.029574

The more parallelism we had, the *slower* the program ran! The synchronization overhead was just too much to be compensated by the parallel computation.

However, parallelization did bring benefits on larger problems:

nv	nth	time
25000	1	2.861814
25000	2	1.710665
25000	4	1.453052

4.8.2 Some Fine Tuning

How could we make our Dijkstra code faster? One idea would be to eliminate the critical section. Recall that in each iteration, the threads compute their local minimum distance values **md** and **mv**, and then update the global values **md** and **mv**. Since the update must be atomic, this causes some serialization of the program. Instead, we could have the threads store their values **mymd** and **mymv** in a global array **mymins**, with each thread using a separate pair of locations within that array, and then at the end of the iteration we could have just one task scan through **mymins** and update **md** and **mv**.

Here is the resulting code:

```

1 // Dijkstra.c
2
3 // OpenMP example program: Dijkstra shortest-path finder in a
4 // bidirectional graph; finds the shortest path from vertex 0 to all
5 // others
6
7 // **** in this version, instead of having a critical section in which
8 // each thread updates md and mv, the threads record their mymd and mymv
9 // values in a global array mymins, which one thread then later uses to
10 // update md and mv
11
12 // usage: dijkstra nv print
13
14 // where nv is the size of the graph, and print is 1 if graph and min
15 // distances are to be printed out, 0 otherwise
16
17 #include <omp.h>
18
19 // global variables, shared by all threads by default
20
21 int nv, // number of vertices
22     *notdone, // vertices not checked yet
23     nth, // number of threads
24     chunk, // number of vertices handled by each thread

```

```

25     md, // current min over all threads
26     mv, // vertex which achieves that min
27     largeint = -1; // max possible unsigned int
28
29 int *mymins; // (mymd,mymv) for each thread; see dowork()
30
31 unsigned *ohd, // 1-hop distances between vertices; "ohd[i][j]" is
32             // ohd[i*nv+j]
33             *mind; // min distances found so far
34
35 void init(int ac, char **av)
36 { int i,j,tmp;
37   nv = atoi(av[1]);
38   ohd = malloc(nv*nv*sizeof(int));
39   mind = malloc(nv*sizeof(int));
40   notdone = malloc(nv*sizeof(int));
41   // random graph
42   for (i = 0; i < nv; i++)
43     for (j = i; j < nv; j++) {
44       if (j == i) ohd[i*nv+i] = 0;
45       else {
46         ohd[nv*i+j] = rand() % 20;
47         ohd[nv*j+i] = ohd[nv*i+j];
48       }
49     }
50   for (i = 1; i < nv; i++) {
51     notdone[i] = 1;
52     mind[i] = ohd[i];
53   }
54 }
55
56 // finds closest to 0 among notdone, among s through e
57 void findmymin(int s, int e, unsigned *d, int *v)
58 { int i;
59   *d = largeint;
60   for (i = s; i <= e; i++)
61     if (notdone[i] && mind[i] < *d) {
62       *d = ohd[i];
63       *v = i;
64     }
65 }
66
67 // for each i in [s,e], ask whether a shorter path to i exists, through
68 // mv
69 void updatemind(int s, int e)
70 { int i;
71   for (i = s; i <= e; i++)
72     if (mind[mv] + ohd[mv*nv+i] < mind[i])
73       mind[i] = mind[mv] + ohd[mv*nv+i];
74 }
75
76 void dowork()
77 {
78   #pragma omp parallel
79   { int startv,endv, // start, end vertices for my thread
80     step, // whole procedure goes nv steps
81     me,
82     mymv; // vertex which attains the min value in my chunk

```

```

83     unsigned mymd; // min value found by this thread
84     int i;
85     me = omp_get_thread_num();
86     #pragma omp single
87     { nth = omp_get_num_threads();
88       if (nv % nth != 0) {
89         printf("nv must be divisible by nth\n");
90         exit(1);
91       }
92       chunk = nv/nth;
93       mymins = malloc(2*nth*sizeof(int));
94     }
95     startv = me * chunk;
96     endv = startv + chunk - 1;
97     for (step = 0; step < nv; step++) {
98       // find closest vertex to 0 among notdone; each thread finds
99       // closest in its group, then we find overall closest
100      findmymin(startv,endv,&mymd,&mymv);
101      mymins[2*me] = mymd;
102      mymins[2*me+1] = mymv;
103      #pragma omp barrier
104      // mark new vertex as done
105      #pragma omp single
106      { md = largeint; mv = 0;
107        for (i = 1; i < nth; i++)
108          if (mymins[2*i] < md) {
109            md = mymins[2*i];
110            mv = mymins[2*i+1];
111          }
112        notdone[mv] = 0;
113      }
114      // now update my section of mind
115      updatemind(startv,endv);
116      #pragma omp barrier
117    }
118  }
119 }
120
121 int main(int argc, char **argv)
122 { int i,j,print;
123   double starttime,endtime;
124   init(argc,argv);
125   starttime = omp_get_wtime();
126   // parallel
127   dowork();
128   // back to single thread
129   endtime = omp_get_wtime();
130   printf("elapsed time: %f\n",endtime-starttime);
131   print = atoi(argv[2]);
132   if (print) {
133     printf("graph weights:\n");
134     for (i = 0; i < nv; i++) {
135       for (j = 0; j < nv; j++)
136         printf("%u ",ohd[nv*i+j]);
137       printf("\n");
138     }
139     printf("minimum distances:\n");
140     for (i = 1; i < nv; i++)

```

```

141         printf("%u\n",mind[i]);
142     }
143 }

```

Let's take a look at the latter part of the code for one iteration;

```

1     findmymin(startv, endv, &mynd, &mymv);
2     mymins[2*me] = mynd;
3     mymins[2*me+1] = mymv;
4     #pragma omp barrier
5     // mark new vertex as done
6     #pragma omp single
7     { notdone[mv] = 0;
8       for (i = 1; i < nth; i++)
9         if (mymins[2*i] < md) {
10            md = mymins[2*i];
11            mv = mymins[2*i+1];
12        }
13    }
14    // now update my section of mind
15    updatemind(startv, endv);
16    #pragma omp barrier

```

The call to **findmymin()** is as before; this thread finds the closest vertex to 0 among this thread's range of vertices. But instead of comparing the result to **md** and possibly updating it and **mv**, the thread simply stores its **mynd** and **mymv** in the global array **mymins**. After all threads have done this and then waited at the barrier, we have just one thread update **md** and **mv**.

Let's see how well this tack worked:

nv	nth	time
25000	1	2.546335
25000	2	1.449387
25000	4	1.411387

This brought us about a 15% speedup in the two-thread case, though less for four threads.

What else could we do? Here are a few ideas:

- False sharing could be a problem here. To address it, we could make **mymins** much longer, changing the places at which the threads write their data, leaving most of the array as padding.
- We could try the modification of our program in Section 4.3.1, in which we use the OpenMP **for** pragma, as well as the refinements stated there, such as **schedule**.
- We could try combining all of the ideas here.

4.8.3 OpenMP Internals

We may be able to write faster code if we know a bit about how OpenMP works inside.

You can get some idea of this from your compiler. For example, if you use the **-t** option with the Omni compiler, or **-k** with Ompi, you can inspect the result of the preprocessing of the OpenMP pragmas.

Here for instance is the code produced by Omni from the call to **findmymin()** in our Dijkstra program:

```
# 93 "Dijkstra.c"
findmymin(startv, endv, &(mynd), &(mymv)); {
__ompc_enter_critical(&__ompc_lock_critical);
# 96 "Dijkstra.c"
if((mynd)<(((unsigned )(md)))){

# 97 "Dijkstra.c"
(md)=(((int )(mymd)));
# 97 "Dijkstra.c"
(mv)=(mymv);
}__ompc_exit_critical(&__ompc_lock_critical);
```

Fortunately Omni saves the line numbers from our original source file, but the pragmas have been replaced by calls to OpenMP library functions.

The document, *The GNU OpenMP Implementation*, <http://pl.postech.ac.kr/~gla/cs700-07f/ref/openMp/libgomp.pdf>, includes good outline of how the pragmas are translated.

4.9 The Rest of OpenMP

There is much, much more to OpenMP than what we have seen here. Here are a couple of examples:

- The **sections** pragma: Suppose at a certain point in the code, there are several different code blocks which you want different threads to execute simultaneously. You can use this pragma to specify this.
- The **ordered** pragma: This is used to ensure that a subblock within a loop is executed in sequential order.

To see the details of these and other OpenMP constructs, there are many Web pages you can check, and there is also the excellent book, *Using OpenMP: Portable Shared Memory Parallel Programming*, by Barbara Chapman, Gabriele Jost and Ruud Van Der Pas, MIT Press, 2008.

4.10 Further Examples

There are additional OpenMP examples in later sections of this book, such as:⁵

- Jacobi algorithm for solving systems of linear equations, with a good example of the OpenMP **reduction** clause, Section 8.6
- another implementation of Quicksort, Section 10.1.1
- matrix multiplication, Section 8.4.2.1

⁵If you are reading this presentation on OpenMP separately from the book, the book is at <http://heather.cs.ucdavis.edu/~matloff/158/PLN/ParProcBook.pdf>

Chapter 5

Introduction to GPU Programming with CUDA

Even if you don't play video games, you can be grateful to the game players, as their numbers have given rise to a class of highly powerful parallel processing devices—**graphics processing units** (GPUs). Yes, you program right on the video card in your computer, even though your program may have nothing to do with graphics or games.

5.1 Overview

The video game market is so lucrative that the industry has developed ever-faster GPUs, in order to handle ever-faster and ever-more visually detailed video games. These actually are parallel processing hardware devices, so around 2003 some people began to wonder if one might use them for parallel processing of nongraphics applications.

Originally this was cumbersome. One needed to figure out clever ways of mapping one's application to some kind of graphics problem, i.e. of way of disguising one's problem so that it appeared to be doing graphics computations. Though some high-level interfaces were developed to automate this transformation, effective coding required some understanding of graphics principles.

But current-generation GPUs separate out the graphics operations, and now consist of multiprocessor elements that run under the familiar shared-memory threads model. Thus they are easily programmable. Granted, effective coding still requires an intimate knowledge of the hardware, but at least it's (more or less) familiar hardware, not requiring knowledge of graphics.

Moreover, unlike a multicore machine, with the ability to run just a few threads at one time, e.g. four

threads on a quad core machine, GPUs can run *hundreds or thousands* of threads at once. There are various restrictions that come with this, but you can see that there is fantastic potential for speed here.

NVIDIA has developed the CUDA language as a vehicle for programming on their GPUs. It's basically just a slight extension of C, and has become very popular. More recently, the OpenCL language has been developed by Apple, AMD and others (including NVIDIA). It too is a slight extension of C, and it aims to provide a uniform interface that works with multicore machines in addition to GPUs. OpenCL is not yet in as broad use as CUDA, so our discussion here focuses on CUDA and NVIDIA GPUs. (Thus we will be implicitly discussing CUDA/NVIDIA without reminding the reader of that fact.)

Some terminology:

- A CUDA program consists of code to be run on the **host**, i.e. the CPU, and code to run on the **device**, i.e. the GPU.
- A function that is called by the host to execute on the device is called a **kernel**.
- Threads in an application are grouped into **blocks**. The entirety of blocks is called the **grid** of that application.

5.2 Sample Program

Here's a sample program. And I've kept the sample simple: It just finds the sums of all the rows of a matrix.

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <cuda.h>
4
5  // CUDA example:  finds row sums of an integer matrix m
6
7  // findlelt() finds the rowsum of one row of the nxn matrix m, storing the
8  // result in the corresponding position in the rowsum array rs; matrix
9  // stored as 1-dimensional, row-major order
10
11 __global__ void findlelt(int *m, int *rs, int n)
12 {
13     int rownum = blockIdx.x; // this thread will handle row # rownum
14     int sum = 0;
15     for (int k = 0; k < n; k++)
16         sum += m[rownum*n+k];
17     rs[rownum] = sum;
18 }
19
20 int main(int argc, char **argv)
21 {
22     int n = atoi(argv[1]); // number of matrix rows/cols
23     int *hm, // host matrix

```

```

24     *dm, // device matrix
25     *hrs, // host rowsums
26     *drs; // device rowsums
27     int msize = n * n * sizeof(int); // size of matrix in bytes
28     // allocate space for host matrix
29     hm = (int *) malloc(msize);
30     // as a test, fill matrix with consecutive integers
31     int t = 0,i,j;
32     for (i = 0; i < n; i++) {
33         for (j = 0; j < n; j++) {
34             hm[i*n+j] = t++;
35         }
36     }
37     // allocate space for device matrix
38     cudaMalloc((void **)&dm,msize);
39     // copy host matrix to device matrix
40     cudaMemcpy(dm,hm,msize,cudaMemcpyHostToDevice);
41     // allocate host, device rowsum arrays
42     int rssize = n * sizeof(int);
43     hrs = (int *) malloc(rssize);
44     cudaMalloc((void **)&drs,rssize);
45     // set up parameters for threads structure
46     dim3 dimGrid(n,1); // n blocks
47     dim3 dimBlock(1,1,1); // 1 thread per block
48     // invoke the kernel
49     findlelt<<<dimGrid,dimBlock>>>(dm,drs,n);
50     // wait for kernel to finish
51     cudaThreadSynchronize();
52     // copy row vector from device to host
53     cudaMemcpy(hrs,drs,rssize,cudaMemcpyDeviceToHost);
54     // check results
55     if (n < 10) for(int i=0; i<n; i++) printf("%d\n",hrs[i]);
56     // clean up
57     free(hm);
58     cudaFree(dm);
59     free(hrs);
60     cudaFree(drs);
61 }

```

This is mostly C, with a bit of CUDA added here and there. Here's how the program works:

- Our **main()** runs on the host.
- Kernel functions are identified by **__global__ void**. They are called by the host and run on the device, thus serving as entries to the device.

We have only one kernel invocation here, but could have many, say with the output of one serving as input to the next.

- Other functions that will run on the device, called by functions running on the device, must be identified by **__device__**, e.g.

```
__device__ int sumvector(float *x, int n)
```

- When a kernel is called, each thread runs it. Each thread receives the same arguments.
- Each block and thread has an ID, stored in programmer-accessible structs **blockIdx** and **threadIdx**. We'll discuss the details later, but for now, we'll just note that here the statement

```
int rownum = blockIdx.x;
```

picks up the block number, which our code in this example uses to determine which row to sum.

- One calls **cudaMalloc()** to dynamically allocate space on the device's memory. Execution of the statement

```
cudaMalloc((void **) &drs, rssize);
```

allocates space *on the device*, pointed to by **drs**, a variable in the *host's* address space.

The space allocated by a **cudaMalloc()** call on the device is global to all kernels, and resides in the global memory of the device (details on memory types later).

One can also allocate device memory statically. For example, the statement

```
__device int z[100];
```

appearing outside any function definition would allocate space on device global memory, with scope global to all kernels. However, it is not accessible to the host.

- Data is transferred to and from the host and device memories via **cudaMemcpy()**. The fourth argument specifies the direction, e.g. `cudaMemcpyHostToDevice`, `cudaMemcpyDeviceToHost` or `cudaMemcpyDeviceToDevice`.
- Kernels return **void** values, so values are returned via a kernel's arguments.
- Device functions (which we don't have here) can return values. They are called only by kernel functions or other device functions.
- Note carefully that a call to the kernel doesn't block; it returns immediately. For that reason, the code above has a host barrier call, to avoid copying the results back to the host from the device before they're ready:

```
cudaThreadSynchronize();
```

On the other hand, if our code were to have another kernel call, say on the next line after

```
findIelt<<<dimGrid, dimBlock>>>(dm, drs, n);
```

and if some of the second call's input arguments were the outputs of the first call, there would be an implied barrier between the two calls; the second would not start execution before the first finished.

Calls like `cudaMemcpy()` do block until the operation completes.

There is also a thread barrier available for the threads themselves, at the block level. The call is

```
__syncthreads();
```

This can only be invoked by threads within a block, not across blocks. In other words, this is barrier synchronization within blocks.

- I've written the program so that each thread will handle one row of the matrix. Since (a) I've chosen to store the matrix in one-dimensional form, (b) the matrix is of size $n \times n$ and (c) in the C language, two-dimensional arrays are stored in column-major order, the loop

```
for (int k = 0; k < n; k++)
    sum += m[rownum*n+k];
```

will indeed traverse the n elements of row number **rownum**, and compute their sum. That sum is then placed in the proper element of the output array:

```
rs[rownum] = sum;
```

- After the kernel returns, the host must copy the result back from the device memory to the host memory, in order to access the results of the call.

5.3 Understanding the Hardware Structure

Scorecards, get your scorecards here! You can't tell the players without a scorecard—classic cry of vendors at baseball games

Know thy enemy—Sun Tzu, *The Art of War*

The enormous computational potential of GPUs cannot be unlocked without an intimate understanding of the hardware. This of course is a fundamental truism in the parallel processing world, but it is acutely important for GPU programming. This section presents an overview of the hardware.

5.3.1 Processing Units

A GPU consists of a large set of **streaming multiprocessors** (SMs); you might say it's a multi-multiprocessor machine. Each SM consists of a number of **streaming processors** (SPs). It is important to understand the motivation for this hierarchy: Two threads located in different SMs cannot synchronize with each other in

the barrier sense. Though this sounds like a negative at first, it is actually a great advantage, as the independence of threads in separate SMs means that the hardware can run faster. So, if the CUDA application programmer can write his/her algorithm so as to have certain independent chunks, and those chunks can be assigned to different SMs (we'll see how, shortly), then that's a "win."

Note that at present, word size is 32 bits. Thus for instance floating-point operations in hardware were originally in single precision only, though newer devices are capable of double precision.

5.3.2 Thread Operation

GPU operation is highly threaded, and again, understanding of the details of thread operation is key to good performance.

5.3.2.1 SIMT Architecture

The threads running within an SM *can* synchronize with each other, but there is further hierarchy.

When you write a CUDA application program, you partition the threads into groups called **blocks**. The hardware will assign an entire block to a single SM, though several blocks can run in the same SM. The hardware will then divide a block into **warps**, 32 threads to a warp. Knowing that the hardware works this way, the programmer controls the block size and the number of blocks, and in general writes the code to take advantage of how the hardware works.

A key point is that *all the threads in a warp run the code in lockstep*. During the machine instruction fetch cycle, the same instruction will be fetched for all of the threads in the warp. Then in the execution cycle, each thread will either execute that particular instruction or execute nothing. The execute-nothing case occurs in the case of branches; see below. This is the classical **single instruction, multiple data** (SIMD) pattern used in some early special-purpose computers such as the ILLIAC; here it is called **single instruction, multiple thread** (SIMT).

The syntactic details of grid and block configuration will be presented in Section 5.3.4.

5.3.2.2 The Problem of Thread Divergence

The SIMT nature of thread execution has major implications for performance. Consider what happens with if/then/else code. If some threads in a warp take the "then" branch and others go in the "else" direction, they cannot operate in lockstep. That means that some threads must wait while others execute. This renders the code at that point serial rather than parallel, a situation called **thread divergence**. As one CUDA Web tutorial points out, this can be a "performance killer." (On the other hand, threads in the same block but in different warps can diverge with no problem.)

In turn, the implication for writing CUDA code is that if you have an “embarrassingly parallel” application—parallelizable into independent chunks with very little communication between them—or your application has components with that property, you should arrange to have the chunks run in different blocks.

5.3.2.3 “OS in Hardware”

Each SM runs the threads on a timesharing basis, just like an operating system (OS). This timesharing is implemented in the hardware, though, not in software as in the OS case. Just as a process in an OS is given a fixed-length timeslice, so that processes take turns running, in a GPU’s hardware OS, warps take turns running, with fixed-length timeslices.

Also, with an OS, if a process reaches an input/output operation, the OS suspends the process while I/O is pending, even if its turn is not up. The OS then runs some other process instead, so as to avoid wasting CPU cycles during the long period of time needed for the I/O. With an SM, the analogous situation is a long memory operation, to global memory; if a warp of threads needs to access global memory (including local memory), the SM will schedule some other warp while the memory access is pending.

The hardware support for threads is extremely good; a context switch takes very little time, quite a contrast to the OS case. Moreover, as noted above, the long latency of global memory may be solvable by having a lot of threads that the hardware can timeshare to hide that latency; while one warp is fetching data from memory, another warp can be executing, thus not losing time due to the long fetch delay. For these reasons, CUDA programmers typically employ a large number of threads, each of which does only a small amount of work—again, quite a contrast to something like OpenMP.

5.3.3 Memory Structure

In addition to the GPU/SM/SP and block/thread/warp hierarchies discussed above, another key hierarchy is memory structure. Let’s discuss the most important two types of memory first—shared and global. Here is a summary.

5.3.3.1 Shared and Global Memory

Here is a summary:

type	shared	global
scope	gbl. to block	gbl. to app.
size	small	large
loc.	on-chip	off-chip
speed	blinding	molasses
lifetime	kernel	application
host access?	no	yes
cached?	no	no

In prose form:

- **Shared memory:** All the threads in an SM share this memory, and use it to communicate among themselves, just as is the case with threads in CPUs. Access is very fast, as this memory is on-chip. It is declared inside the kernel, or in the kernel call (details below).

On the other hand, shared memory is small, currently 16K bytes per SM, and the data stored in it are valid only for the life of the currently-executing kernel. Also, shared memory cannot be accessed by the host.

- **Global memory:** This is shared by all the threads in an entire application, and is persistent across kernel calls, throughout the life of the application, i.e. until the program running on the host exits. It is usually much larger than shared memory. It is accessible from the host. Pointers to global memory can (but do not have to) be declared outside the kernel.

On the other hand, global memory is off-chip and very slow, taking hundreds of clock cycles per access instead of just a few. As noted earlier, this can be ameliorated by exploiting latency hiding; we will elaborate on this in Section 5.3.3.2.

The reader should pause here and reread the above comparison between shared and global memories. *The key implication is that shared memory is used essentially as a programmer-managed cache.* Data will start out in global memory, but if a variable is to be accessed multiple times by the GPU code, it's probably better for the programmer to write code that copies it to shared memory, and then access the copy instead of the original. If the variable is changed and is to be eventually transmitted back to the host, the programmer must include code to copy it back to global memory.

Neither memory type is (hardware) cached.

Accesses to global and shared memory are done via half-warps, i.e. an attempt is made to do all memory accesses in a half-warp simultaneously. In that sense, only threads in a half-warp run simultaneously, but the full warp is *scheduled* simultaneously by the hardware OS.

The host can access global memory via `cudaMemcpy()`, as seen earlier. It cannot access shared memory. Here is a typical pattern:

```
__global__ void abckernel(int *abcglobalmem)
{
```

```

__shared__ int abcsharedmem[100];
// ... code to copy some of abcglobalmem to some of abcsharedmem
// ... code for computation
// ... code to copy some of abcsharedmem to some of abcglobalmem
}

```

Typically you would write the code so that each thread deals with its own portion of the shared data, e.g. its own portion of **abcsharedmem** and **abcglobalmem** above. However, all the threads in that block can read/write any element in **abcsharedmem**.

Shared memory consistency (recall Section 2.6) is sequential within a thread, but **relaxed** among threads in a block. A write by one thread is not guaranteed to be visible to the others in a block until **__syncthreads()** is called. On the other hand, writes by a thread *will* be visible to that same thread in subsequent reads. Among the implications of this is that if each thread writes only to portions of shared memory that are not read by other threads in the block, then **__syncthreads()** need not be called.

In the code fragment above, we allocated the shared memory through a C-style declaration:

```
__shared__ int abcsharedmem[100];
```

It is also possible to allocate shared memory in the kernel call, along with the block and thread configuration. Here is an example:

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <cuda.h>
4
5  // CUDA example: illustrates kernel-allocated shared memory; does
6  // nothing useful, just copying an array from host to device global,
7  // then to device shared, doubling it there, then copying back to device
8  // global then host
9
10 __global__ void doubleit(int *dv, int n)
11 { extern __shared__ int sv[];
12   int me = threadIdx.x;
13   // threads share in copying dv to sv, with each thread copying one
14   // element
15   sv[me] = 2 * dv[me];
16   __syncthreads(); // probably not needed in this case
17   dv[me] = sv[me];
18 }
19
20 int main(int argc, char **argv)
21 {
22   int n = atoi(argv[1]); // number of matrix rows/cols
23   int *hv, // host array
24       *dv; // device array
25   int vsize = n * sizeof(int); // size of array in bytes
26   // allocate space for host array
27   hv = (int *) malloc(vsize);

```

```

28     // fill test array with consecutive integers
29     int t = 0,i;
30     for (i = 0; i < n; i++)
31         hv[i] = t++;
32     // allocate space for device array
33     cudaMalloc((void **)&dv,vsize);
34     // copy host array to device array
35     cudaMemcpy(dv,hv,vsize,cudaMemcpyHostToDevice);
36     // set up parameters for threads structure
37     dim3 dimGrid(1,1);
38     dim3 dimBlock(n,1,1); // all n threads in the same block
39     // invoke the kernel; third argument is amount of shared memory
40     doubleit<<<dimGrid,dimBlock,vsize>>>(dv,n);
41     // wait for kernel to finish
42     cudaThreadSynchronize();
43     // copy row array from device to host
44     cudaMemcpy(hv,dv,vsize,cudaMemcpyDeviceToHost);
45     // check results
46     if (n < 10) for(int i=0; i<n; i++) printf("%d\n",hv[i]);
47     // clean up
48     free(hv);
49     cudaFree(dv);
50 }

```

Here the variable `sv` is kernel allocated. It's declared in the statement

```
extern __shared__ int sv[];
```

but actually allocated during the kernel invocation

```
doubleit<<<dimGrid,dimBlock,vsize>>>(dv,n);
```

in that third argument within the chevrons, **`vsize`**.

Note that one can only directly declare one region of space in this manner. This has two implications:

- Suppose we have two **`__device__`** functions, each declared an **`extern __shared__`** array like this. Those two arrays will occupy the same place in memory!
- Suppose within one **`__device__`** function, we wish to have two **`extern __shared__`** arrays. We cannot do that literally, but we can share the space via subarrays, e.g.:

```
int *x = &sv[120];
```

would set up `x` as a subarray of `sv` above, starting at element 120.

One can also set up shared arrays of fixed length in the same code. Declare them before the variable-length one.

In our example above, the array `sv` is syntactically local to the function `doubleit()`, but is shared by all invocations of that function in the block, thus acting “global” to them in a sense. But the point is that it is not accessible from within *other* functions running in that block. In order to achieve the latter situation, a shared array can be declared outside any function.

5.3.3.2 Global-Memory Performance Issues

As noted, the **latency**—time to access a single word—for global memory is quite high, on the order of hundreds of clock cycles. However, the hardware attempts to ameliorate this problem in a couple of ways.

First, as mentioned earlier, if a warp has requested a global memory access that will take a long time, the hardware will schedule another warp to run while the first is waiting. This is an example of a common parallel processing technique called **latency hiding**.

Second, the **bandwidth** to global memory—the number of words accessed per unit time—can be high, due to hardware actions called **coalescing**. This simply means that if the hardware sees that the threads in this half-warp (or at least the ones currently accessing global memory) are accessing consecutive words, the hardware can execute the memory requests in groups of up to 32 words at a time. This is true for both reads and writes.

The newer GPUs go even further, coalescing much more general access patterns, not just to consecutive words.

5.3.3.3 Shared-Memory Performance Issues

Shared memory is divided into **banks**, in a low-order interleaved manner (recall Section 2.2): Words with consecutive addresses are stored in consecutive banks, mod the number of banks, i.e. wrapping back to 0 when hitting the last bank. If for instance there are 8 banks, addresses 0, 8, 16,... will be in bank 0, addresses 1, 9, 17,... will be in bank 1 and so on. (Actually, older devices have 16 banks, while newer ones have 32.)

The fact that all memory accesses in a half-warp are attempted simultaneously implies that the best access to shared memory arises when the accesses are to different banks. For this reason, if one needs to access an entire array in order, one should write one’s code so that consecutive threads access consecutive array elements, so as to avoid bank conflicts, which inhibit speed.

An exception occurs in **broadcast**. If all threads in the block wish to read from the same word in the same bank, the word will be sent to all the requestors simultaneously without conflict. However, if only some threads try to read the same word, there may or may not be a conflict, as the hardware chooses a bank for

broadcast in some unspecified way.

5.3.3.4 Host/Device Memory Transfer Performance Issues

Copying data between host and device can be a major bottleneck. One way to ameliorate this is to use `cudaMallocHost()` instead of `malloc()` when allocating memory on the host. This sets up page-locked memory, meaning that it cannot be swapped out by the OS' virtual memory system. This allows the use of DMA hardware to do the memory copy, said to make `cudaMemcpy()` twice as fast.

5.3.3.5 Other Types of Memory

There are also other types of memory. Again, let's start with a summary:

type	registers	local	constant
scope	single thread	single thread	glbl. to app.
loc.	on-chip	off-chip	off-chip
speed	blinding	molasses	molasses
lifetime	kernel	kernel	application
host access?	no	no	yes
cached?	no	no	yes

- Registers: Each SM has a set of registers. They are much more numerous than in a CPU. Access to them is very fast, said to be slightly faster than to shared memory.

The compiler normally stores the local variables for a device function in registers, but there are exceptions. An array won't be placed in registers if the array is too large, or if the array has variable index values, such as

```
int z[20], i;
...
y = z[i];
```

Since registers are not indexable, the compiler cannot allocate `z` to registers in this case. If on the other hand, the only code accessing `z` has constant indices, e.g. `z[8]`, the compiler may put `z` in registers.

- Local memory: This is physically part of global memory, but is an area within that memory that is allocated by the compiler for a given thread. As such, it is slow, and accessible only by that thread. The compiler allocates this memory for local variables in a device function if the compiler cannot store them in registers. This is called **register spill**.

- **Constant memory:** As the name implies, it's read-only from the device (read/write by the host), for storing values that will not change. It is off-chip, thus potentially slow, but has a cache on the chip. At present, the size is 64K.

One designates this memory with `__constant__`, as a global variable in the source file. One sets its contents from the host via `cudaMemcpyToSymbol()`, For example:

```
__constant__ int x;
int y = 3;
// host code
cudaMemcpyToSymbol("x", &y, sizeof(int));
...
// device code
int z;
z = x;
```

Visible to all threads.

- **Texture memory:** This memory is closer to graphics applications, and is essentially two-dimensional. It is read-only. It also is off-chip but has an on-chip cache.

5.3.4 Threads Hierarchy

Following the hardware, threads in CUDA software follow a hierarchy:

- The entirety of threads for an application is called a **grid**.
- A grid consists of one or more **blocks** of threads.
- Each block has its own ID within the grid, consisting of an “x coordinate” and a “y coordinate.”
- Likewise each thread has x, y and z coordinates within whichever block it belongs to.
- Just as an ordinary CPU thread needs to be able to sense its ID, e.g. by calling `omp_get_thread_num()` in OpenMP, CUDA threads need to do the same. A CUDA thread can access its block ID via the built-in variables `blockIdx.x` and `blockIdx.y`, and can access its thread ID within its block via `threadIdx.x`, `threadIdx.y` and `threadIdx.z`.
- The programmer specifies the grid size (the numbers of rows and columns of blocks within a grid) and the block size (numbers of rows, columns and layers of threads within a block). In the first example above, this was done by the code

```
dim3 dimGrid(n,1);
dim3 dimBlock(1,1,1);
find1elt<<<dimGrid,dimBlock>>>(dm,drs,n);
```

Here the grid is specified to consist of n ($n \times 1$) blocks, and each block consists of just one ($1 \times 1 \times 1$) thread.

That last line is of course the call to the kernel. As you can see, CUDA extends C syntax to allow specifying the grid and block sizes. CUDA will store this information in structs of type **dim3**, in this case **gridDim** and **blockDim**, accessible to the programmer, again with member variables for the various dimensions, e.g. **blockDim.x** for the size of the X dimension for the number of threads per block.

- All threads in a block run in the same SM, though more than one block might be on the same SM.
- The “coordinates” of a block within the grid, and of a thread within a block, are merely abstractions. If for instance one is programming computation of heat flow across a two-dimensional slab, the programmer may find it clearer to use two-dimensional IDs for the threads. *But this does not correspond to any physical arrangement in the hardware.*

As noted, the motivation for the two-dimensional block arrangement is to make coding conceptually simpler for the programmer if he/she is working an application that is two-dimensional in nature.

For example, in a matrix application one’s parallel algorithm might be based on partitioning the matrix into rectangular submatrices (**tiles**), as we’ll do in Section 8.3. In a small example there, the matrix

$$A = \begin{pmatrix} 1 & 5 & 12 \\ 0 & 3 & 6 \\ 4 & 8 & 2 \end{pmatrix} \quad (5.1)$$

is partitioned as

$$A = \begin{pmatrix} A_{00} & A_{01} \\ A_{10} & A_{11} \end{pmatrix}, \quad (5.2)$$

where

$$A_{00} = \begin{pmatrix} 1 & 5 \\ 0 & 3 \end{pmatrix}, \quad (5.3)$$

$$A_{01} = \begin{pmatrix} 12 \\ 6 \end{pmatrix}, \quad (5.4)$$

$$A_{10} = (4 \ 8) \quad (5.5)$$

and

$$A_{11} = \begin{pmatrix} 2 \end{pmatrix}. \quad (5.6)$$

We might then have one block of threads handle A_{00} , another block handle A_{01} and so on. CUDA's two-dimensional ID system for blocks makes life easier for programmers in such situations.

5.3.5 What's NOT There

We're not in Kansas anymore, Toto—character Dorothy Gale in *The Wizard of Oz*

It looks like C, it feels like C, and for the most part, it *is* C. But in many ways, it's quite different from what you're used to:

- You don't have access to the C library, e.g. **printf()** (the library consists of host machine language, after all). There are special versions of math functions, however, e.g. **__sin()**.
- No recursion.
- No stack. Functions are essentially inlined, rather than their calls being handled by pushes onto a stack.
- No pointers to functions.

5.4 Synchronization

As mentioned earlier, a barrier for the threads in the same block is available by calling **__syncthreads()**. Note carefully that if one thread writes a variable to shared memory and another then reads that variable, one must call this function in order to get the latest value. Remember, threads across blocks cannot sync with each other in this manner.

Several **atomic** operations—read/modify/write actions that a thread can execute without **pre-emption**, i.e. without interruption—are available on both global and shared memory. For example, **atomicAdd()** performs a fetch-and-add operation, as described in Section 2.7 of this book. The call is

```
atomicAdd(address of integer variable,inc);
```

where **address of integer variable** is the address of the (device) variable to add to, and **inc** is the amount to be added. The return value of the function is the value originally at that address before the operation.

There are also **atomicExch()** (exchange the two operands), **atomicCAS()** (if the first operand equals the second, replace the first by the third), **atomicMin()**, **atomicMax()**, **atomicAnd()**, **atomicOr()**, and so on.

Use **-arch=sm_11** when compiling, e.g.

```
nvcc -g -G yoursrc.cu -arch=sm_11
```

5.5 Hardware Requirements, Installation, Compilation, Debugging

You do need what is currently (December 2010) a high-end NVIDIA video card. There is a list at http://www.nvidia.com/object/cuda_gpus.html. If you have a Linux system, run **lspci** to determine what kind you have.

Download the CUDA toolkit from NVIDIA. Just plug “CUDA download” into a Web search engine to find the site. Install as directed.

You’ll need to set your search and library paths to include the CUDA **bin** and **lib** directories.

To compile **x.cu** (and yes, use the **.cu** suffix), type

```
$ nvcc -g -G x.cu
```

The **-g -G** options are for setting up debugging, the first for host code, the second for device code. You may also need to specify

```
-I/your_CUDA_include_path
```

to pick up the file **cuda.h**. Run the code as you normally would.

You may need to take special action to set your library path properly. For example, on Linux machines, set the environment variable **LD_LIBRARY_PATH** to include the CUDA library.

To determine the limits, e.g. maximum number of threads, for your device, use code like this:

```
cudaDeviceProp Props;  
cudaGetDeviceProperties(&Props, 0);
```

The 0 is for device 0, assuming you only have one device. The return value of **cudaGetDeviceProperties()** is a complex C struct whose components are listed at http://developer.download.nvidia.com/compute/cuda/2_3/toolkit/docs/online/group__CUDART__DEVICE_g5aa4f47938af8276f08074c

html. But I recommend printing it from within GDB to see the values. One of the fields gives clock speed, which is typically slower than that of the host.

Under older versions of CUDA, such as 2.3, one can debug using GDB as usual. You must compile your program in emulation mode, using the **-deviceemu** command-line option. This is no longer available as of version 3.2. CUDA also includes a special version of GDB, CUDA-GDB (invoked as **cuda-gdb**) for real-time debugging. However, on Unix-family platforms it runs only if X11 is not running. Short of dedicating a machine for debugging, you may find it useful to install a version 2.3 in addition to the most recent one to use for debugging.

5.6 Improving the Sample Program

The issues involving coalescing in Section 5.3.3.2 would suggest that our rowsum code might run faster with column sums. As two threads in the same half-warp march down adjoining columns in lockstep, they will always be accessing adjoining words in memory.

So, I modified the program accordingly (not shown), and compiled the two versions, as **rs** and **cs**, the row- and column-sum versions of the code, respectively.

This did produce a small improvement (confirmed in subsequent runs, needed in any timing experiment):

```
pc5:~/CUDA% time rs 20000
2.585u 1.753s 0:04.54 95.3%    0+0k 7104+0io 54pf+0w
pc5:~/CUDA% time cs 20000
2.518u 1.814s 0:04.40 98.1%    0+0k 536+0io 5pf+0w
```

But let's compare it to a version running only on the CPU,

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 // non-CUDA example: finds col sums of an integer matrix m
5
6 // findlelt() finds the colsum of one col of the nxn matrix m, storing the
7 // result in the corresponding position in the colsum array cs; matrix
8 // stored as 1-dimensional, row-major order
9
10 void findlelt(int *m, int *cs, int n)
11 {
12     int sum=0;
13     int topofcol;
14     int col,k;
15     for (col = 0; col < n; col++) {
16         topofcol = col;
17         sum = 0;
18         for (k = 0; k < n; k++)
```

```

19         sum += m[topofcol+k*n];
20         cs[col] = sum;
21     }
22 }
23
24 int main(int argc, char **argv)
25 {
26     int n = atoi(argv[1]); // number of matrix cols/cols
27     int *hm, // host matrix
28         *hcs; // host colsums
29     int msize = n * n * sizeof(int); // size of matrix in bytes
30     // allocate space for host matrix
31     hm = (int *) malloc(msize);
32     // as a test, fill matrix with consecutive integers
33     int t = 0,i,j;
34     for (i = 0; i < n; i++) {
35         for (j = 0; j < n; j++) {
36             hm[i*n+j] = t++;
37         }
38     }
39     int cssize = n * sizeof(int);
40     hcs = (int *) malloc(cssize);
41     findlelt(hm,hcs,n);
42     if (n < 10) for(i=0; i<n; i++) printf("%d\n",hcs[i]);
43     // clean up
44     free(hm);
45     free(hcs);
46 }

```

How fast does this non-CUDA version run?

```

pc5:~/CUDA% time csc 20000
61.110u 1.719s 1:02.86 99.9%    0+0k 0+0io 0pf+0w

```

Very impressive! No wonder people talk of CUDA in terms like “a supercomputer on our desktop.” And remember, this includes the time to copy the matrix from the host to the device (and to copy the output array back). And we didn’t even try to optimize thread configuration, memory coalescing and bank usage, making good use of memory hierarchy, etc.

On the other hand, remember that this is an “embarrassingly parallel” application, and in many applications we may have to settle for a much more modest increase, and work harder to get it.

5.7 More Examples

5.7.1 Finding the Mean Number of Mutual Outlinks

Consider a network graph of some kind, such as Web links. For any two vertices, say any two Web sites, we might be interested in mutual outlinks, i.e. outbound links that are common to two Web sites. The CUDA

code below finds the mean number of mutual outlinks, among all pairs of sites in a set of Web sites.

```

1 #include <cuda.h>
2 #include <stdio.h>
3
4 // CUDA example: finds mean number of mutual outlinks, among all pairs
5 // of Web sites in our set
6
7 // for a given thread number tn, determines the pair of rows to be
8 // processed by that thread, in an nxn matrix; returns pointer to the
9 // pair of row numbers
10 __device__ void findpair(int tn, int n, int *pair)
11 {
12     int sum=0,oldsum=0,i;
13     for(i=0; ;i++) {
14         sum += n - i - 1;
15         if (tn <= sum-1) {
16             pair[0] = i;
17             pair[1] = tn - oldsum + i + 1;
18             return;
19         }
20         oldsum = sum;
21     }
22 }
23 // proclpair() processes one pair of Web sites, i.e. one pair of rows in
24 // the nxn adjacency matrix m; the number of mutual outlinks is added to
25 // tot
26 __global__ void proclpair(int *m, int *tot, int n)
27 {
28     // find (i,j) pair to assess for mutuality
29     int pair[2];
30     findpair(threadIdx.x,n,pair);
31     int sum=0;
32     int startrowa = pair[0] * n,
33         startrowb = pair[1] * n;
34     for (int k = 0; k < n; k++)
35         sum += m[startrowa + k] * m[startrowb + k];
36     atomicAdd(tot,sum);
37 }
38
39 int main(int argc, char **argv)
40 {
41     int n = atoi(argv[1]); // number of matrix rows/cols
42     int *hm, // host matrix
43         *dm, // device matrix
44         htot, // host grand total
45         *dtot; // device grand total
46     int msize = n * n * sizeof(int); // size of matrix in bytes
47     // allocate space for host matrix
48     hm = (int *) malloc(msize);
49     // as a test, fill matrix with random 1s and 0s
50     int i,j;
51     for (i = 0; i < n; i++) {
52         hm[n*i+i] = 0;
53         for (j = 0; j < n; j++) {
54             if (j != i) hm[i*n+j] = rand() % 2;
55         }
56     }
57 }

```

```

56     }
57     // allocate space for device matrix
58     cudaMalloc((void **)&dm,msize);
59     // copy host matrix to device matrix
60     cudaMemcpy(dm,hm,msize,cudaMemcpyHostToDevice);
61     htot = 0;
62     // set up device total and initialize it
63     cudaMalloc((void **)&dtot,sizeof(int));
64     cudaMemcpy(dtot,&htot,sizeof(int),cudaMemcpyHostToDevice);
65     // set up parameters for threads structure
66     dim3 dimGrid(1,1);
67     int npairs = n*(n-1)/2;
68     dim3 dimBlock(npairs,1,1);
69     // invoke the kernel
70     proclpair<<<dimGrid,dimBlock>>>(dm,dtot,n);
71     // wait for kernel to finish
72     cudaThreadSynchronize();
73     // copy total from device to host
74     cudaMemcpy(&htot,dtot,sizeof(int),cudaMemcpyDeviceToHost);
75     // check results
76     if (n <= 15) {
77         for (i = 0; i < n; i++) {
78             for (j = 0; j < n; j++)
79                 printf("%d ",hm[n*i+j]);
80             printf("\n");
81         }
82     }
83     printf("mean = %f\n",htot/float(npairs));
84     // clean up
85     free(hm);
86     cudaFree(dm);
87     cudaFree(dtot);
88 }

```

The main programming issue here is finding a way to partition the various pairs (i,j) to the different threads. The function **findpair()** here does that.

Note the use of **atomicAdd()**.

5.7.2 Finding Prime Numbers

The code below finds all the prime numbers from 2 to **n**.

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <cuda.h>
4
5  // CUDA example: illustration of shared memory allocation at run time;
6  // finds primes using classical Sieve of Erathosthenes: make list of
7  // numbers 2 to n, then cross out all multiples of 2 (but not 2 itself),
8  // then all multiples of 3, etc.; whatever is left over is prime; in our
9  // array, 1 will mean "not crossed out" and 0 will mean "crossed out"

```

```

10
11 // IMPORTANT NOTE: uses shared memory, in a single block, without
12 // rotating parts of array in and out of shared memory; thus limited to
13 // n <= 4000 if have 16K shared memory
14
15 // initialize sprimes, 1s for the odds, 0s for the evens; see sieve()
16 // for the nature of the arguments
17 __device__ void initsp(int *sprimes, int n, int nth, int me)
18 {
19     int chunk, startsetsp, endsetsp, val, i;
20     sprimes[2] = 1;
21     // determine sprimes chunk for this thread to init
22     chunk = (n-1) / nth;
23     startsetsp = 2 + me*chunk;
24     if (me < nth-1) endsetsp = startsetsp + chunk - 1;
25     else endsetsp = n;
26     // now do the init
27     val = startsetsp % 2;
28     for (i = startsetsp; i <= endsetsp; i++) {
29         sprimes[i] = val;
30         val = 1 - val;
31     }
32     // make sure sprimes up to date for all
33     __syncthreads();
34 }
35
36 // copy sprimes back to device global memory; see sieve() for the nature
37 // of the arguments
38 __device__ void cpytogl(int *dprimes, int *sprimes, int n, int nth, int me)
39 {
40     int startcpy, endcpy, chunk, i;
41     chunk = (n-1) / nth;
42     startcpy = 2 + me*chunk;
43     if (me < nth-1) endcpy = startcpy + chunk - 1;
44     else endcpy = n;
45     for (i = startcpy; i <= endcpy; i++) dprimes[i] = sprimes[i];
46     __syncthreads();
47 }
48
49 // finds primes from 2 to n, storing the information in dprimes, with
50 // dprimes[i] being 1 if i is prime, 0 if composite; nth is the number
51 // of threads (threadDim somehow not recognized)
52 __global__ void sieve(int *dprimes, int n, int nth)
53 {
54     extern __shared__ int sprimes[];
55     int me = threadIdx.x;
56     int nth1 = nth - 1;
57     // initialize sprimes array, 1s for odds, 0 for evens
58     initsp(sprimes, n, nth, me);
59     // "cross out" multiples of various numbers m, with each thread doing
60     // a chunk of m's; always check first to determine whether m has
61     // already been found to be composite; finish when m*m > n
62     int maxmult, m, startmult, endmult, chunk, i;
63     for (m = 3; m*m <= n; m++) {
64         if (sprimes[m] != 0) {
65             // find largest multiple of m that is <= n
66             maxmult = n / m;
67             // now partition 2,3,...,maxmult among the threads

```

```

68     chunk = (maxmult - 1) / nth;
69     startmult = 2 + me*chunk;
70     if (me < nth-1) endmult = startmult + chunk - 1;
71     else endmult = maxmult;
72     }
73     // OK, cross out my chunk
74     for (i = startmult; i <= endmult; i++) sprimes[i*m] = 0;
75     }
76     __syncthreads();
77     // copy back to device global memory for return to host
78     cpytogl(dprimes, sprimes, n, nth, me);
79 }
80
81 int main(int argc, char **argv)
82 {
83     int n = atoi(argv[1]), // will find primes among 1, ..., n
84         nth = atoi(argv[2]); // number of threads
85     int *hprimes, // host primes list
86         *dprimes; // device primes list
87     int psize = (n+1) * sizeof(int); // size of primes lists in bytes
88     // allocate space for host list
89     hprimes = (int *) malloc(psize);
90     // allocate space for device list
91     cudaMalloc((void **)&dprimes, psize);
92     dim3 dimGrid(1,1);
93     dim3 dimBlock(nth,1,1);
94     // invoke the kernel, including a request to allocate shared memory
95     sieve<<<dimGrid, dimBlock, psize>>>(dprimes, n, nth);
96     // check whether we asked for too much shared memory
97     cudaError_t err = cudaGetLastError();
98     if(err != cudaSuccess) printf("%s\n", cudaGetErrorString(err));
99     // wait for kernel to finish
100    cudaThreadSynchronize();
101    // copy list from device to host
102    cudaMemcpy(hprimes, dprimes, psize, cudaMemcpyDeviceToHost);
103    // check results
104    if (n <= 1000) for(int i=2; i<=n; i++)
105        if (hprimes[i] == 1) printf("%d\n", i);
106    // clean up
107    free(hprimes);
108    cudaFree(dprimes);
109 }

```

This code has been designed with some thought as to memory speed and thread divergence. Ideally, we would like to use device shared memory if possible, and to exploit the lockstep, SIMD nature of the hardware.

The code uses the classical Sieve of Erathosthenes, “crossing out” multiples of 2, 3, 5, 7 and so on to get rid of all the composite numbers. However, the code here differs from that in Section 1.3.1.2, even though both programs use the Sieve of Erathosthenes.

Say we have just two threads, A and B. In the earlier version, thread A might cross out all multiples of 19 while B handles multiples of 23. In this new version, thread A deals with only some multiples of 19 and B handles the others for 19. Then they both handle their own portions of multiples of 23, and so on.

The thinking here is that the second version will be more amenable to lockstep execution, thus causing less thread divergence.

Thus in this new version, each thread handles a chunk of multiples of the given prime. Note the contrast of this with many CUDA examples, in which each thread does only a small amount of work, such as computing a single element in the product of two matrices.

In order to enhance memory performance, this code uses device shared memory. All the “crossing out” is done in the shared memory array **sprimes**, and then when we are all done, that is copied to the device global memory array **dprimes**, which is in turn copies to host memory. By the way, note that the amount of shared memory here is determined dynamically.

However, device shared memory consists only of 16K bytes, which would limit us here to values of **n** up to about 4000. Extending the program to work for larger values of **n** would require some careful planning if we still wish to use shared memory.

5.8 CUBLAS

CUDA includes some parallel linear algebra routines callable from straight C code. In other words, you can get the benefit of GPU in linear algebra contexts without using CUDA.¹ Note, though, that in fact you *are* using CUDA, behind the scenes.

And indeed, you can mix CUDA and CUBLAS code. Your program might have multiple kernel invocations, some CUDA and others CUBLAS, with each using data in device global memory that was written by earlier kernels. Again, remember, the contents of device global memory (including the bindings of variable names) are persistent across kernel calls in the same application.

Below is an example **RowSumsCB.c**, the matrix row sums example again, this time using CUBLAS. We can find the vector of row sums of the matrix A by post-multiplying A by a column vector of all 1s.

I compiled the code by typing

```
gcc -g -I/usr/local/cuda/include -L/usr/local/cuda/lib RowSumsCB.c -lcublas -lcudart
```

You should modify for your own CUDA locations accordingly. Users who merely wish to use CUBLAS will find the above more convenient, but if you are mixing CUDA and CUBLAS, you would use **nvcc**:

```
nvcc -g -G RowSumsCB.c -lcublas
```

Here is the code:

¹CUDA also includes a similar package for Fast Fourier Transform computation, CUFFT.

```

1 #include <stdio.h>
2 #include <cublas.h> // required include
3
4 int main(int argc, char **argv)
5 {
6     int n = atoi(argv[1]); // number of matrix rows/cols
7     float *hm, // host matrix
8           *hrs, // host rowsums vector
9           *ones, // 1s vector for multiply
10          *dm, // device matrix
11          *drs; // device rowsums vector
12     // allocate space on host
13     hm = (float *) malloc(n*n*sizeof(float));
14     hrs = (float *) malloc(n*sizeof(float));
15     ones = (float *) malloc(n*sizeof(float));
16     // as a test, fill hm with consecutive integers, but in column-major
17     // order for CUBLAS; also put 1s in ones
18     int i, j;
19     float t = 0.0;
20     for (i = 0; i < n; i++) {
21         ones[i] = 1.0;
22         for (j = 0; j < n; j++)
23             hm[j*n+i] = t++;
24     }
25     cublasInit(); // required init
26     // set up space on the device
27     cublasAlloc(n*n, sizeof(float), (void**)&dm);
28     cublasAlloc(n, sizeof(float), (void**)&drs);
29     // copy data from host to device
30     cublasSetMatrix(n, n, sizeof(float), hm, n, dm, n);
31     cublasSetVector(n, sizeof(float), ones, 1, drs, 1);
32     // matrix times vector
33     cublasSgemv('n', n, n, 1.0, dm, n, drs, 1, 0.0, drs, 1);
34     // copy result back to host
35     cublasGetVector(n, sizeof(float), drs, 1, hrs, 1);
36     // check results
37     if (n < 20) for (i = 0; i < n; i++) printf("%f\n", hrs[i]);
38     // clean up on device (should call free() on host too)
39     cublasFree(dm);
40     cublasFree(drs);
41     cublasShutdown();
42 }

```

As noted in the comments, CUBLAS assumes FORTRAN-style, i.e. column-major order, for matrices.

Now that you know the basic format of CUDA calls, the CUBLAS versions will look similar. In the call

```
cublasAlloc(n*n, sizeof(float), (void**)&dm);
```

for instance, we are allocating space on the device for an $n \times n$ matrix of **floats**, on the device.

The call

```
cublasSetMatrix(n, n, sizeof(float), hm, n, dm, n);
```

is slightly more complicated. Here we are saying that we are copying **hm**, an $n \times n$ matrix of **floats** on the host, to **dm** on the host. The **n** arguments in the last and third-to-last positions again say that the two matrices each have **n** dimensioned rows. This seems redundant, but this is needed in cases of matrix tiling, where the number of rows of a tile would be less than the number of rows of the matrix as a whole.

The 1s in the call

```
cublasSetVector(n, sizeof(float), ones, 1, drs, 1);
```

are needed for similar reasons. We are saying that in our source vector **ones**, for example, the elements of interest are spaced 1 elements apart, i.e. they are contiguous. But if we wanted our vector to be some row in a matrix with, say, 500 rows, the elements of interesting would be spaced 500 elements apart, again keeping in mind that column-major order is assumed.

The actual matrix multiplication is done here:

```
cublasSgemv('n', n, n, 1.0, dm, n, drs, 1, 0.0, drs, 1);
```

The “mv” in “cublasSgemv” stands for “matrix times vector.” Here the call says: no (‘n’), we do not want the matrix to be transposed; the matrix has **n** rows and **n** columns; we wish the matrix to be multiplied by 1.0 (if 0, the multiplication is not actually performed, which we could have here); the matrix is at **dm**; the number of dimensioned rows of the matrix is **n**; the vector is at **drs**; the elements of the vector are spaced 1 word apart; we wish the vector to not be multiplied by a scalar (see note above); the resulting vector will be stored at **drs**, 1 word apart.

Further information is available in the CUBLAS manual.

5.9 Error Checking

Every CUDA call (except for kernel invocations) returns an error code of type **cudaError_t**. One can view the nature of the error by calling **cudaGetErrorString()** and printing its output.

For kernel invocations, one can call **cudaGetLastError()**, which does what its name implies. A call would typically have the form

```
cudaError_t err = cudaGetLastError();
if(err != cudaSuccess) printf("%s\n", cudaGetErrorString(err));
```

You may also wish to **cudaSafeCall()**, which is used by wrapping your regular CUDA call. It automatically prints out error messages as above.

Each CUBLAS call returns a potential error code, of type **cublasStatus**, not checked here.

5.10 The New Generation

The latest GPU architecture from NVIDIA is called Fermi. Many of the advances are of the “bigger and faster than before” type. These are important, but be sure to note the significant architectural changes, including:

- Host memory, device global memory and device shared memory share a unified address space.
- On-chip memory can be apportioned to both shared memory and cache memory. Since shared memory is in essence a programmer-managed cache, this gives the programmer access to a real cache.

5.11 Further Examples

There are additional CUDA examples in later sections of this book. These include:²

- Prof. Richard Edgar’s matrix-multiply code, optimized for use of shared memory, Section 8.4.2.2
- odd/even transposition sort, Section 10.3.3, showing a typical CUDA pattern for iterative algorithms

²If you are reading this presentation on CUDA separately from the book, the book is at <http://heather.cs.ucdavis.edu/~matloff/158/PLN/ParProcBook.pdf>

Chapter 6

Message Passing Systems

Message passing systems are probably the most common platforms for parallel processing today.

6.1 Overview

Traditionally, shared-memory hardware has been extremely expensive, with a typical system costing hundreds of thousands of dollars. Accordingly, the main users were for very large corporations or government agencies, with the machines being used for heavy-duty server applications, such as for large databases and World Wide Web sites. The conventional wisdom is that these applications require the efficiency that good shared-memory hardware can provide.

But the huge expense of shared-memory machines led to a quest for high-performance message-passing alternatives, first in hypercubes and then in networks of workstations (NOWs).

The situation changed radically around 2005, when “shared-memory hardware for the masses” became available in dual-core commodity PCs. Chips of higher core multiplicity are commercially available, with a decline of price being inevitable. Ordinary users will soon be able to afford shared-memory machines featuring dozens of processors.

Yet the message-passing paradigm continues to thrive. Many people believe it is more amenable to writing really fast code, and the advent of **cloud computing** has given message-passing a big boost. In addition, many of the world’s very fastest systems (see www.top500.org for the latest list) are in fact of the message-passing type.

In this chapter, we take a closer look at this approach to parallel processing.

6.2 A Historical Example: Hypercubes


A popular class of parallel machines in the 1980s and early 90s was that of **hypercubes**. Intel sold them, for example, as did a subsidiary of Oracle, nCube. A hypercube would consist of some number of ordinary Intel processors, with each processor having some memory and serial I/O hardware for connection to its “neighbor” processors.

Hypercubes proved to be too expensive for the type of performance they could achieve, and the market was small anyway. Thus they are not common today, but they are still important, both for historical reasons (in the computer field, old techniques are often recycled decades later), and because the algorithms developed for them have become quite popular for use on general machines. In this section we will discuss architecture, algorithms and software for such machines.

6.2.0.1 Definitions A **hypercube** of dimension d consists of $D = 2^d$ **processing elements** (PEs), i.e. processor-memory pairs, with fast serial I/O connections between neighboring PEs. We refer to such a cube as a **d-cube**.

The PEs in a d -cube will have numbers 0 through $D-1$. Let (c_{d-1}, \dots, c_0) be the base-2 representation of a PE’s number. The PE has fast point-to-point links to d other PEs, which we will call its **neighbors**. Its i th neighbor has number $(c_{d-1}, \dots, 1 - c_{i-1}, \dots, c_0)$.¹

For example, consider a hypercube having $D = 16$, i.e. $d = 4$. The PE numbered 1011, for instance, would have four neighbors, 0011, 1111, 1001 and 1010.




It is sometimes helpful to build up a cube from the lower-dimensional cases. To build a $(d+1)$ -dimensional cube from two d -dimensional cubes, just follow this recipe:

¹Note that we number the digits from right to left, with the rightmost digit being digit 0.

- (a) Take a d -dimensional cube and duplicate it. Call these two cubes subcube 0 and subcube 1.
- (b) For each pair of same-numbered PEs in the two subcubes, add a binary digit 0 to the front of the number for the PE in subcube 0, and add a 1 in the case of subcube 1. Add a link between them.

The following figure shows how a 4-cube can be constructed in this way from two 3-cubes:



Given a PE of number (c_{d-1}, \dots, c_0) in a d -cube, we will discuss the i -cube to which this PE belongs, meaning all PEs whose first $d-i$ digits match this PE's.² Of all these PEs, the one whose last i digits are all 0s is called the **root** of this i -cube.

For the 4-cube and PE 1011 mentioned above, for instance, the 2-cube to which that PE belongs consists of 1000, 1001, 1010 and 1011—i.e. all PEs whose first two digits are 10—and the root is 1000.

Given a PE, we can split the i -cube to which it belongs into two $(i-1)$ -subcubes, one consisting of those PEs whose digit $i-1$ is 0 (to be called subcube 0), and the other consisting of those PEs whose digit $i-1$ is 1 (to be called subcube 1). Each given PE in subcube 0 has as its **partner** the PE in subcube 1 whose digits match those of the given PE, except for digit $i-1$.

To illustrate this, again consider the 4-cube and the PE 1011. As an example, let us look at how the 3-cube it belongs to will split into two 2-cubes. The 3-cube to which 1011 belongs consists of 1000, 1001, 1010, 1011, 1100, 1101, 1110 and 1111. This 3-cube can be split into two 2-cubes, one being 1000, 1001, 1010

²Note that this is indeed an i -dimensional cube, because the last i digits are free to vary.

and 1011, and the other being 1100, 1101, 1110 and 1111. Then PE 1000 is partners with PE 1100, PE 1001 is partners with PE 1101, and so on.

Each link between two PEs is a dedicated connection, much preferable to the shared link we have when we run, say, MPI, on a collection of workstations on an Ethernet. On the other hand, if one PE needs to communicate with a non-neighbor PE, multiple links (as many as d of them) will need to be traversed. Thus the nature of the communications costs here is much different than for a network of workstations, and this must be borne in mind when developing programs.

6.3 Networks of Workstations (NOWs)

The idea here is simple: Take a bunch of commodity PCs and network them for use as parallel processing systems. They are of course individual machines, capable of the usual uniprocessor, nonparallel applications, but by networking them together and using message-passing software environments such as MPI, we can form very powerful parallel systems.

The networking does result in a significant loss of performance, but the price/performance ratio in NOW can be much superior in many applications to that of shared-memory or hypercube hardware of comparable number of CPUs.

6.3.1 The Network Is Literally the Weakest Link

Still, one factor which can be key to the success of a NOW is to use a fast network, both in terms of hardware and network protocol. Ordinary Ethernet and TCP/IP are fine for the applications envisioned by the original designers of the Internet, e.g. e-mail and file transfer, but they are slow in the NOW context.

A popular network for a NOW today is Infiniband (IB) (www.infinibandta.org). It features low latency, about 1.0-3.0 microseconds, high bandwidth, about 1.0-2.0 gigaBytes per second), and uses a low amount of the CPU's cycles, around 5-10%.

The basic building block of IB is a switch, with many inputs and outputs, similar in concept to Ω -net. You can build arbitrarily large and complex topologies from these switches.

A central point is that IB, as with other high-performance networks designed for NOWs, uses RDMA (Remote Direct Memory Access) read/write, which eliminates the extra copying of data between the application program's address space to that of the operating system.

IB has high performance and scalable³ implementations of distributed locks, semaphores, collective communication operations. An atomic operation takes about 3-5 microseconds.

³The term *scalable* arises frequently in conversations on parallel processing. It means that this particular method of dealing with some aspect of parallel processing continues to work well as the system size increases. We say that the method *scales*.

IB implements true **multicast**, i.e. the simultaneous sending of messages to many nodes. Note carefully that even though MPI has its `MPI_Bcast()` function, it will send things out one at a time unless your network hardware is capable of multicast, and the MPI implementation you use is configured specifically for that hardware.

For information on network protocols, e.g. for example www.rdmaconsortium.org. A research paper evaluating a tuned implementation of MPI on IB is available at nowlab.cse.ohio-state.edu/publications/journal-papers/2004/liuj-ijpp04.pdf.

6.3.2 Other Issues

Increasingly today, the workstations themselves are multiprocessor machines, so a NOW really is a hybrid arrangement. They can be programmed either purely in a message-passing manner—e.g. running eight MPI processes on four dual-core machines—or in a mixed way, with a shared-memory approach being used within a workstation but message-passing used between them.

NOWs have become so popular that there are now “recipes” on how to build them for the specific purpose of parallel processing. The term **Beowulf** come to mean a NOW, usually with a fast network connecting them, used for parallel processing. The term *NOW* itself is no longer in use, replaced by *cluster*. Software packages such as ROCKS (<http://www.rockclusters.org/wordpress/>) have been developed to make it easy to set up and administer such systems.

6.4 Systems Using Nonexplicit Message-Passing

Writing message-passing code is a lot of work, as the programmer must explicitly arrange for transfer of data. Contrast that, for instance, to shared-memory machines, in which cache coherency transactions will cause data transfers, but which are not arranged by the programmer and not even seen by him/her.

In order to make coding on message-passing machines easier, higher-level systems have been devised. These basically operate in the **scatter/gather** paradigm, in which a “manager” node sends out chunks of work to the other nodes, serving as “workers,” and then collects and assembles the results sent back the workers.

One example of this is R’s **snow** package, which will be discussed in Section 12.5.2. But the most common approach today—and the one attracting the most attention—is MapReduce, to be discussed below.

6.4.1 MapReduce

MapReduce was developed as part of a recently-popularized computational approach known as **cloud computing**. The idea is that a large corporation that has many computers could sell time on them, thus mak-

ing profitable use of excess capacity. The typical customer would have occasional need for large-scale computing—and often large-scale data storage. The customer would submit a program to the cloud computing vendor, who would run it in parallel on the vendor’s many machines (unseen, thus forming the “cloud”), then return the output to the customer.

Google, Yahoo! and Amazon, among others, have recently gotten into the cloud computing business. The open-source application of choice for this is Hadoop, an implementation of MapReduce.

The key issue, of course, is the parallelizability of the inherently serial code. But all the user need do is provide code to break the data into chunks, code to work on a chunk, and code to collect the outputs from the chunks back into the overall output of the program.

For this to work, the program’s data usage pattern must have a simple, regular structure, as in these examples:

Example 1: Suppose we wish to list all the words used in a file, together with the counts of the numbers of instances of the words. If we have 100000 lines in the file and 10 processors, we could divide the file into chunks of 10000 lines each, have each processor run code to do the word counts in its chunk, and then combine the results.

Example 2: Suppose we wish to multiply an $n \times 1$ vector X by an $n \times n$ matrix A . Say $n = 100000$, and again we have 10 processors. We could divide A into chunks of 10000 rows each, have each processor multiply X by its chunk, and then combine the outputs.

To illustrate this, here is a pseudocode summary of a word-count program written in Python by Michael Noll; see http://www.michael-noll.com/wiki/Writing_An_Hadoop_MapReduce_Program_In_Python. Actually Hadoop is really written for Java applications. However, Hadoop can work with programs in any language under Hadoop’s Streaming option, by reading from STDIN and writing to STDOUT. This does cause some slowdown in numeric programs, for the conversion of strings to numbers and vice versa.⁴

mapper.py:

```

1 for each line in STDIN
2     break line into words, placed in wordarray
3     for each word in wordarray
4         print word, '1' to STDOUT # we have found 1 instance of the word

```

reducer.py:

```

1 # dictionary will consist of (word,count) pairs
2 dictionary = empty
3 for each line in STDIN
4     split line into word, thiscount

```

⁴In the case of Python, we could also run Jython, a Python interpreter that produces Java byte code. Hadoop also offers communication via Unix pipes.

```

5     if word not in dictionary:
6         add (word,thiscount) to dictionary
7     else
8         change (word,count) entry to (word,count+thiscount)
9 print dictionary to STDOUT

```

Note that these two user programs have nothing in them at all regarding parallelism. Instead, the process works as follows:

- the user provides Hadoop the original data file, by copying the file to Hadoop’s own file system, the Hadoop Distributed File System (HDFS)
- the user provides Hadoop with the mapper and reducer programs; Hadoop runs several instances of each
- Hadoop forms chunks by forming groups of lines in the file
- Hadoop has each instance of the mapper program work on a chunk:

```

mapper.py < chunk > outputchunk
# output is replicated and sent to the various instances of reducer

```

- Hadoop runs

```

reducer.py < outputchunk > myfinalchunk
# in this way final output is distributed to the nodes in HDFS

```

In the matrix-multiply model, the mapper program would produce chunks of X , together with the corresponding row numbers. Then the reducer program would sort the rows by row number, and place the result in X .

Note too that by having the file in HDFS, we minimize communications costs in shipping the data. “Moving computation is cheaper than moving data.”

Hadoop also incorporates rather sophisticated fault tolerance mechanisms. If a node goes down, the show goes on.

Note again that this works well only on problems of a certain structure. Also, some say that the idea has been overpromoted; see for instance “MapReduce: A Major Step Backwards,” *The Database Column*, by Professor David DeWitt, <http://www.databasecolumn.com/2008/01/mapreduce-a-major-step-back.html>

Chapter 7

Introduction to MPI

MPI is the *de facto* standard for message-passing software.

7.1 Overview

7.1.1 History

Though (small) shared-memory machines have come down radically in price, to the point at which a dual-core PC is affordable in the home, historically shared-memory machines were available only to the “very rich”—large banks, national research labs and so on.

The first “affordable” message-machine type was the Hypercube, developed by a physics professor at Cal Tech. It consisted of a number of **processing elements** (PEs) connected by fast serial I/O cards. This was in the range of university departmental research labs. It was later commercialized by Intel and NCube.

Later, the notion of **networks of workstations** (NOWs) became popular. Here the PEs were entirely independent PCs, connected via a standard network. This was refined a bit, by the use of more suitable network hardware and protocols, with the new term being **clusters**.

All of this necessitated the development of standardized software tools based on a message-passing paradigm. The first popular such tool was Parallel Virtual Machine (PVM). It still has its adherents today, but has largely been supplanted by the Message Passing Interface (MPI).

MPI itself later became MPI 2. Our document here is intended mainly for the original.

7.1.2 Structure and Execution

MPI is merely a set of Application Programmer Interfaces (APIs), called from user programs written in C, C++ and other languages. It has many implementations, with some being open source and generic, while others are proprietary and fine-tuned for specific commercial hardware.

Suppose we have written an MPI program *x*, and will run it on four machines in a cluster. Each machine will be running its own copy of *x*. Official MPI terminology refers to this as four **processes**. Now that multicore machines are commonplace, one might indeed run two or more cooperating MPI processes—where now we use the term *processes* in the real OS sense—on the same multicore machine. In this document, we will tend to refer to the various MPI processes as **nodes**, with an eye to the cluster setting.

Though the nodes are all running the same program, they will likely be working on different parts of the program's data. This is called the Single Program Multiple Data (SPMD) model. This is the typical approach, but there could be different programs running on different nodes. Most of the APIs involve a node sending information to, or receiving information from, other nodes.

7.1.3 Implementations

Two of the most popular implementations of MPI are MPICH and LAM. MPICH offers more tailoring to various networks and other platforms, while LAM runs on networks. Introductions to MPICH and LAM can be found, for example, at <http://heather.cs.ucdavis.edu/~matloff/MPI/NotesMPICH.NM.html> and <http://heather.cs.ucdavis.edu/~matloff/MPI/NotesLAM.NM.html>, respectively.

7.1.4 Performance Issues

Mere usage of a parallel language on a parallel platform does not guarantee a performance improvement over a serial version of your program. The central issue here is the overhead involved in internode communication.

As of 2008, Infiniband, one of the fastest cluster networks commercially available, has a **latency** of about 1.0-3.0 microseconds, meaning that it takes the first bit of a packet that long to get from one node on an Infiniband switch to another. Comparing that to the nanosecond time scale of CPU speeds, one can see that the communications overhead can destroy a program's performance. And Ethernet is quite a bit slower than Infiniband.

Note carefully that latency is a major problem even if the **bandwidth**—the number of bits per second which are sent—is high. For this reason, it is quite possible that your parallel program may actually run more slowly than its serial version.

Of course, if your platform is a shared-memory multiprocessor (especially a multicore one, where communication between cores is particularly fast) and you are running all your MPI processor on that machine, the problem is less severe.

7.2 Running Example

7.2.1 The Algorithm

The code implements the Dijkstra algorithm for finding the shortest paths in an undirected graph. Pseudocode for the algorithm is

```

1 Done = {0}
2 NonDone = {1,2,...,N-1}
3 for J = 1 to N-1 Dist[J] = infinity`
4 Dist[0] = 0
5 for Step = 1 to N-1
6     find J such that Dist[J] is min among all J in NonDone
7     transfer J from NonDone to Done
8     NewDone = J
9     for K = 1 to N-1
10        if K is in NonDone
11            Dist[K] = min(Dist[K],Dist[NewDone]+G[NewDone,K])

```

At each iteration, the algorithm finds the closest vertex J to 0 among all those not yet processed, and then updates the list of minimum distances to each vertex from 0 by considering paths that go through J. Two obvious potential candidate part of the algorithm for parallelization are the “find J” and “for K” lines, and the above OpenMP code takes this approach.

7.2.2 The Code

```

1 // Dijkstra.c
2
3 // MPI example program: Dijkstra shortest-path finder in a
4 // bidirectional graph; finds the shortest path from vertex 0 to all
5 // others
6
7 // command line arguments: nv print dbg
8
9 // where: nv is the size of the graph; print is 1 if graph and min
10 // distances are to be printed out, 0 otherwise; and dbg is 1 or 0, 1
11 // for debug
12
13 // node 0 will both participate in the computation and serve as a
14 // "manager"
15
16 #include <stdio.h>

```

```

17 #include <mpi.h>
18
19 #define MYMIN_MSG 0
20 #define OVRLMIN_MSG 1
21 #define COLLECT_MSG 2
22
23 // global variables (but of course not shared across nodes)
24
25 int nv, // number of vertices
26     *notdone, // vertices not checked yet
27     nnodes, // number of MPI nodes in the computation
28     chunk, // number of vertices handled by each node
29     startv, endv, // start, end vertices for this node
30     me, // my node number
31     dbg;
32 unsigned largeint, // max possible unsigned int
33     mymin[2], // mymin[0] is min for my chunk,
34              // mymin[1] is vertex which achieves that min
35     othermin[2], // othermin[0] is min over the other chunks
36                // (used by node 0 only)
37                // othermin[1] is vertex which achieves that min
38     overallmin[2], // overallmin[0] is current min over all nodes,
39                  // overallmin[1] is vertex which achieves that min
40     *ohd, // 1-hop distances between vertices; "ohd[i][j]" is
41           // ohd[i*nv+j]
42     *mind; // min distances found so far
43
44 double T1, T2; // start and finish times
45
46 void init(int ac, char **av)
47 { int i, j, tmp; unsigned u;
48   nv = atoi(av[1]);
49   dbg = atoi(av[3]);
50   MPI_Init(&ac, &av);
51   MPI_Comm_size(MPI_COMM_WORLD, &nnodes);
52   MPI_Comm_rank(MPI_COMM_WORLD, &me);
53   chunk = nv/nnodes;
54   startv = me * chunk;
55   endv = startv + chunk - 1;
56   u = -1;
57   largeint = u >> 1;
58   ohd = malloc(nv*nv*sizeof(int));
59   mind = malloc(nv*sizeof(int));
60   notdone = malloc(nv*sizeof(int));
61   // random graph
62   // note that this will be generated at all nodes; could generate just
63   // at node 0 and then send to others, but faster this way
64   srand(9999);
65   for (i = 0; i < nv; i++)
66     for (j = i; j < nv; j++) {
67       if (j == i) ohd[i*nv+i] = 0;
68       else {
69         ohd[nv*i+j] = rand() % 20;
70         ohd[nv*j+i] = ohd[nv*i+j];
71       }
72     }
73   for (i = 0; i < nv; i++) {
74     notdone[i] = 1;

```



```

75     mind[i] = largeint;
76     }
77     mind[0] = 0;
78     while (dbg) ; // stalling so can attach debugger
79 }
80
81 // finds closest to 0 among notdone, among startv through endv
82 void findmymin()
83 { int i;
84   mymin[0] = largeint;
85   for (i = startv; i <= endv; i++)
86     if (notdone[i] && mind[i] < mymin[0]) {
87       mymin[0] = mind[i];
88       mymin[1] = i;
89     }
90 }
91
92 void findoverallmin()
93 { int i;
94   MPI_Status status; // describes result of MPI_Recv() call
95   // nodes other than 0 report their mins to node 0, which receives
96   // them and updates its value for the global min
97   if (me > 0)
98     MPI_Send(mymin,2,MPI_INT,0,MYMIN_MSG,MPI_COMM_WORLD);
99   else {
100    // check my own first
101    overallmin[0] = mymin[0];
102    overallmin[1] = mymin[1];
103    // check the others
104    for (i = 1; i < nnodes; i++) {
105      MPI_Recv(overallmin,2,MPI_INT,i,MYMIN_MSG,MPI_COMM_WORLD,&status);
106      if (overallmin[0] < overallmin[0]) {
107        overallmin[0] = overallmin[0];
108        overallmin[1] = overallmin[1];
109      }
110    }
111  }
112 }
113
114 void updatemymin() // update my mind segment
115 { // for each i in [startv,endv], ask whether a shorter path to i
116   // exists, through mv
117   int i, mv = overallmin[1];
118   unsigned md = overallmin[0];
119   for (i = startv; i <= endv; i++)
120     if (md + ohd[mv*nv+i] < mind[i])
121       mind[i] = md + ohd[mv*nv+i];
122 }
123
124 void disseminateoverallmin()
125 { int i;
126   MPI_Status status;
127   if (me == 0)
128     for (i = 1; i < nnodes; i++)
129       MPI_Send(overallmin,2,MPI_INT,i,OVRMLIN_MSG,MPI_COMM_WORLD);
130   else
131     MPI_Recv(overallmin,2,MPI_INT,0,OVRMLIN_MSG,MPI_COMM_WORLD,&status);
132 }

```

```

133
134 void updateallmind() // collects all the mind segments at node 0
135 { int i;
136   MPI_Status status;
137   if (me > 0)
138     MPI_Send(mind+startv, chunk, MPI_INT, 0, COLLECT_MSG, MPI_COMM_WORLD);
139   else
140     for (i = 1; i < nnodes; i++)
141       MPI_Recv(mind+i*chunk, chunk, MPI_INT, i, COLLECT_MSG, MPI_COMM_WORLD,
142              &status);
143 }
144
145 void printmind() // partly for debugging (call from GDB)
146 { int i;
147   printf("minimum distances:\n");
148   for (i = 1; i < nv; i++)
149     printf("%u\n", mind[i]);
150 }
151
152 void dowork()
153 { int step, // index for loop of nv steps
154   i;
155   if (me == 0) T1 = MPI_Wtime();
156   for (step = 0; step < nv; step++) {
157     findmymin();
158     findoverallmin();
159     disseminateoverallmin();
160     // mark new vertex as done
161     notdone[overallmin[1]] = 0;
162     updatemymin(startv, endv);
163   }
164   updateallmind();
165   T2 = MPI_Wtime();
166 }
167
168 int main(int ac, char **av)
169 { int i, j, print;
170   init(ac, av);
171   dowork();
172   print = atoi(av[2]);
173   if (print && me == 0) {
174     printf("graph weights:\n");
175     for (i = 0; i < nv; i++) {
176       for (j = 0; j < nv; j++)
177         printf("%u ", ohd[nv*i+j]);
178       printf("\n");
179     }
180     printmind();
181   }
182   if (me == 0) printf("time at node 0: %f\n", (float) (T2-T1));
183   MPI_Finalize();
184 }
185

```

The various MPI functions will be explained in the next section.

7.2.3 Introduction to MPI APIs

7.2.3.1 MPI_Init() and MPI_Finalize()

These are required for starting and ending execution of an MPI program. Their actions may be implementation-dependent. For instance, if our platform is an Ethernet-based cluster, **MPI_Init()** will probably set up the TCP/IP sockets via which the various nodes communicate with each other. On an Infiniband-based cluster, connections in the special Infiniband network protocol will be established. On a shared-memory multiprocessor, an implementation of MPI that is tailored to that platform would take very different actions.

7.2.3.2 MPI_Comm_size() and MPI_Comm_rank()

In our function **init()** above, note the calls

```
MPI_Comm_size(MPI_COMM_WORLD, &nnodes);  
MPI_Comm_rank(MPI_COMM_WORLD, &me);
```

The first call determines how many nodes are participating in our computation, placing the result in our variable **nnodes**. Here **MPI_COMM_WORLD** is our node group, termed a **communicator** in MPI parlance. MPI allows the programmer to subdivide the nodes into groups, to facilitate performance and clarity of code. Note that for some operations, such as barriers, the only way to apply the operation to a proper subset of all nodes is to form a group. The totality of all groups is denoted by **MPI_COMM_WORLD**. In our program here, we are not subdividing into groups.

The second call determines this node's ID number, called its **rank**, within its group. As mentioned earlier, even though the nodes are all running the same program, they are typically working on different parts of the program's data. So, the program needs to be able to sense which node it is running on, so as to access the appropriate data. Here we record that information in our variable **me**.

7.2.3.3 MPI_Send()

To see how MPI's basic send function works, consider our line above,

```
MPI_Send(mymin, 2, MPI_INT, 0, MYMIN_MSG, MPI_COMM_WORLD);
```

Let's look at the arguments:

mymin: We are sending a set of bytes. This argument states the address at which these bytes begin.

2, MPI_INT: This says that our set of bytes to be sent consists of 2 objects of type **MPI_INT**. That means 8 bytes on 32-bit machines, so why not just collapse these two arguments to one, namely the number 8? Why did the designers of MPI bother to define data types? The answer is that we want to be able to run MPI on a heterogeneous set of machines, with MPI serving as the “broker” between them in case different architectures among those machines handle data differently.

First of all, there is the issue of **endianness**. Intel machines, for instance, are **little-endian**, which means that the least significant byte of a memory word has the smallest address among bytes of the word. Sun SPARC chips, on the other hand, are **big-endian**, with the opposite storage scheme. If our set of nodes included machines of both types, straight transmission of sequences of 8 bytes might mean that some of the machines literally receive the data backwards!

Secondly, these days 64-bit machines are becoming more and more common. Again, if our set of nodes were to include both 32-bit and 64-bit words, some major problems would occur if no conversion were done.

0: We are sending to node 0.

MYMIN_MSG: This is the message type, programmer-defined in our line

```
#define MYMIN_MSG 0
```

Receive calls, described in the next section, can ask to receive only messages of a certain type.

COMM_WORLD: This is the node group to which the message is to be sent. Above, where we said we are sending to node 0, we technically should say we are sending to node 0 within the group **COMM_WORLD**.

7.2.3.4 MPI_Recv()

Let’s now look at the arguments for a basic receive:

```
MPI_Recv(othermin, 2, MPI_INT, i, MYMIN_MSG, COMM_WORLD, &status);
```

othermin: The received message is to be placed at our location **othermin**.

2, MPI_INT: Two objects of **MPI_INT** type are to be received.

i: Receive only messages of from node **i**. If we did not care what node we received a message from, we could specify the value **MPI_ANY_SOURCE**.

MYMIN_MSG: Receive only messages of type **MYMIN_MSG**. If we did not care what type of message we received, we would specify the value **MPI_ANY_TAG**.

COMM_WORLD: Group name.

status: Recall our line

```
MPI_Status status; // describes result of MPI_Recv() call
```

The type is an MPI **struct** containing information about the received message. Its primary fields of interest are **MPI_SOURCE**, which contains the identity of the sending node, and **MPI_TAG**, which contains the message type. These would be useful if the receive had been done with **MPI_ANY_SOURCE** or **MPI_ANY_TAG**; the status argument would then tell us which node sent the message and what type the message was.

7.3 Collective Communications

MPI features a number of **collective communication** capabilities, a number of which are used in the following refinement of our Dijkstra program:

7.3.1 Example

```
1 // Dijkstra.coll1.c
2
3 // MPI example program: Dijkstra shortest-path finder in a
4 // bidirectional graph; finds the shortest path from vertex 0 to all
5 // others; this version uses collective communication
6
7 // command line arguments: nv print dbg
8
9 // where: nv is the size of the graph; print is 1 if graph and min
10 // distances are to be printed out, 0 otherwise; and dbg is 1 or 0, 1
11 // for debug
12
13 // node 0 will both participate in the computation and serve as a
14 // "manager"
15
16 #include <stdio.h>
17 #include <mpi.h>
18
19 // global variables (but of course not shared across nodes)
20
21 int nv, // number of vertices
22     *notdone, // vertices not checked yet
23     nnodes, // number of MPI nodes in the computation
24     chunk, // number of vertices handled by each node
25     startv, endv, // start, end vertices for this node
26     me, // my node number
27     dbg;
28 unsigned largeint, // max possible unsigned int
29     mymin[2], // mymin[0] is min for my chunk,
30     // mymin[1] is vertex which achieves that min
31     overallmin[2], // overallmin[0] is current min over all nodes,
```

```

32             // overallmin[1] is vertex which achieves that min
33     *ohd, // 1-hop distances between vertices; "ohd[i][j]" is
34           // ohd[i*nv+j]
35     *mind; // min distances found so far
36
37 double T1,T2; // start and finish times
38
39 void init(int ac, char **av)
40 { int i,j,tmp; unsigned u;
41   nv = atoi(av[1]);
42   dbg = atoi(av[3]);
43   MPI_Init(&ac,&av);
44   MPI_Comm_size(MPI_COMM_WORLD,&nnodes);
45   MPI_Comm_rank(MPI_COMM_WORLD,&me);
46   chunk = nv/nnodes;
47   startv = me * chunk;
48   endv = startv + chunk - 1;
49   u = -1;
50   largeint = u >> 1;
51   ohd = malloc(nv*nv*sizeof(int));
52   mind = malloc(nv*sizeof(int));
53   notdone = malloc(nv*sizeof(int));
54   // random graph
55   // note that this will be generated at all nodes; could generate just
56   // at node 0 and then send to others, but faster this way
57   srand(9999);
58   for (i = 0; i < nv; i++)
59     for (j = i; j < nv; j++) {
60       if (j == i) ohd[i*nv+i] = 0;
61       else {
62         ohd[nv*i+j] = rand() % 20;
63         ohd[nv*j+i] = ohd[nv*i+j];
64       }
65     }
66   for (i = 0; i < nv; i++) {
67     notdone[i] = 1;
68     mind[i] = largeint;
69   }
70   mind[0] = 0;
71   while (dbg) ; // stalling so can attach debugger
72 }
73
74 // finds closest to 0 among notdone, among startv through endv
75 void findmymin()
76 { int i;
77   mymin[0] = largeint;
78   for (i = startv; i <= endv; i++)
79     if (notdone[i] && mind[i] < mymin[0]) {
80       mymin[0] = mind[i];
81       mymin[1] = i;
82     }
83 }
84
85 void updatemymin() // update my mind segment
86 { // for each i in [startv,endv], ask whether a shorter path to i
87   // exists, through mv
88   int i, mv = overallmin[1];
89   unsigned md = overallmin[0];

```

```

90     for (i = startv; i <= endv; i++)
91         if (md + ohd[mv*nv+i] < mind[i])
92             mind[i] = md + ohd[mv*nv+i];
93     }
94
95 void printmind() // partly for debugging (call from GDB)
96 { int i;
97   printf("minimum distances:\n");
98   for (i = 1; i < nv; i++)
99       printf("%u\n",mind[i]);
100 }
101
102 void dowork()
103 { int step, // index for loop of nv steps
104   i;
105   if (me == 0) T1 = MPI_Wtime();
106   for (step = 0; step < nv; step++) {
107       findmymin();
108       MPI_Reduce(mymin,overallmin,1,MPI_2INT,MPI_MINLOC,0,MPI_COMM_WORLD);
109       MPI_Bcast(overallmin,1,MPI_2INT,0,MPI_COMM_WORLD);
110       // mark new vertex as done
111       notdone[overallmin[1]] = 0;
112       updatemymin(startv,endv);
113   }
114   // now need to collect all the mind values from other nodes to node 0
115   MPI_Gather(mind+startv,chunk,MPI_INT,mind,chunk,MPI_INT,0,MPI_COMM_WORLD);
116   T2 = MPI_Wtime();
117 }
118
119 int main(int ac, char **av)
120 { int i,j,print;
121   init(ac,av);
122   dowork();
123   print = atoi(av[2]);
124   if (print && me == 0) {
125       printf("graph weights:\n");
126       for (i = 0; i < nv; i++) {
127           for (j = 0; j < nv; j++)
128               printf("%u ",ohd[nv*i+j]);
129           printf("\n");
130       }
131       printmind();
132   }
133   if (me == 0) printf("time at node 0: %f\n", (float) (T2-T1));
134   MPI_Finalize();
135 }

```

The new calls will be explained in the next section.

7.3.2 MPI_Bcast()

In our original Dijkstra example, we had a loop

```
for (i = 1; i < nnodes; i++)
    MPI_Send(overallmin, 2, MPI_INT, i, OVRLMIN_MSG, MPI_COMM_WORLD);
```

in which node 0 sends to all other nodes. We can replace this by

```
MPI_Bcast(overallmin, 2, MPI_INT, 0, MPI_COMM_WORLD);
```

In English, this call would say,

At this point all nodes participate in a broadcast operation, in which node 0 sends 2 objects of type **MPI_INT**. The source of the data will be located at address **overallmin** at node 0, and the other nodes will receive the data at a location of that name.

Note my word “participate” above. The name of the function is “broadcast,” which makes it sound like only node 0 executes this line of code, which is not the case; all the nodes in the group (in this case that means all nodes in our entire computation) execute this line. The only difference is the action; most nodes participate by receiving, while node 0 participates by sending.

Why might this be preferable to using an explicit loop? First, it would obviously be much clearer. That makes the program easier to write, easier to debug, and easier for others (and ourselves, later) to read.

But even more importantly, using the broadcast may improve performance. We may, for instance, be using an implementation of MPI which is tailored to the platform on which we are running MPI. If for instance we are running on a network designed for parallel computing, such as Myrinet or Infiniband, an optimized broadcast may achieve a much higher performance level than would simply a loop with individual send calls. On a shared-memory multiprocessor system, special machine instructions specific to that platform’s architecture can be exploited, as for instance IBM has done for its shared-memory machines. Even on an ordinary Ethernet, one could exploit Ethernet’s own broadcast mechanism, as had been done for PVM, a system like MPI (G. Davies and N. Matloff, *Network-Specific Performance Enhancements for PVM, Proceedings of the Fourth IEEE International Symposium on High-Performance Distributed Computing*, 1995, 205-210).

7.3.2.1 MPI_Reduce()/MPI_Allreduce()

Look at our call

```
MPI_Reduce(mymin, overallmin, 1, MPI_2INT, MPI_MINLOC, 0, MPI_COMM_WORLD);
```

above. In English, this would say,

At this point all nodes in this group participate in a “reduce” operation. The type of reduce operation is **MPI_MINLOC**, which means that the minimum value among the nodes will be computed, and the index attaining that minimum will be recorded as well. Each node contributes a value to be checked, and an associated index, from a location **mymin** in their programs; the type of the pair is **MPI_2INT**. The overall min value/index will be computed by combining all of these values at node 0, where they will be placed at a location **overallmin**.

MPI also includes a function **MPI_Allreduce()**, which does the same operation, except that instead of just depositing the result at one node, it does so at all nodes. So for instance our code above,

```
MPI_Reduce(mymin,overallmin,1,MPI_2INT,MPI_MINLOC,0,MPI_COMM_WORLD);
MPI_Bcast(overallmin,1,MPI_2INT,0,MPI_COMM_WORLD);
```

could be replaced by

```
MPI_Allreduce(mymin,overallmin,1,MPI_2INT,MPI_MINLOC,MPI_COMM_WORLD);
```

Again, these can be optimized for particular platforms.

7.3.2.2 MPI_Gather()/MPI_Allgather()

A classical approach to parallel computation is to first break the data for the application into chunks, then have each node work on its chunk, and then gather all the processed chunks together at some node. The MPI function **MPI_Gather()** does this.

In our program above, look at the line

```
MPI_Gather(mind+startv,chunk,MPI_INT,mind,chunk,MPI_INT,0,MPI_COMM_WORLD);
```

In English, this says,

At this point all nodes participate in a gather operation, in which each node contributes data, consisting of **chunk** number of MPI integers, from a location **mind+startv** in its program. All that data is strung together and deposited at the location **mind** in the program running at node 0.

There is also **MPI_Allgather()**, which places the result at all nodes, not just one.

7.3.2.3 The MPI_Scatter()

This is the opposite of `MPI_Gather()`, i.e. it breaks long data into chunks which it parcels out to individual nodes.

Here is MPI code to count the number of edges in a directed graph. (A link from i to j does not necessarily imply one from j to i .) In the context here, **me** is the node's rank; **nv** is the number of vertices; **oh** is the one-hop distance matrix; and **nnodes** is the number of MPI processes. At the beginning only the process of rank 0 has a copy of **oh**, but it sends that matrix out in chunks to the other nodes, each of which stores its chunk in an array **ohchunk**.

```

1 MPI_Scatter(oh, nv*nv, MPI_INT, ohchunk, nv/nnodes, MPI_INT, 0,
2 MPI_COMM_WORLD);
3 mycount = 0;
4 for (i = 0; i < nv*nv/nnodes)
5     if (ohchunk[i] != 0) mycount++;
6 MPI_Reduce(&mycount, &numedge, 1, MPI_INT, MPI_SUM, 0, MPI_COMM_WORLD);
7 if (me == 0) printf("there are %d edges\n", numedge);

```

7.3.2.4 The MPI_Barrier()

This implements a barrier for a given communicator. The name of the communicator is the sole argument for the function.

Explicit barriers are less common in message-passing programs than in the shared-memory world.

7.3.3 Creating Communicators

Again, a communicator is a subset (either proper or improper) of all of our nodes. MPI includes a number of functions for use in creating communicators. Some set up a virtual “topology” among the nodes.

For instance, many physics problems consist of solving differential equations in two- or three-dimensional space, via approximation on a grid of points. In two dimensions, groups may consist of rows in the grid.

Here's how we might divide an MPI run into two groups (assumes an even number of MPI processes to begin with):

```

MPI_Comm_size(MPI_COMM_WORLD, &nnodes);
MPI_Comm_rank(MPI_COMM_WORLD, &me);
...
// declare variables to bind to groups
MPI_Group worldgroup, subgroup;
// declare variable to bind to a communicator
MPI_Comm subcomm;
...

```

```

int i,startrank,nn2 = nnodes/2;
int *subranks = malloc(nn2*sizeof(int));
if (me < nn2) start = 0;
else start = nn2;
for (i = 0; i < nn2; i++)
    subranks[i] = i + start;
// bind the world to a group variable
MPI_Comm_group(MPI_COMM_WORLD, &worldgroup);
// take worldgroup the nn2 ranks in "subranks" and form group
// "subgroup" from them
MPI_Group_incl(worldgroup, nn2, subranks, subgroup);
// create a communicator for that new group
MPI_Comm_create(MPI_COMM_WORLD, subgroup, subcomm);
// get my rank in this new group
MPI_Group_rank (subgroup, &subme);

```

You would then use **subcomm** instead of `MPI_COMM_WORLD` whenever you wish to, say, broadcast, only to that group.

7.4 Buffering, Synchrony and Related Issues

As noted several times so far, interprocess communication in parallel systems can be quite expensive in terms of time delay. In this section we will consider some issues which can be extremely important in this regard.

7.4.1 Buffering, Etc.

To understand this point, first consider situations in which MPI is running on some network, under the TCP/IP protocol. Say an MPI program at node A is sending to one at node B.

It is extremely import to keep in mind the levels of abstraction here. The OS's TCP/IP stack is running at the Session, Transport and Network layers of the network. MPI—meaning the MPI internals—is running above the TCP/IP stack, in the Application layers at A and B. And the MPI user-written application could be considered to be running at a “Super-application” layer, since it calls the MPI internals. (From here on, we will refer to the MPI internals as simply “MPI.”)

MPI at node A will have set up a TCP/IP socket to B during the user program's call to **MPI_Init()**. The other end of the socket will be a corresponding one at B. This setting up of this socket pair as establishing a **connection** between A and B. When node A calls **MPI_Send()**, MPI will write to the socket, and the TCP/IP stack will transmit that data to the TCP/IP socket at B. The TCP/IP stack at B will then send whatever bytes come in to MPI at B.

Now, it is important to keep in mind that in TCP/IP the totality of bytes sent by A to B during lifetime of the connection is considered one long message. So for instance if the MPI program at A calls **MPI_Send()** five

times, the MPI internals will write to the socket five times, but the bytes from those five messages will not be perceived by the TCP/IP stack at B as five messages, but rather as just one long message (in fact, only part of one long message, since more may be yet to come).

MPI at B continually reads that “long message” and breaks it back into MPI messages, keeping them ready for calls to `MPI_Recv()` from the MPI application program at B. Note carefully that phrase, *keeping them ready*; it refers to the fact that the order in which the MPI application program requests those messages may be different from the order in which they arrive.

On the other hand, looking again at the TCP/IP level, even though all the bytes sent are considered one long message, it will physically be sent out in pieces. These pieces don’t correspond to the pieces written to the socket, i.e. the MPI messages. Rather, the breaking into pieces is done for the purpose of **flow control**, meaning that the TCP/IP stack at A will not send data to the one at B if the OS at B has no room for it. The **buffer** space the OS at B has set up for receiving data is limited. As A is sending to B, the TCP layer at B is telling its counterpart at A when A is allowed to send more data.

Think of what happens the MPI application at B calls `MPI_Recv()`, requesting to receive from A, with a certain tag T. Say the first argument is named `x`, i.e. the data to be received is to be deposited at `x`. If MPI sees that it already has a message of tag T, it will have its `MPI_Recv()` function return the message to the caller, i.e. to the MPI application at B. **If no such message has arrived yet, MPI won’t return to the caller yet, and thus the caller blocks.**

`MPI_Send()` can block too. If the platform and MPI implementation is that of the TCP/IP network context described above, then the send call will return when its call to the OS’ `write()` (or equivalent, depending on OS) returns, but that could be delayed if the OS’ buffer space is full. On the other hand, another implementation could require a positive response from B before allowing the send call to return.

Note that buffering slows everything down. In our TCP scenario above, `MPI_Recv()` at B must copy messages from the OS’ buffer space to the MPI application program’s program variables, e.g. `x` above. This is definitely a blow to performance. That in fact is why networks developed specially for parallel processing typically include mechanisms to avoid the copying. Infiniband, for example, has a Remote Direct Memory Access capability, meaning that A can write directly to `x` at B. Of course, if our implementation uses **synchronous** communication, with A’s send call not returning until A gets a response from B, we must wait even longer.

Technically, the MPI standard states that `MPI_Send(x,...)` will return only when it is safe for the application program to write over the array which it is using to store its message, i.e. `x`. As we have seen, there are various ways to implement this, with performance implications. Similarly, `MPI_Recv(y,...)` will return only when it is safe to read `y`.

7.4.2 Safety

With **synchronous** communication, deadlock is a real risk. Say A wants to send two messages to B, of types U and V, but that B wants to receive V first. Then A won't even get to send V, because in preparing to send U it must wait for a notice from B that B wants to read U—a notice which will never come, because B sends such a notice for V first. This would not occur if the communication were asynchronous.

But beyond formal deadlock, programs can fail in other ways, even with buffering, as buffer space is always by nature finite. A program can fail if it runs out of buffer space, either at the sender or the receiver. See www.llnl.gov/computing/tutorials/mpi_performance/samples/unsafe.c for an example of a test program which demonstrates this on a certain platform, by deliberately overwhelming the buffers at the receiver.

In MPI terminology, asynchronous communication is considered **unsafe**. The program may run fine on most systems, as most systems are buffered, but fail on some systems. Of course, as long as you know your program won't be run in nonbuffered settings, it's fine, and since there is potentially such a performance penalty for doing things synchronously, most people are willing to go ahead with their “unsafe” code.

7.4.3 Living Dangerously

If one is sure that there will be no problems of buffer overflow and so on, one can use variant send and receive calls provided by MPI, such as **MPI_Isend()** and **MPI_Irecv()**. The key difference between them and **MPI_Send()** and **MPI_Recv()** is that they return immediately, and thus are termed **nonblocking**. Your code can go on and do other things, not having to wait.

This does mean that at A you cannot touch the data you are sending until you determine that it has either been buffered somewhere or has reached **x** at B. Similarly, at B you can't use the data at **x** until you determine that it has arrived. Such determinations can be made via **MPI_Wait()**. In other words, you can do your send or receive, then perform some other computations for a while, and then call **MPI_Wait()** to determine whether you can go on. Or you can call **MPI_Probe()** to ask whether the operation has completed yet.

7.4.4 Safe Exchange Operations

In many applications A and B are swapping data, so both are sending and both are receiving. This too can lead to deadlock. An obvious solution would be, for instance, to have the lower-rank node send first and the higher-rank node receive first.

But a more convenient, safer and possibly faster alternative would be to use MPI's **MPI_Sendrecv()** function. Its prototype is

```
int MPI_Sendrecv_replace(void* buf, int count, MPI_Datatype datatype,
```

```
int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
MPI_Status *status)
```

Note that the sent and received messages can be of different lengths and can use different tags.

7.5 Use of MPI from Other Languages

MPI is a vehicle for parallelizing C/C++, but some clever people have extended the concept to other languages, such as the cases of Python and R that we treat here.

7.5.1 Python: pyMPI

(Important note: As of April 2010, a much more widely used Python/MPI interface is MPI4Py. It works similarly to what is described here.)

A number of interfaces of Python to MPI have been developed.¹ A well-known example is pyMPI, developed by a PhD graduate in computer science in UCD, Patrick Miller.

One writes one’s pyMPI code, say in **x.py**, by calling pyMPI versions of the usual MPI routines. To run the code, one then runs MPI on the program **pyMPI** with **x.py** as a command-line argument.

Python is a very elegant language, and pyMPI does a nice job of elegantly interfacing to MPI. Following is a rendition of Quicksort in pyMPI. Don’t worry if you haven’t worked in Python before; the “non-C-like” Python constructs are explained in comments at the end of the code.

```
1 # a type of quicksort; break array x (actually a Python "list") into
2 # p quicksort-style piles, based # on comparison with the first p-1
3 # elements of x, where p is the number # of MPI nodes; the nodes sort
4 # their piles, then return them to node 0, # which strings them all
5 # together into the final sorted array
6
7 import mpi # load pyMPI module
8
9 # makes npls quicksort-style piles
10 def makepiles(x, npls):
11     pivot = x[:npls] # we'll use the first npls elements of x as pivots,
12                     # i.e. we'll compare all other elements of x to these
13     pivot.sort() # sort() is a member function of the Python list class
14     pls = [] # initialize piles list to empty
15     lp = len(pivot) # length of the pivot array
16     # pls will be a list of lists, with the i-th list in pls storing the
17     # i-th pile; the i-th pile will start with ID i (to enable
```

¹If you are not familiar with Python, I have a quick tutorial at <http://heather.cs.ucdavis.edu/~matloff/python.html>.

7.5.2 R

See the material on Rmpi, in Chapter 12.

Chapter 8

Introduction to Parallel Matrix Operations

8.1 It's Not Just Physics Anymore

In the early days parallel processing was mostly used in physics problems. Typical problems of interest would be grid computations such as the heat equation, matrix multiplication, matrix inversion (or equivalent operations) and so on. These matrices are not those little 3x3 things you worked with in your linear algebra class. In parallel processing applications of matrix algebra, our matrices can have thousands of rows and columns, or even larger.

The range of applications of parallel processing is of course far broader today. In many of these applications, problems which at first glance seem not to involve matrices, actually do have matrix solutions.

Example from graphy theory:

Let n denote the number of vertices in the graph. Define the graph's **adjacency matrix** A to be the $n \times n$ matrix whose element (i,j) is equal to 1 if there is an edge connecting vertices i and j (i.e. i and j are "adjacent"), and 0 otherwise. The corresponding **reachability matrix** R has its (i,j) element equal to 1 if there is some path from i to j , and 0 otherwise.

One can prove that

$$R = b[(I + A)^{n-1}], \quad (8.1)$$

where I is the identity matrix and the function $b()$ ('b' for "boolean") is applied elementwise to its matrix argument, replacing each nonzero element by 1 while leaving the elements which are 0 unchanged. The graph is connected if and only if all elements of R are 1s.

So, the original graph connectivity problem reduces to a matrix problem.

8.2 CUBLAS

Of course, remember that CUDA provides some excellent matrix-operation routines, in CUBLAS.

8.3 Partitioned Matrices

Parallel processing of course relies on finding a way to partition the work to be done. In the matrix algorithm case, this is often done by dividing a matrix into blocks (often called **tiles** these days).

For example, let

$$A = \begin{pmatrix} 1 & 5 & 12 \\ 0 & 3 & 6 \\ 4 & 8 & 2 \end{pmatrix} \quad (8.2)$$

and

$$B = \begin{pmatrix} 0 & 2 & 5 \\ 0 & 9 & 10 \\ 1 & 1 & 2 \end{pmatrix}, \quad (8.3)$$

so that

$$C = AB = \begin{pmatrix} 12 & 59 & 79 \\ 6 & 33 & 42 \\ 2 & 82 & 104 \end{pmatrix}. \quad (8.4)$$

We could partition A as

$$A = \begin{pmatrix} A_{00} & A_{01} \\ A_{10} & A_{11} \end{pmatrix}, \quad (8.5)$$

where

$$A_{00} = \begin{pmatrix} 1 & 5 \\ 0 & 3 \end{pmatrix}, \quad (8.6)$$

$$A_{01} = \begin{pmatrix} 12 \\ 6 \end{pmatrix}, \quad (8.7)$$

$$A_{10} = \begin{pmatrix} 4 & 8 \end{pmatrix} \quad (8.8)$$

and

$$A_{11} = \begin{pmatrix} 2 \end{pmatrix}. \quad (8.9)$$

Similarly we would partition B and C into blocks of the same size as in A,

$$B = \begin{pmatrix} B_{00} & B_{01} \\ B_{10} & B_{11} \end{pmatrix} \quad (8.10)$$

and

$$C = \begin{pmatrix} C_{00} & C_{01} \\ C_{10} & C_{11} \end{pmatrix}, \quad (8.11)$$

so that for example

$$B_{10} = \begin{pmatrix} 1 & 1 \end{pmatrix}. \quad (8.12)$$

The key point is that multiplication still works if we pretend that those submatrices are numbers. For example, pretending like that would give the relation

$$C_{00} = A_{00}B_{00} + A_{01}B_{10}, \quad (8.13)$$

which the reader should verify really is correct as matrices, i.e. the computation on the right side really does yield a matrix equal to C_{00} .

8.4 Matrix Multiplication

Since so many parallel matrix algorithms rely on matrix multiplication, a core issue is how to parallelize that operation.

Let's suppose for the sake of simplicity that each of the matrices to be multiplied is of dimensions $n \times n$. Let p denote the number of "processes," such as shared-memory threads or message-passing nodes.

We assume that the matrices are **dense**, meaning that most of their entries are nonzero. This is in contrast to **sparse** matrices, with many zeros. For instance, in **tridiagonal** matrices, in which the only nonzero elements are either on the diagonal or on subdiagonals just below or above the diagonal, and all other elements are guaranteed to be 0. Or we might just know that most elements are zeros but have no guarantee as to where they are; here we might have a system of pointers to get from one nonzero element to another. Clearly we would use different type of algorithms for sparse matrices than for dense ones.

8.4.1 Message-Passing Case

For concreteness here and in other sections below on message passing, assume we are using MPI.

The obvious plan of attack here is to break the matrices into blocks, and then assign different blocks to different MPI nodes. Assume that \sqrt{p} evenly divides n , and partition each matrix into submatrices of size $n/\sqrt{p} \times n/\sqrt{p}$. In other words, each matrix will be divided into m rows and m columns of blocks, where $m = n/\sqrt{p}$.

One of the conditions assumed here is that the matrices A and B are stored in a distributed manner across the nodes. This situation could arise for several reasons:

- The application is such that it is natural for each node to possess only part of A and B .
- One node, say node 0, originally contains all of A and B , but in order to conserve communication time, it sends each node only parts of those matrices.
- The entire matrix would not fit in the available memory at the individual nodes.

As you'll see, the algorithms then have the nodes passing blocks among themselves.

8.4.1.1 Fox's Algorithm

Consider the node that has the responsibility of calculating block (i,j) of the product C , which it calculates as

$$A_{i0}B_{0j} + A_{i1}B_{1j} + \dots + A_{ii}B_{ij} + \dots + A_{i,m-1}B_{m-1,j} \quad (8.14)$$

Rearrange this as

$$A_{ii}B_{ij} + A_{i,i+1}B_{i+1,j} + \dots + A_{i,m-1}B_{m-1,j} + A_{i0}B_{0j} + A_{i1}B_{1j} + \dots + A_{i,i-1}B_{i-1,j} \quad (8.15)$$

Written more compactly, this is

$$\sum_{k=0}^{m-1} A_{i,(i+k) \bmod m} B_{(i+k) \bmod m, j} \quad (8.16)$$

In other words, start with the A_{ii} term, then go down column i of A , wrapping back up to the top when you reach the bottom. The order of summation in this rearrangement will be the actual order of computation.

The algorithm is then as follows. The node which is handling the computation of C_{ij} does this (in parallel with the other nodes which are working with their own values of i and j):

```

1 iup = i+1 mod m;
2 idown = i-1 mod m;
3 for (k = 0; k < m; k++) {
4     km = (i+k) mod m;
5     broadcast(A[i,km]) to all nodes handling row i of C;
6     C[i,j] = C[i,j] + A[i,km]*B[km,j]
7     send B[km,j] to the node handling C[idown,j]
8     receive new B[km+1 mod m,j] from the node handling C[iup,j]
9 }
```

This is Fox's algorithm. Cannon's algorithm is similar, except that it does cyclical rotation in both rows and columns, compared to Fox's rotation only in columns but broadcast within rows.

The algorithm can be adapted in the obvious way to nonsquare matrices, etc.

8.4.1.2 Performance Issues

Note that in MPI we would probably want to implement this algorithm using communicators. For example, this would make broadcasting within a block row more convenient and efficient.

Note too that there is a lot of opportunity here to overlap computation and communication, which is the best way to solve the communication problem. For instance, we can do the broadcast above at the same time as we do the computation.

Obviously this algorithm is best suited to settings in which we have PEs in a mesh topology. This includes hypercubes, though one needs to be a little more careful about communications costs there.

8.4.2 Shared-Memory Case

8.4.2.1 OpenMP

Since a matrix multiplication in serial form consists of nested loops, a natural way to parallelize the operation in OpenMP is through the **for** pragma, e.g.

```

1 #pragma omp parallel for
2 for (i = 0; i < ncolsa; i++)
3     for (j = 0; i < nrowb; j++) {
4         sum = 0;
5         for (k = 0; i < ncolsa; i++)
6             sum += a[i][k] * b[k][j];
7     }

```

This would parallelize the outer loop, and we could do so at deeper nesting levels if profitable.

8.4.2.2 CUDA

Given that CUDA tends to work better if we use a large number of threads, a natural choice is for each thread to compute one element of the product, like this:

```

1 __global__ void matmul(float *ma, float *mb, float *mc, int nrowa, int ncolsa, int ncolsb)
2 { int k;
3   sum = 0;
4   for (k = 0; i < ncolsa; i++)
5       sum += a[i*ncolsa+k] * b[k*ncols+j];
6 }

```

This should produce a good speedup. But we can do even better, much much better.

The CUBLAS package includes very finely-tuned algorithms for matrix multiplication. The CUBLAS source code is not public, though, so in order to get an idea of how such tuning might be done, let's look at Prof. Richard Edgar's algorithm, which makes use of shared memory (<http://astro.pas.rochester.edu/~aquillen/gpuworkshop/AdvancedCUDA.pdf>):¹

```

1 __global__ void MultiplyOptimise(const float *A, const float *B, float *C) {
2   // Extract block and thread numbers
3   int bx = blockIdx.x; int by = blockIdx.y;
4   int tx = threadIdx.x; int ty = threadIdx.y;
5
6   // Index of first A sub-matrix processed by this block
7   int aBegin = dc_wA * BLOCK_SIZE * by;

```

¹Actually, this may be what CUBLAS uses.

```

8 // Index of last A sub-matrix
9 int aEnd = aBegin + dc_wA - 1;
10 // Stepsize of A sub-matrices
11 int aStep = BLOCK_SIZE;
12 // Index of first B sub-matrix
13 // processed by this block
14 int bBegin = BLOCK_SIZE * bx;
15 // Stepsize for B sub-matrices
16 int bStep = BLOCK_SIZE * dc_wB;
17 // Accumulator for this thread
18 float Csub = 0;
19 for(int a = aBegin, b = bBegin; a <= aEnd; a += aStep, b+= bStep) {
20 // Shared memory for sub-matrices
21 __shared__ float As[BLOCK_SIZE][BLOCK_SIZE];
22 __shared__ float Bs[BLOCK_SIZE][BLOCK_SIZE];
23 // Load matrices from global memory into shared memory
24 // Each thread loads one element of each sub-matrix
25 As[ty][tx] = A[a + (dc_wA * ty) + tx];
26 Bs[ty][tx] = B[b + (dc_wB * ty) + tx];
27 // Synchronise to make sure load is complete
28 __syncthreads();
29 // Perform multiplication on sub-matrices
30 // Each thread computes one element of the C sub-matrix
31 for( int k = 0; k < BLOCK_SIZE; k++ ) {
32 Csub += As[ty][k] * Bs[k][tx];
33 }
34 // Synchronise again
35 __syncthreads();
36 }
37 // Write the C sub-matrix back to global memory
38 // Each thread writes one element
39 int c = (dc_wB * BLOCK_SIZE * by) + (BLOCK_SIZE*bx);
40 C[c + (dc_wB*ty) + tx] = Csub;
41 }

```

Here are the relevant portions of the calling code, including global variables giving the number of columns (“width”) of the multiplier matrix and the number of rows (“height”) of the multiplicand:

```

#define BLOCK_SIZE 16
...
__constant__ int dc_wA;
__constant__ int dc_wB;
...
// Sizes must be multiples of BLOCK_SIZE
dim3 threads(BLOCK_SIZE,BLOCK_SIZE);
dim3 grid(wB/BLOCK_SIZE,ha/BLOCK_SIZE);
MultiplySimple<<<grid,threads>>>(d_A, d_B, d_C);
...

```

(Note the alternative way to configure threads, using the functions **threads()** and **grid()**.)

Here the the term “block” in the defined value **BLOCK_SIZE** refers both to blocks of threads and the partitioning of matrices. In other words, a thread block consists of 256 threads, to be thought of as a 16x16 “array” of threads, and each matrix is partitioned into submatrices of size 16x16.

In addition, in terms of grid configuration, there is again a one-to-one correspondence between thread blocks and submatrices. Each submatrix of the product matrix C will correspond to, and will be computed by, one block in the grid.

We are computing the matrix product $C = AB$. Denote the elements of A by a_{ij} for the element in row i , column j , and do the same for B and C . Row-major storage is used.

Each thread will compute one element of C , i.e. one c_{ij} . It will do so in the usual way, by multiplying column j of B by row i of A . However, the key issue is how this is done in concert with the other threads, and the timing of what portions of A and B are in shared memory at various times.

Concerning the latter, note the code

```
for(int a = aBegin, b = bBegin; a <= aEnd; a += aStep, b+= bStep) {
    // Shared memory for sub-matrices
    __shared__ float As[BLOCK_SIZE][BLOCK_SIZE];
    __shared__ float Bs[BLOCK_SIZE][BLOCK_SIZE];
    // Load matrices from global memory into shared memory
    // Each thread loads one element of each sub-matrix
    As[ty][tx] = A[a + (dc_wA * ty) + tx];
    Bs[ty][tx] = B[b + (dc_wB * ty) + tx];
}
```

Here we loop across a row of submatrices of A , and a column of submatrices of B , calculating one submatrix of C . In each iteration of the loop, we bring into shared memory a new submatrix of A and a new one of B . Note how even this copying from device global memory to device shared memory is shared among the threads.

As an example, suppose

$$A = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 7 & 8 & 9 & 10 & 11 & 12 \end{pmatrix} \quad (8.17)$$

and

$$B = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \\ 9 & 10 & 11 & 12 \\ 13 & 14 & 15 & 16 \\ 17 & 18 & 19 & 20 \\ 21 & 22 & 23 & 24 \end{pmatrix} \quad (8.18)$$

Further suppose that **BLOCK_SIZE** is 2. That's too small for good efficiency—giving only four threads per block rather than 256—but it's good for the purposes of illustration.

Let's see what happens when we compute C_{11} , the 2x2 submatrix of C 's upper-left corner. Due to the fact that partitioned matrices multiply "just like numbers," we have

$$C_{11} = A_{11}B_{11} + A_{12}B_{21} + A_{13}B_{31} \quad (8.19)$$

$$= \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 5 & 6 \end{pmatrix} + \dots \quad (8.20)$$

Now, all this will be handled by thread block number (0,0), i.e. the block whose X and Y "coordinates" are both 0. In the first iteration of the loop, A_{11} and B_{11} are copied to shared memory for that block, then in the next iteration, A_{12} and B_{21} are brought in, and so on.

Consider what is happening with thread number (1,0) within that block. Remember, its ultimate goal is to compute c_{21} (adjusting for the fact that in math, matrix subscripts start at 1). In the first iteration, this thread is computing

$$\begin{pmatrix} 1 & 2 \end{pmatrix} \begin{pmatrix} 1 \\ 5 \end{pmatrix} = 11 \quad (8.21)$$

It saves that 11 in its running total **Csub**, eventually writing it to the corresponding element of C :

```
int c = (dc_wB * BLOCK_SIZE * by) + (BLOCK_SIZE*bx);
C[c + (dc_wB*ty) + tx] = Csub;
```

Professor Edgar found that use of shared device memory resulted a huge improvement, extending the original speedup of 20X to 500X!

8.4.3 Finding Powers of Matrices

In Section (8.1), we saw a special case of matrix multiplication, powers, so that in our context here we have $A = B$. There are some small improvements that we could make in our algorithm for this case, but also there is something big.

Suppose for instance we need to find A^{32} . We could apply the above algorithm 31 times. But a much faster approach would be to first calculate A^2 , then square that result to get A^4 , then square it to get A^8 and so on. That would get us A^{32} by applying the algorithm in Section 8.4.1.1 only five times, instead of 31.

8.5 Solving Systems of Linear Equations

Suppose we have a system of equations

$$a_{i0}x_0 + \dots + a_{i,n-1}x_{n-1} = b_i, i = 0, 1, \dots, n - 1, \quad (8.22)$$

where the x_i are the unknowns to be solved for.

As you know, this system can be represented compactly as

$$AX = b, \quad (8.23)$$

where A is $n \times n$ and X and b is $n \times 1$.

8.5.1 Gaussian Elimination

Form the $n \times (n+1)$ matrix $C = (A \mid b)$ by appending the column vector b to the right of A. Then we work on the rows of C, with the pseudocode for the sequential case in the most basic form being

```

1 for ii = 0 to n-1
2   divide row ii by c[ii][ii]
3   for r = ii+1 to n-1 // vacuous if r = n-1
4     replace row r by row r - c[r][ii] times row ii
5 set new b to be column n-1 of C

```

(In the divide operation in the above pseudocode, if c_{ii} is 0, or even close to 0, that row is first swapped with another one further down.)

This transforms C to upper triangular form, i.e. all the elements c_{ij} with $i > j$ are 0. Also, all diagonal elements are equal to 1. This corresponds to a new set of equations,

$$\begin{aligned}
 c_{00}x_0 + c_{11}x_1 + c_{22}x_2 + \dots + c_{0,n-1}x_{n-1} &= b_0 \\
 c_{11}x_1 + c_{22}x_2 + \dots + c_{1,n-1}x_{n-1} &= b_1 \\
 c_{22}x_2 + \dots + c_{2,n-1}x_{n-1} &= b_2 \\
 &\dots \\
 c_{n-1,n-1}x_{n-1} &= b_{n-1}
 \end{aligned}$$

We then find the x_i via **back substitution**:

```

1  x[n-1] = b[n-1] / c[n-1,n-1]
2  for i = n-2 downto 0
3    x[i] = (b[i] - c[i][n-1] * x[n-1] - ... - c[i][i+1] * x[i+1]) / c[i][i]

```

An obvious parallelization of this algorithm would be to assign each process one contiguous group of rows. Then each process would do

```

1  for ii = 0 to n-1
2    if ii is in my group of rows
3      pivot = c[i][i]
4      divide row ii by pivot
5      broadcast row ii
6    else receive row ii
7    for r = ii+1 to n-1 in my group
8      subtract c[r][ii] times row ii from row r
9  set new b to be column n-1 of C

```

One problem with this is that in the outer loop, when ii gets past a given process' group of column indices, that process becomes idle. This can be solved by giving each process several groups of rows, in cyclic order. For example, say we have four processes. Then process 0 could take rows 0-99, 400-499, 800-899 and so on, process 1 would take rows 100-199, 500-599 etc.

8.5.2 The Jacobi Algorithm

One can rewrite (8.22) as

$$x_i = \frac{1}{a_{ii}} [b_i - (a_{i0}x_0 + \dots + a_{i,i-1}x_{i-1} + a_{i,i+1}x_{i+1} + \dots + a_{i,n-1}x_{n-1})], i = 0, 1, \dots, n-1. \quad (8.24)$$

This suggests a natural iterative algorithm for solving the equations. We start with our guess being, say, $x_i = b_i$ for all i . At our k^{th} iteration, we find our $(k+1)^{st}$ guess by plugging in our k^{th} guess into the right-hand side of (8.24). We keep iterating until the difference between successive guesses is small enough to indicate convergence.

This algorithm is guaranteed to converge if each diagonal element of A is larger in absolute value than the sum of the absolute values of the other elements in its row.

Parallelization of this algorithm is easy: Just assign each process to handle a block of X . Note that this means that each process must make sure that all other processes get the new value of this block after every iteration.

Note too that in matrix terms (8.24) can be expressed as

$$x^{(k+1)} = D^{-1}(b - O x^{(k)}) \quad (8.25)$$

where D is the diagonal matrix consisting of the diagonal elements of A (so its inverse is just the diagonal matrix consisting of the reciprocals of those elements), O is the square matrix obtained by replacing A 's diagonal elements by 0s, and $x^{(i)}$ is our guess for x in the i^{th} iteration. This reduces the problem to one of matrix multiplication, and thus we can parallelize the Jacobi algorithm by utilizing a method for doing parallel matrix multiplication.

8.6 OpenMP Implementation of the Jacobi Algorithm

```

1 #include <omp.h>
2
3 // partitions s..e into nc chunks, placing the ith in first and last (i
4 // = 0,...,nc-1)
5 void chunker(int s, int e, int nc, int i, int *first, int *last)
6 { int chunksize = (e-s+1) / nc;
7   *first = s + i * chunksize;
8   if (i < nc-1) *last = *first + chunksize - 1;
9   else *last = e;
10 }
11
12 // returns the "dot product" of vectors u and v
13 float innerprod(float *u, float *v, int n)
14 { float sum = 0.0; int i;
15   for (i = 0; i < n; i++)
16     sum += u[i] * v[i];
17   return sum;
18 }
19
20 // solves AX = Y, A nxn; stops iteration when total change is < n*eps
21 void jacobi(float *a, float *x, float *y, int n, float eps)
22 {
23   float *oldx = malloc(n*sizeof(float));
24   float se;
25   #pragma omp parallel
26   { int i;
27     int thn = omp_get_thread_num();
28     int nth = omp_get_num_threads();
29     int first, last;
30     chunker(0, n-1, nth, thn, &first, &last);
31     for (i = first; i <= last; i++) oldx[i] = x[i] = 1.0;
32     float tmp;
33     while (1) {
34       for (i = first; i <= last; i++) {
35         tmp = innerprod(&a[n*i], oldx, n);
36         tmp -= a[n*i+i] * oldx[i];
37         x[i] = (y[i] - tmp) / a[n*i+i];
38       }
39       #pragma omp barrier
40       #pragma omp for reduction(+:se)
41       for (i = first; i <= last; i++)
42         se += abs(x[i]-oldx[i]);
43       #pragma omp barrier
44       if (se < n*eps) break;
45       for (i = first; i <= last; i++)

```

```

46         oldx[i] = x[i];
47     }
48 }
49 }

```

Note the use of the OpenMP **reduction** clause.

8.7 Matrix Inversion

Many applications make use of A^{-1} for an $n \times n$ square matrix A . In many cases, it is not computed directly, but here we address methods for direct computation.

8.7.1 Using the Methods for Solving Systems of Linear Equations

Let B_i be the vector $(0, 0, \dots, 0, 1, 0, \dots, 0)'$, where the 1 is in position i , $i = 0, 1, \dots, n-1$. The B_i are then the columns of the $n \times n$ identity matrix I . So, we can find A^{-1} by using the methods in Section 8.5, setting $b = B_i$, $i = 0, 1, \dots, n-1$; the n solutions x we obtain this way will then be the columns of A^{-1} .

This can be formalized by replacing the x and b vectors in (8.25) by the corresponding matrices.

Again, this is easily parallelized. For instance, each process could be responsible for a given group of columns.

8.7.2 Power Series Method

Recall that for numbers x that are smaller than 1 in absolute value,

$$\frac{1}{1-x} = 1 + x + x^2 + \dots \quad (8.26)$$

In algebraic terms, this would be that for an $n \times n$ matrix C ,

$$(I - C)^{-1} = I + C + C^2 + \dots \quad (8.27)$$

This can be shown to converge if

$$\max_{i,j} |c_{ij}| < 1 \quad (8.28)$$

To invert our matrix A , then, we can set $C = I - A$, giving us

$$A^{-1} = (I - C)^{-1} = I + C + C^2 + \dots = I + (I - A) + (I - A)^2 + \dots \quad (8.29)$$

To meet the convergence condition, we could set $\tilde{A} = dA$, where d is small enough so that (8.28) holds for $I - \tilde{A}$. This will be possible, if all the elements of A are nonnegative.

Chapter 9

Parallel Combinatorial Algorithms

9.1 Overview

In Chapter 1, we saw Dijkstra's algorithm for finding the shortest path in a graph. In Chapter 8, we saw an algorithm for finding bridges within a graph. Both of these are **combinatorial search algorithms**. Such algorithms generally have exponential time complexity, and thus are natural candidates for parallel computation. This chapter will present a few more examples.

9.2 The 8 Queens Problem

A famous example is the 8 Queens Problem, in which one wishes to place eight queens on a standard 8x8 chessboard in such a way that no queen is attacking any other. (The generalization, of course, would involve n queens on an $n \times n$ board.) Suppose our goal is to find all possible solutions.

To start a solution to this problem, we first note in any solution will have the property that no row will contain more than one queen. This suggests building up a solution row by row: Suppose we have successfully placed queens so far in rows 0, 1, ..., $k-1$ (row 0 being the top row of the board). Where can we place a queen in row k ? Well, since we cannot use any column already occupied by the preceding k queens, that means we have a choice of $8-k$ columns. But even among those k columns, there will be j of them, for some $0 \leq j \leq 8-k$ that are in the diagonal attack path of some preceding queen. Then we can extend our tentative k -row solution to $8-k-j$ new $(k+1)$ -row solutions.

We will define our solution here for the shared-memory paradigm, though it would be easy to change this for the message-passing paradigm.¹ Define

¹The main point would be to change linked lists and pointers to arrays and array indices.

```

struct TentSoln {
    int RowsSoFar;
    int Cols[8];
    struct TentSoln *Next;
}

```

Each such **struct** contains a partial solution, up through row number **RowsSoFar**. The array **Cols** has the interpretation that **Col[I] == J** tells us which column the queen in row **I** occupies.

Each **struct** is a task showing one partial solution. The node which obtains this task will then extend this partial solution to several new, longer partial solutions.

The tasks are all placed into a linked list. **Next** points to the next item in the work pool.

A parallel solution based on this idea would like something like this:

```

1 while (work pool nonempty or at least one nonidle processor) {
2     get a TentSoln struct from the work pool, and point P to it;
3     I = P->RowsSoFar;
4     for (J = 0; J < 8; J++) {
5         if (a queen at row I, column J would not attack the previous queens) {
6             Q = malloc(sizeof(struct TentSoln));
7             Q->RowsSoFar = I+1;
8             add the struct pointed to by Q to the work pool;
9         }
10    }
11 }

```

There of course would also be code in the case $I = 8$ to check and see if we have found a solution, and if so, to record it, etc.

Note that any rotation of a solution—interchanging rows and columns—is also a solution. Similarly, any reflection across one of the two main diagonals of the board is also a solution. This information could be used to speed up computation, though at the expense of additional complexity of the code.

9.3 The 8-Square Puzzle Problem

This game was invented more than 100 years ago. Here is what a typical board position looks like:

0	5	3
1	4	
7	2	6

(The real puzzle has numbering from 1 to 8, but we use 0 to 7.)

Each number is on a little movable square, which can be moved up, down, left and right as long as the spot in the given direction is unoccupied. In the example above, the square 3, for instance, could be moved downward, producing an empty spot at the top right of the puzzle. The object of the game is arrange the squares in ascending numerical order, with square 0 at the upper left of the puzzle (which in this example happens to be the case already).

We again solve this by setting up a work pool, in this case a pool of board positions. Each board position would be implemented in something like this:

```
struct BoardPos {
    int Row[9];
    int Col[9];
    struct BoardPos *Next;
}
```

Here **Row[I]** and **Col[I]** would be the position of the square numbered **I**. For convenience, we also store the location of the blank position, in **Row[8]** and **Col[8]**.

Suppose a processor goes to the work pool and gets the board position depicted above. In the simplest form of the algorithm, the processor would check each of the three possible moves (4 right, 3 down, 6 up) to see if the resulting board position would duplicate one that had already been checked. All moves that lead to new positions would be added to the work pool. Each processor would loop around, pulling items from the work pool, until some processor somewhere finds a solution to the game (in which case that processor would add termination messages to the work pool, so that the other processors knew to stop). An outline of the algorithm would be as follows:

```
1 while (work pool nonempty or at least one nonidle processor) {
2     get a BoardPos struct from the work pool, and point P to it;
3     for (I = 0; I < 8; I++) {
4         for all possible moves of square I do {
5             Q = malloc(sizeof(struct BoardPos));
6             fill in *Q according to this move;
7             if *Q has not already been checked
8                 add this board to the work pool;
9         }
10    }
11 }
```

Again, code would need to be included for checking to see if a solution has been found, whether we have found that no solution exists, and so on.

Note the operation

```
    if *Q has not already been checked
        add this board to the work pool;
```

Clearly this is needed, to avoid endless cycling. But it is not as innocuous as it looks. If the set of all previously-checked board positions is to be made available to all processors, this may produce substantial increases in contention for memory and interprocessor interconnects. On the other hand, we could arrange the code such that only certain processors have to know about certain subsets of the set of previously-checked board positions, but this makes the code more complex and may produce load-balancing problems.

A more sophisticated version of the algorithm would use a **branch-and-bound** technique. The idea here is to reduce computation by giving priority in the work pool to those board positions which appear “promising” by some reasonable measure. For example, we could take as our measure the “distance” between a given board position and the goal board position, as defined by the sum of the distances from each numbered square to its place in the winning position. In the example above, for instance, the square numbered 5 is a distance of 2 from its ultimate place (2 meaning, one square to the right, one square down, so $1+1 = 2$). The board above is a distance 15 from the winning board.

The idea, then would be that we implement the work pool as an ordered linked list (or other ordered structure), and when a board position is added to the work pool, we insert it according to its distance from the winning board. This way the processors will usually work on the more promising boards, and thus hopefully reach the solution faster.

9.4 Itemset Analysis in Data Mining

9.4.1 What Is It?

The term **data mining** is a buzzword, but all it means is the process of finding relationships among a set of variables. In other words, it would seem to simply be a good old-fashioned statistics problem.

Well, in fact it *is* simply a statistics problem—but writ large. Instead of the tiny sample sizes of 25 you likely saw in your statistics courses, typical sample sizes in the data mining arena run in the hundreds of thousands or even hundreds of millions. And there may be hundreds of variables, in contrast to the, say, half dozen you might see in a statistics course.

Major, Major Warning: With so many variables, the chances of picking up spurious relations between variables is large. And although many books and tutorials on data mining will at least pay lip service to this issue (referring to it as **overfitting**, they don’t emphasize it enough.²

Putting the overfitting problem aside, though, by now the reader’s reaction should be, “This calls for parallel processing,” and he/she is correct. Here we’ll look at parallelizing a particular problem, called **itemset analysis**, the most famous example of which is the **market basket problem**:

²Some writers recommend splitting one’s data into a **training set**, which is used to discover relationships, and a **validation set**, which is used to confirm those relationships. However, overfitting can still occur even with this precaution.

9.4.2 The Market Basket Problem

Consider an online bookstore has records of every sale on the store's site. Those sales may be represented as a matrix S , whose (i,j) th element S_{ij} is equal to either 1 or 0, depending on whether the i th sale included book j , $i = 0, 1, \dots, s-1$, $j = 0, 1, \dots, t-1$. So each row of S represents one sale, with the 1s in that row showing which titles were bought. Each column of S represents one book title, with the 1s showing which sales transactions included that book.

Let's denote the entire line of book titles by T_0, \dots, T_{t-1} . An **itemset** is just a subset of this. A **frequent** itemset is one which appears in many of sales transactions. But there is more to it than that. The store wants to choose some books for special ads, of the form "We see you bought books X and Y . We think you may be interested in Z ."

Though we are using marketing as a running example here (which is the typical way that this subject is introduced), we will usually just refer to "items" instead of books, and to "database records" rather than sales transactions.

We have the following terminology:

- An **association rule** $I \rightarrow J$ is simply an ordered pair of disjoint itemsets I and J .
- The **support** of an association rule $I \rightarrow J$ is the proportion of records which include both I and J .
- The **confidence** of an association rule $I \rightarrow J$ is the proportion of records which include J , *among those records which include I* .

Note that in probability terms, the support is basically $P(I \text{ and } J)$ while the confidence is $P(J|I)$. If the confidence in the book business, it means that buyers of the books in set I also tend to buy those in J . But this information is not very useful if the support is low, because it means that the combination occurs so rarely that it's not worth our time to deal with it.

So, the user—let's call him/her the "data miner"—will first set thresholds for support and confidence, and then set out to find all association rules for which support and confidence exceed their respective thresholds.

9.4.3 Serial Algorithms

Various algorithms have been developed to find frequent itemsets and association rules. The most famous one for the former task is the **Apriori** algorithm. Even it has many forms. We will discuss one of the simplest forms here.

The algorithm is basically a breadth-first tree search. At the root we find the frequent 1-item itemsets. At the second level, we find the frequent 2-item itemsets, and so on. After we finish with level i , we then generate new candidate itemsets of size $i+1$ from the frequent itemsets we found of size i , by

The key point in the latter operation is that if an itemset is not frequent, i.e. has support less than the threshold, then adding further items to it will make it even less frequent. That itemset is then pruned from the tree, and the branch ends.

Here is the pseudocode:

```

set  $F_1$  to the set of 1-item itemsets whose support exceeds the threshold
for  $i = 2$  to  $b$ 
   $F_i = \phi$ 
  for each  $I$  in  $F_{i-1}$ 
    for each  $K$  in  $F_1$ 
       $Q = I \cup K$ 
      if support( $Q$ ) exceeds support threshold
        add  $Q$  to  $F_i$ 
  if  $F_i$  is empty break
return  $\cup_i F_i$ 

```

Again, there are many refinements of this, which shave off work to be done and thus increase speed. For example, we should avoid checking the same itemsets twice, e.g. first $\{1,2\}$ then $\{2,1\}$. This can be accomplished by keeping itemsets in lexicographical order. We will not pursue any refinements here.

9.4.4 Parallelizing the Apriori Algorithm

Clearly there is lots of opportunity for parallelizing the serial algorithm above. Both of the inner **for** loops can be parallelized in straightforward ways; they are “embarrassingly parallel.” There are of course critical sections to worry about in the shared-memory setting, and in the message-passing setting one must designate a manager node in which to store the F_i .

However, as more and more refinements are made in the serial algorithm, then the parallelism in this algorithm become less and less “embarrassing.” And things become more challenging if the storage needs of the F_i , and of their associated “accounting materials” such as a directory showing the current tree structure (done via hash trees), become greater than what can be stored in the memory of one node.

In other words, parallelizing the market basket problem can be very challenging. The interested reader is referred to the considerable literature which has developed on this topic.

Chapter 10

Introduction to Parallel Sorting

Sorting is one of the most common operations in parallel processing applications. For example, it is central to many parallel database operations, and important in areas such as image processing, statistical methodology and so on. A number of different types of parallel sorting schemes have been developed. Here we look at some of these schemes.

10.1 Quicksort

You are probably familiar with the idea of quicksort: First break the original array into a “small-element” pile and a “large-element” pile, by comparing to a **pivot** element. In a naive implementation, the first element of the array serves as the pivot, but better performance can be obtained by taking, say, the median of the first three elements. Then “recurse” on each of the two piles, and then string the results back together again.

This is an example of the **divide and conquer** approach seen in so many serial algorithms. It is easily parallelized (though load-balancing issues may arise). Here, for instance, we might assign one pile to one thread and the other pile to another thread.

Suppose the array to be sorted is named **x**, and consists of **n** elements.

In a naive implementation, the piles would be put into new arrays, but this is bad in two senses: It wastes memory space, and wastes time, since much copying of arrays needs to be done. A better implementation places the two piles back into the original array **x**. The following C code does that.

The function **separate()** is intended to be used in a recursive quicksort operation. It operates on **x[l]** through **x[h]**, a subarray of **x** that itself may have been formed at an earlier stage of the recursion. It forms two piles from those elements, and placing the piles back in the same region **x[l]** through **x[h]**. It also has a return value, showing where the first pile ends.

```

int separate(int l, int h)
{ int ref,i,j,k,tmp;
  ref = x[h]; i = l-1; j = h;
  do {
    do i++; while (x[i] < ref && i < h);
    do j--; while (x[j] > ref && j > l);
    tmp = x[i]; x[i] = x[j]; x[j] = tmp;
  } while (j > i);
  x[j] = x[i]; x[i] = x[h]; x[h] = tmp;
  return i;
}

```

The function `separate()` rearranges the subarray, returning a value `m`, so that:

- `x[l]` through `x[m-1]` are less than `x[m]`,
- `x[m+1]` through `x[h]` are greater than `x[m]`, and
- `x[m]` is in its “final resting place,” meaning that `x[m]` will never move again for the remainder of the sorting process. (Another way of saying this is that the current `x[m]` is the `m`-th smallest of all the original `x[i]`, `i = 0,1,...,n-1`.)

By the way, `x[l]` through `x[m-1]` will also be in their final resting places as a group. They may be exchanging places with each other from now on, but they will never again leave the range `i` though `m-1` within the `x` array as a whole. A similar statement holds for `x[m+1]` through `x[n-1]`.

10.1.1 Shared-Memory Quicksort

Here is OpenMP code which performs quicksort in the shared-memory paradigm (adapted from code in the OpenMP Source Code Repository, <http://www.pcg.uill.es/ompscr/>):¹

```

1 void qs(int *x, int l, int h)
2 { int newl[2], newh[2], i, m;
3   m = separate(x,l,h);
4   newl[0] = l; newh[0] = m-1;
5   newl[1] = m+1; newh[1] = h;
6   #pragma omp parallel
7   {
8     #pragma omp for nowait
9     for (i = 0; i < 2; i++)
10      qs(newl[i],newh[i]);
11  }
12 }

```

¹Recall that another implementation, using the `task` directive, was given earlier in Section 4.4.

Note the **nowait** clause. Since different threads are operating on different portions of the array, they need not be synchronized.

A variant on this which might achieve better load balancing would set up a **task pool**, consisting of an array of **(l, h)** pairs. Initially the pool consists of just [0,n-1]. The function **qs()** would then become iterative instead of recursive, with its main loop looking something like this for an array of length n:

```

fetch an (l,h) pair from the task pool
while not done
  call separate() on x[l] through x[h], yielding m
  if m < h
    add (m+1,h) to the task pool
  h = m-1
  if l == h
    fetch [l,h] from the task pool

```

This pseudocode is missing important details. For example, How does the iteration within a thread stop? The key lies in pairs of the form (i,i) , which I'll call *singletons*. The sort is done when the number of singletons reaches n.

10.1.2 Hyperquicksort

This algorithm was originally developed for hypercubes, but can be used on any message-passing system having a power of 2 for the number of nodes.²

It is assumed that at the beginning each PE contains some chunk of the array to be sorted. After sorting, each PE will contain some chunk of the sorted array, meaning that:

- each chunk is itself in sorted form
- for all cases of $i < j$, the elements at PE i are less than the elements at PE j

If the sorted array itself were our end, rather than our means to something else, we could now collect it at some node, say node 0. If, as is more likely, the sorting is merely an intermediate step in a larger distributed computation, we may just leave the chunks at the nodes and go to the next phase of work.

Say we are on a d-cube. The intuition behind the algorithm is quite simple:

```

for i = d downto 1
  for each i-cube:
    root of the i-cube broadcasts its median to all in the i-cube,
    to serve as pivot

```

²See Chapter 6 for definitions of hypercube terms.

```

consider the two (i-1)-subcubes of this i-cube
each pair of partners in the (i-1)-subcubes exchanges data:
  low-numbered PE gives its partner its data larger than pivot
  high-numbered PE gives its partner its data smaller than pivot

```

To avoid deadlock, have the lower-numbered partner send then receive, and vice versa for the higher-numbered one. Better, in MPI, use `MPI_SendRcv()`.

After the first iteration, all elements in the lower (d-1)-cube are less than all elements in higher (d-1)-cube. After d such steps, the array will be sorted.

10.2 Mergesorts

10.2.1 Sequential Form

In its serial form, mergesort has the following pseudocode:

```

1 // initially called with l = 0 and h = n-1, where n is the length of the
2 // array and is assumed here to be a power of 2
3 void seqmergesort(int *x, int l, int h)
4 {   seqmergesort(x,0,h/2-1);
5     seqmergesort(x,h/2,h);
6     merge(x,l,h);
7 }

```

The function `merge()` should be done in-place, i.e. without using an auxiliary array. It basically codes the operation shown in pseudocode for the message-passing case in Section 10.2.3.

10.2.2 Shared-Memory Mergesort

This is similar to the patterns for shared-memory quicksort in Section 10.1.1 above.

10.2.3 Message Passing Mergesort on a Tree Topology

First, we organize the processing nodes into a binary tree. This is simply from the point of view of the software, rather than a physical grouping of the nodes. We will assume, though, that the number of nodes is one less than a power of 2.

To illustrate the plan, say we have seven nodes in all. We could label node 0 as the root of the tree, label nodes 1 and 2 to be its two children, label nodes 3 and 4 to be node 1's children, and finally label nodes 5 and 6 to be node 2's children.

It is assumed that the array to be sorted is initially distributed in the leaf nodes (recall a similar situation for hyperquicksort), i.e. nodes 3-6 in the above example. The algorithm works best if there are approximately the same number of array elements in the various leaves.

In the first stage of the algorithm, each leaf node applies a regular sequential sort to its current holdings. Then each node begins sending its now-sorted array elements to its parent, one at a time, in ascending numerical order.

Each nonleaf node then will merge the lists handed to it by its two children. Eventually the root node will have the entire sorted array. Specifically, each nonleaf node does the following:

```
do
  if my left-child datum < my right-child datum
    pass my left-child datum to my parent
  else
    pass my right-child datum to my parent
until receive the "no more data" signal from both children
```

There is quite a load balancing issue here. On the one hand, due to network latency and the like, one may get better performance if each node accumulates a chunk of data before sending to the parent, rather than sending just one datum at a time. Otherwise, “upstream” nodes will frequently have no work to do.

On the other hand, the larger the chunk size, the earlier the leaf nodes will have no work to do. So for any particular platform, there will be some optimal chunk size, which would need to be determined by experimentation.

10.2.4 Compare-Exchange Operations

These are key to many sorting algorithms.

A **compare-exchange**, also known as **compare-split**, simply means in English, “Let’s pool our data, and then I’ll take the lower half and you take the upper half.” Each node executes the following pseudocode:

```
send all my data to partner
receive all my partner’s data
if I have a lower id than my partner
  I keep the lower half of the pooled data
else
  I keep the upper half of the pooled data
```

10.2.5 Bitonic Mergesort

Definition: A sequence $(a_0, a_1, \dots, a_{k-1})$ is called **bitonic** if either of the following conditions holds:

(a) The sequence is first nondecreasing then nonincreasing, meaning that for some r

$$(a_0 \leq a_1 \leq \dots \leq a_r \geq a_{r+1} \geq a_{n-1})$$

(b) The sequence can be converted to the form in (a) by rotation, i.e. by moving the last k elements from the right end to the left end, for some k .

As an example of (b), the sequence (3,8,12,15,14,5,1,2) can be rotated rightward by two element positions to form (1,2,3,8,12,15,14,5). Or we could just rotate by one element, moving the 2 to forming (2,3,8,12,15,14,5,1).

Note that the definition includes the cases in which the sequence is purely nondecreasing ($r = n-1$) or purely nonincreasing ($r = 0$).

Also included are “V-shape” sequences, in which the numbers first decrease then increase, such as (12,5,2,8,20). By (b), these can be rotated to form (a), with (12,5,2,8,20) being rotated to form (2,8,20,12,5), an “A-shape” sequence.

(For convenience, from here on I will use the terms *increasing* and *decreasing* instead of *nonincreasing* and *nondecreasing*.)

Suppose we have bitonic sequence $(a_0, a_1, \dots, a_{k-1})$, where k is a power of 2. Rearrange the sequence by doing compare-exchange operations between a_i and $a_{n/2+i}$, $i = 0, 1, \dots, n/2-1$. Then it is not hard to prove that the new $(a_0, a_1, \dots, a_{k/2-1})$ and $(a_{k/2}, a_{k/2+1}, \dots, a_{k-1})$ are bitonic, and every element of that first subarray is less than or equal to every element in the second one.

So, we have set things up for yet another divide-and-conquer attack:

```

1 // x is bitonic of length n, n a power of 2
2 void sortbitonic(int *x, int n)
3 { do the pairwise compare-exchange operations
4   if (n > 2) {
5       sortbitonic(x, n/2);
6       sortbitonic(x+n/2, n/2);
7   }
8 }
```

This can be parallelized in the same ways we saw for Quicksort earlier.

So much for sorting bitonic sequences. But what about general sequences?

We can proceed as follows, using our function `sortbitonic()` above:

1. For each $i = 0, 2, 4, \dots, n-2$:

- Each of the pairs (a_i, a_{i+1}) , $i = 0, 2, \dots, n-2$ is bitonic, since *any* 2-element array is bitonic!
 - Apply **sortbitonic()** to (a_i, a_{i+1}) . In this case, we are simply doing a compare-exchange.
 - If $i/2$ is odd, reverse the pair, so that this pair and the pair immediately preceding it now form a 4-element bitonic sequence.
2. For each $i = 0, 4, 8, \dots, n-4$:
- Apply **sortbitonic()** to $(a_i, a_{i+1}, a_{i+2}, a_{i+3})$.
 - If $i/4$ is odd, reverse the quartet, so that this quartet and the quartet immediately preceding it now form an 8-element bitonic sequence.
3. Keep building in this manner, until get to a single sorted n -element list.

There are many ways to parallelize this. In the hypercube case, the algorithm consists of doing compare-exchange operations with all neighbors, pretty much in the same pattern as hyperquicksort.

10.3 The Bubble Sort and Its Cousins

10.3.1 The Much-Maligned Bubble Sort

Recall the **bubble sort**:

```

1 void bubblesort(int *x, int n)
2 { for i = n-1 downto 1
3     for j = 0 to i
4         compare-exchange(x, i, j, n)
5 }
```

Here the function **compare-exchange()** is as in Section 10.2.4 above. In the context here, it boils down to

```

if x[i] > x[j]
    swap x[i] and x[j]
```

In the first i iteration, the largest element “bubbles” all the way to the right end of the array. In the second iteration, the second-largest element bubbles to the next-to-right-end position, and so on.

You learned in your algorithms class that this is a very inefficient algorithm—when used serially. But it’s actually rather usable in parallel systems.

For example, in the shared-memory setting, suppose we have one thread for each value of *i*. Then those threads can work in parallel, as long as a thread with a larger value of *i* does not overtake a thread with a smaller *i*, where “overtake” means working on a larger *j* value.

Once again, it probably pays to chunk the data. In this case, `compare-exchange()` fully takes on the meaning it had in Section 10.2.4.

10.3.2 A Popular Variant: Odd-Even Transposition

A popular variant of this is the **odd-even transposition sort**. The pseudocode for a shared-memory version is:

```

1 // the argument "me" is this thread's ID
2 void oddevensort(int *x, int n, int me)
3 { for i = 1 to n
4     if i is odd
5         if me is even
6             compare-exchange(x,me,me+1,n)
7         else // me is odd
8             compare-exchange(x,me,me-1,n)
9     else // i is even
10        if me is even
11            compare-exchange(x,me,me-1,n)
12        else // me is odd
13            compare-exchange(x,me,me+1,n)

```

If the second or third argument of `compare-exchange()` is less than 0 or greater than *n*-1, the function has no action.

This looks a bit complicated, but all it's saying is that, from the point of view of an even-numbered element of *x*, it trades with its right neighbor during odd phases of the procedure and with its left neighbor during even phases.

Again, this is usually much more effective if done in chunks.

10.3.3 CUDA Implementation of Odd/Even Transposition Sort

```

1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <cuda.h>
4
5 // compare and swap; copies from the f to t, swapping f[i] and
6 // f[j] if the higher-index value is smaller; it is required that i < j
7 __device__ void cas(int *f,int *t,int i,int j, int n, int me)
8 {
9     if (i < 0 || j >= n) return;
10    if (me == i) {

```

```

11     if (f[i] > f[j]) t[me] = f[j];
12     else t[me] = f[i];
13 } else { // me == j
14     if (f[i] > f[j]) t[me] = f[i];
15     else t[me] = f[j];
16 }
17 }
18
19 // does one iteration of the sort
20 __global__ void oekern(int *da, int *daaux, int n, int iter)
21 { int bix = blockIdx.x; // block number within grid
22   if (iter % 2) {
23     if (bix % 2) cas(da, daaux, bix-1, bix, n, bix);
24     else cas(da, daaux, bix, bix+1, n, bix);
25   } else {
26     if (bix % 2) cas(da, daaux, bix, bix+1, n, bix);
27     else cas(da, daaux, bix-1, bix, n, bix);
28   }
29 }
30
31 // sorts the array ha, length n, using odd/even transp. sort;
32 // kept simple for illustration, no optimization
33 void oddeven(int *ha, int n)
34 {
35   int *da;
36   int dasize = n * sizeof(int);
37   cudaMalloc((void **)&da, dasize);
38   cudaMemcpy(da, ha, dasize, cudaMemcpyHostToDevice);
39   // the array daaux will serve as "scratch space"
40   int *daaux;
41   cudaMalloc((void **)&daaux, dasize);
42   dim3 dimGrid(n, 1);
43   dim3 dimBlock(1, 1, 1);
44   int *tmp;
45   for (int iter = 1; iter <= n; iter++) {
46     oekern<<<dimGrid, dimBlock>>>(da, daaux, n, iter);
47     cudaThreadSynchronize();
48     if (iter < n) {
49       // swap pointers
50       tmp = da;
51       da = daaux;
52       daaux = tmp;
53     } else
54       cudaMemcpy(ha, daaux, dasize, cudaMemcpyDeviceToHost);
55   }
56 }

```

Recall that in CUDA code, separate blocks of threads cannot synchronize with each other. Unless we deal with just a single block, this necessitates limiting the kernel to a single iteration of the algorithm, so that as iterations progress, execution alternates between the device and the host.

Moreover, we do not take advantage of shared memory. One possible solution would be to use `__syncthreads()` within each block for most of the compare-and-exchange operations, and then having the host take care of the operations on the boundaries between blocks.

10.4 Shearsort

In some contexts, our hardware consists of a two-dimensional mesh of PEs. A number of methods have been developed for such settings, one of the most well known being Shearsort, developed by Sen, Shamir and the eponymous Isaac Scherson of UC Irvine. Again, the data is assumed to be initially distributed among the PEs. Here is the pseudocode:

```

1  for i = 1 to ceiling(log2(n)) + 1
2      if i is odd
3          sort each even row in descending order
4          sort each odd row in ascending order
5      else
6          sort each column in ascending order

```

At the end, the numbers are sorted in a “snakelike” manner.

For example:

6	12
5	9

6	12
9	5

6	5
9	12

5	6 ↓
12	← 9

No matter what kind of system we have, a natural domain decomposition for this problem would be for each process to be responsible for a group of rows. There then is the question about what to do during the even-numbered iterations, in which column operations are done. This can be handled via a parallel matrix transpose operation. In MPI, the function `MPI_Alltoall()` may be useful.

10.5 Bucket Sort with Sampling

For concreteness, suppose we are using MPI on message-passing hardware, say with 10 PEs. As usual in such a setting, suppose our data is initially distributed among the PEs.

Suppose we knew that our array to be sorted is a random sample from the uniform distribution on (0,1). In other words, about 20% of our array will be in (0,0.2), 38% will be in (0.45,0.83) and so on.

What we could do is assign PE0 to the interval (0,0.1), PE1 to (0.1,0.2) etc. Each PE would look at its local data, and distribute it to the other PEs according to this interval scheme. Then each PE would do a local sort.

In general, we don't know what distribution our data comes from. We solve this problem by doing sampling. In our example here, each PE would sample some of its local data, and send the sample to PE0. From all of these samples, PE0 would find the decile values, i.e. 10th percentile, 20th percentile, ..., 90th percentile. These values, called **splitters** would then be broadcast to all the PEs, and they would then distribute their local data to the other PEs according to these intervals.

10.6 Enumeration Sort

This one is really simple. Take for instance the array (12,5,13,18,6). There are 2 elements less than 12, so in the end, it should go in position 2 of the sorted array, (5,6,12,13,18).

Say we wish to sort x , which for convenience we assume contains no tied values. Then the pseudocode for this algorithm, placing the results in y , is

```
for all i in 0...n-1:
  count = 0
  elt = x[i]
  for all j in 0...n-1:
    if x[j] < elt then count++
  y[count] = elt
```

The outer (or inner) loop is easily parallelized.

Chapter 11

Parallel Computation of Fourier Series, with an Introduction to Parallel Imaging


Mathematical computations involving sounds and images, for example for voice and pattern recognition are often performed using **Fourier** analysis.

11.1 General Principles

11.1.1 One-Dimensional Fourier Series

A sound **wave form** graphs volume of the sound against time. Here, for instance, is the wave form for a vibrating reed:¹

¹Reproduced here by permission of Prof. Peter Hamburger, Indiana-Purdue University, Fort Wayne. See <http://www.ipfw.edu/math/Workshop/PBC.html>



Recall that we say a function of time $g(t)$ is **periodic** (“repeating,” in our casual wording above) with period T if $g(u+T) = g(u)$ for all u . The **fundamental frequency** of $g()$ is then defined to be the number of periods per unit time,

$$f_0 = \frac{1}{T} \quad (11.1)$$

Recall also from calculus that we can write a function $g(t)$ (not necessarily periodic) as a Taylor series, which is an “infinite polynomial”:

$$g(t) = \sum_{n=0}^{\infty} c_n t^n. \quad (11.2)$$

The specific values of the c_n may be derived by differentiating both sides of (11.2) and evaluating at $t = 0$, yielding

$$c_n = \frac{g^{(n)}(0)}{n!}, \quad (11.3)$$

where $g^{(j)}$ denotes the j th derivative of $g()$.

For instance, for e^t ,

$$e^t = \sum_{n=0}^{\infty} \frac{1}{n!} t^n \quad (11.4)$$

In the case of a repeating function, it is more convenient to use another kind of series representation, an “infinite trig polynomial,” called a **Fourier series**. This is just a fancy name for a weighted sum of sines and

cosines of different frequencies. More precisely, we can write any repeating function $g(t)$ with period T and fundamental frequency f_0 as

$$g(t) = \sum_{n=0}^{\infty} a_n \cos(2\pi n f_0 t) + \sum_{n=1}^{\infty} b_n \sin(2\pi n f_0 t) \quad (11.5)$$

for some set of weights a_n and b_n . Here, instead of having a weighted sum of terms

$$1, t, t^2, t^3, \dots \quad (11.6)$$

as in a Taylor series, we have a weighted sum of terms

$$1, \cos(2\pi f_0 t), \cos(4\pi f_0 t), \cos(6\pi f_0 t), \dots \quad (11.7)$$

and of similar sine terms. Note that the frequencies $n f_0$, in those sines and cosines are integer multiples of the fundamental frequency of x , f_0 , called **harmonics**.

The weights a_n and b_n , $n = 0, 1, 2, \dots$ are called the **frequency spectrum** of $g()$. The coefficients are calculated as follows:²

$$a_0 = \frac{1}{T} \int_0^T g(t) dt \quad (11.8)$$

$$a_n = \frac{2}{T} \int_0^T g(t) \cos(2\pi n f_0 t) dt \quad (11.9)$$

$$b_n = \frac{2}{T} \int_0^T g(t) \sin(2\pi n f_0 t) dt \quad (11.10)$$

By analyzing these weights, we can do things like machine-based voice recognition (distinguishing one person's voice from another) and speech recognition (determining what a person is saying). If for example one person's voice is higher-pitched than that of another, the first person's weights will be concentrated more on the higher-frequency sines and cosines than will the weights of the second.


Since $g(t)$ is a graph of loudness against time, this representation of the sound is called the **time domain**. When we find the Fourier series of the sound, the set of weights a_n and b_n is said to be a representation of

²The get an idea as to how these formulas arise, see Section 11.8. But for now, if you integrate both sides of (11.5), you will at least verify that the formulas below do work.

the sound in the **frequency domain**. One can recover the original time-domain representation from that of the frequency domain, and vice versa, as seen in Equations (11.8), (11.9), (11.10) and (11.5).

In other words, the transformations between the two domains are inverses of each other, and there is a one-to-one correspondence between them. Every $g()$ corresponds to a unique set of weights and vice versa.

Now here is the frequency-domain version of the reed sound:



Note that this graph is very “spiky.” In other words, even though the reed’s waveform includes all frequencies, most of the power of the signal is at a few frequencies which arise from the physical properties of the reed.

Fourier series are often expressed in terms of complex numbers, making use of the relation

$$e^{i\theta} = \cos(\theta) + i \sin(\theta), \quad (11.11)$$

where $i = \sqrt{-1}$.³

³There is basically no physical interpretation of complex numbers. Instead, they are just mathematical abstractions. However, they are highly useful abstractions, with the complex form of Fourier series, beginning with (11.12), being a case in point.

The complex form of (11.5) is

$$g(t) = \sum_{j=-\infty}^{\infty} c_j e^{2\pi i j \frac{t}{T}}. \quad (11.12)$$

The c_j are now generally complex numbers. They are functions of the a_j and b_j , and thus comprise the frequency spectrum.

Equation (11.12) has a simpler, more compact form than (11.5). Do you now see why I referred to Fourier series as trig polynomials? The series (11.12) involves the j^{th} powers of $e^{2\pi i \frac{t}{T}}$.

11.1.2 Two-Dimensional Fourier Series

Let's now move from sounds images. Here $g()$ is a function of two variables, $g(u,v)$, where u and v are the horizontal and vertical coordinates of a pixel in the image; $g(u,v)$ is the intensity of the image at that pixel. If it is a gray-scale image, the intensity is whiteness of the image at that pixel, typically with 0 being pure black and 255 being pure white. If it is a color image, a typical graphics format is to store three intensity values at a point, one for each of red, green and blue. The various colors come from combining three colors at various intensities.

Since images are two-dimensional instead of one-dimensional like a sound wave form, the Fourier series for an image is a sum of sines and cosines in two variables, i.e. a double sum $\sum_j \sum_k \dots$ instead of $\sum_j \dots$

The terminology changes a bit. Our original data is now referred to as being in the **spatial domain**, rather than the time domain. But the Fourier series coefficients are still said to be in the frequency domain.

11.2 Discrete Fourier Transforms

In sound and image applications, we seldom if ever know the exact form of the repeating function $g()$. All we have is a **sampling** from $g()$, i.e. we only have values of $g(t)$ for a set of discrete values of t .

In the sound example above, a typical sampling rate is 8000 samples per second.⁴ So, we may have $g(0)$, $g(0.000125)$, $g(0.000250)$, $g(0.000375)$, and so on. In the image case, we sample the image pixel by pixel.

Thus we can't calculate integrals like (11.8). So, how do we approximate the Fourier transform based on the sample data?

⁴See Section 11.9 for the reasons behind this.

11.2.1 One-Dimensional Data

Let $X = (x_0, \dots, x_{n-1})$ denote the sampled values, i.e. the time-domain representation of $g()$ based on our sample data. These are interpreted as data from one period of $g()$, with the period being n and the fundamental frequency being $1/n$. The frequency-domain representation will also consist of n numbers, c_0, \dots, c_{n-1} , defined as follows:⁵

$$c_k = \frac{1}{n} \sum_{j=0}^{n-1} x_j e^{-2\pi i j k / n} = \frac{1}{n} \sum_{j=0}^{n-1} x_j q^{j k} \quad (11.13)$$

where

$$q = e^{-2\pi i / n} \quad (11.14)$$

again with $i = \sqrt{-1}$. The array C of complex numbers c_k is called the **discrete Fourier transform (DFT)** of X .

Note that (11.13) is basically a discrete analog of (11.9) and (11.10).

As in the continuous case, we can recover each domain from the other. So, while (11.13) shows how to go to the frequency domain from the time domain, we can go from the frequency domain to the time domain via the inverse transform, whose equation is

$$x_k = \sum_{j=0}^{n-1} c_j e^{2\pi i j k / n} = \sum_{j=0}^{n-1} c_j q^{-j k} \quad (11.15)$$

Note that (11.15) is basically a discrete analog of (11.5).

Note too that instead of having infinitely many harmonics, we can only have n of them: $1, 1/n, 2/n, \dots, (n-1)/n$. It would be impossible to have more than n , as can be seen by reasoning as follows: The x_k are given, q is a constant, and we are solving for the c_j . So, we have n equations in n unknowns. If we had more than n unknowns, the system would be indeterminate.

⁵It should be noted that there are many variant definitions of these transforms. One common variation is to include/exclude a scale factor, such as our $1/n$ in (11.13). Another type of variations involve changing only c_0 , in order to make certain matrices have more convenient forms.

11.2.2 Two-Dimensional Data

The spectrum numbers c_{rs} are double-subscripted, like the original data x_{uv} , the latter being the pixel intensity in row u , column v of the image, $u = 0, 1, \dots, n-1$, $v = 0, 1, \dots, m-1$. Equation (11.13) becomes

$$c_{rs} = \frac{1}{n} \frac{1}{m} \sum_{j=0}^{n-1} \sum_{k=0}^{m-1} x_{jk} e^{-2\pi i \left(\frac{jr}{n} + \frac{ks}{m} \right)} \quad (11.16)$$

where $r = 0, 1, \dots, n-1$, $s = 0, 1, \dots, m-1$.

Its inverse is

$$x_{rs} = \sum_{j=0}^{n-1} \sum_{k=0}^{m-1} c_{jk} e^{2\pi i \left(\frac{jr}{n} + \frac{ks}{m} \right)} \quad (11.17)$$

11.3 Parallel Computation of Discrete Fourier Transforms

11.3.1 CUFFT

Remember that CUDA includes some excellent FFT routines, in CUFFT.

11.3.2 The Fast Fourier Transform

Speedy computation of a discrete Fourier transform was developed by Cooley and Tukey in their famous Fast Fourier Transform (FFT), which takes a “divide and conquer” approach:

Equation (11.13) can be rewritten as

$$c_k = \frac{1}{n} \left[\sum_{j=0}^{m-1} x_{2j} q^{2jk} + \sum_{j=0}^{m-1} x_{2j+1} q^{(2j+1)k} \right] \quad (11.18)$$

where $m = n/2$.

After some algebraic manipulation, this becomes

$$c_k = \frac{1}{2} \left[\frac{1}{m} \sum_{j=0}^{m-1} x_{2j} z^{jk} + q^k \frac{1}{m} \sum_{j=0}^{m-1} x_{2j+1} z^{jk} \right] \quad (11.19)$$

where $z = e^{-2\pi i/m}$.

A look at Equation (11.19) shows that the two sums within the brackets have the same form as Equation (11.13). In other words, Equation (11.19) shows how we can compute an n -point FFT from two $\frac{n}{2}$ -point FFTs. That means that a DFT can be computed recursively, cutting the sample size in half at each recursive step.

In a shared-memory setting such as OpenMP, we could implement this recursive algorithm in the manners of Quicksort in Chapter 10.

In a message-passing setting, again because this is a divide-and-conquer algorithm, we can use the pattern of Hyperquicksort, also in 10.

Some digital signal processing chips implement this in hardware, with a special interconnection network to implement this algorithm.

11.3.3 A Matrix Approach

The matrix form of (11.13) is

$$C = \frac{1}{n} AX \quad (11.20)$$

where A is $n \times n$. Element (j,k) of A is q^{jk} , while element j of X is x_j . This formulation of the problem then naturally leads one to use parallel methods for matrix multiplication, as in Chapter 8.

Divide-and-conquer tends not to work too well in shared-memory settings, because after some point, fewer and fewer threads will have work to do. Thus this matrix formulation is quite valuable.

11.3.4 Parallelizing Computation of the Inverse Transform

The form of the DFT (11.13) and its inverse (11.15) are very similar. For example, the inverse transform is again of a matrix form as in (11.20); even the new matrix looks a lot like the old one.⁶

⁶In fact, one can obtain the new matrix easily from the old, as explained in Section 11.8.

Thus the methods mentioned above, e.g. FFT and the matrix approach, apply to calculation of the inverse transforms too.

11.3.5 Parallelizing Computation of the Two-Dimensional Transform

Regroup (11.16) as:

$$c_{rs} = \frac{1}{n} \sum_{j=0}^{n-1} \left(\frac{1}{m} \sum_{k=0}^{m-1} x_{jk} e^{-2\pi i \left(\frac{ks}{m}\right)} \right) e^{-2\pi i \left(\frac{jr}{n}\right)} \quad (11.21)$$

$$= \frac{1}{n} \sum_{j=0}^{n-1} y_{js} e^{-2\pi i \left(\frac{jr}{n}\right)} \quad (11.22)$$

Note that y_{js} , i.e. the expression between the large parentheses, is the s^{th} component of the DFT of the j^{th} row of our data. And hey, the last expression (11.22) above is in the same form as (11.13)! Of course, this means we are taking the DFT of the spectral coefficients rather than observed data, but numbers are numbers.

In other words: To get the two-dimensional DFT of our data, we first get the one-dimensional DFTs of each row of the data, place these in rows, and then find the DFTs of each column. This property is called **separability**.

This certainly opens possibilities for parallelization. Each thread (shared memory case) or node (message passing case) could handle groups of rows of the original data, and in the second stage each thread could handle columns.

Or, we could interchange rows and columns in this process, i.e. put the j sum inside and k sum outside in the above derivation.

11.4 Applications to Image Processing

In image processing, there are a number of different operations which we wish to perform. We will consider two of them here.

11.4.1 Smoothing

An image may be too “rough.” There may be some pixels which are noise, accidental values that don’t fit smoothly with the neighboring points in the image.

One way to smooth things out would be to replace each pixel intensity value⁷ by the mean or median among the pixels neighbors. These could be the four immediate neighbors if just a little smoothing is needed, or we could go further out for a higher amount of smoothing. There are many variants of this.

But another way would be to apply a **low-pass filter** to the DFT of our image. This means that after we compute the DFT, we simply delete the higher harmonics, i.e. set c_{rs} to 0 for the larger values of r and s . We then take the inverse transform back to the spatial domain. Remember, the sine and cosine functions of higher harmonics are “wigglier,” so you can see that all this will have the effect of removing some of the wiggleness in our image—exactly what we wanted.

We can control the amount of smoothing by the number of harmonics we remove.

The term *low-pass filter* obviously alludes to the fact that the low frequencies “pass” through the filter but the high frequencies are blocked. Since we’ve removed the high-oscillatory components, the effect is a smoother image.⁸

To do smoothing in parallel, if we just average neighbors, this is easily parallelized. If we try a low-pass filter, then we use the parallelization methods shown here earlier.

11.4.2 Edge Detection

In computer vision applications, we need to have a machine-automated way to deduce which pixels in an image form an edge of an object.

Again, edge-detection can be done in primitive ways. Since an edge is a place in the image in which there is a sharp change in the intensities at the pixels, we can calculate slopes of the intensities, in the horizontal and vertical directions. (This is really calculating the approximate values of the partial derivatives in those directions.)

But the Fourier approach would be to apply a high-pass filter. Since an edge is a set of pixels which are abruptly different from their neighbors, we want to keep the high-frequency components and block out the low ones.

Again, this means first taking the Fourier transform of the original, then deleting the low-frequency terms, then taking the inverse transform to go back to the spatial domain.

Below we have “before and after” pictures, first of original data and then the picture after an edge-detection process has been applied.⁹

⁷Remember, there may be three intensity values per pixel, for red, green and blue.

⁸Note that we may do more smoothing in some parts of the image than in others.

⁹These pictures are courtesy of Bill Green of the Robotics Laboratory at Drexel University. In this case he is using a Sobel process instead of Fourier analysis, but the result would have been similar for the latter. See his Web tutorial at www.pages.drexel.edu/~weg22/edge.html, including the original pictures, which may not show up well in our printed book here.



The second picture looks like a charcoal sketch! But it was derived mathematically from the original picture, using edge-detection methods.

Note that edge detection methods also may be used to determine where sounds (“ah,” “ee”) begin and end in speech-recognition applications. In the image case, edge detection is useful for face recognition, etc.

Parallelization here is similar to that of the smoothing case.

11.5 The Cosine Transform

It’s inconvenient, to say the least, to work with all those complex numbers. But an alternative exists in the form of the **cosine transform**, which is a linear combination of cosines in the one-dimensional case, and of products of cosines in the two-dimensional case.

$$d_{uv} = \frac{2}{\sqrt{mn}} Y(u)Y(v) \sum_{j=0}^{n-1} \sum_{k=0}^{m-1} x_{jk} \cos \frac{(2j+1)u\pi}{2n} \cos \frac{(2k+1)v\pi}{2m}, \quad (11.23)$$

where $Y(0) = 1/\sqrt{2}$ and $Y(t) = 1$ for $t > 0$.

$$x_{jk} = \frac{2}{\sqrt{mn}} \sum_{u=0}^{n-1} \sum_{v=0}^{m-1} Y(u)Y(v)d_{uv} \cos \frac{(2j+1)u\pi}{2n} \cos \frac{(2k+1)v\pi}{2m}, \quad (11.24)$$

11.6 Keeping the Pixel Intensities in the Proper Range

Normally pixel intensities are stored as integers between 0 and 255, inclusive. With many of the operations mentioned above, both Fourier-based and otherwise, we can get negative intensity values, or values higher than 255. We may wish to discard the negative values and scale down the positive ones so that most or all are smaller than 256.

Furthermore, even if most or all of our values are in the range 0 to 255, they may be near 0, i.e. too faint. If so, we may wish to multiply them by a constant.

11.7 Does the Function $g()$ Really Have to Be Repeating?

It is clear that in the case of a vibrating reed, our loudness function $g(t)$ really is periodic. What about other cases?

A graph of your voice would look “locally periodic.” One difference would be that the graph would exhibit more change through time as you make various sounds in speaking, compared to the one repeating sound for the reed. Even in this case, though, your voice *is* repeating within short time intervals, each interval corresponding to a different sound. If you say the word *eye*, for instance, you make an “ah” sound and then an “ee” sound. The graph of your voice would show one repeating pattern during the time you are saying “ah,” and another repeating pattern during the time you are saying “ee.” So, even for voices, we do have repeating patterns over short time intervals.

On the other hand, in the image case, the function may be nearly constant for long distances (horizontally or vertically), so a local periodicity argument doesn’t seem to work there.

The fact is, though, that it really doesn’t matter in the applications we are considering here. Even though mathematically our work here has tacitly assumed that our image is duplicated infinitely times (horizontally and vertically),¹⁰ we don’t care about this. We just want to get a measure of “wiggleness,” and fitting linear combinations of trig functions does this for us.

¹⁰And in the case of the cosine transform, implicitly we are assuming that the image flips itself on every adjacent copy of the image, first right-side up, then upside-down, then right-side up again, etc.

11.8 Vector Space Issues (optional section)

The theory of Fourier series (and of other similar transforms), relies on vector spaces. It actually is helpful to look at some of that here. Let's first discuss the derivation of (11.13).

Define X and C as in Section 11.2. X 's components are real, but it is also a member of the vector space V of all n -component arrays of complex numbers.

For any complex number $a+bi$, define its **conjugate**, $\overline{a+bi} = a - bi$. Note that

$$\overline{e^{i\theta}} = \cos \theta - i \sin \theta = \cos(-\theta) + i \sin(-\theta) = e^{-i\theta} \quad (11.25)$$

Define an inner product ("dot product"),

$$[u, w] = \frac{1}{n} \sum_{j=0}^{n-1} u_j \bar{w}_j. \quad (11.26)$$

Define

$$v_h = (1, q^{-h}, q^{-2h}, \dots, q^{-(n-1)h}), h = 0, 1, \dots, n-1. \quad (11.27)$$

Then it turns out that the v_h form an orthonormal basis for V .¹¹ For example, to show orthonogonality, observe that for $r \neq s$

$$[v_r, v_s] = \frac{1}{n} \sum_{j=0}^{n-1} v_{rj} \bar{v}_{sj} \quad (11.28)$$

$$= \frac{1}{n} \sum_{j=0}^{n-1} q^{j(-r+s)} \quad (11.29)$$

$$= \frac{1 - q^{(-r+s)n}}{n(1 - q)} \quad (11.30)$$

$$= 0, \quad (11.31)$$

due to the identity $1 + y + y^2 + \dots + y^k = \frac{1-y^{k+1}}{1-y}$ and the fact that $q^n = 1$. In the case $r = s$, the above computation shows that $[v_r, v_s] = 1$.

¹¹Recall that this means that these vectors are orthogonal to each other, and have length 1, and that they span V .

The DFT of X , which we called C , can be considered the “coordinates” of X in V , relative to this orthonormal basis. The k th coordinate is then $[X, v_k]$, which by definition is (11.13).

The fact that we have an orthonormal basis for V here means that the matrix A/n in (11.20) is an orthogonal matrix. For real numbers, this means that this matrix’s inverse is its transpose. In the complex case, instead of a straight transpose, we do a conjugate transpose, $B = \overline{A/n}^t$, where t means transpose. So, B is the inverse of A/n . In other words, in (11.20), we can easily get back to X from C , via

$$X = BC = \frac{1}{n} \overline{A}^t C. \quad (11.32)$$

It’s really the same for the nondiscrete case. Here the vector space consists of all the possible periodic functions $g()$ (with reasonable conditions placed regarding continuity etc.) forms the vector space, and the sine and cosine functions form an orthonormal basis. The a_n and b_n are then the “coordinates” of $g()$ when the latter is viewed as an element of that space.

11.9 Bandwidth: How to Read the *San Francisco Chronicle* Business Page (optional section)

The popular press, especially business or technical sections, often uses the term **bandwidth**. What does this mean?

Any transmission medium has a natural range $[f_{min}, f_{max}]$ of frequencies that it can handle well. For example, an ordinary voice-grade telephone line can do a good job of transmitting signals of frequencies in the range 0 Hz to 4000 Hz, where “Hz” means cycles per second. Signals of frequencies outside this range suffer fade in strength, i.e. are **attenuated**, as they pass through the phone line.¹²

We call the frequency interval $[0, 4000]$ the **effective bandwidth** (or just the **bandwidth**) of the phone line.

In addition to the bandwidth of a **medium**, we also speak of the bandwidth of a **signal**. For instance, although your voice is a mixture of many different frequencies, represented in the Fourier series for your voice’s waveform, the really low and really high frequency components, outside the range $[340, 3400]$, have very low power, i.e. their a_n and b_n coefficients are small. Most of the power of your voice signal is in that range of frequencies, which we would call the effective bandwidth of your voice waveform. This is also the reason why digitized speech is sampled at the rate of 8,000 samples per second. A famous theorem, due to Nyquist, shows that the sampling rate should be double the maximum frequency. Here the number 3,400 is “rounded up” to 4,000, and after doubling we get 8,000.

Obviously, in order for your voice to be heard well on the other end of your phone connection, the bandwidth of the phone line must be at least as broad as that of your voice signal, and that is the case.

¹²And in fact will probably be deliberately filtered out.

11.9. BANDWIDTH: HOW TO READ THE SAN FRANCISCO CHRONICLE BUSINESS PAGE (OPTIONAL SECTION)2

However, the phone line's bandwidth is not much broader than that of your voice signal. So, some of the frequencies in your voice will fade out before they reach the other person, and thus some degree of distortion will occur. It is common, for example, for the letter 'f' spoken on one end to be mis-heard as 's' on the other end. This also explains why your voice sounds a little different on the phone than in person. Still, most frequencies are reproduced well and phone conversations work well.

We often use the term "bandwidth" to literally refer to width, i.e. the width of the interval $[f_{min}, f_{max}]$.

There is huge variation in bandwidth among transmission media. As we have seen, phone lines have bandwidth intervals covering values on the order of 10^3 . For optical fibers, these numbers are more on the order of 10^{15} .

The radio and TV frequency ranges are large also, which is why, for example, we can have many AM radio stations in a given city. The AM frequency range is divided into subranges, called **channels**. The width of these channels is on the order of the 4000 we need for a voice conversation. That means that the transmitter at a station needs to shift its content, which is something like in the $[0,4000]$ range, to its channel range. It does that by multiplying its content times a sine wave of frequency equal to the center of the channel. If one applies a few trig identities, one finds that the product signal falls into the proper channel!

Accordingly, an optical fiber could also carry many simultaneous phone conversations.

Bandwidth also determines how fast we can set digital bits. Think of sending the sequence 10101010... If we graph this over time, we get a "squarewave" shape. Since it is repeating, it has a Fourier series. What happens if we double the bit rate? We get the same graph, only horizontally compressed by a factor of two. The effect of this on this graph's Fourier series is that, for example, our former a_3 will now be our new a_6 , i.e. the $2\pi \cdot 3f_0$ frequency cosine wave component of the graph now has the double the old frequency, i.e. is now $2\pi \cdot 6f_0$. That in turn means that the effective bandwidth of our 10101010... signal has doubled too.

In other words: To send high bit rates, we need media with large bandwidths.

Chapter 12

Applications to Statistics/Data Mining

How did the word *statistics* get supplanted by *data mining*? In a word, it is a matter of scale.

In the old days of statistics, a data set of 300 observations on 3 or 4 variables was considered large. Today, the widespread use of computers and the Web yield data sets with numbers of observations that are easily in the tens of thousands range, and in a number of cases even tens of millions. The numbers of variables can also be in the thousands or more.

In addition, the methods have become much more combinatorial in nature. In a classification problem, for instance, the old discriminant analysis involved only matrix computation, whereas a nearest-neighbor analysis requires far more computer cycles to complete.

In short, this calls for parallel methods of computation.

12.1 Itemset Analysis

12.1.1 What Is It?

The term **data mining** is a buzzword, but all it means is the process of finding relationships among a set of variables. In other words, it would seem to simply be a good old-fashioned statistics problem.

Well, in fact it *is* simply a statistics problem—but writ large, as mentioned earlier.

Major, Major Warning: With so many variables, the chances of picking up spurious relations between variables is large. And although many books and tutorials on data mining will at least pay lip service to this issue (referring to it as **overfitting**, they don't emphasize it enough.¹

¹Some writers recommend splitting one's data into a **training set**, which is used to discover relationships, and a **validation set**,

Putting the overfitting problem aside, though, by now the reader's reaction should be, "This calls for parallel processing," and he/she is correct. Here we'll look at parallelizing a particular problem, called **itemset analysis**, the most famous example of which is the **market basket problem**:

12.1.2 The Market Basket Problem

Consider an online bookstore has records of every sale on the store's site. Those sales may be represented as a matrix S , whose (i,j) th element S_{ij} is equal to either 1 or 0, depending on whether the i th sale included book j , $i = 0, 1, \dots, s-1$, $j = 0, 1, \dots, t-1$. So each row of S represents one sale, with the 1s in that row showing which titles were bought. Each column of S represents one book title, with the 1s showing which sales transactions included that book.

Let's denote the entire line of book titles by T_0, \dots, T_{b-1} . An **itemset** is just a subset of this. A **frequent itemset** is one which appears in many of sales transactions. But there is more to it than that. The store wants to choose some books for special ads, of the form "We see you bought books X and Y . We think you may be interested in Z ."

Though we are using marketing as a running example here (which is the typical way that this subject is introduced), we will usually just refer to "items" instead of books, and to "database records" rather than sales transactions.

We have the following terminology:

- An **association rule** $I \rightarrow J$ is simply an ordered pair of disjoint itemsets I and J .
- The **support** of an association rule $I \rightarrow J$ is the proportion of records which include both I and J .
- The **confidence** of an association rule $I \rightarrow J$ is the proportion of records which include J , *among those records which include I* .

Note that in probability terms, the support is basically $P(I \text{ and } J)$ while the confidence is $P(J|I)$. If the confidence in the book business, it means that buyers of the books in set I also tend to buy those in J . But this information is not very useful if the support is low, because it means that the combination occurs so rarely that it's not worth our time to deal with it.

So, the user—let's call him/her the "data miner"—will first set thresholds for support and confidence, and then set out to find all association rules for which support and confidence exceed their respective thresholds.

which is used to confirm those relationships. However, overfitting can still occur even with this precaution.

12.1.3 Serial Algorithms

Various algorithms have been developed to find frequent itemsets and association rules. The most famous one for the former task is the **Apriori** algorithm. Even it has many forms. We will discuss one of the simplest forms here.

The algorithm is basically a breadth-first tree search. At the root we find the frequent 1-item itemsets. In the online bookstore, for instance, this would mean finding all individual books that appear in at least r of our sales transaction records, where r is our threshold.

At the second level, we find the frequent 2-item itemsets, e.g. all pairs of books that appear in at least r sales records, and so on. After we finish with level i , we then generate new candidate itemsets of size $i+1$ from the frequent itemsets we found of size i .

The key point in the latter operation is that if an itemset is not frequent, i.e. has support less than the threshold, then adding further items to it will make it even less frequent. That itemset is then pruned from the tree, and the branch ends.

Here is the pseudocode:

```

set  $F_1$  to the set of 1-item itemsets whose support exceeds the threshold
for  $i = 2$  to  $b$ 
   $F_i = \phi$ 
  for each  $I$  in  $F_{i-1}$ 
    for each  $K$  in  $F_1$ 
       $Q = I \cup K$ 
      if support( $Q$ ) exceeds support threshold
        add  $Q$  to  $F_i$ 
  if  $F_i$  is empty break
return  $\cup_i F_i$ 

```

Again, there are many refinements of this, which shave off work to be done and thus increase speed. For example, we should avoid checking the same itemsets twice, e.g. first $\{1,2\}$ then $\{2,1\}$. This can be accomplished by keeping itemsets in lexicographical order. We will not pursue any refinements here.

12.1.4 Parallelizing the Apriori Algorithm

Clearly there is lots of opportunity for parallelizing the serial algorithm above. Both of the inner **for** loops can be parallelized in straightforward ways; they are “embarrassingly parallel.” There are of course critical sections to worry about in the shared-memory setting, and in the message-passing setting one must designate a manager node in which to store the F_i .

However, as more and more refinements are made in the serial algorithm, then the parallelism in this algorithm become less and less “embarrassing.” And things become more challenging if the storage needs of the F_i , and of their associated “accounting materials” such as a directory showing the current tree structure (done via hash trees), become greater than what can be stored in the memory of one node.

In other words, parallelizing the market basket problem can be very challenging. The interested reader is referred to the considerable literature which has developed on this topic.

12.2 Probability Density Estimation

Let X denote some quantity of interest in a given population, say people’s heights. Technically, the **probability density function** of X , typically denoted by f , is a function on the real line with the following properties:

- $f(t) \geq 0$ for all t
- for any $r < s$,

$$P(r < X < s) = \int_r^s f(t) dt \quad (12.1)$$

(Note that this implies that f integrates to 1.)

This seems abstract, but it’s really very simple: Say we have data on X , n sample values X_1, \dots, X_n , and we plot a histogram from this data. Then f is what the histogram is estimating. If we have more and more data, the histogram gets closer and closer to the true f .²

So, how do we estimate f , and how do we use parallel computing to reduce the time needed?

12.2.1 Kernel-Based Density Estimation

Histogram computation breaks the real down into intervals, and then counts how many X_i fall into each interval. This is fine as a crude method, but one can do better.

No matter what the interval width is, the histogram will consist of a bunch of rectangles, rather than a smooth curve. This problem basically stems from a lack of weighting on the data.

For example, suppose we are estimating $f(25.8)$, and suppose our histogram interval is $[24.0, 26.0]$, with 54 points falling into that interval. Intuitively, we can do better if we give the points closer to 25.8 more weight.

²The histogram must be scaled to have total area 1. Most statistical programs have options for this.

One way to do this is called **kernel-based** density estimation, which for instance in R is handled by the function **density()**.

We need a set of weights, more precisely a weight function k , called the **kernel**. Any nonnegative function which integrates to 1—i.e. a density function in its own right—will work. Typically k is taken to be the Gaussian or normal density function,

$$k(u) = \frac{1}{\sqrt{2\pi}} e^{-0.5u^2} \quad (12.2)$$

Our estimator is then

$$\hat{f}(t) = \frac{1}{nh} \sum_{i=1}^n k\left(\frac{t - X_i}{h}\right) \quad (12.3)$$

In statistics, it is customary to use the $\hat{}$ symbol (pronounced “hat”) to mean “estimate of.” Here \hat{f} means the estimate of f .

Note carefully that we are estimating an entire function! There are infinitely many possible values of t , thus infinitely many values of $f(t)$ to be estimated. This is reflected in (12.3), as $\hat{f}(t)$ does indeed give a (potentially) different value for each t .



Here h , called the *bandwidth*, is playing a role analogous to the interval width in the case of histograms.

Again, this looks very abstract, but all it is doing is assigning weights to the data. Consider our example above in which we wish to estimate $f(25.8)$, i.e. $t = 25.8$ and $h = 6.0$. If say, X_{88} is 1209.1, very far as away from 25.8, we don’t want this data point to have much weight in our estimation of $f(25.8)$. Well, it won’t have much weight at all, because the quantity

$$u = \frac{25.8 - 88}{6} \quad (12.4)$$

will be very large, and (12.2) will be tiny, as u will be way, way out in the left tail.

Now, keep all this in perspective. In the end, we will be plotting a curve, *just like we do with a histogram*. We simply have a more sophisticated way to do this than plotting a histogram. Following are the graphs generated first by the histogram method, then by the kernel method, on the same data:



There are many ways to parallelize this computation, such as:

- Remember, we are going to compute (12.3) for many values of t . So, we can just have each process compute a block of those values.
- We may wish to try several different values of h , just as we might try several different interval widths for a histogram. We could have each process compute using its own values of h .
- It can be shown that (12.3) has the form of something called a **convolution**. The theory of convolution

would take us too far afield,³ but this fact is useful here, as the Fourier transform of a convolution can be shown to be the product of the Fourier transforms of the two convolved components.⁴ In other words, *this reduces the problem to that of parallelizing Fourier transforms*—something we know how to do, from Chapter 11.

12.2.2 Histogram Computation for Images

In image processing, histograms are used to find tallies of how many pixels there are of each intensity. (Note that there is thus no interval width issue, as there is a separate “interval” value for each possible intensity level.) The serial pseudocode is:

```
for i = 1, ..., numintenslevels:
    count = 0
    for row = 1, ..., numRows:
        for col = 1, ..., numcols:
            if image[row][col] == i: count++
    hist[i] = count
```

On the surface, this is certainly an “embarrassingly parallel” problem. In OpenMP, for instance, we might have each thread handle a block of rows of the image, i.e. parallelize the **for row** loop. In CUDA, we might have each thread handle an individual pixel, thus parallelizing the nested **for row/col** loops.

However, to make this go fast is a challenge, say in CUDA, due to issues of what to store in shared memory, when to swap it out, etc. A very nice account of fine-tuning this computation in CUDA is given in *Histogram Calculation in CUDA*, by Victor Podlozhnyuk of NVIDIA, 2007, http://developer.download.nvidia.com/compute/cuda/1_1/Website/projects/histogram256/doc/histogram.pdf. The actual code is at http://developer.download.nvidia.com/compute/cuda/sdk/website/Data-Parallel_Algorithms.html#histogram. A summary follows:

3

If you’ve seen the term before and are curious as to how this is a convolution, read on:
Write (12.3) as

$$\hat{f}(t) = \sum_{i=1}^n \frac{1}{h} k \left(\frac{t - X_i}{h} \right) \cdot \frac{1}{n} \quad (12.5)$$

Now consider two artificial random variables U and V, created just for the purpose of facilitating computation, defined as follows.

The random variable U takes on the values ih with probability $g \cdot \frac{1}{h} k(i)$, $i = -c, -c+1, \dots, 0, 1, \dots, c$ for some value of c that we choose to cover most of the area under k, with g chosen so that the probabilities sum to 1. The random variable V takes on the values X_1, \dots, X_n (considered fixed here), with probability 1/n each. U and V are set to be independent.

Then (g times) (12.5) becomes $P(U+V=t)$, exactly what convolution is about, the probability mass function (or density, in the continuous case) of a random variable arising as the sum of two independent nonnegative random variables.

⁴Again, if you have some background in probability and have seen characteristic functions, this fact comes from the fact that the characteristic function of the sum of two independent random variables is equal to the product of the characteristic functions of the two variables.

Podlozhnyuk's overall plan is to have the threads compute subhistograms for various chunks of the image, then merge the subhistograms to create the histogram for the entire data set. Each thread will handle $1/k$ of the image's pixels, where k is the total number of threads in the grid, i.e. across all blocks.

Since the subhistograms are accessed repeatedly, we want to store them in shared memory. In Podlozhnyuk's first cut at the problem, he maintains a separate subhistogram for each thread. He calls this version of the code **histogram64**. The name stems from the fact that only 64 intensity levels are used, i.e. the more significant 6 bits of each pixel's data byte. The reason for this restriction will be discussed later.

The subhistograms in this code are one-dimensional arrays of length 64 bytes, one count per image intensity level. However, these arrays must be placed judiciously within the shared memory for a block, so that the threads get involved in as few bank conflicts as possible. Podlozhnyuk devised a clever way to do this, which in fact generates no bank conflicts at all.


In the end, the many subhistograms within a block must be merged, and those merged counts must in turn be merged across all blocks. The former operation is done again by careful ordering to avoid any bank conflicts, and then (if the hardware has the capability) via calls to **atomicAdd()**.

Now, why does **histogram64** tabulate image intensities at only 6-bit granularity? It's simply a matter of resource limitations. Podlozhnyuk notes that NVIDIA says that for best efficiency, there should be between 128 and 256 threads per block. He takes the middle ground, 192. With 16K of shared memory per block, $16K/192$ works out to about 85 bytes per thread. That eliminates computing a histogram for the full 8-bit image data, with 256 intensity levels; all we can accommodate is 6 bits, for 64 levels.

Accordingly, Podlozhnyuk offers **histogram256**, which refines the process, by having one subhistogram per warp, instead of per thread. This allows the full 8-bit data, 256 levels, to be tabulated, one word devoted to each count, rather than just one byte. A subhistogram is now a table, 256 rows by 32 columns (one for each thread in the warp), with each table entry being 4 bytes (1 byte is not sufficient, as 32 threads are tabulating with it).

12.3 Clustering

Suppose you have data consisting of (X,Y) pairs, which when plotted look like this:



It looks like there may be two or three groups here. What clustering algorithms do is to form groups, both their number and their membership, i.e. which data points belong to which groups. (Note carefully that *there is no “correct” answer here. This is merely an exploratory data analysis tool.*

Clustering is used in many diverse fields. For instance, it is used in image processing for segmentation and edge detection.

Here we have two variables, say people’s heights and weights. In general we have many variables, say p of them, so whatever clustering we find will be in p -dimensional space. No, we can’t picture it very easily if p is larger than (or even equal to) 3, but we can at least identify membership, i.e. John and Mary are in group 1, Jenny is in group 2, etc. We may derive some insight from this.

There are many, many types of clustering algorithms. Here we will discuss the famous **k-means** algorithm, developed by Prof. Jim MacQueen of the UCLA business school.

The method couldn’t be simpler. Choose k , the number of groups you want to form, and then run this:

```

1 # form initial groups from the first k data points (or choose randomly)
2 for i = 1,...,k:
3   group[i] = (x[i],y[i])
4   center[i] = (x[i],y[i])
5 do:
6   for j = 1,...,n:
7     find the closest center[i] to (x[j],y[j])
8     cl[j] = the i you got in the previous line

```

```

9   for i = 1, ..., k:
10      group[i] = all (x[j], y[j]) such that cl[j] = i
11      center[i] = average of all (x, y) in group[i]
12 until group memberships do not change from one iteration to the next

```

Definitions of terms:

- *Closest* means in p-dimensional space, with the usual Euclidean distance: The distance from (a_1, \dots, a_p) to (b_1, \dots, b_p) is

$$\sqrt{(b_1 - a_1)^2 + \dots + (b_p - a_p)^2} \quad (12.6)$$

- The *center* of a group is its **centroid**, which is a fancy name for taking the average value in each component of the data points in the group. If $p = 2$, for example, the center consists of the point whose X coordinate is the average X value among members of the group, and whose Y coordinate is the average Y value in the group.

In terms of parallelization, again we have an embarrassingly parallel problem.

12.4 Principal Component Analysis (PCA)

Consider data consisting of (X,Y) pairs as we saw in Section 12.3. Suppose X and Y are highly correlated with each other. Then for some constants c and d,

$$Y \approx c + dX \quad (12.7)$$

Then in a sense there is really just one random variable here, as the second is nearly equal to some linear combination of the first. The second provides us with almost no new information, once we have the first. In other words, even though the vector (X,Y) roams in *two*-dimensional space, it usually sticks close to a *one*-dimensional object, namely the line (12.7).

Now think again of p variables. It may be the case that there exist $r < p$ variables, consisting of linear combinations of the p variables, that carry most of the information of the full set of p variables. If r is much less than p, we would prefer to work with those r variables. In data mining, this is called **dimension reduction**.

It can be shown that we can find these r variables by finding the r eigenvectors corresponding to the r largest eigenvalues of a certain matrix. We will not pursue that here, but the point is that again we have a matrix formulation, and thus parallelizing the problem can be done easily by using methods for parallel matrix operations.

12.5 Parallel Processing in R

R is a widely used programming language for statistics and data manipulation. Given that huge statistical problems have become commonplace today, a number of “parallel R” packages have been developed.

Here are a few such packages:

- Message-passing or quasi-message passing: **Rmpi**, **snow**, **foreach**
- Shared-memory: **Rdsm**, **bigmemory**
- GPU: **gputools**, **rgpu**

A far more extensive list is at <http://cran.r-project.org/web/views/HighPerformanceComputing.html>.

Some of these packages will be covered in the following sections.

12.5.1 Rmpi

The **Rmpi** package provides an interface from R to MPI, much like that of **pyMPI**. So, we run **Rmpi** on top of MPI.

Installation:

Say you want to install in the directory */a/b/c/*. The easiest way to do so is use R’s **install.packages()** function, say:

```
> install.packages("{\bf Rmpi}", "/a/b/c/")
```

This will install **Rmpi** in the directory */a/b/c/Rmpi*.

You’ll need to arrange for the directory */a/b/c* (not */a/b/c/Rmpi*) to be added to your R library search path. I recommend placing a line

```
.libPaths("/a/b/c/")
```

in a file **.Rprofile** in your home directory.

Usage:

Fire up MPI, and then in R load in **Rmpi**, by typing

```
> library({\bf Rmpi})
```

Then start **Rmpi**:

```
> mpi.spawn.Rslaves()
```

On some systems, the call to **mpi.spawn.Rslaves()** may encounter problems. An alternate method of launching the worker processes is to copy the **Rprofile** file in the **Rmpi** distribution to **.Rprofile** in, say, your home directory. (You may wish to add a line to the file that calls **.libPaths()** so that your **Rmpi** directory becomes visible to R.) Then start R, say for two workers and a manager, by running something like (for the LAM case)

```
mpirun -c 3 R --no-save -q
```

This will start R on all machines in the group you started MPI on.

The available functions are similar to (though much more extensive than) those of **pyMPI**. Here are just a few:

- **mpi.comm.size():**
Returns the number of MPI processes, including the master that spawned the other processes.
- **mpi.comm.rank():**
Returns the rank of the process that executes it.
- **mpi.send(), mpi.recv():**
The usual send/receive operations.
- **mpi.bcast(), mpi.scatter(), mpi.gather():**
The usual broadcast, scatter and gather operations.
- Etc.

12.5.2 The R snow Package

The **snow** package runs on top of **Rmpi** (or directly via sockets), allowing the programmer to more conveniently express the parallel disposition of work. It operates under a scatter/gather model.

For instance, just as the ordinary R function **apply()** applies the same function to all rows of a matrix, the **snow** function **parApply()** does that in parallel, across multiple machines; different machines will work on different rows.

Installation:

Use **install.packages()**, as described above for **Rmpi**. If you plan to have **snow** run on top of **Rmpi**, you'll of course need the latter too.

Usage:

Make sure **snow** is in your library path (see material on **Rmpi** above).

Load **snow**:

```
> library(snow)
```

One then sets up a cluster, by calling the **snow** function **makeCluster()**. The named argument **type** of that function indicates the networking platform, e.g. "MPI," "PVM" or "SOCK." The last indicates that you wish **snow** to run on TCP/IP sockets that it creates itself, rather than going through MPI. In the examples here, I used "SOCK," on machines named **pc48** and **pc49**, setting up the cluster this way:

```
> cls <- makeCluster(type="SOCK",c("pc48","pc49"))
```

For MPI or PVM, one specifies the number of nodes to create, rather than specifying the nodes themselves.

Note that the above R code sets up worker nodes at the machines named **pc48** and **pc49**; these are in addition to the master node, which is the machine on which that R code is executed

There are various other optional arguments. One you may find useful is **outfile**, which records the result of the call in the file **outfile**. This can be helpful if the call fails.

Let's look at a simple example of multiplication of a vector by a matrix. We set up a test matrix:

```
> a <- matrix(c(1,2,3,4,5,6,7,8,9,10,11,12),nrow=6)
> a
      [,1] [,2]
[1,]    1    7
[2,]    2    8
[3,]    3    9
[4,]    4   10
[5,]    5   11
[6,]    6   12
```

We will multiply the vector $(1, 1)^T$ (T meaning transpose) by our matrix **a**, by defining a dot product function:

```
> dot <- function(x,y) {return(x**y)}
```

Let's test it using the ordinary **apply()**:

```
> apply(a,1,dot,c(1,1))
[1] 8 10 12 14 16 18
```

To review your R, note that this applies the function **dot()** to each row (indicated by the 1, with 2 meaning column) of **a** playing the role of the first argument to **dot()**, and with **c(1,1)** playing the role of the second argument.

Now let's do this in parallel, across our two machines in our cluster **cls**:

```
> parApply(cls,a,1,dot,c(1,1))
[1] 8 10 12 14 16 18
```

The function **clusterCall(cls,f,args)** applies the given function **f()** at each worker node in the cluster **cls**, using the arguments provided in **args**.

The function **clusterExport(cls,varlist)** copies the variables in the list **varlist** to each worker in the cluster **cls**. You can use this to avoid constant shipping of large data sets from the master to the workers; you just do so once, using **clusterExport()** on the corresponding variables, and then access those variables as global. For instance:

```
> z <- function() return(x)
> x <- 5
> y <- 12
> clusterExport(cls,list("x","y"))
> clusterCall(cls,z)
[[1]]
[1] 5

[[2]]
[1] 5
```

The function **clusterEvalQ(cls,expression)** runs **expression** at each worker node in **cls**. Continuing the above example, we have

```
> clusterEvalQ(cls,x <- x+1)
[[1]]
[1] 6

[[2]]
[1] 6
```

```
> clusterCall(cls,z)
[[1]]
[1] 6

[[2]]
[1] 6

> x
[1] 5
```

Note that **x** still has its original version back at the master.

The function **clusterApply(cls,individualargs,f,commonargsgohere)** runs **f()** at each worker node in **cls**, with arguments as follows. The first argument to **f()** for worker *i* is the *i*th element of the list **individualargs**, i.e. **individualargs[[i]]**, and optionally one can give additional arguments for **f()** following **f()** in the argument list for **clusterApply()**.

Here for instance is how we can assign an ID to each worker node, like MPI **rank**:⁵

```
> myid <- 0
> clusterExport(cls,"myid")
> setid <- function(i) {myid <<- i} # note superassignment operator
> clusterApply(cls,1:2,setid)
[[1]]
[1] 1

[[2]]
[1] 2

> clusterCall(cls,function() {return(myid)})
[[1]]
[1] 1

[[2]]
[1] 2
```

Don't forget to stop your clusters before exiting R, by calling **stopCluster(clustername)**.

There are various other useful **snow** functions. See the user's manual for details.

12.5.3 Rdsm

My **Rdsm** package can be used as a threads system regardless of whether you are on a NOW or a multicore machine. It is an extension of a similar package I wrote in 2002 for Perl, called PerlDSM. It gives the R programmer a shared memory view, but the objects are not physically shared. Instead, they are stored in a

⁵I don't see a provision in **snow** itself that does this.

server, and accessed through network sockets, thus enabling a threads-like view for R programmers even on NOWs. There is no manager/worker structure here. All of the R processes execute the same code, as peers.

Shared objects in **Rdsm** can be numerical vectors or matrices, via the classes **dsmv** and **dsmm**, or R lists, using the class **dsm1**. Communication with the server in the vector and matrix cases is done in binary form for efficiency, while serialization is used for lists. There is as a built-in variable **myinfo** that gives a process' ID number and the total number of processes, analogous to the information obtained in **Rmpi** from the functions **mpi.comm.rank()** and **mpi.comm.size()**.

To install, again use **install.packages()** as above. There is built-in documentation, but it's best to read through the code **MatMul.R** in the **examples** directory of the **Rdsm** distribution first. It is heavily commented, with the goal of serving as an introduction to the package.

The example below repeatedly cycles through a large list of Web sites, taking measurements on the time to access each one. The data are stored in a shared variable **acesstimes**; the **n** most recent access times are stored. Each process works on one Web site at a time.

An unusual feature here is that one of the processes immediately exits, returning to the R interactive command line. This allows the user to monitor the data that is being collected. Remember, the shared variables are still accessible to that process. Thus while the other processes are continually adding data to **acesstimes** (and deleted one item for each one added), the user can give commands to the exited process to analyze the data, say with histograms, as the collection progresses.

Note the use of lock/unlock operations here, with the **Rdsm** variables of the same names.

```
# if the variable accesstimes is length n, then the Rdsm vector
# accesstimes stores the n most recent probed access times, with element
# i being the i-th oldest

# arguments:
#   sitefile: IPs, one Web site per line
#   ww: window width, desired length of accesstimes
webprobe <- function(sitefile,ww) {
  if(myinfo$myid == 1) {
    newdsm("acesstimes","dsmv","double",val=rep(0,ww))
  } else {
    newdsm("acesstimes","dsmv","double",size=ww)
  }
  barr()
  # last process is intended simply to provide access to humans, who
  # can do analyses on the data, typing commands, so have it exit this
  # function and return to the R command prompt
  if (myinfo$myid == myinfo$nc1nt) {
    return()
  } else { # the other processes continually probe the Web:
    sites <- scan(sitefile,what="")
    nsites <- length(sites)
    repeat {
      # choose random site to probe
      site <- sites[sample(1:nsites,1)]
      # now probe it
```



```

    acc <- system.time(system(paste("wget --spider -q",site)))[3]
    # add to accesstimes, in sliding-window fashion
    lock("acclock")
    newvec <- c(accesstimes[-1],acc)
    accesstimes[] <- newvec
    unlock("acclock")
  }
}

```

12.5.4 R with GPUs

The blinding speed of GPUs is sure to be of interest to more and more R users in the coming years. As of now, there are two main packages for GPU programming with R, **gputools** and **rgpu**. Both deal mainly with linear algebra operations.

Note that, due to issues involving linking to the CUDA libraries, in the cases of these two packages, you probably will *not* be able to install them by merely calling **install.packages()**. The alternative I recommend works as follows:

- Download the package in **.tar.gz** form.
- Unpack the package, producing a directory that we'll call **x**.
- Modify some files within **x**.
- Then run

```
R CMD INSTALL -l your_installation_directory x
```

12.5.4.1 The gputools Package

In installing **gputools**, I downloaded the source from the CRAN R repository site, and unpacked as above. I then removed the subcommand

```
-gencode arch=compute_20,code=sm_20
```

from the file **Makefile.in** in the **src** directory. I also changed the **LIBS** line in that file to read

```
LIBS := @LIBS@ -L/usr/lib/nvidia
```

in order to pick up the file **libcuda.so** needed for a **-lcuda** link.

The package offers various linear algebra routines, such as matrix multiplication, solution of $Ax = b$ (and thus matrix inversion), and singular value decomposition, as well as some computation-intensive operations such as linear/generalized linear model estimation and hierarchical clustering.

Here for instance is how to find the square of a matrix **m**:

```
> m2 <- gpuMatMult(m,m)
```

12.5.4.2 The rgpu Package

In installing **rgpu**, I downloaded the source code from https://gforge.nbic.nl/frs/?group_id=38 and unpacked as above. I then changed the file **Makefile**, with the modified lines being

```
LIBS = -L/usr/lib/nvidia -lcuda -lcudart -lcublas
CUDA_INC_PATH ?= /home/matloff/NVIDIA_GPU_Computing_SDK/C/common/inc
R_INC_PATH ?= /usr/include/R
```

The first line was needed to pick up **-lcuda**, as with **gputools**. The second line was needed to acquire the file **cutil.h** in the NVIDIA SDK, which I had installed earlier at the location see above.

For the third line, I made a file **z.c** consisting solely of the line

```
#include <R.h>
```

and ran

```
R CMD SHLIB z.c
```

just to see whether the R include file was.

As of May 2010, the routines in **rgpu** are much less extensive than those of **gputools**. However, one very nice feature of **rgpu** is that one can compute matrix expressions without bringing intermediate results back from the device memory to the host memory, which would be a big slowdown. Here for instance is how to compute the square of the matrix **m**, plus itself:

```
> m2m <- evalgpu(m %*% m + m)
```

12.5.4.3 Debugging R Applications

The built-in debugging facilities in R are primitive. However, if you are a Vim editor fan, I've developed a tool that greatly enhances the power of R's debugger. Download **edtdbg** from R's CRAN repository.

REvolution Analytics, a firm that offers R consulting and develops souped-up versions of R, offers an IDE for R that includes nice debugging facilities. At this stage, it is only available on Windows, but they are developing a platform-independent version for Eclipse.

Packages such as **Rmpi**, **snow**, **foreach** and so on do not set up a terminal for each process, thus making it impossible to use R's debugger on the workers. You are then forced to simply print out trace information, e.g. values of variables. Note that you should use **message()** for this purpose, as **print()** won't work in the worker processes.

Appendix A

Review of Matrix Algebra

This book assumes the reader has had a course in linear algebra (or has self-studied it, always the better approach). This appendix is intended as a review of matrix algebra, rather than a detailed treatment of it.

A.1 Terminology and Notation

A **matrix** is a rectangular array of numbers. A **vector** is a matrix with only one row (a **row vector** or only one column (a **column vector**).

The expression, “the (i,j) element of a matrix,” will mean its element in row i, column j.

Please note the following conventions:

- Capital letters, e.g. A and X, will be used to denote matrices and vectors.
- Lower-case letters with subscripts, e.g. $a_{2,15}$ and x_8 , will be used to denote their elements.
- Capital letters with subscripts, e.g. A_{13} , will be used to denote submatrices and subvectors.

If A is a **square** matrix, i.e. one with equal numbers n of rows and columns, then its **diagonal** elements are a_{ii} , $i = 1, \dots, n$.

The **norm** (or **length**) of an n-element vector **X** is

$$\| X \| = \sqrt{\sum_{i=1}^n x_i^2} \quad (\text{A.1})$$

A.1.1 Matrix Addition and Multiplication

- For two matrices have the same numbers of rows and same numbers of columns, addition is defined elementwise, e.g.

$$\begin{pmatrix} 1 & 5 \\ 0 & 3 \\ 4 & 8 \end{pmatrix} + \begin{pmatrix} 6 & 2 \\ 0 & 1 \\ 4 & 0 \end{pmatrix} = \begin{pmatrix} 7 & 7 \\ 0 & 4 \\ 8 & 8 \end{pmatrix} \quad (\text{A.2})$$

- Multiplication of a matrix by a **scalar**, i.e. a number, is also defined elementwise, e.g.

$$0.4 \begin{pmatrix} 7 & 7 \\ 0 & 4 \\ 8 & 8 \end{pmatrix} = \begin{pmatrix} 2.8 & 2.8 \\ 0 & 1.6 \\ 3.2 & 3.2 \end{pmatrix} \quad (\text{A.3})$$

- The **inner product** or **dot product** of equal-length vectors X and Y is defined to be

$$\sum_{k=1}^n x_k y_k \quad (\text{A.4})$$

- The product of matrices A and B is defined if the number of rows of B equals the number of columns of A (A and B are said to be **conformable**). In that case, the (i,j) element of the product C is defined to be

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj} \quad (\text{A.5})$$

For instance,

$$\begin{pmatrix} 7 & 6 \\ 0 & 4 \\ 8 & 8 \end{pmatrix} \begin{pmatrix} 1 & 6 \\ 2 & 4 \end{pmatrix} = \begin{pmatrix} 19 & 66 \\ 8 & 16 \\ 24 & 80 \end{pmatrix} \quad (\text{A.6})$$

It is helpful to visualize c_{ij} as the inner product of row i of A and column j of B, e.g. as shown in bold face here:

$$\begin{pmatrix} \mathbf{7} & \mathbf{6} \\ 0 & 4 \\ 8 & 8 \end{pmatrix} \begin{pmatrix} \mathbf{1} & 6 \\ \mathbf{2} & 4 \end{pmatrix} = \begin{pmatrix} \mathbf{7} & 70 \\ 8 & 16 \\ 8 & 80 \end{pmatrix} \quad (\text{A.7})$$

- Matrix multiplication is associative and distributive, but in general not commutative:

$$A(BC) = (AB)C \quad (\text{A.8})$$

$$A(B + C) = AB + AC \quad (\text{A.9})$$

$$AB \neq BA \quad (\text{A.10})$$

A.2 Matrix Transpose

- The transpose of a matrix A , denoted A' or A^T , is obtained by exchanging the rows and columns of A , e.g.

$$\begin{pmatrix} 7 & 70 \\ 8 & 16 \\ 8 & 80 \end{pmatrix}' = \begin{pmatrix} 7 & 8 & 8 \\ 70 & 16 & 80 \end{pmatrix} \quad (\text{A.11})$$

- If $A + B$ is defined, then

$$(A + B)' = A' + B' \quad (\text{A.12})$$

- If A and B are conformable, then

$$(AB)' = B'A' \quad (\text{A.13})$$

A.3 Linear Independence

Equal-length vectors X_1, \dots, X_k are said to be **linearly independent** if it is impossible for

$$a_1X_1 + \dots + a_kX_k = 0 \quad (\text{A.14})$$

unless all the a_i are 0.

A.4 Determinants

Let A be an $n \times n$ matrix. The definition of the determinant of A , $\det(A)$, involves an abstract formula featuring permutations. It will be omitted here, in favor of the following computational method.

Let $A_{-(i,j)}$ denote the submatrix of A obtained by deleting its i^{th} row and j^{th} column. Then the determinant can be computed recursively across the k^{th} row of A as

$$\det(A) = \sum_{m=1}^n (-1)^{k+m} \det(A_{-(k,m)}) \quad (\text{A.15})$$

where

$$\det \begin{pmatrix} s & t \\ u & v \end{pmatrix} = sv - tu \quad (\text{A.16})$$

A.5 Matrix Inverse

- The **identity** matrix I of size n has 1s in all of its diagonal elements but 0s in all off-diagonal elements. It has the property that $AI = A$ and $IA = A$ whenever those products are defined.
- The A is a square matrix and $AB = I$, then B is said to be the **inverse** of A , denoted A^{-1} . Then $BA = I$ will hold as well.
- A^{-1} exists if and only if its rows (or columns) are linearly independent.
- A^{-1} exists if and only if $\det(A) \neq 0$.
- If A and B are square, conformable and invertible, then AB is also invertible, and

$$(AB)^{-1} = B^{-1}A^{-1} \quad (\text{A.17})$$

A.6 Eigenvalues and Eigenvectors

Let A be a square matrix.¹

¹For nonsquare matrices, the discussion here would generalize to the topic of **singular value decomposition**.

- A scalar λ and a nonzero vector X that satisfy

$$AX = \lambda X \quad (\text{A.18})$$

are called an **eigenvalue** and **eigenvector** of A , respectively.

- A matrix U is said to be **orthogonal** if its rows have norm 1 and are orthogonal to each other, i.e. their inner product is 0. U thus has the property that $UU' = I$ i.e. $U^{-1} = U'$.
- If A is symmetric and real, then it is **diagonalizable**, i.e. there exists an orthogonal matrix U such that

$$U'AU = D \quad (\text{A.19})$$

for a diagonal matrix D . The elements of D are the eigenvalues of A , and the columns of U are the eigenvectors of A .