Chapter 2

"Why Is My Program So Slow?": Obstacles to Speed

Here is an all-too-common scenario: An analyst acquires a brand new multicore machine, capable of wondrous things. With great excitement, he codes up his favorite large problem on the new machine—only to find that the parallel version runs more slowly than the serial one. What a disappointment!

Though you are no doubt eager to get to some more code, a firm grounding in the infrastructural issues will prove to be quite valuable indeed, hence the need for this chapter. These issues will arise repeatedly in the rest of the book. If you wish, you could skip ahead to the other chapters now, and come back to this one as the need arises, but it's better if you go through it now. So, let's see what factors can lead to such a situation in which our hapless analyst above sees his wonderful plans go awry.

2.1 Obstacles to Speed

Let's refer to the computational entities as *processes*, such as the workers in the case of **snow**. There are two main performance issues in parallel programming:

- *Communications overhead:* Typically data must be transferred back and forth between processes. This takes time, which can take quite a toll on performance.

  In addition, the processes can get in each other's way if they all try to access the same data at once. They can collide when trying to access the same communications channel, the same memory module, and so on. This is another sap on speed.

  The term *granularity* is used to refer, roughly, to the ratio of computation to overhead. *Large-grained* or *coarse-grained* algorithms involve large enough chunks of computation that the overhead isn't much of a problem. In *fine-grained* algorithms, we really need to avoid overhead as much as possible.

- *Load balance:* As noted in the last chapter, if we are not careful in the way in which we assign work to processes, we risk assigning much more work to some than to others. This compromises performance, as it leaves some processes unproductive at the end of the run, while there is still work to be done.

There are a number of issues of this sort that occur generally enough to be collected into this chapter, as an "early warning" of issues that can arise. This is just an overview, with details coming in subsequent chapters, but being forewarned of the problems will make it easier to recognize them as they are encountered.

2.2 Performance and Hardware Structures

Scorecards, scorecards! You can't tell the players without the scorecards!— old chant of scorecard vendors at baseball games

The foot bone connected to the ankle bone, The ankle bone connected to the shin bone...— from the children's song, "Dem Bones"

The reason our unfortunate analyst in the preceding section was surprised that his code ran more slowly on the parallel machine was almost certainly due to a lack of understanding of the underlying hardware and systems software. While one certainly need not understand the hardware on an electronics level, a basic knowledge of "what is connected to what" is essential.

In this section, we'll present overviews of the major hardware issues, and of the two parallel hardware technologies the reader is mostly likely to

encounter, *multiprocessors* and *clusters*:[1]

- A multiprocessor system has, as the name implies, two or more processors, i.e., two or more CPUs, so that two or more programs (or parts of the same program) can be doing computation at the same time. A *multicore* system, common in the home, is essentially a low-end multiprocessor, as we will see later. Multiprocessors are also known as *shared-memory* systems, since they indeed share the same physical RAM.

    These days, almost any home PC or laptop is at least dual core. If you own such a machine, congratulations, you own a multiprocessor system!

    You are also to be congratulated for owning a multiprocessor system if you have a fairly sophisticated video card in your computer, one that can serve as a *graphics processing unit*. GPUs are specialized shared-memory systems.

- A cluster consists of multiple computers, each capable of running independently, that are networked together, enabling their engaging in a concerted effort to solve a big numerical problem.

    If you have a network at home, say with a wireless or wired router, then congratulations, you own a cluster![2]

I emphasize the "household item" aspect above, to stress that these are not esoteric architectures, though of course scale can vary widely from what you have at home to far more sophisticated and expensive systems, with quite a bit in between.

The terms *shared-memory* and *networked* above give clues as to the obstacles to computational speed that arise, which are key. So, we will first discuss the high-level workings of these two hardware structures, in Sections 2.3 and 2.4.

We'll then explain how they apply to the overhead issue with our two basic platform types, multicore (Section 2.5.1.1) and cluster (Section 2.5.1.2). We'll cover just enough details to illustrate the performance issues discussed later in this chapter, and return for further details in later chapters.

---

[1]What about clouds? A cloud consists of multicore machines and clusters too, but operating behind the scenes.

[2]It should be noted that in the case of large clusters used for intensive computation, one generally must install software for the purpose of controlling which program runs on which machines. But your two-node home system is still a cluster.

2.3   Memory Basics

Slowness of memory access is one of the most common issues arising in high-performance computing. Thus a basic understanding of memory is vital.

Consider an ordinary assignment statement, copying one variable (a single integer, say) to another:

y = x

Typically, both **x** and **y** will be stored somewhere in memory, i.e., RAM (Random Access Memory). Memory is broken down into *bytes*, designed to hold one character, and *words*, usually designed to contain one number. A byte consists of eight bits, i.e., eight 0s and 1s. On typical computers today, the word size is 64 bits, or eight bytes.

Each word has an ID number, called an *address*. (Individual bytes have addresses too, but this will not concern us here.) So the compiler (in the case of C/C++/FORTRAN) or the interpreter (in the case of a language like R), will assign specific addresses in memory at which **x** and **y** are to be stored. The above assignment will be executed by the machine's copying one word to the other.

A vector will typically be stored in a set of consecutive words. This will be the case for matrices too, but there is a question as to whether this storage will be row-by-row or column-by-column. C/C++ uses *row-major order*: First all of the first row (called row 0) is stored, then all of the second row, and so on. R and FORTRAN use *column-major order*, storing all of the first column (named column 1) etc. So, for instance, if **z** is a 5×8 matrix in R, then **z[2,3]** will be in the 12^{th} word (5+5+2) in the portion of memory occupied by **z**. These considerations will affect performance, as we will see later.

Memory access time, even though measured in tens of nanoseconds—billionths of a second—is slow relative to CPU speeds. This is due not only to electronic delays within the memory chips themselves, but also due to the fact that the pathway to memory is often a bottleneck. More on this below.

2.3.1   Caches

A device commonly used to deal with slow memory access is a *cache*. This is a small but fast chunk of memory that is located on or near the processor

chip. For this purpose, memory is divided into *blocks*, say of 64 bytes each. Memory address 1200, for instance, would be in block 18, since 1200/64 is equal to 18 plus a fraction. (The first block is called Block 0.)

The cache is divided into *lines*, each the size of a memory block. At any given time, the cache contains local copies of some blocks of memory, with the specific choice of blocks being dynamic—at some times the cache will contain copies of some memory blocks, while a bit later it may contain copies of some other blocks.[3]

If we are lucky, in most cases, the memory word that the processor wishes to access (i.e., the variable in the programmer's code she wishes to access) already has a copy in its cache—a *cache hit*. If this is a read access (of **x** in our little example above), then it's great—we avoid the slow memory access.

On the other hand, in the case of a write access (to **y** above), if the requested word is currently in the cache, that's nice too, as it saves us the long trip to memory (if we do not "write through" and update memory right away, as we are assuming here). But it does produce a discrepancy between the given word in memory and its copy in the cache. In the cache architecture we are discussing here, that discrepancy is tolerated, and eventually resolved when the block in question is "evicted," as we will see below. (With a multicore machine, cache operation becomes more complicated, as typically each core will have its own cache, thus potentially causing severe discrepancies. This will be discussed in Section 2.5.1.1.)

If in a read or write access the desired memory word is not currently in the cache, this is termed a *cache miss*. This is fairly expensive. When it occurs, the entire block containing the requested word must be brought into the cache. In other words, we must access many words of memory, not just one. Moreover, usually a block currently in the cache must be *evicted* to make room for the new one being brought in. If the old block had been written to at all, we must now write that entire block back to memory, to update the latter.[4]

So, though we save memory access time when we have a cache hit, we incur a substantial penalty at a miss. Good cache design can make it so that the penalty is incurred only rarely. When a read miss occurs, the hardware makes "educated guesses" as to which blocks are least likely to be needed again in the near future, and evicts one of these. It usually guesses well, so

---

[3]What follows below is a description of a common cache design. There are many variations, not discussed here.

[4]There is a *dirty bit* that records whether we've written to the block, but not which particular words were affected. Thus the entire block must be written.

that cache hit rates are typically well above 90%. Note carefully, though, that this can be affected by the way we code. This will be discussed in future chapters.

A machine will typically have two or more levels of cache. The one in or next to the CPU is called the L1, or Level 1 cache. Then there may be an L2 cache, a "cache for the cache." If the desired item is not found in the L1 cache, the CPU will then search the L2 cache before resorting to accessing the item in memory.

2.3.2 Virtual Memory

Though it won't arise much in our context, we should at least briefly discuss *virtual memory*. Consider our example above, in which our program contained variables **x** and **y**. Say these are assigned to addresses 200 and 8888, respectively. Fine, but what if another program is also running on the machine? The compiler/interpreter may have assigned one of its variables, say **g**, to address 200. How do we resolve this?

The standard solution is to make the address 200 (and all others) only "virtual." It may be, for instance, that **x** from the first program is actually stored in physical address 7260. The program will still say **x** is at word 200, but the hardware will translate 200 to 7260 as the program executes. If **g** in the second program is actually in word 6548, the hardware will replace 200 by 6548 every time the program requests access to word 200. The hardware has a table to do these lookups, one table for each program currently running on the machine, with the table being maintained by the operating system.

Virtual memory systems break memory into *pages*, say of 4096 bytes each, analogous to cache blocks. Usually, only some of your program's pages are *resident* in memory at any given time, with the remainder of the pages out on disk. If your program needs some memory word not currently resident— a *page fault*, analogous to a cache miss—the hardware senses this, and transfers control to the operating system. The OS must bring in the requested page from disk, an extremely expensive operation in terms of time, due to the fact that a disk drive is mechanical rather than electronic like RAM.[5] Thus page faults can really slow down program speed, and again as with the cache case, you may be able to reduce page faults through careful design of your code.

---

[5]Some more expensive drives, known as Solid State Drives (SSDs), are in fact electronic.

2.3.3   Monitoring Cache Misses and Page Faults

Both cache misses and page faults are enemies of good performance, so it would be nice to monitor them.

This actually can be done in the case of page faults. As noted, a page fault triggers a jump to the OS, which can thus record it. In Unix-family systems, the **time** command gives not only run time but also a count of page faults.

By contrast, cache misses are handled purely in hardware, thus not recordable by the OS. But one might try to gauge the cache behavior of a program by using the number of page faults as a proxy. There are also simulators, such as **valgrind**, which can be used to measure cache performance.

2.3.4   Locality of Reference

Clearly, the effectiveness of caches and virtual memory depend on repeatedly using items in the same blocks (*spatial locality*) within short time periods (*temporal locality*). As mentioned earlier, this in turn can be affected to some degree by the way the programmer codes things.

Say we wish to find the sum of all elements in a matrix. Should our code traverse the matrix row-by-row or column-by-column? In R, for instance, which as mentioned stores matrices in column-major order, we should go column-by-column, to get better locality.

A detailed case study on cache behavior will be presented in Section 5.8.

2.4   Network Basics

A single Ethernet (or other similar system), say within a building, is called a *network*. The *Internet* is simply the interconnection of many networks–millions of them.

Say you direct the browser on your computer to go to the Cable Network News (CNN) home page, and you are located in San Francisco. Since CNN is headquartered in Atlanta, *packets* of information will go from San Francisco to Atlanta. (Actually, they may not go that far, since Internet service providers (ISPs) often cache Web pages, but let's suppose that doesn't occur.) Actually, a packet's journey will be rather complicated:

- Your browser program will write your Web request to a *socket.* The latter is not a physical object, but rather a software interface from your program to the network.

- The socket software will form a packet from your request, which will then go through several layers of the *network protocol stack* in your OS. Along the way, the packet will grow, as more information is being added, but also it will split into multiple, smaller packets.

- Eventually the packets will reach your computer's network interface hardware, from which they go onto the network.

- A *gateway* on the network will notice that the ultimate destination is external to this network, so the packets will be transferred to another network that the gateway is also attached to.

- Your packets will wend their way across the country, being sent from one network to the next.[6]

- When your packets reach a CNN computer, they will now work their way *up* the levels of the OS, finally reaching the Web server program.

2.5   Latency and Bandwidth

Getting there is half the fun—old saying, regarding the pleasures of traveling

The speed of a communications channel—whether between processor cores and memory in shared-memory platforms, or between network nodes in a cluster of machines—is measured in terms of *latency*, the end-to-end travel time for a single bit, and *bandwidth*, the number of bits per second that we can pump onto the channel.

To make the notions a little more concrete, consider the San Francisco Bay Bridge, a long, multilane structure for which westbound drivers pay a toll. The notion of latency would describe the time it takes for a car to drive from one end of the bridge to the other. (For simplicity, assume they all go the same speed.) By contrast, the bandwidth would be the number of cars exiting from the toll booths per unit time. We can reduce the latency by raising the speed limit on the bridge, while we could increase the bandwidth by adding more lanes and more toll booths.

---

[6]Run the **traceroute** command on your machine to see the exact path, though this can change over time.

The network time in seconds to send an n-byte message, with a latency of l seconds and a bandwidth of b bytes/second, is clearly

$$l + n/b \tag{2.1}$$

Of course, this assumes that there are no other messages contending for the communication channel.

Clearly there are numerous delays in networks, including the less-obvious ones incurred in traversing the layers of the OS. Such traversal involves copying the packet from layer to layer, and in cases of interest in this book, such copying can involve huge matrices and thus take a lot of time.

Though parallel computation is typically done within a network rather than across networks as above, many of those delays are still there. So, network speeds are much, much lower than processor speeds, both in terms of latency and bandwidth.

The latency in even a fast network such as Infiniband is on the order of microseconds, i.e., millionths of a second, which is eons compared to the nanosecond level of execution time for a machine instruction in a processor. (Beware of a network that is said to be fast but turns out only to have high bandwidth, not also low latency.)

Latency and bandwidth issues arise in shared-memory systems too. Consider GPUs, for instance. In most applications, there is a lot of data transfer between the CPU and the GPU, with attendant potential for slowdown. Latency, for example, is the time for a single bit to go from the CPU to the GPU, or vice versa.

One way to ameliorate the slowdown from long latency delays is *latency hiding*. The basic idea is to try to do other useful work while a communication having long latency is pending. This approach is used, for instance, in the use of nonblocking I/O in message-passing systems (Section 8.7.1) to deal with network latency, and in GPUs (Chapter 6) to deal with memory latency.

2.5.1 Two Representative Hardware Platforms: Multicore Machines and Clusters

Multicore machines have become standard on the desktop (even in the cell phone!), and many data scientists have access to computer clusters. What are the performance issues on these platforms? The next two sections

bus

Figure 2.1: Symmetric Multiprocsssor System

provide an overview.

2.5.1.1   Multicore

A *symmetric multiprocessor system* looks something like Figure 2.1 in terms of components and, most importantly, their interconnection. What do we see?

- There are *processors*, depicted by the Ps, in which your program is physically executed.

- There are *memory banks*, the Ms, in which your program and data reside during execution.[7]

- The processors and memory banks are connected to a *bus*, a set of parallel wires used for communication between these computer components.

Your input/output hardware—disk drives, keyboards and so on—are also connected to the bus, and there may actually be more than one bus, but our focus will be mainly on the processors and memory.

A *threaded* program will have several instantiations of itself, called *threads*, that are working in concert to achieve parallelism. They run independently, except that they share the data of the program in common. If your program is threaded, it will be running on several of the processors at once, each thread on a different core. A key point, as we will see, is that the shared

---

[7]These were called *banks* in the old days. Later the term *modules* became more popular, but with the recent popularity of GPUs, the word *banks* has come back into favor.

memory becomes the vehicle for communication between the various processes.

Your program consists of a number of machine language instructions. (If you write in an interpreted language such as R, the interpreter itself consists of such instructions.) As the processors execute your program, they will fetch the instructions from memory.

As noted earlier, your data—the variables in your program—is stored in memory. The machine instructions fetch the data from memory as needed, so that it can be processed, e.g., summed, in the processors.

Until recently, ordinary PCs sold at your local electronics store followed the model in Figure 2.1 but with only one P. Multiprocessor systems enabled parallel computation, but cost hundreds of thousands of dollars. But then it became standard for systems to have a *multicore* form. This means that there are multiple Ps, but with the important distinction that they are all on a single chip (each P is one core), making for inexpensive systems.[8] Whether on a single chip or not, having multiple Ps sets up parallel computation, and is known as the *shared-memory* paradigm, for obvious reasons.

By the way, why are there multiple Ms in Figure 2.1? To improve memory performance, the system is set up so that memory is partitioned into several banks (typically there are the same number of Ms as Ps). This enables us to not only do *computation* on a parallel basis—several Ps working on different pieces of a problem in parallel—but also to do *memory access* in parallel—several memory accesses being active in parallel, in different banks. This amortizes the memory access penalty. Of course, if more than one P happens to need to access the same M at about the same time, we lose this parallelism.

As you can see, a potential bottleneck is the bus. When more than one P needs to access memory at a time, even if to different banks, attempting to place memory access requests on the bus, all but one of them will need to wait. This *bus contention* can cause significant slowdown. Much more elaborate systems, featuring multiple communications channels to memory rather than just a bus, have also been developed and serve to ameliorate the bottleneck issue. Most readers of this book, however, are more likely to use a multicore system on a single memory bus.

You can see now why efficient memory access is such a crucial factor in achieving high performance. There is one more tool to handle this that is

---

[8]Terminology is not standardized, unfortunately. It is common to refer to that chip as "the" processor, even though there actually are multiple processors inside.

vital to discuss here: Use of caches. Note the plural; in Figure 2.1, there is usually a C in between each P and the bus.

As with uniprocessor systems, caching can bring a big win in performance. In fact, the potential is even greater with a multiprocessor system, since caching will now bring the additional benefit of reducing bus contention. Unfortunately, it also produces a new problem, *cache coherency*, as follows.[9]

Consider what happens upon a *write hit*, i.e., a write to a location for which a local cache copy exists. For instance, consider code such as

x = 28;

with **x** having address 200. This code might be executed at a time when there is a copy of word 200 in that processor's cache. The problem is that other caches may also have a copy of this word, so they are now invalid for that block. (Recall that validity is defined only at the block level; if all words in a block but one are valid, the whole block is considered invalid.) The hardware must now inform those other caches that their copies of this block are invalid.

The hardware does so via the bus, thus incurring an expensive bus operation. Moreover, the next time this word (or for that matter, any word in this block) is requested at one of the other caches, there will be a cache miss, again an expensive event.

Once again, proper coding on the programmer's part can sometimes ameliorate the cache coherency problem.

A final point on multicore structure: Even on a uniprocessor machine, one generally has multiple programs running concurrently. You might have your browser busy downloading a file, say, while at the same time you are using a photo processing application. With just a single processor, these programs will actually take turns running; each one will run for a short time, say 50 milliseconds, then hand off the processor to the next program, in a cyclic manner. (You as the user probably won't be aware of this directly, but you may notice the system as a whole slowing down.) Note by the way that if a program is doing a lot of input/output (e.g., file access), it is effectively idle during I/O times; as soon as it starts an I/O operation, it will relinquish the processor.

By contrast, on a multicore machine, you can have multiple programs running physically simultaneously (though of course they will still take turns if there are more of them than there are cores).

---

[9]As noted earlier, there are variations of the structure described here, but this one is typical.

2.5.1.2 Clusters

These are much simpler to describe, though with equally thorny performance obstacles.

The term *cluster* simply refers to a set of independent *processing elements* (PEs) or *nodes* that are connected by a local area network, such as the common Ethernet or the high-performance Infiniband. Each PE consists of a CPU and some RAM. The PE could be a full desktop computer, including keyboard, disk drive and monitor, but if it is used primarily for parallel computation, then just one monitor, keyboard and so on suffice for the entire system. A cluster may also have a special operating system, to coordinate assigning of user programs to PEs.

We will have one computational process per PE (unless each PE is a multi-core system, as is common). Communication between the processes occurs via the network. The latter aspect, of course, is where the major problems occur.

2.5.2 The Principle of "Just Leave It There"

All of these considerations regarding latency and bandwidth means, among other things, that data copying is often the enemy of speed. This is particularly true for high-latency platforms such as clusters and GPUs.

In such a situation, it may be crucial to design the algorithm to minimize copying. With an iterative algorithm, for instance, be sure to leave intermediate results on the remote nodes in the cluster case, and in the GPU memory in that setting, if possible.

2.6 Thread Scheduling

Say you have a threaded program, for example with four threads and a machine with four cores. Then the four threads will run physically simultaneously (if there are no other programs competing with them). That of course is the entire point, to achieve parallelism, but there is more to it than that.

Modern operating systems for general-purpose computers use *timesharing*: Unseen by the users, programs are taking turns (*timeslices*) using the cores of the machine. Say for instance that Manny and Moe are using a university

computer named Jack, with Manny sitting at the console and Moe logged in remotely. For concreteness, say Manny is running an R program, and Moe is running something in Python, with both currently having long-running computations in progress.

Assume first that this is a single-core machine. Then only one program can run at a time. Manny's program will run for a while, but after a set amount of time, the hardware timer will issue an *interrupt*, causing a jump to another program. That program has been configured to be the operating system. The OS will look on its *process table* to find another program in *ready* state, meaning runnable (as opposed to say, suspended while awaiting keyboard input). Assuming there are no other processes, Moe's program will now get a turn. This transition from Manny to Moe is called a *context switch*. Moe's program will run for a while, then another interrupt comes, and Manny will get another turn, and so on.

Now suppose it is a dual-core machine. Here Manny and Moe's programs will run more or less continuously, in parallel, though with periodic downtimes due to the interrupts and attendant brief OS runs.

But suppose Moe's code is threaded, running two threads. Now we will have three threads—Moe's two and Manny's one (even a non-threaded program consists of one thread)—competing to use three cores. Moe's two threads will sometimes run in parallel with each other but sometimes not. Instead of a 2X speedup, Moe is getting about 1.5X.[10]

There are also possible cache issues. When a thread starts a new turn, it may be on a different core than that used in the last turn. If there is a separate cache for each core, the cache as the new core probably contains little if anything useful to this thread. Thus there will be a lot of cache misses for a while in this timeslice. There may be a remedy in the form of setting *processor affinity*; see Section 5.9.4.

By the way, what happens when one of those programs finishes its computation and returns to the user prompt, e.g., > in the case of Manny's R program? R will then be waiting for Manny's keyboard input. But the OS won't wait, and the OS does in fact get involved. R is trying to read from the keyboard, and to do this it calls a C library function, which in turn makes a call to a function in the OS. The OS, realizing that it may be quite a while before Manny types, will mark his entry in the process table as being in *sleep* state. When he finally does hit a key, the keyboard sends an interrupt,[11] causing the OS to run, and the latter will mark his program as now being back in ready state, and it will eventually get another turn.

---

[10]Even the 2X figure assumes that Moe's code was load balanced in the first place, which may not be the case.

[11]In Moe's case, the interrupt will come from Jack's network card.

2.7 How Many Processes/Threads?

As mentioned earlier, it is customary in the R world to refer to each worker in a **snow** program as a process. A question that then arises is, how many processes should we run?

Say for instance we have a cluster of 16 nodes. Should we set up 16 workers for our **snow** program? The same issues arise with threaded programs, say with **Rdsm** or OpenMP (Chapters 4 and 5). On a quadcore machine, should we run 4 threads?

The answer is *not* automatically Yes to these questions. With a fine-grained program, using too many processes/threads may actually degrade performance, as the overhead may overwhelm the presumed advantage of throwing more hardware at the problem. So, one might actually use fewer cluster nodes or fewer cores than one has available.

On the other hand, one might try to *oversubscribe* the resources. As discussed earlier, a cache miss causes a considerable delay, and a page fault even more. This is time during which one of the nodes/cores will not be doing any computation, exacting an opportunity cost from performance. It may pay, then, to have "extra" threads for the program available to run.

2.8 Example: Mutual Outlink Problem

To make this concrete, let's measure times for the mutual outlinks problem (Section 1.4), with larger and larger numbers of processes.

Here I ran on a shared-memory machine consisting of four processor chips, each of which has eight cores. This gives us a 32-core system, and I ran the mutual outlinks problem with values of **nc**, the number of cores, equal to 2, 4, 6, 8, 10, 12, 16, 24, 28 and 32. The problem size was 1000 rows by 1000 columns. The times are plotted in Figure 2.2.

Here we see a classical U-shaped pattern: As we throw more and more processes on the problem, it helps in the early stages, but performance actually degrades after a certain point. The latter phenomenon is probably due to the communications overhead we discussed earlier, in this case bus contention and the like.[12]

[12]Though the processes are independent and do not share memory, they do share the bus.

Figure 2.2: Run Time Versus Number of Cores

By the way, for each of our **nc** workers, we had one invocation of R running on the machine. There was also an additional invocation, for the manager. However, this is not a performance issue in this case, as the manager spends most of its time idle, waiting for the workers.

2.9   "Big O" Notation

With all this talk of physical obstacles to overcome, such as memory access time, it's important also to raise the question as to whether the application itself is very parallelizable in the first place. One measure of that is "big O" notation.

In our mutual outlinks example with an $n \times n$ adjacency matrix, we need to do on average $n/2$ sum operations per row, with n rows, thus $n \cdot n/2$ operations in all. In parallel processing circles, the key question asked about hardware, software, algorithms and so on is, "Does it scale?", meaning, Does the run time grow manageably as the problem size grows?

We see above that the run time of the mutual outlinks problem grows proportionally to the *square* of the problem size, in this case the number of websites. (Dividing by 2 doesn't affect this growth rate.) We write this as $O(n^2)$, known colloquially as "big O" notation. When applied to analysis of run time, we say that it measures the *time complexity*.

Ironically, applications that *are* manageable often are poor candidates for parallel processing, due to overhead playing a greater role in such problems. An application with $O(n)$ time complexity, for instance, may present a challenge. We will return to this notion at various points in this book.

2.10 Data Serialization

Some parallel R packages, e.g., **snow**, that send data through a network *serialize* the data, meaning to convert it to ASCII form. The data must then be unserialized on the receiving end. This creates a delay, which may or may not be serious but must be taken into consideration.

2.11 "Embarrassingly Parallel" Applications

The term *embarrassingly parallel* is heard often in talk about parallel programming. It is a central topic, hence deserving of having a separate section devoted to it.

2.11.1 What People Mean by "Embarrassingly Parallel"

It's no shame to be poor...but it's no great honor either—the character Tevye in *Fiddler on the Roof*

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 to 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 problem is too easy, i.e., 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, less obvious 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.

In recent years, the term *embarrassingly parallel* has drifted to a somewhat different meaning. Algorithms that are embarrassingly parallel in the above sense of simplicity tend to have very low communication between processes, key to good performance. That latter trait is the center of attention nowadays, so the term *embarrassingly parallel* generally refers to an algorithm with low communication needs.

2.11.2   Suitable Platforms for Non-Embarrassingly Parallel Applications

The only general-purpose parallel computing platform suitable for non-embarrassingly parallel applications is that of the multicore/multiprocessor system. This is due to the fact that processor/memory copies have the least communication overhead. Note carefully that this does not mean there is NO overhead—if a cache coherency transaction occurs, we pay a heavy price. But at least the "base" overhead is small.

Still, non-embarrassingly parallel problems are generally tough nuts to crack. A good, commonplace example is linear regression analysis. Here a matrix inversion or equivalent such as QR factorization, is tough to parallelize. We'll return to this issue frequently in this book.

2.12   Further Reading

For more information on memory addresses and structures, timesharing operating systems and so on, see my open source book on computer systems, *An Introduction to Computer Systems*, N. Matloff, `http://heather.cs.ucdavis.edu/~matloff/50/PLN/CompSystsBook.pdf`.