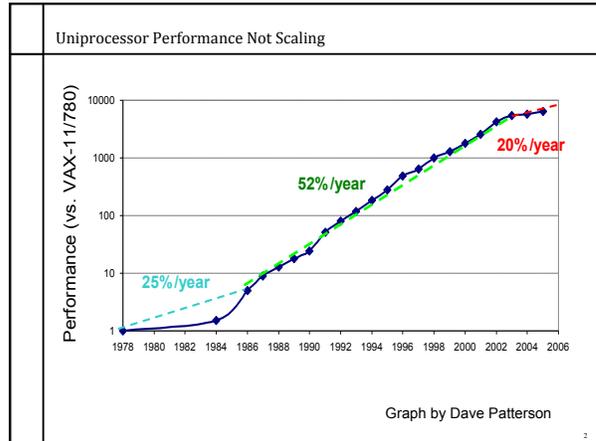
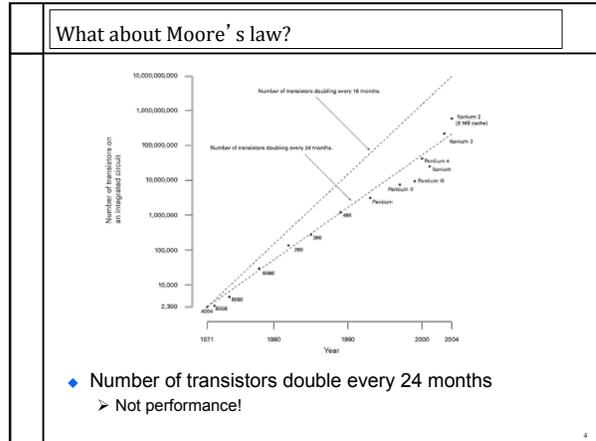


*Concurrent Programming:  
 Why you should care, deeply*  
  
*Don Porter*  
*Portions courtesy Emmett Witchel*

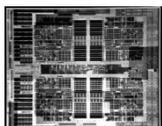
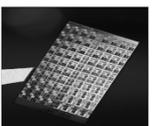


- Power and heat lay waste to processor makers
- ◆ Intel P4 (2000-2007)
    - 1.3GHz to 3.8GHz, 31 stage pipeline
    - “Prescott” in 02/04 was too hot. Needed 5.2GHz to beat 2.6GHz Athalon
  - ◆ Intel Pentium Core, (2006-)
    - 1.06GHz to 3GHz, 14 stage pipeline
    - Based on mobile (Pentium M) micro-architecture
      - ❖ Power efficient
  - ◆ 2% of electricity in the U.S. feeds computers
    - Doubled in last 5 years



- Architectural trends that favor multicore
- ◆ Power is a first class design constraint
    - Performance per watt the important metric
  - ◆ Leakage power significant with small transistors
    - Chip dissipates power even when idle!
  - ◆ Small transistors fail more frequently
    - Lower yield, or CPUs that fail?
  - ◆ Wires are slow
    - Light in vacuum can travel ~1m in 1 cycle at 3GHz
    - Motivates multicore designs (simpler, lower-power cores)
  - ◆ Quantum effects
  - ◆ Motivates multicore designs (simpler, lower-power cores)

**Multicores are here, and coming fast!**

4 cores in 2007	16 cores in 2009	80 cores in 20??
		
AMD Quad Core	Sun Rock	Intel TeraFLOP

“[AMD] quad-core processors ... are just the beginning....”  
<http://www.amd.com>  
 “Intel has more than 15 multi-core related projects underway”  
<http://www.intel.com>

Multicore programming will be in demand	
	<ul style="list-style-type: none"> <li>◆ Hardware manufacturers betting big on multicore</li> <li>◆ Software developers are needed</li> <li>◆ Writing concurrent programs is not easy</li> <li>◆ You will learn how to do it in this class</li> </ul>

Concurrency Problem									
	<ul style="list-style-type: none"> <li>◆ Order of thread execution is non-deterministic <ul style="list-style-type: none"> <li>➢ Multiprocessing <ul style="list-style-type: none"> <li>❖ A system may contain multiple processors → cooperating threads/processes can execute simultaneously</li> </ul> </li> <li>➢ Multi-programming <ul style="list-style-type: none"> <li>❖ Thread/process execution can be interleaved because of time-slicing</li> </ul> </li> </ul> </li> <li>◆ Operations often consist of multiple, visible steps <ul style="list-style-type: none"> <li>➢ Example: <math>x = x + 1</math> is not a single operation <table style="margin-left: 20px;"> <tr> <td>❖ read x from memory into a register</td> <td>Thread 2</td> </tr> <tr> <td>❖ increment register</td> <td>read</td> </tr> <tr> <td>❖ store register back to memory</td> <td>increment</td> </tr> <tr> <td></td> <td>store</td> </tr> </table> </li> </ul> </li> <li>◆ Goal: <ul style="list-style-type: none"> <li>➢ Ensure that your concurrent program works under ALL possible interleaving</li> </ul> </li> </ul>	❖ read x from memory into a register	Thread 2	❖ increment register	read	❖ store register back to memory	increment		store
❖ read x from memory into a register	Thread 2								
❖ increment register	read								
❖ store register back to memory	increment								
	store								

Questions	
	<ul style="list-style-type: none"> <li>◆ Do the following either completely succeed or completely fail?</li> <li>◆ Writing an 8-bit byte to memory <ul style="list-style-type: none"> <li>➢ A. Yes B. No</li> </ul> </li> <li>◆ Creating a file <ul style="list-style-type: none"> <li>➢ A. Yes B. No</li> </ul> </li> <li>◆ Writing a 512-byte disk sector <ul style="list-style-type: none"> <li>➢ A. Yes B. No</li> </ul> </li> </ul>

Sharing among threads increases performance...	
	<pre> int a = 1, b = 2; main() {     CreateThread(fn1, 4);     CreateThread(fn2, 5); } fn1(int arg1) {     if(a) b++; } fn2(int arg1) {     a = arg1; } </pre> <p style="text-align: right;">What are the values of a &amp; b at the end of execution?</p>

Sharing among threads increases performance, but can lead to problems!!	
	<pre> int a = 1, b = 2; main() {     CreateThread(fn1, 4);     CreateThread(fn2, 5); } fn1(int arg1) {     if(a) b++; } fn2(int arg1) {     a = 0; } </pre> <p style="text-align: right;">What are the values of a &amp; b at the end of execution?</p>

Some More Examples	
	<ul style="list-style-type: none"> <li>◆ What are the possible values of x in these cases?</li> </ul> <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Thread1: <math>x = 1</math>;                      Thread2: <math>x = 2</math>;</p> </div> <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Initially <math>y = 10</math>; Thread1: <math>x = y + 1</math>;                      Thread2: <math>y = y * 2</math>;</p> </div> <div style="border: 1px solid black; padding: 5px;"> <p>Initially <math>x = 0</math>; Thread1: <math>x = x + 1</math>;                      Thread2: <math>x = x + 2</math>;</p> </div>

Critical Sections	
<ul style="list-style-type: none"> <li>◆ A critical section is an abstraction               <ul style="list-style-type: none"> <li>➢ Consists of a number of consecutive program instructions</li> <li>➢ Usually, crit sec are mutually exclusive and can wait/signal                   <ul style="list-style-type: none"> <li>✦ Later, we will talk about atomicity and isolation</li> </ul> </li> </ul> </li> <li>◆ Critical sections are used frequently in an OS to protect data structures (e.g., queues, shared variables, lists, ...)</li> <li>◆ A critical section implementation must be:               <ul style="list-style-type: none"> <li>➢ Correct: the system behaves as if only 1 thread can execute in the critical section at any given time</li> <li>➢ Efficient: getting into and out of critical section must be fast. Critical sections should be as short as possible.</li> <li>➢ Concurrency control: a good implementation allows maximum concurrency while preserving correctness</li> <li>➢ Flexible: a good implementation must have as few restrictions as practically possible</li> </ul> </li> </ul>	13

The Need For Mutual Exclusion	
<ul style="list-style-type: none"> <li>◆ Running multiple processes/threads in parallel increases performance</li> <li>◆ Some computer resources cannot be accessed by multiple threads at the same time               <ul style="list-style-type: none"> <li>➢ E.g., a printer can't print two documents at once</li> </ul> </li> <li>◆ Mutual exclusion is the term to indicate that some resource can only be used by one thread at a time               <ul style="list-style-type: none"> <li>➢ Active thread excludes its peers</li> </ul> </li> <li>◆ For shared memory architectures, data structures are often mutually exclusive               <ul style="list-style-type: none"> <li>➢ Two threads adding to a linked list can corrupt the list</li> </ul> </li> </ul>	14

Exclusion Problems, Real Life Example	
<ul style="list-style-type: none"> <li>◆ Imagine multiple chefs in the same kitchen               <ul style="list-style-type: none"> <li>➢ Each chef follows a different recipe</li> </ul> </li> <li>◆ Chef 1               <ul style="list-style-type: none"> <li>➢ Grab butter, grab salt, do other stuff</li> </ul> </li> <li>◆ Chef 2               <ul style="list-style-type: none"> <li>➢ Grab salt, grab butter, do other stuff</li> </ul> </li> <li>◆ What if Chef 1 grabs the butter and Chef 2 grabs the salt?               <ul style="list-style-type: none"> <li>➢ Yell at each other (not a computer science solution)</li> <li>➢ Chef 1 grabs salt from Chef 2 (preempt resource)</li> <li>➢ Chefs all grab ingredients in the same order                   <ul style="list-style-type: none"> <li>✦ Current best solution, but difficult as recipes get complex</li> <li>✦ Ingredient like cheese might be sans refrigeration for a while</li> </ul> </li> </ul> </li> </ul>	15

The Need To Wait	
<ul style="list-style-type: none"> <li>◆ Very often, synchronization consists of one thread waiting for another to make a condition true               <ul style="list-style-type: none"> <li>➢ Master tells worker a request has arrived</li> <li>➢ Cleaning thread waits until all lanes are colored</li> </ul> </li> <li>◆ Until condition is true, thread can sleep               <ul style="list-style-type: none"> <li>➢ Ties synchronization to scheduling</li> </ul> </li> <li>◆ Mutual exclusion for data structure               <ul style="list-style-type: none"> <li>➢ Code can wait (await)</li> <li>➢ Another thread signals (notify)</li> </ul> </li> </ul>	16

Example 2: Traverse a singly-linked list	
<ul style="list-style-type: none"> <li>◆ Suppose we want to find an element in a singly linked list, and move it to the head</li> <li>◆ Visual intuition:               <div style="text-align: center; margin-top: 10px;"> </div> </li> </ul>	17

Example 2: Traverse a singly-linked list	
<ul style="list-style-type: none"> <li>◆ Suppose we want to find an element in a singly linked list, and move it to the head</li> <li>◆ Visual intuition:               <div style="text-align: center; margin-top: 10px;"> </div> </li> </ul>	18

Even more real life, linked lists	
<pre> lprev = NULL; for(lpitr = lhead; lpitr; lpitr = lpitr-&gt;next) {     if(lpitr-&gt;val == target){         // Already head?, break         if(lprev == NULL) break;         // Move cell to head         lprev-&gt;next = lpitr-&gt;next;         lpitr-&gt;next = lhead;         lhead = lpitr;         break;     }     lprev = lpitr; } </pre>	<p>◆ Where is the critical section?</p>

Even more real life, linked lists	
<pre> Thread 1 // Move cell to head lprev-&gt;next = lpitr-&gt;next; lpitr-&gt;next = lhead; lhead = lpitr; </pre>	<pre> Thread 2 lprev-&gt;next = lpitr-&gt;next; lpitr-&gt;next = lhead; lhead = lpitr; </pre>
<p>◆ A critical section often needs to be larger than it first appears</p> <p>➢ The 3 key lines are not enough of a critical section</p>	

Even more real life, linked lists	
<pre> Thread 1 if(lpitr-&gt;val == target){     elt = lpitr;     // Already head?, break     if(lprev == NULL) break;     // Move cell to head     lprev-&gt;next = lpitr-&gt;next;     // lpitr no longer in list }  Thread 2 for(lpitr = lhead; lpitr;     lpitr = lpitr-&gt;next) {     if(lpitr-&gt;val == target){ </pre>	<p>◆ Putting entire search in a critical section reduces concurrency, but it is safe.</p>

Safety and Liveness	
<p>◆ <b>Safety property</b>: "nothing bad happens"</p> <p>➢ holds in every finite execution prefix</p> <ul style="list-style-type: none"> <li>◆ Windows™ never crashes</li> <li>◆ a program never terminates with a wrong answer</li> </ul> <p>◆ <b>Liveness property</b>: "something good eventually happens"</p> <p>➢ no partial execution is irremediable</p> <ul style="list-style-type: none"> <li>◆ Windows™ always reboots</li> <li>◆ a program eventually terminates</li> </ul> <p>◆ Every property is a combination of a safety property and a liveness property - (Alpern and Schneider)</p>	

Safety and liveness for critical sections	
<p>◆ At most k threads are concurrently in the critical section</p> <ul style="list-style-type: none"> <li>➢ A. Safety</li> <li>➢ B. Liveness</li> <li>➢ C. Both</li> </ul> <p>◆ A thread that wants to enter the critical section will eventually succeed</p> <ul style="list-style-type: none"> <li>➢ A. Safety</li> <li>➢ B. Liveness</li> <li>➢ C. Both</li> </ul> <p>◆ Bounded waiting: If a thread <i>i</i> is in entry section, then there is a bound on the number of times that other threads are allowed to enter the critical section (only 1 thread is allowed in at a time) before thread <i>i</i>'s request is granted.</p> <ul style="list-style-type: none"> <li>➢ A. Safety B. Liveness C. Both</li> </ul>	