From Serial to Parallel

Stephen Blair-Chappell Intel Compiler Labs <u>www.intel.com</u>



Agenda

- Why Parallel?
- Optimising Applications
- Steps to move from Serial to Parallel



6/18/2010

2

Intel® Software Development Products Overview



Copyright s 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

A Message From Steve Lionel Intel Compiler Labs

Congratulations BCS FIG.wmv

BCS Fortran Interest Group 40th Anniversary



The Chartered Institute for IT Fortran Specialist Group

Software & Services Group



Moving to Parallel - a view from some developers

- Top 5 challenges
 - -Legacy
 - -Education
 - -Tools
 - -Fear of many cores
 - -Maintainability



6/18/2010

4

Intel® Software Development Products Overview



Copyright C 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Why Parallel?

Section 1



Why is everyone going multicore? Power Density Race



Intel® Software Development Products Overview



Copyright C 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

intel

Software

Products



- Speed no longer increasing
- Num transistors still growing
- Num Cores rather than clock speed is doubling every 18 months



Intel® Software Development Products Overview



Copyright S 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Theoretical growth of cores



Intel® Software Development Products Overview



Copyright C 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

intel

Software

Products

Future: Multicore and Manycore

All Large Core Mixed Large and Small Core All Small Core Connections to memory bank(s), connections between processors, memory coherency models - all come into play. Diversity!

Note: the above pictures don't necessarily represent any current or future Intel products



Intel® Software Development Products Overview



Copyright S 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Multi-core : beating the *power\performance* barrier



Relative single-core frequency and Vcc



Intel® Software Development Products Overview



Copyright s 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Industry's First 45 nm High-K + Metal Gate Transistor Technology

Improved Transistor Density~2xImproved Transistor Switching Speed>20%Reduced Transistor Switching Power~30%Reduction in gate oxide leakage power>10x



6-transistor 65 nm S-RAM Cell 0.570 µm2

*Other brands and names are the property of their respective owners

Copyright



45 nm S-RAM Cell 0.346 μm2



65 nm Transistor



45 nm HK + MG



Enables New Features, Higher Performance, Greater Energy Efficiency



Intel's Teraflops Research Chip



Podtech Intel Research Day Terascale.flv



Intel® Software Development Products Overview



Copyright s 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Intel's Teraflops Research Chip

			and the second se
Speed	Power	Perf.	
GHz	Watts	Teraflops	
3.16	62	1.01	
5.1	175	1.63	
5.7	265	1.81	





Intel® Software Development Products Overview



Copyright S 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Larrabee

Products



Copyright O 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners





Intel® Software Development Products Overview



Copyright \circledast 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

What can we do with faster Computers?

- Solve problems faster
- Reduce turn-around time of big jobs
- Increase responsiveness of interactive apps

- Get better solutions in the same amount of time
 - Increase resolution of models
 - Make model more sophisticated





Intel® Software Development Products Overview



Optimising Code

Section 2



Two points to address before you start parallelising

•Will buying a faster computer solve your problem?



Dr Yann Golanski, York Just buy a faster machine!

First look at how much it will cost you to make your program parallel. If it will take say 2 months of coding, can you just buy a faster machine that will give you the speedup you want? Of course

Tip 1

faster machine that will give you the speedup you want? Of course once you reach the limits of a machines speed, you are going to have to then do some parallelisation.

Maybe Serial Optimisation will be sufficient.



6/18/2010

20

Intel® Software Development Products Overview



Copyright C 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

A Three-Tiered Tuning Model

Tuning Level	Question being asked	Examples of issues
System wide	Can my system be 'tuned' to improve the performance of my application	Network, disk and memory performance. Intrusion by 3 rd party programs such as virus scanners.
Application Heuristics	Can my application code or heuristics to improve performance?	Code redundancy. Inefficient program algorithms. Poor memory allocation strategies. Bad \ missing threading implementation.
Architectural Bottlenecks	Is the CPU architecture being used at its best?	Stalls in CPU pipeline. Data alignment . Cache misses. Using expensive instructions. Failing to use latest generation optimised instructions.



21 6/18/2010

Intel® Software Development Products Overview



Copyright © 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Compiler generated Optimisations





Global Compiler Options

Inter-procedural Optimisations

Profile Guided Optimisations

Vectorisation

Parallelisation



22 6/18/2010

From Serial to Parallel



Copyright C 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Example of hand-crafted SSE instructions

	Time Taken	Speedup
No SSE	4.55 sec	1
Vith SSE	0.19 sec	24



Intel® Software Development Products Overview



Copyright o 2007, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners



FP/SIMD execution cluster



Intel[®] VTune[™] Performance Analyzer

(Sampling)	Will Status Status See TE SS # @ TE SS # &	· · · · · · · · · · · · · · · · · · ·	🗣 🗣 💊	Activity	/1 (Sampling)	•	I II X	🔶 🖉 🐰	章 🧏	1					
15 mode@l 01 105 02 2000.05 seet 01 105 02 105 02 100 105 100 100 100 100 100 100 100 100	P) Event LX_UNRALED CK Ims1_REFIRED ANY LX_UNRALED CK Ims1_REFIRED ANY CalcSum Ims1_REFIRED ANY CalcSum Ims1_REFIRED ANY CalcSum Ims1_REFIRED ANY Set Event Activy ID Set Event Activy ID Set Set Set <th>8</th> <th>1 2 3 E</th> <th>5 5 8</th> <th>A 2017 7 20</th> <th>0 🔳 {}</th> <th></th> <th></th> <th></th> <th></th> <th></th> <th></th> <th></th> <th></th> <th></th>	8	1 2 3 E	5 5 8	A 2017 7 20	0 🔳 {}									
Calc Calc <th< th=""><th>Edit of Callsburger RTC_CheckE Image: Callsburger Image: Callsbu</th><th>1 (Sampling)</th><th></th><th></th><th></th><th></th><th></th><th></th><th></th><th></th><th></th><th></th><th>Events</th><th></th><th>Te</th></th<>	Edit of Callsburger RTC_CheckE Image: Callsburger Image: Callsbu	1 (Sampling)											Events		Te
CalcPi CalcSum CalcSum CalcSum CalcSum Solt Event Actively Scale Solt Event Actively Scale Solt Event Actively Scale Solt Solt Intel 4 Media Solt Solt Solt Solt Solt Solt Solt Solt	CalcPi CalcSum Calc	CPU_CLK_UNHALTED.CC	RTC_Check	:kE									CPU_CLK_UNHALTE INST_RETIRED AND CPU_CLK_UNHALTE INST_RETIRED AND CPU_CLK_UNHALTE INST_RETIRED AND	ED.CORE % ED.CORE % 7 % ED.CORE Y events	17 6 6 7 35 13
CalcPi CalcSum Calc	CalcPi CalcSum CalcSum CalcSum CalcSum Soft Event Soft Event Modulet			_											
CalcSum CalcSum Sort Evert ActivityID Scale Sample ActivityID Sc	CalcSum CalcSum 0 5 10 15 20 25 30 76 Sott Event ActivityID Scale Sample Alter Value Tald Sample Duarkon Fing 0 Fing 3 CPUID Ide time V Sott Event ActivityID Scale Sample Alter Value Tald Sample Duarkon Fing 0 Fing 3 CPUID Ide time V Sott Event ActivityID Scale Sample Alter Value Tald Sample Duarkon Fing 0 Fing 3 CPUID Ide time V INST_IDETINED.ARY S4 1e007 2000.000 245 11.484 4267 1222 1 2.35% V Processes Threads Modules Hutipode V		CalcPi												
CalcSum Solt Event ActivityID Scale Scale CPUID Ide time 0 5 10 15 20 25 30 25 Soit Event ActivityID Scale Sample Alter Value Table Samples Duarkon Ping 0 Ping 3 CPU ID Ide time ✓ Strip (PL_GL_U)IN=Alt (PD CODE 54 1e-002 2000,000 52485 11.484 4566 7222 1 2.85% ✓ Disclar per functions Resed-CP 54 1e-002 2.000,000 2463 11.484 1267 1222 1 2.35% ✓ Disclar per functions Resed-CP 54 1e-00 0 0 0 0 0	CalcSum Event CalcSum 0 5 10 15 20 25 20 25 Set Event ActivityID Scale Sample Attent Value Tatal Samples Duration Fing 0 Fing											Ľ.			
0 5 10 15 20 25 30 50 Server Atter Value Ted Server 1 444 446 122 1 2265 WST_RETRED_AVY WST_RETRED_AVY Solds aper Instructions Released OP 54 10 1 0 0 0 0 11 444 1247 1222 1 2265 Clocks aper Instructions Released OP 54 10 1 0 0 0 0 0 11 445 1247 1222 1 2265 Exerver 1 444 1247 1222 1 2265 Exerver 1 445 1247 1247 1247 Exerver 1 445 1247 1247 1247 Exerver 1 445 1247 1247 1247 Exerver 1 445 1247 1247 1247 1247 Exerver 1 445 1247 1247 1247 1247 1247 Exerver 1 445 1247 1247 1247 1247 1247 1247 1247 1247	0 5 10 15 20 25 30 35 Series Atter Value Ted Series Data Reg 0 Fing 3 (PVI U) Ide time		▶ CalcSum												
Soft Evert Activity D Sould Sample Aller Value Tourkson Ring D Ring 3 CPU ID Add traine V CPU LD, UNALTED CORE 54 TeO/07 2000,000 2459 11.464 4467 1220 1 2.855 V Dods per Instructions Relies-CP 54 1 0	Sort Event Active/ID Soude Soude/Attr-VAte Total Soudes Total Soudes Prog 0			ů.	5 10	 1	5	20	25	30		35	1		
✓ ✓ OPU_CU_UNReLID: UDHE 54 1e007 2000.000 57.66 11.484 1427 1222 1 2.3852 ✓ Clock per Instruction Retied - CPI 54 10 1 0 0 2 2 1 2.3852	V CPU_CU_QUEWELED.UDHE 54 14:007 2000.000 57/85 11:144 1427 U 24/25 V Dods per Instruction Relead-CPI 54 16:007 2000.000 24/85 11:144 1247 12 23/85 V Dods per Instruction Relead-CPI 54 10 1 0 0 0 0 2 V Processes Treeds Modules Hatapats Hatapats Hatapats			Sort	Event	Activity ID	Scale	Sample After Value	Total Samples	Duration	Ring 0	Ring 3	CPU ID	Idle ti	me
Clocks per Instructions Retired - CPI 54 10 1 0 0 0 0	Cooks per instruction Reted-CPI 54 10 1 0 0 0 0 Processes Threads Modules Holipoids			~	INST_RETIRED.ANY	54	1e-007	2,000,000	2469	11.404	1247	1222	1	2.36%	
	Horizon Standard College			•	Clocks per Instructions Retired - CPI	54	10	1	0	0	0	0			
	Processes Treeds Modules Holdpols		1	1								•	<u> </u>		_
	ADDENT	>													
		<u> </u>													

Graphical tool

Helps characterise runtime performance

System-wide View of application environment

Use to tune serial and parallel code

Use to identify Hot Spots in Code

Use to generate a call graph



25 6/18/2010

From Serial to Parallel



Copyright C 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Call graph: Application workflow



The life of a program instruction





Hardware Performance Events



Demo 0 – Using Intel[®] VTune[™] Performance Analyzer

From 1 to 1,000,000





From Serial to Parallel

Copyright C 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Notice the System Wide View – Can you see any problems?

mcshield.exe		1,146
cidaemon.exe		1,067
02_Hotspot Analysis.exe		838
osrss.exe		807
pid_0x0	Process=02_	Hotspot Ana
cisvc.exe		178
rundll32.exe		155
pid_0x4		111
services.exe		97



30 6/18/2010

From Serial to Parallel



Copyright O 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

What is the problem here?

• VTune Sample-over-time view

ProcessName	ProcessID	CPU_CLK_UNHALT samples	Time, sec 0.5 1 1.5 2 2.5
moshield.exe	2596	1146	
cidaemon.exe	18296	1067	alla and a set and a state of a set of the s
02_Hotspot Analysis.exe	12464	838	an a
csrss.exe	2020	807	الم المالة، بلا من محمد الله من معاد الله المعالية المعالية المعالية المعالية، المحمد في المعالية الم
pid_0x0	0	214	and the second



31 6/18/2010

From Serial to Parallel



Copyright s 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

The Hotspot



Copyright C 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Products

Four Steps in Moving to Parallel

Section 2



34 6/18/2010

From Serial to Parallel



Copyright S 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Steps in moving from Serial to Parallel





35 6/18/2010

From Serial to Parallel



Copyright O 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Key Questions

Design

- Is my program parallel?
- Where is the best place to parallelise my program?
- How can I get my program to run faster?
- What's the expected speedup?

Code & Debug

- How?
- How difficult?
- Is my code still working?

Verify

- Is the parallelism correct?
- Do I have deadlocks or data races?
- Do I have memory errors?
- Does my program still work as intended?

Tune

- Do my tasks do equal amounts of work?
- Is my application scalable?
- Is the threading running efficiently?



36 6/18/2010

From Serial to Parallel



Copyright S 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners





37 6/18/2010

From Serial to Parallel



Copyright O 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

Tools

Serial Architectural Analysis Introducing Parallelism Validating Correctness Performance Tuning Parallel

Existing Intel Software	Intel Parallel Studio
Intel® VTune [™] Performance Analyzer	Advisor/Amplifier
Intel Compilers Parallel Libraries	Composer
Intel® Thread Checker	Inspector
Intel® Thread Profiler	Amplifier



38 6/18/2010

From Serial to Parallel



Copyright O 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

For Microsoft Visual Studio^{*} C++ architects, developers, and software innovators creating parallel Windows^{*} applications.



Microsoft Visual Studio* plug-in End-to-end product suite for parallelism Forward scaling to many core



Intel[®] Parallel Studio includes:

- Intel® Parallel Advisor Lite **
- Intel[®] Parallel **Composer**
- Intel[®] Parallel **Inspector**
- Intel[®] Parallel Amplifier

** Beta – from whatif.intel.com



39 6/18/2010

From Serial to Parallel

(intel)

Copyright S 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners


Key Questions - Design



Is my program parallel?

Where is the best place to parallelise my program?

How can I get my program to run faster?

What's the expected speedup?



From Serial to Parallel



Identifying best parts to Parallelize





43 6/18/2010

From Serial to Parallel



Why use parallelism? Amdahl's Law

Describes the upper bound of parallel execution speedup

Serial code limits speedup





Some code is not worth making parallel...

Don't parallelise code

- just because it's clever
- With low CPU utilisation
- I/O bound



- Do parallelise code that
 - Eats significant CPU cycles
- You need to get visibility of the runtime behaviour





45 6/18/2010

From Serial to Parallel



Architectural Extensions can speed up your code

•Always optimise your code

•Even if you don't go parallel, some architectural features can still give significant speed-up

•Example, SSE extensions



46 6/18/2010

From Serial to Parallel



Our Application - Prime Number Generator

i.	factor
61	357
63	3
65	35
67	357
69	3
71	357
73	3579
75	35
77	357
79	3579

```
bool TestForPrime(int val)
   // let's start checking from 3
    int limit, factor = 3;
    limit = (long)(sqrtf((float)val)+0.5f);
    while( (factor <= limit) && (val % factor) )</pre>
            factor ++;
    return (factor > limit);
void FindPrimes(int start, int end)
    int range = end - start + 1;
    for( int i = start; i <= end; i += 2 )</pre>
    {
        if( TestForPrime(i) )
            globalPrimes[gPrimesFound++] = i;
        ShowProgress(i, range);
```



6/18/2010

47

From Serial to Parallel



Copyright © 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

ł

}

Demo 1 – Getting the Benchmark

From 1 to 1,000,000



Our Application - Prime Number Generator



Optimise the serial code first

- Using Intel Compiler to automatically generate SSE instructions.
 - Code ran twice as fast
 - No change made to original code

Calculating Pi			
	Time Secs	Speedup	
No SSE	1.29	1.00	
With SSE	0.66	1.95	

004018D9	movaps	xmmword ptr [esp],xmm0
004018DD	paddd	xmm5,xmm6
004018E1	addpd	xmm7,xmm3
004018E5	mulpd	xmm7,xmm2
004018E9	add	eax,8
004018EC	mulpd	xmm7,xmm7
004018F0	movaps	xmm0,xmmword ptr ds:[406770h]
004018F7	addpd	xmm7,xmm1
004018FB	divpd	xmm0,xmm7
004018FF	cvtdq2pd	xmm7,xmm5
00401903	paddd	xmm5,xmm6
00401907	addpd	xmm4,xmm0
0040190B	movaps	xmm0,xmmword ptr ds:[406770h]
00401912	addpd	xmm7,xmm3
00401916	mulpd	xmm7,xmm2
0040191A	mulpd	xmm7,xmm7
0040191E	addpd	xmm7,xmm1
00401922	divpd	xmm0,xmm7
00401926	movaps	xmm7,xmmword ptr [esp]
0040192A	addpd	xmm7,xmm0
0040192E	cvtdq2pd	xmm0,xmm5
00401932	movaps	xmmword ptr [esp],xmm7

Example of SSE compiler-generated instructions



50 6/18/2010

From Serial to Parallel

(intel)

Demo 2 - Using the Intel Compiler

From 1 to 1,000,000





From Serial to Parallel

Swapping compilers.

- From solution drop-down menu
- Action is reversible







Program built and run with Intel compiler



*Other brands and names are the property of their respective owners



Identifying Hotspots

Pinpointing places where an application could be parallelised



54 6/18/2010

From Serial to Parallel



The Big Question



"How can I make my code run faster?"





55 6/18/2010

From Serial to Parallel



Today's Question



"Where do I split up my code to take advantage of multiple CPU cores?"





6/18/2010

56

From Serial to Parallel



The task, Identifying the Hot Spot...







57 6/18/2010

From Serial to Parallel



... and Splitting up the Work.







58 6/18/2010

From Serial to Parallel



Demo 2 - Finding the Hotspots



59 6/18/2010

From Serial to Parallel



Finding a Hot Spot

(intel)

Software

Products





Where to Parallelise - Amplifier

Call Stack 💌 👻	CPU Time:Total	CPU Time:Self	*
🖃 😒 Total	100.0%	0s	
⊟ ≥ BaseProcessStart	100.0%	Os	
□ \u2212 _tmainCRTStartup	100.0%	Os	
🖃 🛛 main	100.0%	Os	
🖃 🏼 FindPrimes	100.0%	Os	
□ TestForPrime	24.4%	0.328s 📃	
ShowProgress	75.6%	1.016s	





Design: What's the expected speedup?

allea Amdhale Law	🖃 🗵 main 🛛 👘 Os
• USE AIHUIIdis Law	🖃 🛛 FindPrimes 👘 Os
Speedup = $1/[s+(1-s)/n + H(n)]$	□ TestForPrime 0.328s
s is serial part (fraction of 1)	ShowProgress 1.016s
H is parallel overhead (ignore) n is number of cores	11
	Elapsed Time: 3.347s
	CPU Time: 1.344s
	Unused CPU Time: 5.351s
	Core Count: 2
	Threads Created: 1
S = 0 Speedup = 1 / [0 + (1 - 0) / 2] = 1 / [0 + 0.5]	
Speeaup = 2 (i.e. ne	w speed \sim 0.672 seconds)



62 6/18/2010

From Serial to Parallel



Alternate Calculation

Speedup =
$$1/[s+(1-s)/n + H(n)]$$

s is serial part (fraction of 1) H is parallel overhead (ignore) n is number of cores

Function - Caller Function Tree	×	CP	U Time:S	elf▼
∎ShowProgress		1.336s		
± TestForPrime		0.340s		
±[Import thunk sqrt]		0.012s)	
Summary	Elapsed Tim CPU Time: Unused CPI Core Count Threads Cr	ne: 4. 1. U Time: 6. :: eated:	• 7 305s 688s 923s 2 1	×

$$S = 1 - (1.688 - 0.012)/1.688 = .007$$

Speedup = 1 / [.007 + (1 - .007) / 2]
= 1 / [0007 + 0.4965]
Speedup = 1.986 (i.e. CPU Time ~ 0.850 seconds)



63 6/18/2010

From Serial to Parallel





Key Questions – Code & Debug



How?

How difficult?

Is my code still working?





From Serial to Parallel

Common types of parallelism

Functional or Task Parallelism
Data Parallelism
Software Pipelining



67 6/18/2010

From Serial to Parallel



Task and Data Parallelism

- Different job for each thread
- e.g. one thread prints, another reads keyboard

- Splitting workload between multiple identical threads
- e.g. three identical threads perform calculations on data array



Task parallelism

68

Data parallelism



6/18/2010

From Serial to Parallel



Software Pipeline





69 6/18/2010

From Serial to Parallel





•How many different ways can you think of to implement parallelism?

-*E.g OpenMP, ..., ...*



70 6/18/2010

From Serial to Parallel



Intel's Family of Parallel Models



Auto Parallelism

Loop-level parallelism automatically supplied by the compiler



72 6/18/2010

From Serial to Parallel



Auto-parallelization



• Auto-parallelization: Automatic threading of loops without having to manually insert OpenMP* directives.

Windows*	Linux*	Mac*
/Qparallel	-parallel	-parallel
/Qpar_report[n]	-par_report[n]	-par_report[n]

• Compiler can identify "easy" candidates for parallelization, but large applications are difficult to analyze.



73 6/18/2010

From Serial to Parallel



Optimisation Results – pi application

Optimisation	Time Taken (secs)	Speedup
default	0.938	1
auto-vectorisation	0.375	2.5
auto-parallelism	0.516	1.8
auto-vec. & auto-par.	0.203	4.6



74 6/18/2010

From Serial to Parallel





OpenMP Architecture



- Fork-Join Model
- Worksharing constructs
- Synchronization constructs
- Directive/pragma-based parallelism
- Extensive API for finer control



75 6/18/2010

From Serial to Parallel





OpenMP Runtime





OpenMP Programming Model:



Fork-Join Parallelism:

Master thread spawns a team of threads as needed.

 Parallelism added incrementally until performance are met: i.e. the sequential program evolves into a parallel program.


Introducing Parallelism





78 6/18/2010

From Serial to Parallel



5-1 -2 -3 -4 -7

Demo 3 : Adding parallelism using #pragma omp for



79 6/18/2010

From Serial to Parallel



Results - Open MP





80 6/18/2010

From Serial to Parallel



Code: Is my code still working? Running 01_Converted To Intel.exe 100% 78498 primes found between 0 and 1000000 in 2.28 secs Running 03_Naive OpenMP.exe 50% 36964 primes found between 0 and 1000000 in 1.45 secs Number of primes is wrong



81 6/18/2010

From Serial to Parallel





Are the results right?

Was the run quicker?



82 6/18/2010

From Serial to Parallel





Key Questions - Verify



Is the parallelism correct?

Do I have deadlocks or data races?

Do I have memory errors?

Does my program still work as intended?





From Serial to Parallel

New paradigm requires new tools

- Using traditional debugging tools is difficult /impossible
 - Printf not re-entrant
 - Debugging several threads is notoriously hard
 - Many debuggers \ profilers are not multi-core enabled
- Multi-core tools are available



86 6/18/2010

From Serial to Parallel



Non deterministic Error Sources in parallel Applications

• Shared Resources require locks

- Locks can
 - 'serialize' a program
 - lead to <u>Deadlocks</u>





87 6/18/2010

From Serial to Parallel

(intel)

Demo 4 – Checking for threading errors



88 6/18/2010

From Serial to Parallel



Checking for Errors with Parallel Inspector

0	verv	view 🔶 🌮 Sources 🔶 🤄	🔒 Details
Prob	lem	Sets	
ID 🔺	0	Problem	Sources
P1	8	Data race	PrimeOpenMP.cpp
P2	-	Data race	PrimeOpenMP.cpp
P3	⚠	Potential privacy infringement	PrimeOpenMP.cpp
P4	蛊	Potential privacy infringement	PrimeOpenMP.cpp
P5	◬	Potential privacy infringement	PrimeOpenMP.cpp



89 6/18/2010

From Serial to Parallel



The Offending Sources

```
Focus Observation: PrimeOpenMP.cpp:116 - Read
      111
      112
                #pragma omp parallel for
      113
                for ( int i = start; i <= end; i += 2 )
      114
                £
      115
                    if( TestForPrime(i) )
                        globalPrimes[gPrimesFound++] = i;
      116
      117
      118
                    ShowProgress(i, range);
      119
               ł
      120
          ÷
      121
      122
           int main(int argc, char **argv)
          <
      Related Observation: PrimeOpenMP.cpp:116 - Write
      111
      112
                #pragma omp parallel for
      113
                for ( int i = start; i <= end; i \neq 2 )
      114
                £
      115
                    if( TestForPrime(i) )
      116
                        globalPrimes[gPrimesFound++] = i;
      117
      118
                    ShowProgress(i, range);
      119
                ł
      120
           }
      121
      122
          int main(int argc, char **argv)
          <
                                      From Serial to Parallel
        6/18/2010
90
```

Copyright C 2008, Intel Corporation. All rights reserved. *Other brands and names are the property of their respective owners

intel

Software

Products



Demo 5 - Fixing the threading errors



92 6/18/2010

From Serial to Parallel



Data Races Have Disappeared!

Problem Sets					
ID 🔺	0	Problem	Sources	Modules	State
P1	≜	Potential privacy infringement	PrimeOpenMP.cpp	05_Correct Threading Errors.exe	┡ Not fixed
P2	⚠	Potential privacy infringement	PrimeOpenMP.cpp	05_Correct Threading Errors.exe	훢 Not fixed
P3	≜	Potential privacy infringement	PrimeOpenMP.cpp	05_Correct Threading Errors.exe	ॺ Not fixed



93 6/18/2010

From Serial to Parallel



Number of Primes is Correct





94 6/18/2010

From Serial to Parallel





Step 4

Tune for best performance



Key Questions - Tune



Is the threading running efficiently?

Do my tasks do equal amounts of work?

Is my application scalable?





From Serial to Parallel

Performance Issues

Load Balancing Synchronisation Overhead Scalability

Difficult to examine without the right tools



98 6/18/2010

From Serial to Parallel



A Reminder – Where are we?





Number of

primes is correct



Almost as slow as

the serial version

Demo 6 – Find the Threading Performance Issues



Hotspot Analysis

💯 Hotspots				
😪 Bottom-up				
Call Stack 🛛	CPU Time:Total 👻	CPU Time:Self	*	
🖃 🗁 Total	100.0%	Os		
🖃 🖂 BaseThreadStart	50.8%	Os	ŀ	
□ \u2264 _kmp_launch_worker	50.8%	Os	I	
□ \u2224 _kmpc_invoke_task_func	50.8%	Os	1	
□ \u224 _kmpc_invoke_task_func	50.8%	Os	1	
🖃 🛛 _kmp_invoke_microtask	50.8%	Os	1	
□ \> L_?FindPrimes@@YAXHH	50.8%	0.023s 🕽	(
□ ShowProgress	42.1%	1.231s 🦲 👘	(
□ TestForPrime	7.9%	0.231s 📒	(
⊟ ≥ BaseProcessStart	49.2%	Os	ŀ	
\u2014 _tmainCRTStartup	49.2%	Os	(
🖃 🏼 main	49.2%	Os	(
FindPrimes	49.2%	0.016s	(
🖃 ↘ L_?FindPrimes@@YAXHH@Z	48.7%	0.012s	(
ShowProgress	43.9%	1.283s	(
□ TestForPrime	4.3%	0.127s 🛿	(



101 6/18/2010

From Serial to Parallel

(intel)

Source View

Ami	otspots	tel [®] Parallel Amplifier		
🚷 Bo	🗞 Bottom-up 🔣 Top-down Tree 🔹 🚺 PrimeOpenMP.cpp 🛛 🛛			
79 (1) 🖑 🥸 🗉			
Line	Source	CPU Time:Self 🛛 😒		
81	}			
82				
83	void ShowProgress(int val, int range)			
84	(
85	int percentDone = 0;			
86				
87	#pragma omp critical	0.016s		
88	(
89	gProgress++;	0.070s 🛿		
90				
91	percentDone = (int)((float)gProgress/(float)range *200.0f +	0.036s		
92	}			
93				
94				
95	if(percentDone % 10 == 0)			
96	<pre>printf("\b\b\b%3d%%", percentDone);</pre>	2.393s		
97				
- 98	}			
99				



102 6/18/2010

From Serial to Parallel

(intel)

Improving the Performance

```
void ShowProgress( int val, int range )
{
    int percentDone;
    static int lastPercentDone = 0;
#pragma omp critical
    {
        gProgress++;
        percentDone = (int)((float)gProgress/(float)range*200.0f+0.5f);
    }
    if( percentDone % 10 == 0 && lastPercentDone < percentDone / 10){
        printf("\b\b\b\83d%%", percentDone);
        lastPercentDone++;
    }
}</pre>
```

The algorithm has many more updates than the 10 needed for showing progress

This change should fix the contention issue



103 6/18/2010

From Serial to Parallel



Demo 7 – Fixing the Synchronisation issues



Superb Speedup ... ???





105 6/18/2010

(intel)

Software

Products

From Serial to Parallel



Demo 8 – Getting New Serial Benchmark



That's better (but disappointing)...





107 6/18/2010

From Serial to Parallel



Demo 9 – Correcting the Synchronisation Issue



Hotspots

Function CPU Time:Self * • Caller Function Tree 388.224ms 05 • TestForPrime 388.224ms 05 • L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 388.224ms 05 • L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 388.224ms 05 • L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 388.224ms 05 • FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 140.555ms 05 • FindPrimes 52.530ms 05 • Main ← _tmainCRTStartup ← BaseProcessStart 52.530ms 05 • L_?FindPrimes@@VAXHH@Z_119_par_loop0_2.61 47.070ms 05 • L_?FindPrimes@@VAXHH@Z_119_par_loop0_2.61 47.070ms 05 • L_?FindPrimes@@VAXHH@Z_119_par_loop0_2.61 47.070ms 05 • L_?FindPrimes@@VAXHH@Z_119_par_loop0_2.61 47.070ms 05 • L_?FindPrimes@@VAXHH@Z_119_par_loop0_2.61 46.924ms 05 • ShowProgress 46.924ms 05 • ShowProgress 46.924ms 05 • L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 46.924ms 05 • ShowProgress 46.924ms 05 • ShowProgress 46.924ms	r003hs PrimeOpenMP.cpp		
Function CPU Time:Self CPU Time:Self 388.224ms TestForPrime 388.224ms K_L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 388.224ms K_kmp_invoke_microtask < kmpc_invoke_task_func < kmpc_in 247.670ms FindPrimes < main < tmainCRTStartup < BaseProcessStart 140.555ms FindPrimes 52.530ms 09 K_main < tmainCRTStartup < BaseProcessStart 52.530ms 09 K_main < tmainCRTStartup < BaseProcessStart 52.530ms 09 K_main < tmainCRTStartup < BaseProcessStart 52.530ms 09 K_mp_invoke_microtask < kmpc_invoke_task_func < kmpc_invok 23.560ms 09 K_sindPrimes@@YAXHH@Z_119_par_loop0_2.61 46.924ms 09 K_sindPrimes@@YAXHH@Z_119_par_loop0_2.61 46.924ms 09 K_kmp_invoke_microtask < kmpc_invoke_task_func < kmpc_invoke 09 09 K_kmp_invoke_microtask < kmpc_invoke_task_func < kmpc_invok 23.488ms 09 K_kmp_invoke_microtask < kmpc_invoke_task_func < kmpc_in 23.488ms 09 K_kmp_invoke_microtask < kmpc_invoke_task_func < kmpc_in 23.488ms 09 K_kmp_invoke_microtask < kmpc_invoke_task_func < kmpc_in 23.488ms 09	Hotspots Bottom-up Top-down Tree	Intel® Pa	ага
□ TestForPrime 388.224ms 09 □ ▷ L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 388.224ms 09 ▷ kmp_invoke_microtask ← kmpc_invoke_task_func ← kmpc_in 247.670ms 09 ▷ FindPrimes ← main ← tmainCRTStartup ← BaseProcessStart 140.555ms 09 □ FindPrimes 52.530ms 09 □ K main ← tmainCRTStartup ← BaseProcessStart 52.530ms 09 □ L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 47.070ms 09 □ K kmp_invoke_microtask ← kmpc_invoke_task_func ← kmpc_invol 23.560ms 09 □ ShowProgress 46.924ms 09 □ ▷ L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 46.924ms 09 □ ▷ kmp_invoke_microtask ← kmpc_invoke_task_func ← kmpc_in 23.488ms 09 □ ▷ kmp_invoke_microtask ← kmpc_invoke_task_func ← kmpc_in 23.488ms 09	Function - Caller Function Tree	☆ CPU Time:Self ▼	•
□ \[\] \[\	TestForPrime	388.224ms	09
\[\begin{aligned} & & & & & & & & & & & & & & & & & & &	L_?FindPrimes@@YAXHH@Z_119par_loop0_2.61	388.224ms	09
└ FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 140.555ms 09 □ FindPrimes 52.530ms 09 □ main ← _tmainCRTStartup ← BaseProcessStart 52.530ms 09 □ L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 47.070ms 09 □ kmp_invoke_microtask ← _kmpc_invoke_task_func ← _kmpc_invok 23.560ms 09 □ FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 23.510ms 09 □ ShowProgress 46.924ms 09 □ └.?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 46.924ms 09 □ └	\land _kmp_invoke_microtask \leftarrow _kmpc_invoke_task_func \leftarrow _kmpc_in	247.670ms	09
■ FindPrimes52.530ms09「 main ← _tmainCRTStartup ← BaseProcessStart52.530ms09□ L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.6147.070ms09「 _kmp_invoke_microtask ← _kmpc_invoke_task_func ← _kmpc_invol23.560ms09「 FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart23.510ms09■ ShowProgress46.924ms09「 _kmp_invoke_microtask ← _kmpc_invoke_task_func ← _kmpc_in09「 _ FindPrimes@@YAXHH@Z_119_par_loop0_2.6146.924ms09「 _ FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart23.488ms09「 _ FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart23.435ms09	$\land FindPrimes \leftarrow main \leftarrow _tmainCRTStartup \leftarrow BaseProcessStart$	140.555ms 📃 📃	09
<pre></pre> <pre></pre> <pre>SeeProcessStart</pre>	□ FindPrimes	52.530ms 📃	09
□L_?FindPrimes@@YAXHH@Z_119par_loop0_2.61 47.070ms 09 □<_kmp_invoke_microtask ← _kmpc_invoke_task_func ← _kmpc_invok	\land main \leftarrow _tmainCRTStartup \leftarrow BaseProcessStart	52.530ms 📒	09
_kmp_invoke_microtask ← _kmpc_invoke_task_func ← _kmpc_invol 23.560ms 09 FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 23.510ms 09 ShowProgress 46.924ms 09 KPrindPrimes@@YAXHH@Z_119par_loop0_2.61 46.924ms 09 69 69 69 69 69 69	L_?FindPrimes@@YAXHH@Z_119par_loop0_2.61	47.070ms 🔲	09
FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 23.510ms 09 ■ ShowProgress 46.924ms 09 ■ K L_?FindPrimes@@YAXHH@Z_119_par_loop0_2.61 46.924ms 09 K _kmp_invoke_microtask ← _kmpc_invoke_task_func ← _kmpc_in 23.488ms 09 K FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 23.435ms 09	Kmp_invoke_microtask	23.560ms 📒	09
ShowProgress 46.924ms 09 Image: ShowProgress 50 23.488ms 09 Image: ShowProgress 50 23.435ms 09	\frown FindPrimes \leftarrow main \leftarrow _tmainCRTStartup \leftarrow BaseProcessStart	23.510ms 📒	09
□ \[\] L_?FindPrimes@@YAXHH@Z_119par_loop0_2.61 46.924ms 09 □ \[\]	ShowProgress	46.924ms 🗖	09
23.488ms 09 23.488ms 09 09 09 09	L_?FindPrimes@@YAXHH@Z_119par_loop0_2.61	46.924ms 🔲	09
▷ FindPrimes ← main ← _tmainCRTStartup ← BaseProcessStart 23.435ms	Kmp_invoke_microtask	23.488ms 📒	09
	\land FindPrimes \leftarrow main \leftarrow _tmainCRTStartup \leftarrow BaseProcessStart	23.435ms 📒	09
□ main 9.970ms 09	Emain	9.970ms 🖡	09
_tmainCRTStartup ← BaseProcessStart 9.970ms 09	\frown _tmainCRTStartup \leftarrow BaseProcessStart	9.970ms 🖡	09



109 6/18/2010

From Serial to Parallel



Locks & Waits

r003hs r004lw PrimeOpenMP.cpp				
💯 Locks and Wait	ts		2	
🚱 Bottom-up	o-down Tree 🛛 🚯 PrimeOpenMP.cpp			
Sync Object Name - Wait Function 🛛 💙	Wait Time by Utilization:Self	Wait Count:Self		
- Wait Function Tree	🔲 Idle 📕 Poor 📙 Ok 📕 Ideal 📘 Over			
GMP Critical FindPrimes:	25.070s	479605	0	
■ ShowProgress	21.982s	414666	0	
🖃 🖺 L_?FindPrimes@	21.982s	414666	0	
⊼_kmp_invoka	11.160s	208993	0	
Note: No	10.822s	205673	0	
■L_?FindPrimes@@YA>	3.088s 🔲	64939	0	



110 6/18/2010

From Serial to Parallel

(intel)

Source Code View

1		
84	(
85	int percentDone = 0;	
86	<pre>static int lastPercentDone = 0;</pre>	
87		
88	#pragma omp critical	
89	í	
90	gProgress++;	21.982s
91		
92	<pre>percentDone = (int)((float)gProgress;</pre>	
93	}	
94		
95		
96	if(percentDone % 10 == 0 44 lastPerc	
97	printf("\b\b\b%3d%%", percentDo	0.004s
98	lastPercentDone++;	
99		
100	}	
101	· · · · · · · · · · · · · · · · · · ·	



111 6/18/2010

From Serial to Parallel

(intel)

Fixing the synchronisation issue -1

This fix removes the need for a critical section

```
void FindPrimes(int start, int end)
{
    // start is always odd
    int range = end - start + 1;

#pragma omp parallel for
    for( int i = start; i <= end; i += 2 )
    {
        if( TestForPrime(i) )
            globalPrimes[InterlockedIncrement(&gPrimesFound)] = i;
        ShowProgress(i, range);
    }
}</pre>
```



112 6/18/2010

From Serial to Parallel

(intel)

Fixing the synchronisation issue - 2

This fix removes the need for a critical section

```
void ShowProgress( int val, int range )
{
    long percentDone, localProgress;
    static int lastPercentDone = 0;
    localProgress = InterlockedIncrement(&gProgress);
    percentDone = (int)((float)localProgress/
        (float)range*200.0f+0.5f);
    if( percentDone % 10 == 0 && lastPercentDone < percentDone / 10)
        printf("\b\b\b\b\83d%%", percentDone);
        lastPercentDone++;
    }
}</pre>
```



113 6/18/2010

From Serial to Parallel



That's better (but still disappointing)...



Running 10_Fixing Synchronisation Issue.exe 100% 78498 primes found between 0 and 1000000 in 0.30 secs





114 6/18/2010

From Serial to Parallel



Demo 9 – Improving the Load Balancing


Threads are not doing equal work





117 6/18/2010

From Serial to Parallel



Fixing a Load Imbalance

Distribute the work more evenly

```
void FindPrimes(int start, int end)
{
    // start is always odd
    int range = end - start + 1;
#pragma omp parallel for schedule(static, 8)
    for( int i = start; i <= end; i += 2 )
    {
        if( TestForPrime(i) )
           globalPrimes[InterlockedIncrement(&gPrimesFound)] = i;
        ShowProgress(i, range);
    }
}</pre>
```



118 6/18/2010

From Serial to Parallel



That's better







119 6/18/2010

From Serial to Parallel



A Finely Balanced threaded Program!





120 6/18/2010

From Serial to Parallel



Scalability



121 6/18/2010

From Serial to Parallel







Is the threading running efficiently?

Do my tasks do equal amounts of work?

Is my application scalable?





From Serial to Parallel

Moving to Parallel – a view from some developers

Top 5 challenges

- •Legacy
- Education
- Tools
- Fear of many coresMaintainability



123 6/18/2010

From Serial to Parallel



Scalability http://paralleluniverse.intel.com

Intel® Parallel Universe Portal



	Step 3: Summary
Session Name	prime
Upload zip file	11_Load Balancing.zip
Command Line	11_Load Balancing.exe 1 1000000



The Results



2.08 1.85

1.62

1.39

1.16

0.92 0.69

0.46

0.23

0.00

Throughput Scalability 😨

2 4 8 16

Cores

Throughout Scalabity

Your job has completed!

oparallel.universe@intel.com

Sent: Mon 07/06/2010 06:05

To: OBlair-chappell, Stephen

Your job has completed!

We're sending you this email because your job "prime" has successfully completed Please click the link below to access to the results page <u>http://paralleluniverse.intel.com/Job/Summary/878388be-a821-46f9-a911-0592309db43b</u> Thanks for using Intel® Parallel Universe Portal!



0.19

0.17 0.15

0.12

0.06

0.04

0.02

0.00

Real Time

1 2 4 8 16

💻 Optimal Time

μ^ω 0.10

8 0.08

125 6/18/2010

Performance Improvement 🕜

Cores

From Serial to Parallel



Without the printfs

			-			
ores	Elapsed Time (secs)	Number of Executions	Ø	Cores	Average Concurrency	Number of Executions
1	0.17	3		1	1.00	3
2	0.17	3		2	1.85	3
4	0.12	3		4	2.69	3
8	0.10	3		8	3.59	3
16	0.11	3		16	3.31	3





126 6/18/2010

From Serial to Parallel



A run of 10 million

Command line: Report:		11_Load Balancing.exe 1 10000000 Download Intel® Parallel Amplifier report			Service Pack 2 Intel® Xeon® processor, CPU 2.80 G Hyperffreade 8-Core 12 GB RAM			
Cores	Elapse	d Time (secs)	Number of Executions	Con	es	Average Concurrency	Number of Executions	
1		3.68	3	1		1.00	3	
2		3.58	3	2		1.98	3	
4		1.83	3	4		3.88	3	
8		0.95	3	8		7.49	3	
16		0.60	3	1	5	12.14	3	





127 6/18/2010

From Serial to Parallel



Tools

Serial Architectural Analysis Introducing Parallelism Validating Correctness Performance Tuning Parallel

Existing Intel Software	Intel Parallel Studio
Intel® VTune [™] Performance Analyzer	Advisor/Amplifier
Intel Compilers Parallel Libraries	Composer
Intel® Thread Checker	Inspector
Intel® Thread Profiler	Amplifier



128 6/18/2010

From Serial to Parallel



thank you

intel.com / go / parallel





thank you

