



Remember when  
the sky was the limit?



Intel, the Intel logo, MMX, and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

# Best Practices for Developing Multi-threaded Programs

# Agenda

Motivation for Threading

Concepts in Parallelism

Threading in Action

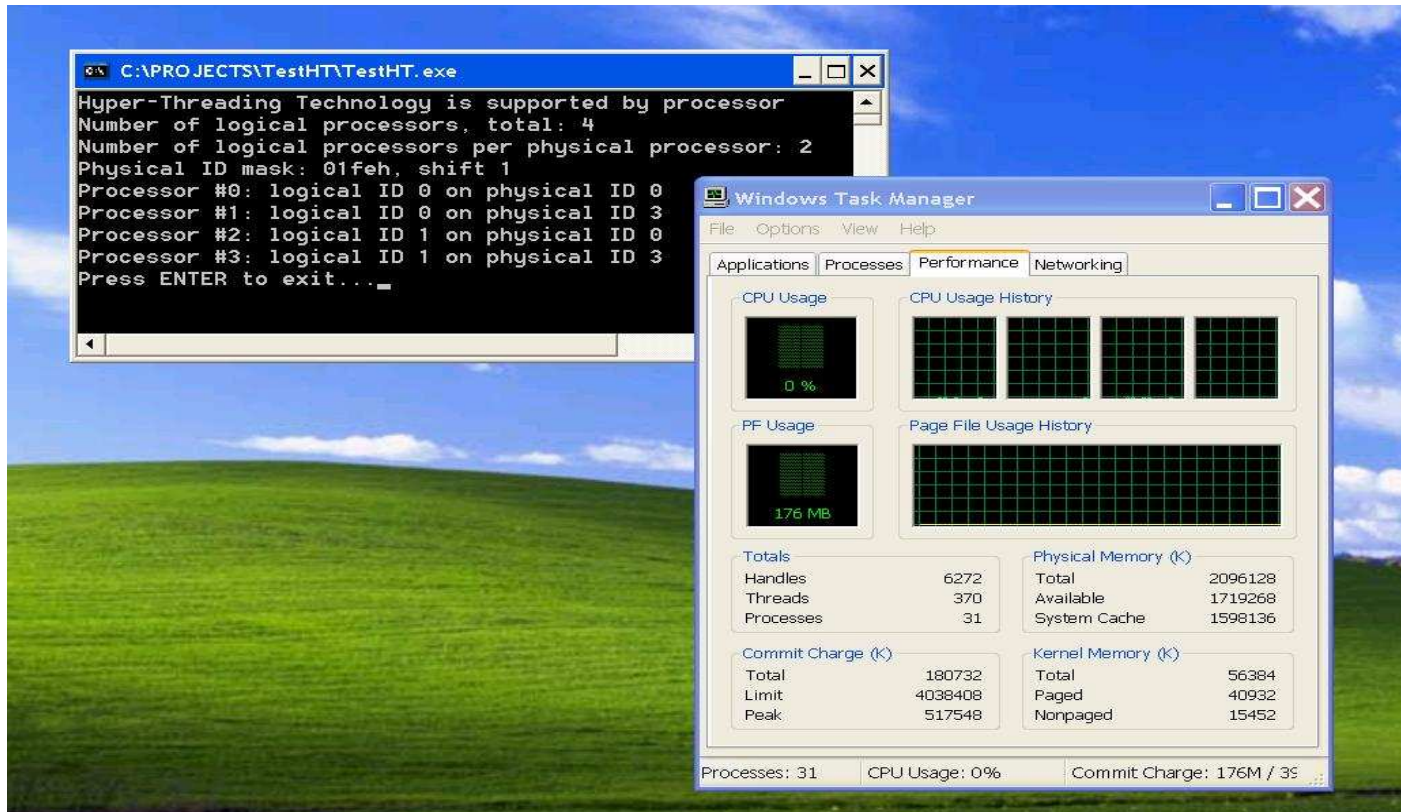
Common Threading Challenges

Intel® Software Tools for Multi-threading

Summary


# Hardware Architecture

The trend toward multi-core mobile, desktop, and server processors is expected to continue into the foreseeable future, and software must be threaded to take full advantage of it.




# Processor Evolution

Intel® Xeon®  
processor




**Single Core**

Dual-Core Intel®  
Xeon® processor  
5100 series




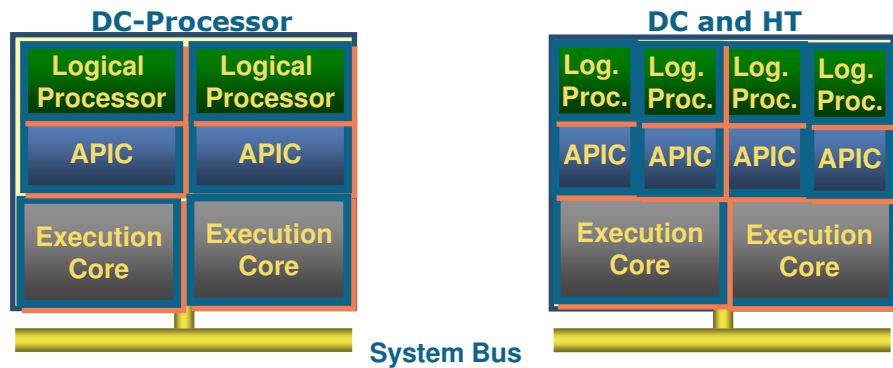
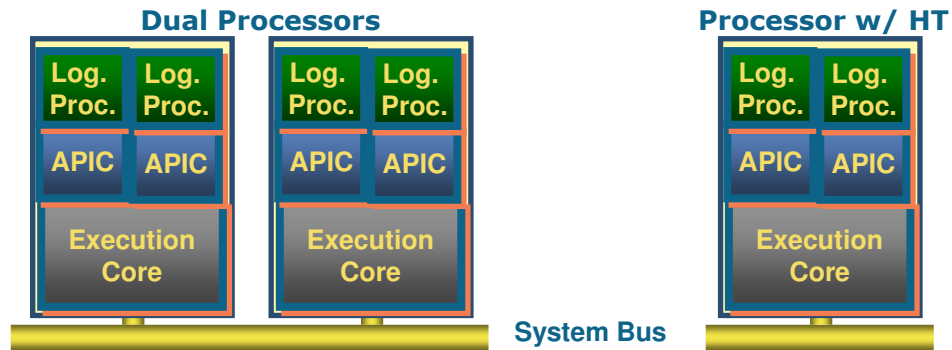
**Dual Core**

New Quad-Core  
Intel® Xeon®  
5300 for 2006

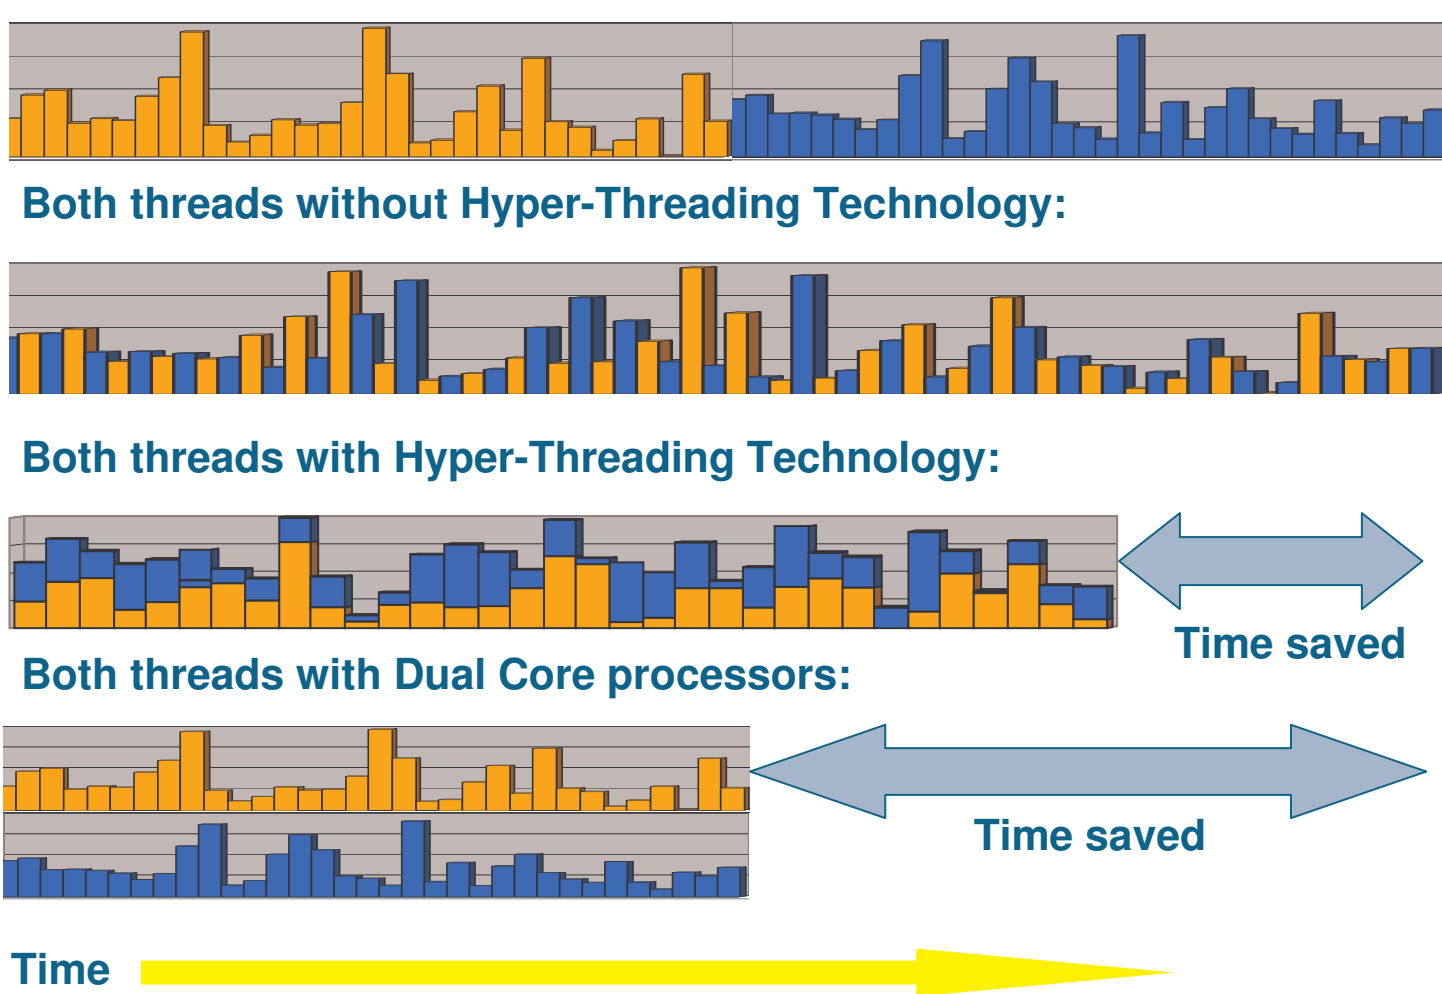


**Quad Core**

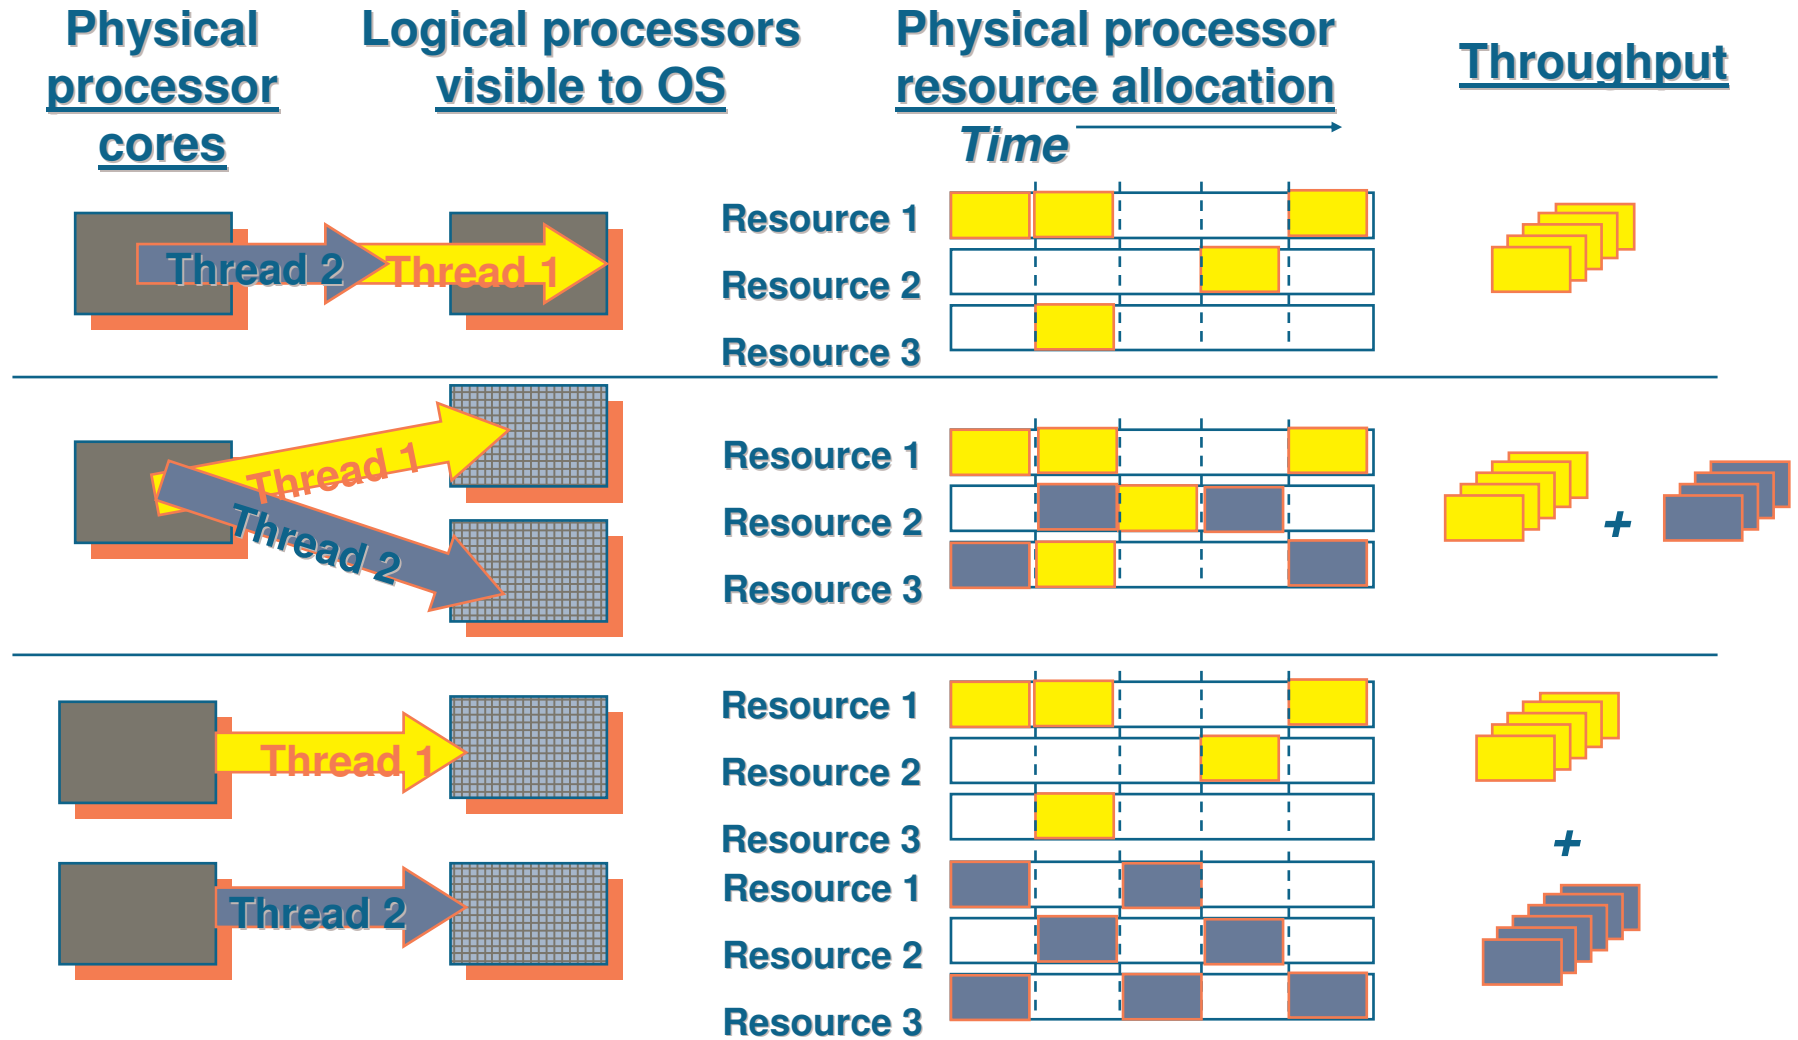
# HT Technology, Dual-Core and Multi-Core



# Hyper-Threading vs. Dual Core



# How HT Technology and Dual-Core Work




# Why use parallelism?

## Amdahl's Law

- Describes the upper bound of parallel execution speedup

### Serial code limits speedup



$$T_{\text{parallel}} = \left\{ (1-P) + \frac{P}{n} \right\} T_{\text{serial}}$$


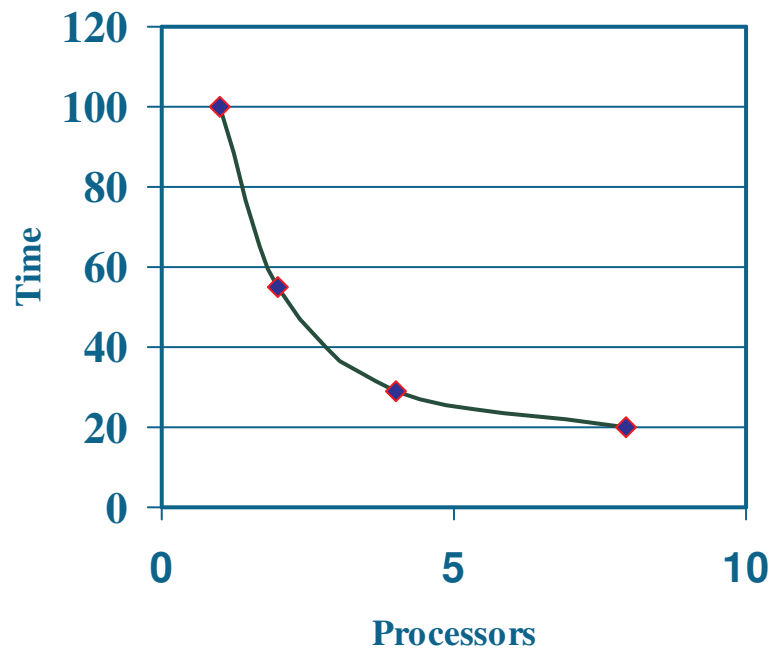
$n = \text{number of processors}$

Speedup =  $T_{\text{serial}} / T_{\text{parallel}} = 1.0 / 0.55 = 2.033$

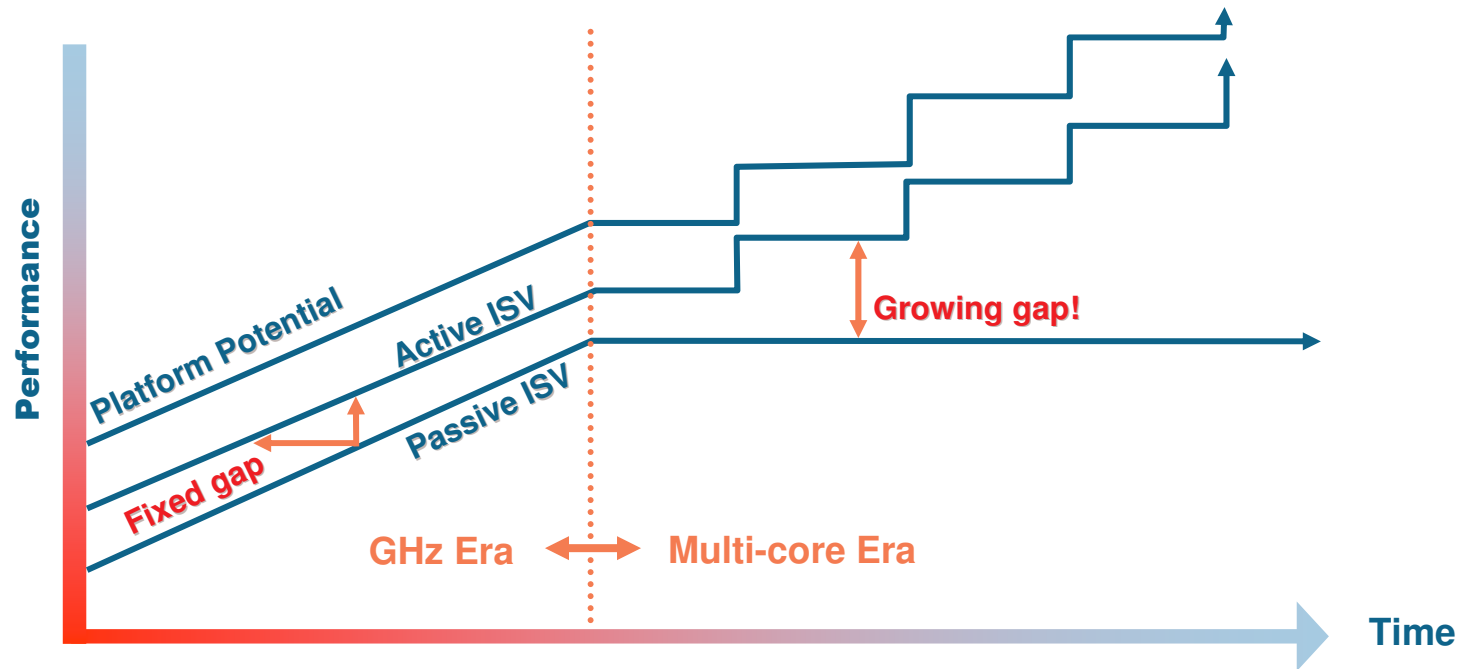


# Why use parallelism ?

- Speed up a given problem
- Scale up to a larger problem with fixed run-time



# Multi-core Processors change the rules



Get your software ready for multi-core using Intel® Tools

# Why Thread Your Application?

## **Increased responsiveness and worker productivity**

- Increased application responsiveness when different tasks run in parallel

## **Improved performance in parallel environments**

- When running computations on multiple processors

## **More computation per cubic foot of data center**

- Web-based apps are often multi-threaded by nature

Taking full advantage of Multi-Core requires multi-threaded software

# Agenda

Motivation for Threading

**Concepts in Parallelism**

Threading in Action

Common Threading Challenges

Intel® Software Tools for Multi-threading

Summary

# What is Parallelism?

Two or more processes or threads execute at the same time

Parallelism for threading architectures

- Multiple processes
  - Communication through Inter-Process Communication (IPC) – Message Passing
- Single process, multiple threads
  - Communication through shared memory

# Different levels of parallelism

- Hardware parallelism
  - Processor pipeline, execution units, SSEx
  - Two logical processors sharing the resources of one physical processor
  - SMP systems, shared memory
- Software parallelism
  - Data parallelism
  - Task parallelism

# Intel® Architecture – Hardware level

## Instruction Level Parallelism (1993)

- Out-of-order instruction pipeline
- Multiple execution units

## Data Level Parallelism (1997)


- MMX™
- Streaming SIMD Extensions (SSE)
- Streaming SIMD Extensions 2 (SSE2)
- Streaming SIMD Extensions 3 (SSE3)

## Thread Level Parallelism (2002)

- Hyper-Threading Technology (HT Technology)
- Multi-core Intel processor architecture


# Partitioning Methods

Grid reprinted with permission of Dr. Phu V. Luong,  
Coastal and Hydraulics Laboratory, ERDC



## Domain Decomposition

- **Data Parallelism:**  
Same operation applied to all data



## Functional Decomposition

- **Task Parallelism**  
Each Thread performs a unique job

**GOAL: Identify independent computations / primitive tasks**



# Most Code Contains Parallelism

## Task parallelism:

Independent subprograms

```
call fluxx(fv, fx)
call fluxy(fv, fy)
call fluxz(fv, fz)
```

## Data parallelism:

Independent loop iterations

```
for (y=0; y<nLines; y++)
    genLine(model, im[y]);
```

# Options for Adding Parallelism

**Explicitly thread your program using Win32\*/POSIX\* threading APIs**

**Use a Compiler to automatically parallelize code**

**Use a Programming Language API (C#\*, Java\*, etc.)**

**Programming Language Extension (OpenMP\*)**

- Use OpenMP directives to tell the compiler how to decompose parts of a serial program for parallel execution

**Use an internally-threaded runtime library for common tasks**

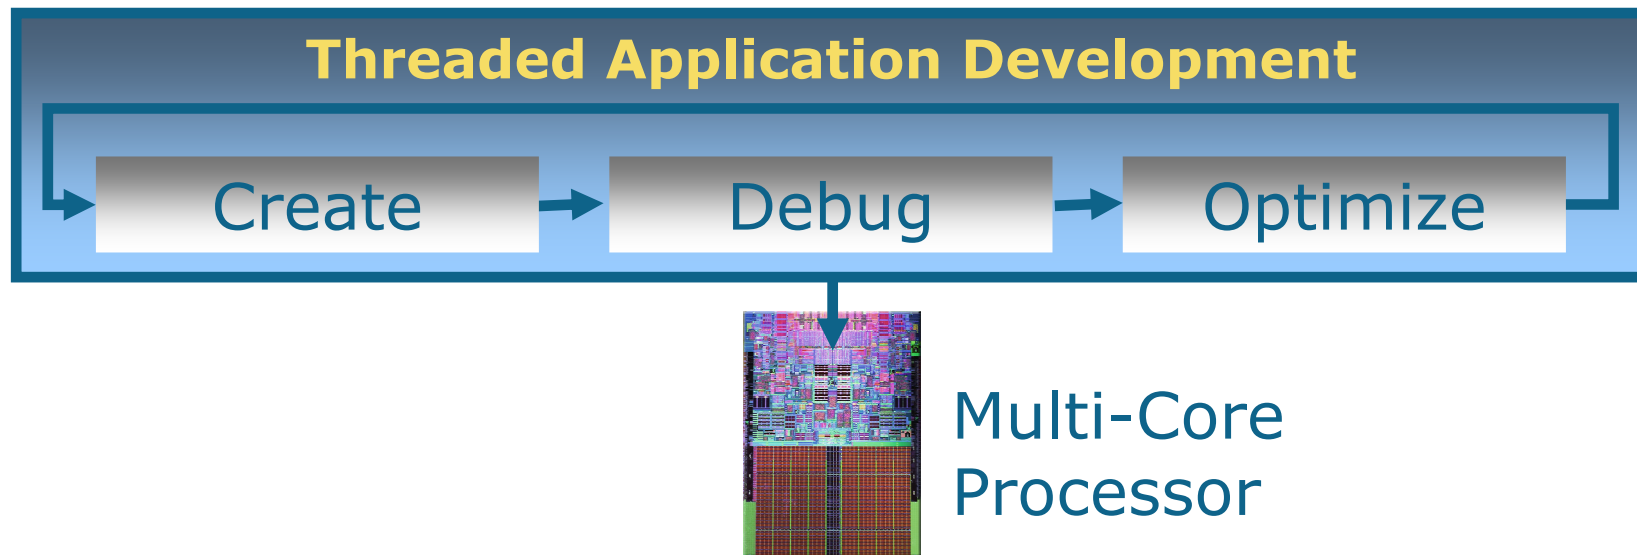
- Intel® Threading Building Blocks (TBB), Intel® Integrated Performance Primitives (Intel® IPP) and Intel® Math Kernel Library (Intel® MKL)
- Parallel memory managers like MicroQuill SmartHeap\* and Hoard\*

# Maximize Multi-Core Performance By Threading Your Software

Software threading needed at the application level

- Breaks problem into pieces that can be solved in parallel
- Performance can scale with number of processors

Need tools to create, debug and optimize multi-threaded applications



# Developer + Tools

**Use your knowledge of the algorithm to identify opportunities for parallelism**

## **Verify with tools**

- Tools will identify dependencies you overlooked
- Tools will help identify regions for greatest ROI
- Tools will improve your productivity

“More computing sins are committed in the name of efficiency (without necessarily achieving it) than for any other single reason ...”

— **William M. Wulf**

**Developer knowledge of algorithm is important!**

# Agenda

Motivation for Threading

Concepts in Parallelism

**Threading in Action**

Common Threading Challenges

Intel® Software Tools for Multi-threading

Summary

# Start Small

## *If full-blown threading looks too hard...*

Look at the code where you spend the most time

Identify code regions that would benefit from parallelism

Try to use the Intel® Compiler to parallelize tight inner loops

- /Qparallel\* and /Qx\* options
- OpenMP\* directives and /Qopenmp\* option

Use an OpenMP\*, Intel® Thread Checker, and Intel® Thread Profiler to prototype possible threading implementations

- Once you have a good algorithm, you can rewrite in a native threading API like Win32\* or pThreads\* as desired.

# Start Small (continued)

## *If full-blown threading looks too hard (continued)*

Replace calls to large common functions with calls to internally parallel libraries such as:

- Intel® Threading Building Blocks
- Intel® Integrated Performance Primitives
- Intel® Math Kernel Library

Make your libraries thread-safe in anticipation of their being called by threaded code

- Use a simple multi-threaded test harness

Consider “functional” parallelism

- Try to separate computation from unrelated tasks (such as the GUI, printing, etc.)

# Threading Methods

For performance hungry apps – parallel threads use all cores

Use OpenMP\* compiler for automated thread creation

	OpenMP	Native Threads	Intel® TBB
Portable	✓	somewhat	Intel only
Performance Oriented	✓		✓
High-Level	✓		✓
No Special Compiler		✓	✓
Serial Code Intact	✓		
Works well with C++	if C-like	✓	✓
Data Parallel	✓		✓
Functional Decomposition		✓	✓
Nested Parallelism Support	some		✓



# Example: Non-Threaded Application

```
public class Report
{
    public void Compute(int x)
    {
        x = x * x; ... // long task
    }
}

public static void main()
{
    Report r = new Report();
    r.Compute(5);
}
```

# Example: Worker Thread

```
public class Report
{
    public void Compute(int x)
    {
        this.x = x;
        mi = new MethodInvoker(AsyncCompute);
        mi.BeginInvoke(null, null);
    }
    void AsyncCompute()
    {
        x = x * x; ... // long task
    }
    MethodInvoker mi;
    int x;
}

public static void main()
{
    Report r = new Report();
    r.Compute(5);
}
```

# POSIX\* pthreads\*: Example

```
#include <stdio.h>
#include <pthread.h>
const int num_threads = 4 ;

void* thread_func(void* arg) {
    do_work();
    return NULL;
}

main() {
    pthread_t threads[num_threads];
    for( int i = 0; i < num_threads; i++)
        pthread_create(&threads[i], NULL, thread_func, NULL);

    for(int j = 0; j < num_threads; j++)
        pthread_join (threads[j], NULL);
}
```

Parallel processing

# OpenMP\* Threads: Implicit

Create  
Threads

Split loop iterations  
among threads

Make local variables  
for each thread

```
#pragma omp parallel for private(pixelX, pixelY)
```

```
for (pixelX = 0; pixelX < imageWidth; pixelX++)
```

```
{
```

```
    for (pixelY = 0; pixelY < imageHeight; pixelY++)
```

```
    {
```

```
        newImage[pixelX, pixelY] =
```

```
            ProcessPixel(pixelX, pixelY, image);
```

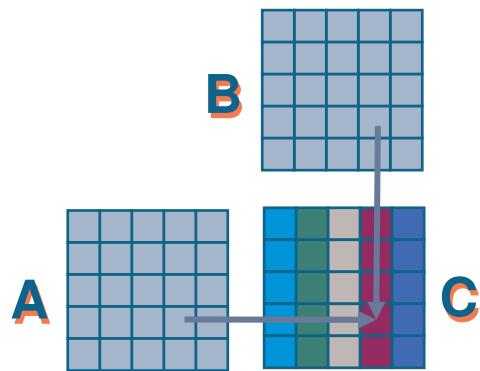
```
    }
```

```
}
```

# Example – OpenMP\* Threads

Number of threads is determined at initialization time (number of processors).

If disabled, code looks and runs like single threaded code



```
...  
  
// Divide work of outer loop between all  
processors on the system  
  
#pragma omp parallel for private(x,y)  
for (x=0; x<width; x++)  
    for (y=0; y<height; y++)  
        C(x,y) = F(A(x,y),B(x,y));  
  
...
```

# Agenda

Motivation for Threading

Concepts in Parallelism

Threading in Action

**Common Threading Challenges**

Intel® Software Tools for Multi-threading

Summary

# Challenges to Implementing Parallelism

## Correctness

- Shared resource identification
- Threading the right code correctly
- Difficult to debug
  - Data Races
  - Deadlocks

## Performance

- Program decomposition (Functional task/Data)
  - Load Imbalance
  - Granularity
- Resource utilization (load balancing)
  - Task Priorities
  - Adequate Memory and I/O bandwidth
- Parallel Overhead (Thread management and synchronization)
  - Due to thread creation, scheduling ...
- Synchronization
  - Excessive use of global data, contention for the same synchronization object
- All these affect scaling

# Program Gives Incorrect Results

## Some possible explanations:

- Race condition or storage conflicts
  - More than one thread accesses memory without synchronization
  - Locking used, but is too local to be effective
- Other components (such as 3<sup>rd</sup>-party APIs) may not be “thread safe” in certain use cases

## Debug via:

- Intel® Thread Checker
- A tool like Rational Purify\* or Compuware BoundsChecker\*



# Data Race Example

Suppose you have global variables  $a=1$ ,  $b=2$

**Thread1**  
 $x = a + b$

**Thread2**  
 $b = 42$

**The end result can be different if**

- Thread**1** runs before Thread**2**:
- Thread**2** runs before Thread**1**:

$x = 3$

$x = 43$

**Execution order is not guaranteed unless updates to “b” and its use are protected by a synchronization mechanism.**

# Threaded Program Hangs

## Possible explanations:

- Deadlocks
  - Actual (A waits on B, which is waiting on A)
  - Potential (will happen under the “right” conditions)
- Thread Stalls and Waits
  - Dangling locks (thread exits holding a lock requested by another thread)
    - Thread exit does not automatically release held locks
    - Must release lock from the same thread where you obtained (entered) it

## Debug using:

- Intel® Thread Checker
- A standard debugger
- Print statements (be sure to identify the thread)

# Deadlock Example

If Thread 1 reached here first

Thread 2 will hang here

ThreadFunc1()

```
{  
    lock(A);  
    globalX++;  
    lock(B);  
    globalY++;  
    unlock(B);  
    unlock(A);  
}
```

ThreadFunc2()

```
{  
    lock(B);  
    globalY++;  
    lock(A);  
    globalX++;  
    unlock(A);  
    unlock(B);  
}
```

**Deadlock - Both threads are now frozen**

**To fix: both functions must acquire and release locks in the same order**

# Agenda

Motivation for Threading

Concepts in Parallelism

Threading in Action

Common Threading Errors

**Intel® Software Tools for Multi-threading**

Summary

# Comprehensive, industry leading solutions for parallelized software development

Analyze your application and identify multi-core performance bottlenecks and hotspots

Highly optimizing compilers delivering scalable solutions

Detect latent programming to address unique challenges

Tune for performance and scalability



Intel has a broad toolset to help develop fast, reliable threaded applications



# Intel® Software Development Products

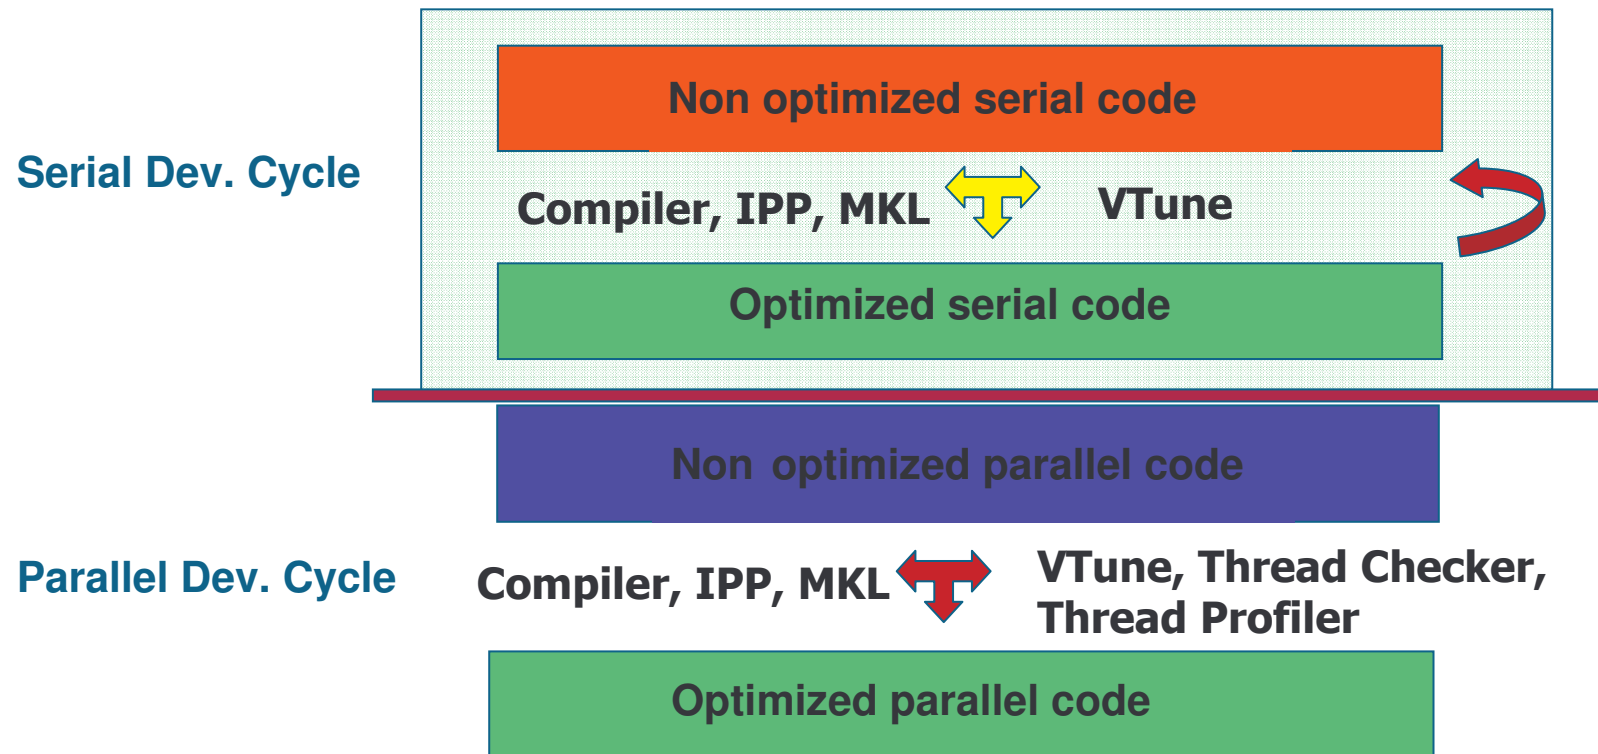
## Intel® VTune™ Performance Analyzer

- Identify “hot spots” of code that may benefit from threading
- Shows callgraph to help identify threading candidates

## Intel® Threading Tools

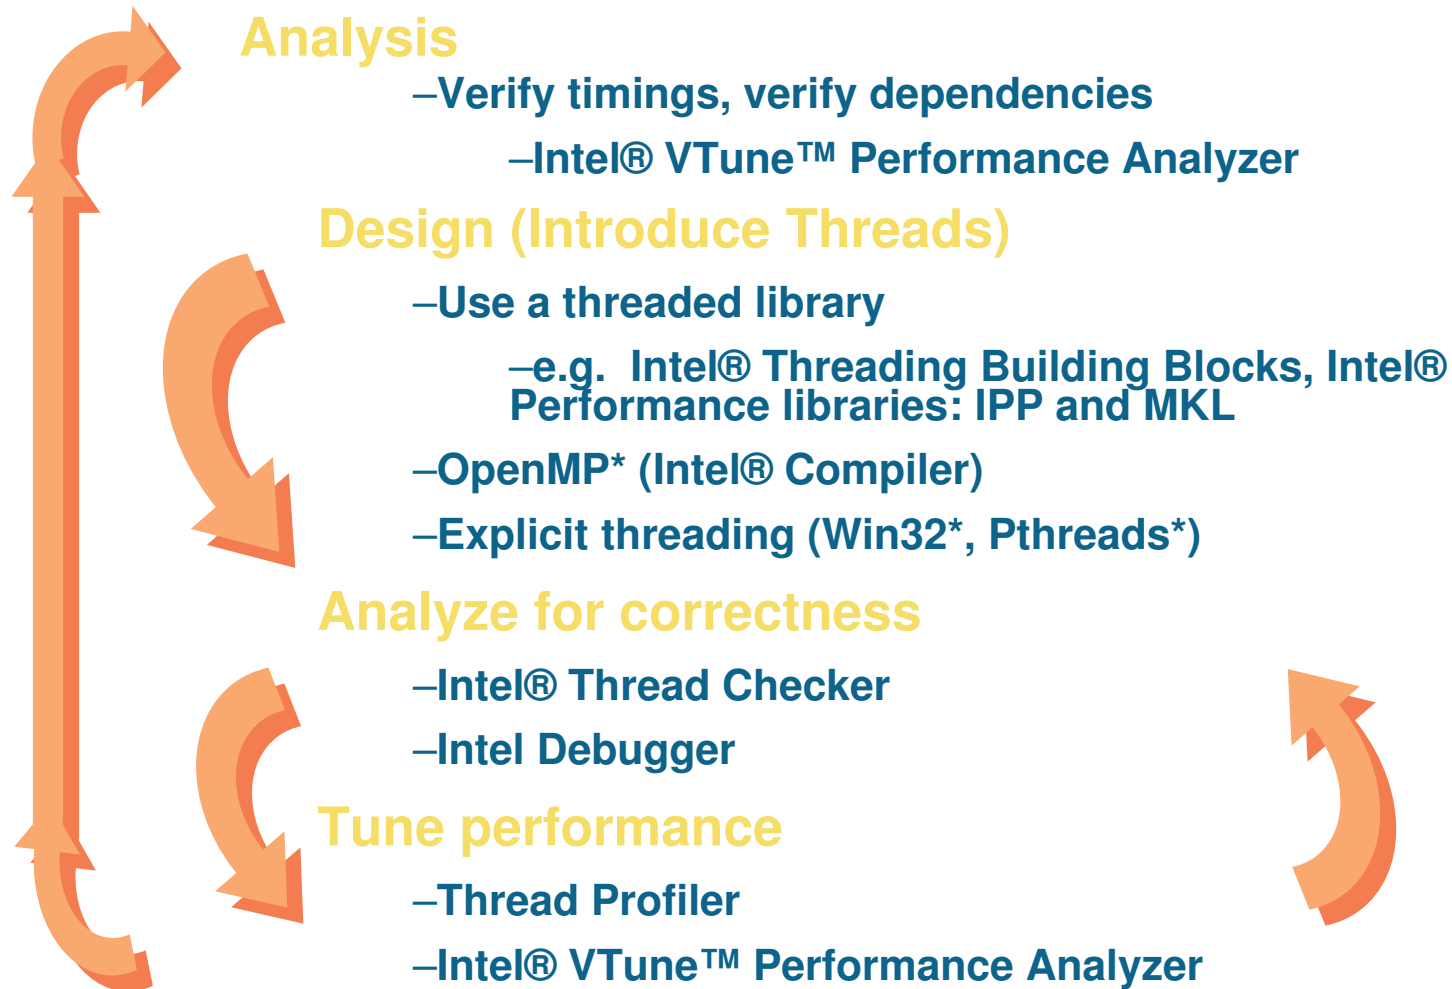
- Locate thread performance bottlenecks
- Estimate achievable/available performance
- Quickly validate designs and create prototypes
- Locate positions of data race conditions (read/write, write/read, write/write)
  - Isolate deadlock
  - Identify inappropriate API arguments

# Development Cycle



**Serial and parallel performance problems are different and need extra tools !**

# Development Cycle





# Example – Prime Number Generation

2  
 3 3  
 5 3,5  
 7 3,5,7  
 9 3  
 11 3,5,7,9,11  
 13 3,5,7  
 15 3  
 17 3,5,7

```

main(int argc, char *argv[])
{
    int i, j, limit;
    clock_t time_1, time_2;
    int start = 1000000;
    int end = 4000000; /* start:end == range of numbers to search */
    int number_of_primes=0; /* number of primes found */
  
```

```

for (i = start; i <= end; i += 2)
{
    limit = (int) sqrt((double)i) + 1;
    prime = 1; /* assume number is prime */

    for (j=3; prime && (j <= limit); j+=2)
        if (i % j == 0)
            prime = 0;
    if (prime)
        number_of_primes++;
}
  
```

```

prime = 1; /* assume number is prime */
  
```

```

for (j=3; prime && (j <= limit); j+=2)
  
```

```

c:\ Intel(R) C++ Compiler 9.0.024 build environment for 32-bit applications
C:\Documents and Settings\hlakyil\My Documents\SPD\Events-Trips\labs\threads>prime_def.exe
Range to check for Primes: 1000001 - 4000000
Found 204648 primes in 2.694 seconds
C:\Documents and Settings\hlakyil\My Documents\SPD\Events-Trips\labs\threads>
  
```

```

number_of_primes,
(double) (time_2-time_1)/CLOCKS_PER_SEC);
(double) (time_2-time_1)/CLOCKS_PER_SEC);
  
```

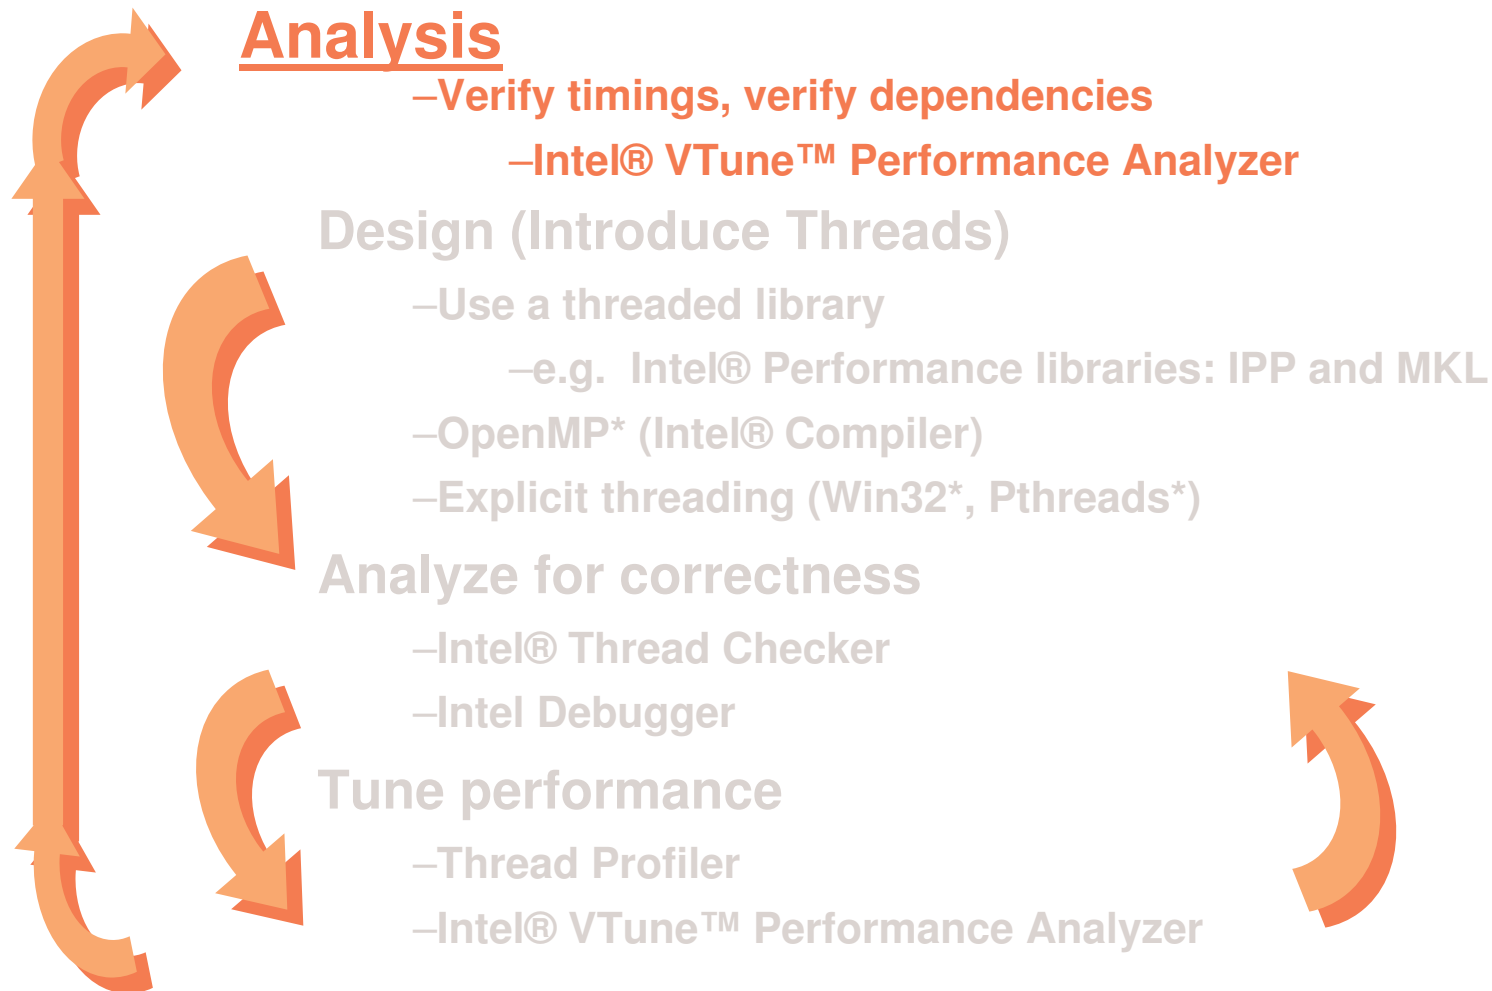
}

# Computing Primes


- Compute the number of primes between *start* and *end*
  - Not the fastest way, really a 'brute force approach'
- Numbers are checked independently of each others

```
for(i = start; i <= end; i += 2){
    limit = (int) sqrt((float)i) + 1;
    prime = 1; /* assume number is prime */
    for ( j=3; (prime && (j <= limit)); j+=2 )
        if (i%j == 0) prime = 0;
    if (prime)
        number_of_primes++;
}
```

# Development Cycle



# VTune™ Performance Analyzer



Helps you identify and characterize performance issues by:

- Collecting performance data from the system running your application
  - Sampling: Event-based or Time-based
  - Call Graph
  - Counter Monitor
- Organizing and displaying the data in a variety of interactive views
  - From system-wide down to source code or processor instruction perspective
  - GUI and CLI
- Identifying potential performance issues and suggesting improvements

# Intel<sup>®</sup> VTune<sup>™</sup> Performance Analyzer


*Identifies hard to find performance bottlenecks*

## Via Call Graph

- Using instrumented application
- Some impact on execution time
- All details related to function calls collected at run-time
- Provides call relationships, frequency of calls, time spent in functions
- Quickly identifies critical path in call-tree

## Via Sampling

- Using statistical sampling approach to analyze where time (CPU cycles) is spent or micro-architectural event like cache misses are occurring
- Processor is interrupted in fixed small intervals; execution context is saved for later analysis
- Almost not intrusive at all – less than 5%
- GUI provide hierarchical display of sampling data in many ways like e.g. 'Sampling over execution time'



**Tuning Browser**

- VTProject1
  - Activity1 (Sampling)
    - Sampling Results [HLAKYIL-MOBL] - Fri Mar
    - Run 1
      - Clockticks
      - Instructions Retired
  - Activity2 (Sampling)
    - Sampling Results [HLAKYIL-MOBL] - Tue Au
    - Run 1
      - L2 Cache Requests
      - L2 Cache Request Misses (highly co

Process	Clockticks sam... (325) - Sam...	Instructions Retired... (325) - Sampling R...	L2 Cache Read Miss... (327) - Sampling Res...	L2 Cache Reads sam... (327) - Sampling Res...	Cycles per Retired Instruction - C... (325) - Sampling Results [HLAK...	Clockticks % (325) - Sampling Results [HLAKYIL-MOBL]	Instr...	Total
matrix.exe	1,549	575	2,516	30,445	2,694	94.97		30,515.00
mcshield.exe	25	23	0	0	1.087	1.53		0.00
VTuneEnv.exe	21	12	9	60	1.750	1.29		3,310.00
PDWERPNT.EXE	9	5	2	1	1.800	0.55		1,176.00
blackd.exe	6	4	4	13	1.500	0.37		152,575,000.00
Explorer.EXE	5	1	2	28	5.000	0.31		96.41
csrss.exe	3	1	3	9	3.000	0.18		0.00
pid_0x4	3	0	2	1	N/A	0.18		0.00
Skype.exe	2	0	1	3	N/A	0.12		2,648,000,000.00
UpdaterUI.exe	1	0	1	2	N/A	0.06		90.34
CopernicDesktopSearch.exe	1	0	2	3	N/A	0.06		940,800,000.00
Alt2evxx.exe	1	1	0	2	1.000	0.06		87.43
1XConfig.exe	1	0	0	0	N/A	0.06		2.81
waatservice.exe	1	0	0	0	N/A	0.06		
lsass.exe	1	0	0	0	N/A	0.06		
services.exe	1	0	0	0	N/A	0.06		
svchost.exe	1	0	1	0	N/A	0.06		
SynTPLpr.exe	0	0	1	0	N/A	0.00		
wmiapsrv.exe	0	0	0	1	N/A	0.00		
clfmom.exe	0	1	0	0	0.000	0.00		
ccmhelp.exe	0	0	1	1	N/A	0.00		
TpShocks.exe	0	0	1	2	N/A	0.00		
EvtEng.exe	0	0	0	1	N/A	0.00		

**Module of Interest Process view**

Legend

Event	Activity ID	Scale	Sample After Value	Total Samples	Duration (s)	Ring 0	Ring 3	Start Time	Machine Name	Processor
L2 Cache Requests	398	0.00000010000x	5000	31650	1.92	656	30994	8/22/2006 4:53:59 PM	HLAKYIL-MOBL	Intel(R) Pentium(R) M processor
L2 Cache Request Misses (highly correlated)	398	1.000000000000x	5000	0	1.92	0	0	8/22/2006 4:53:59 PM	HLAKYIL-MOBL	Intel(R) Pentium(R) M processor
Clockticks	38	0.000000010000x	800000	3664	1.55	241	3423	3/24/2006 11:16:10 AM	HLAKYIL-MOBL	Intel(R) Pentium(R) M processor
Instructions Retired	38	0.00000010000x	800000	1345	1.55	97	1248	3/24/2006 11:16:10 AM	HLAKYIL-MOBL	Intel(R) Pentium(R) M processor
Cycles per Retired Instruction - CPI	38	1.000000000000x	0	0	0.00	0	0	---	---	---

211 items, 5 events, 1 item(s) selected.

**Output**

**General**

Tue Aug 22 16:53:58 2006 <localhost> (Run 1) The Sampling Collector is collecting samples based on the following event(s): L2 Cache Requests, L2 Cache Request Misses (highly correlated).  
 Tue Aug 22 16:54:03 2006 <localhost> (Run 1) Sampling data was successfully collected.

# VTune™ Call graph: Application Workflow

VTune(TM) Performance Analyzer - [Call Graph - [Call Graph Results - [NNWEHARCHEN] - Fri Feb 04 10:58:05 2005]]

File Edit View Activity Configure Window Help

Activity4 (Call Graph)

Module (14)	Thread (14)	Function (14)	Class (14)	Calls (14)	Self Time (14)	Total Time (14)	Callers (14)	Callees (14)	Module Path (14)
VortexMoveme...	Thread_0(EEC)	GetRuntimeClass	CVortex...	1	0	0	1	0	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	NewVortexes	CVortex...	100	2,729	2,729	1	0	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	OnDraw	CVortex...	102	35,704	248,927	2	30	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	OnInitialUpdate	CVortex...	1	1	22	1	3	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	OnMyRepaint	CVortex...	101	1,101	1,899,692	1	9	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	OnRunRun	CVortex...	1	4	13	1	2	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	operator delete	CObject	3	1	3	3	1	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	PreCreateWindow	CVortex...	1	1	1	1	2	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	WinMain		1	1	14,834,245	1	1	s:\dp.SEDUsers\ha..
VortexMoveme...	Thread_0(EEC)	WinMainCRTStartup		1	1,238	14,840,430	1	10	s:\dp.SEDUsers\ha..

Call Graph Results

Activity4 (Call Graph)

Call Graph Results

Graph Call List

Output

Instrumentation Results

Fri Feb 04 10:57:57 2005 Data collection finished...

Fri Feb 04 10:57:57 2005 Updating call graph database...

Fri Feb 04 10:58:05 2005 Done.

# Demo #1

- Analyze the simple Prime application
- Identify hotspots and opportunities for parallelism



# VTune™ Analyzer - Sampling Demo on prime

**VTune(TM) Performance Environment - [Sampling Results]**

File Edit View Activity Configure Window Help

**Tuning Browser**

- FT\_Demo
  - Prime Default [Sample]
    - Sampling Results [EC]
      - Run 1
        - CPU\_CLK\_UNHALTED.CORE
        - INST\_RETIRED.ANY
    - Prime OMP1 Sampling
      - Sampling Results [EC]
        - Run 1
          - CPU\_CLK\_UNHALTED.CORE
          - INST\_RETIRED.ANY

**Source**

Address	Line	Source	CPU_CLK_UNHALTED.CORE	INST_RETIRED.ANY
0x1063	15	int prime;		
0x1063	16	int print_primes=0;		
0x1066	17			
0x1066	18	if (!(start % 2)) start++;		
0x107E	19			
0x107E	20	printf("Range to check for Primes: %d - %d\n\n",start, end);		
0x10A0	21	time_1=clock();		
0x10B0	22			
0x10B0	23	for(i = start; i <= end; i += 2)	7	8
0x10B0	24	{		
0x10D7	25	limit = (int) sqrt((double)i) + 1;	15	1
0x111F	26	prime = 1; /* assume number is prime */		
0x1126	27	for (j=3; prime && (j <= limit); j+=2)	220	46
0x114D	28	if (i%j == 0)	1,863	1,459
0x1160	29	prime = 0;	9	4
0x1169	30	}		
0x1169	31	if (prime)	1	2
0x1174	32	number_of_primes++;		
0x1174	33	}		
0x1174	34	}		
0x117C	35	time_2=clock();		
0x118C	36	printf("\nFound %d primes in %f seconds\n",		
0x119B	37	number_of_primes,		
0x11A2	38	(double) (time_2 - time_1) / (double)CLOCKS_PER_SEC);		
0x11A2	39			
0x11A2	40			
0x11A2	41			

**Call Graph**

- main
  - sqrt

**Call Stack**

- main
  - CPU\_CLK\_UNHALTED.CORE 2,104 x 2,161,000 = 4,546,744,000 (99.858%)
  - INST\_RETIRED.ANY 1,529 x 2,161,000 = 3,304,169,000 (99.935%)

Event	Activity ID	Scale	Sample After Value	Total Samples	Duration (s)	Ring 0	Ring 3	Start Time	Machine Name	
<input checked="" type="checkbox"/>	CPU_CLK_UNHALTED.CORE	8	0.00000001000x	2161000	2339	2.66	107	2232	9/12/2006 2:06:57 PM	EDDNC6400M7400-
<input checked="" type="checkbox"/>	INST_RETIRED.ANY	8	0.00000001000x	2161000	1667	2.66	29	1638	9/12/2006 2:06:57 PM	EDDNC6400M7400-

2 items, 2 events, 1 item(s) selected.

Sampling Modules - [Sampling Results [EDDNC6400M7400-] - Tue Sep 12 14:07:00 2006]    Sampling Hotspots - [Sampling Results [EDDNC6400M7400-] - Tue Sep 12 14:07:00 2006]

**Output**

**General**

Tue Sep 12 14:09:33 2006 EDDNC6400M7400- (Run 1) The Sampling Collector is collecting samples based on the following event(s): CPU\_CLK\_UNHALTED.CORE, INST\_RETIRED.ANY.  
 Tue Sep 12 14:09:35 2006 EDDNC6400M7400- (Run 1) Sampling data was successfully collected.

For Help, press F1

# VTune™ Analyzer – Call Graph Demo on prime

VTune(TM) Performance Environment - [Call Graph - [Call Graph Results - [EDDNC6400M7400-] - Tue Sep 12 15:13:22 2006]]

File Edit View Activity Configure Window Help

Prime Default (Sampling)

**Tuning Browser**

- FT\_VT\_Demo
  - Prime Default (Sampling)
    - Prime OMP1 Sampling
    - Prime OMP2 Sampling
    - Prime OMP3 Sampling
    - Prime OMP4 Sampling
    - Prime Default CallGraph
      - Call Graph Results - [EDDNC6400M7400-]
        - Total Time
        - Self Time
        - Number of Calls
        - Call Site Total Time
        - Call Site Number of Calls

Module (40)	Thread (40)	Function (40)	Class (40)	Calls (40)	Self Time ...	Total Time (40)	% in function (40)	Callers (40)
prime_def.exe	Thread_0(AF4)	main		1	2,147,609	2,183,761	98.3%	1
libmmd.dll	Thread_0(AF4)	sqrt		1,500,000	35,613	35,613	100.0%	1
NTDLL.DLL	Thread_0(AF4)	CsrClientCallServer		11	388	388	100.0%	2
msvcrt.dll	Thread_0(AF4)	printf		2	277	478	57.9%	1
prime_def.exe	Thread_0(AF4)	mainCRTStartup		1	264	2,184,117	0.0%	1
NTDLL.DLL	Thread_0(AF4)	RtlAllocateHeap		153	61	61	100.0%	11
NTDLL.DLL	Thread_0(AF4)	NtOpenKey		1	43	43	100.0%	1
KERNEL32.DLL	Thread_0(AF4)	MultiByteToWideChar		602	38	38	100.0%	4

```

    graph LR
      Thread_0(AF4) --> GetFileType
      Thread_0(AF4) --> BaseProcessInitPostIm...
      Thread_0(AF4) --> mainCRTStartup
      mainCRTStartup --> _p_commode
      mainCRTStartup --> _getmainargs
      mainCRTStartup --> _p_fmode
      mainCRTStartup --> _p__initenv
      mainCRTStartup --> _set_app_type
      mainCRTStartup --> _setdefaultprecision
      mainCRTStartup --> main
      main --> _kmp_c_end
      main --> _kmp_c_begin
      main --> printf
      main --> clock
      main --> _kmp_c_global_thread...
      main --> sqrt
    
```

Function: sqrt  
 Module: c:\Documents and Settings\Administrator\My Documents\ToolsDemo\jabs\libmmd.dll  
 Source:  
 Total Time: 35,613 mcs  
 Self Time: 35,613 mcs  
 Total Wait Time: 0 mcs  
 Self Wait Time: 0 mcs  
 Calls: 1,500,000

Last command: Fit in window 18 nodes, 17 edges; (18 and 17)

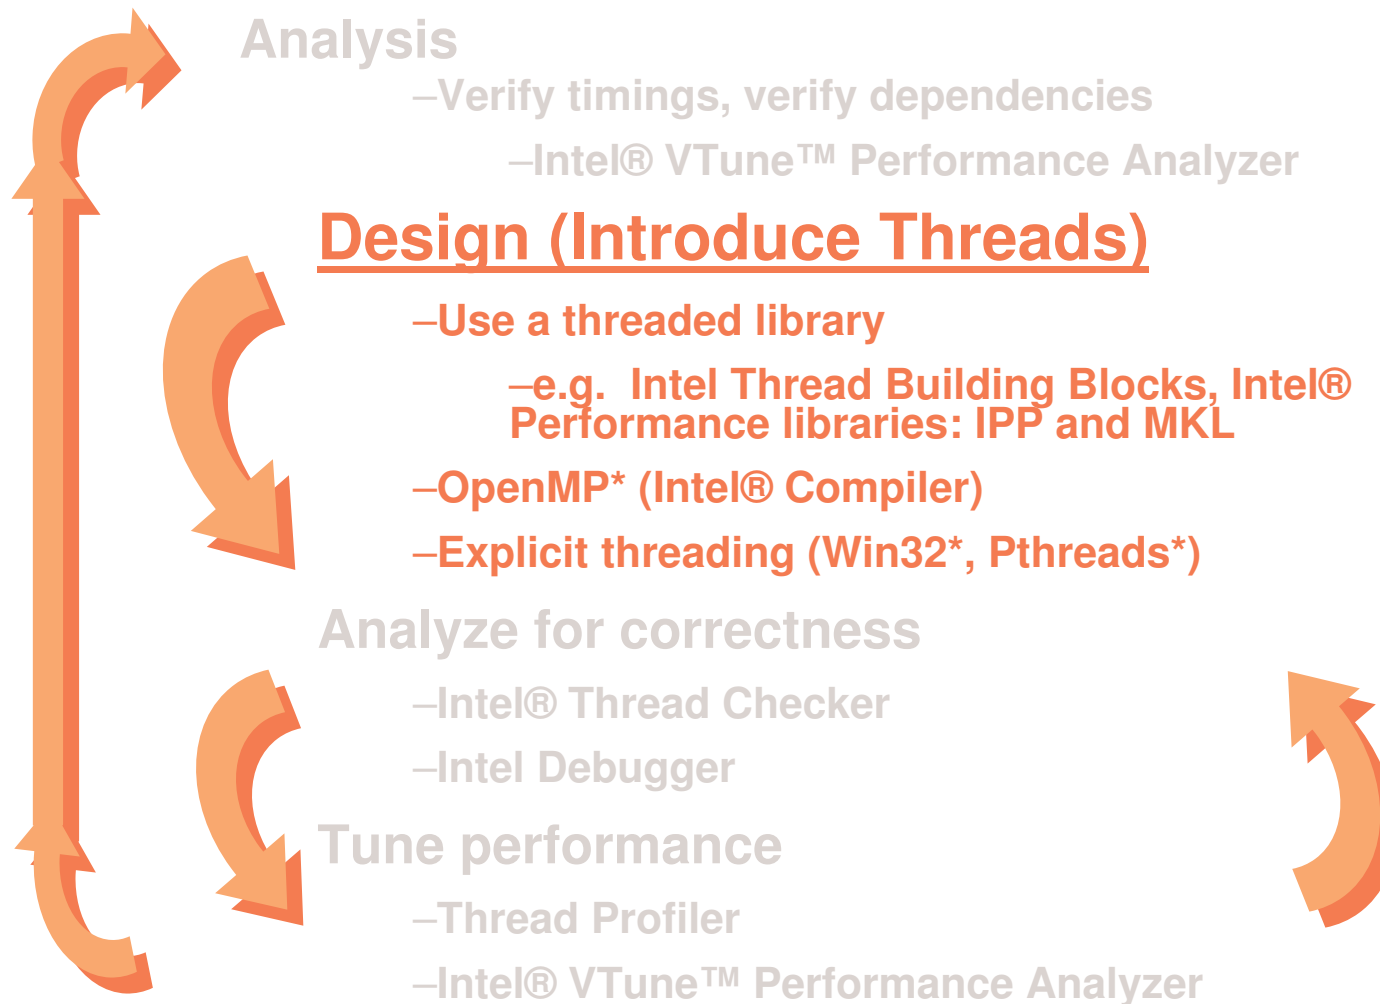
Graph Call List

Output

General

For Help, press F1

# Development Cycle



# Intel® Compilers: OpenMP\* Support

Full support of latest OpenMP\* 2.5 standard

- /Qopenmp
- /Qopenmp\_report{0|1|2|3}
- /Qopenmp\_profile

Automatic parallelization

- /Qparallel
- /Qparthreshold<n> // control over which loops to parallelize

Environment variables to control run time behavior

- E.g. OMP\_NUM\_THREADS




# OpenMP\* Architecture

- Fork-join model
- Work-sharing constructs (for, section, single)
- Data environment constructs (private, shared)
- Synchronization constructs (master, critical, barrier, atomic)
- Extensive Application Program Interface (API) for finer control

# Programming Model

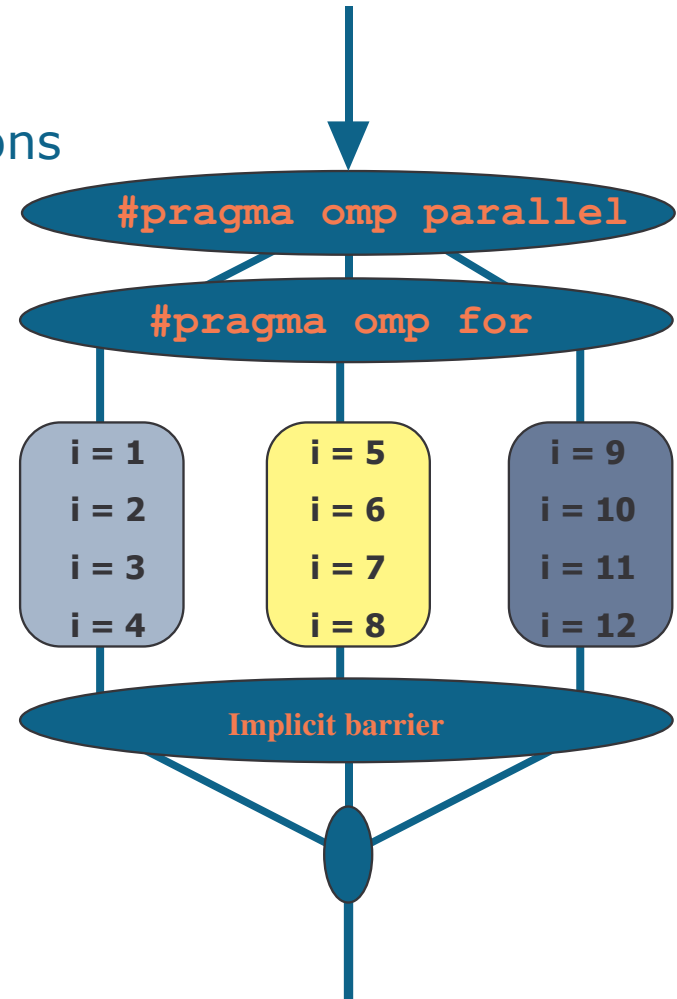
- Master thread spawns a team of threads
- Parallelism is added incrementally: the sequential program evolves into a parallel program



# OpenMP\* Parallel Loop Model

Threads are created – typically one per core  
Each threads executes a subset of all iterations  
Data is classified as shared or private

```
#pragma omp parallel for  
for(i = 1, i < 13, i++)  
    c[i] = a[i] + b[i]
```



\* Other brands and names may be claimed as the property of others.



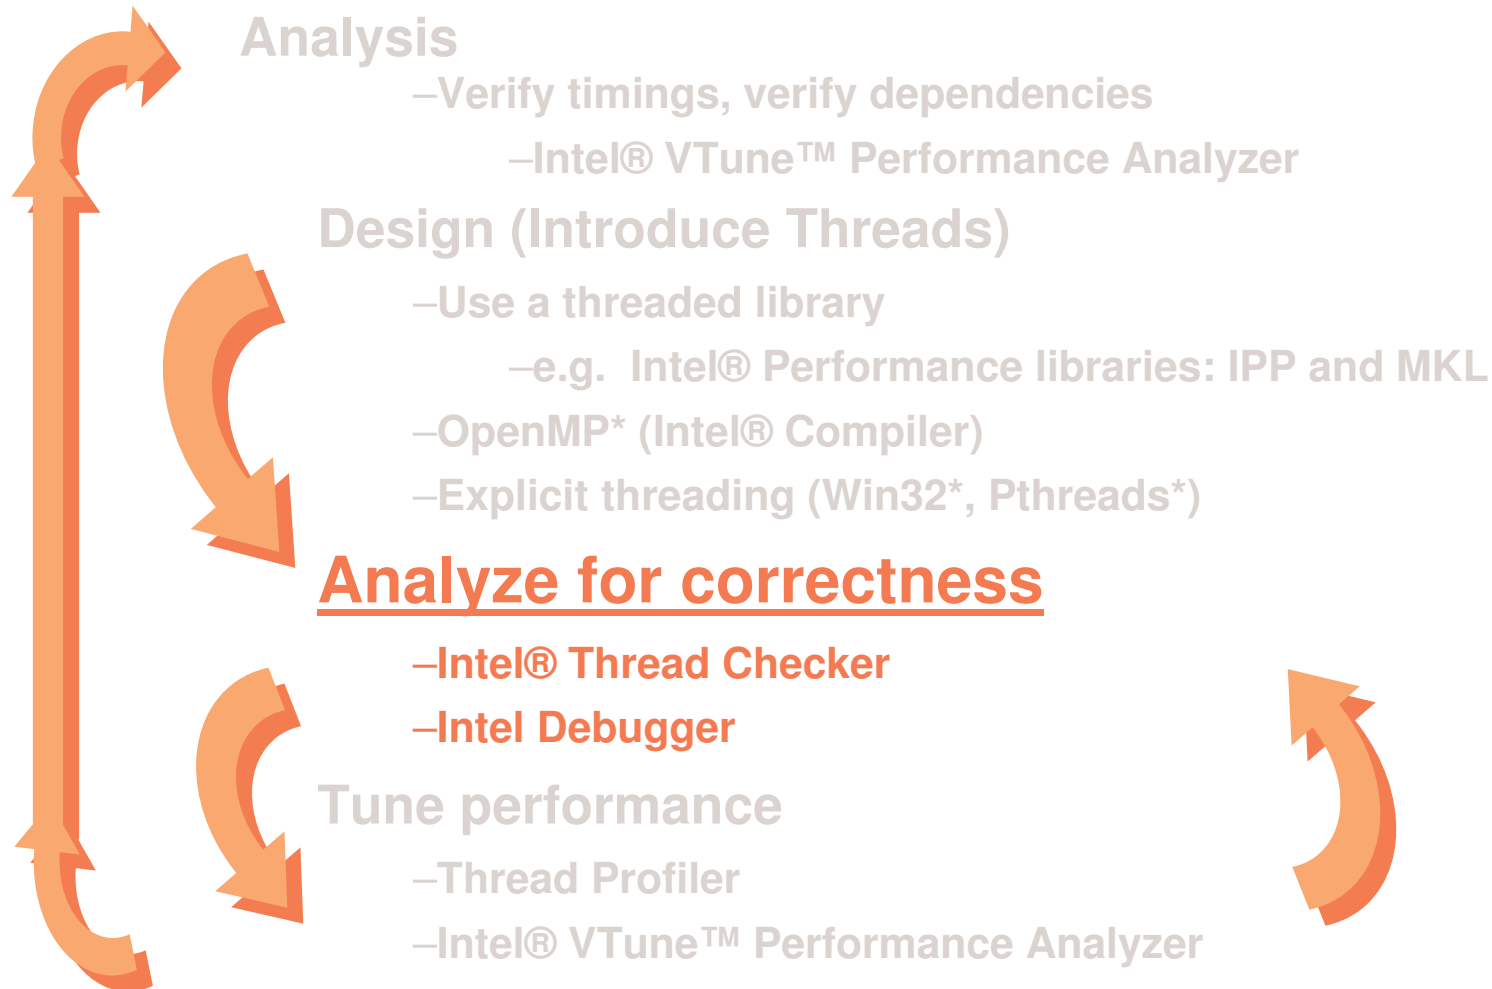
# Demo #2

1. Parallelization, compilation for OpenMP\*
2. Compilation and execution of sequential code
  - `icl /MD /Zi /Qopenmp -o prime_omp1 prime_omp1.c`

```
#pragma omp parallel for private (j, limit, prime)
for(i = start; i <= end; i += 2){
    limit = (int) sqrt((float)i) + 1;
    prime = 1; /* assume number is prime */
    for ( j=3; (prime && (j <= limit); j+=2 )
        if (i%j == 0) prime = 0;
    if (prime)
        number_of_primes++;
}
```



# Development Cycle



# Intel® Thread Checker



Identification of data safety issue:

- Dead-locks, race conditions etc

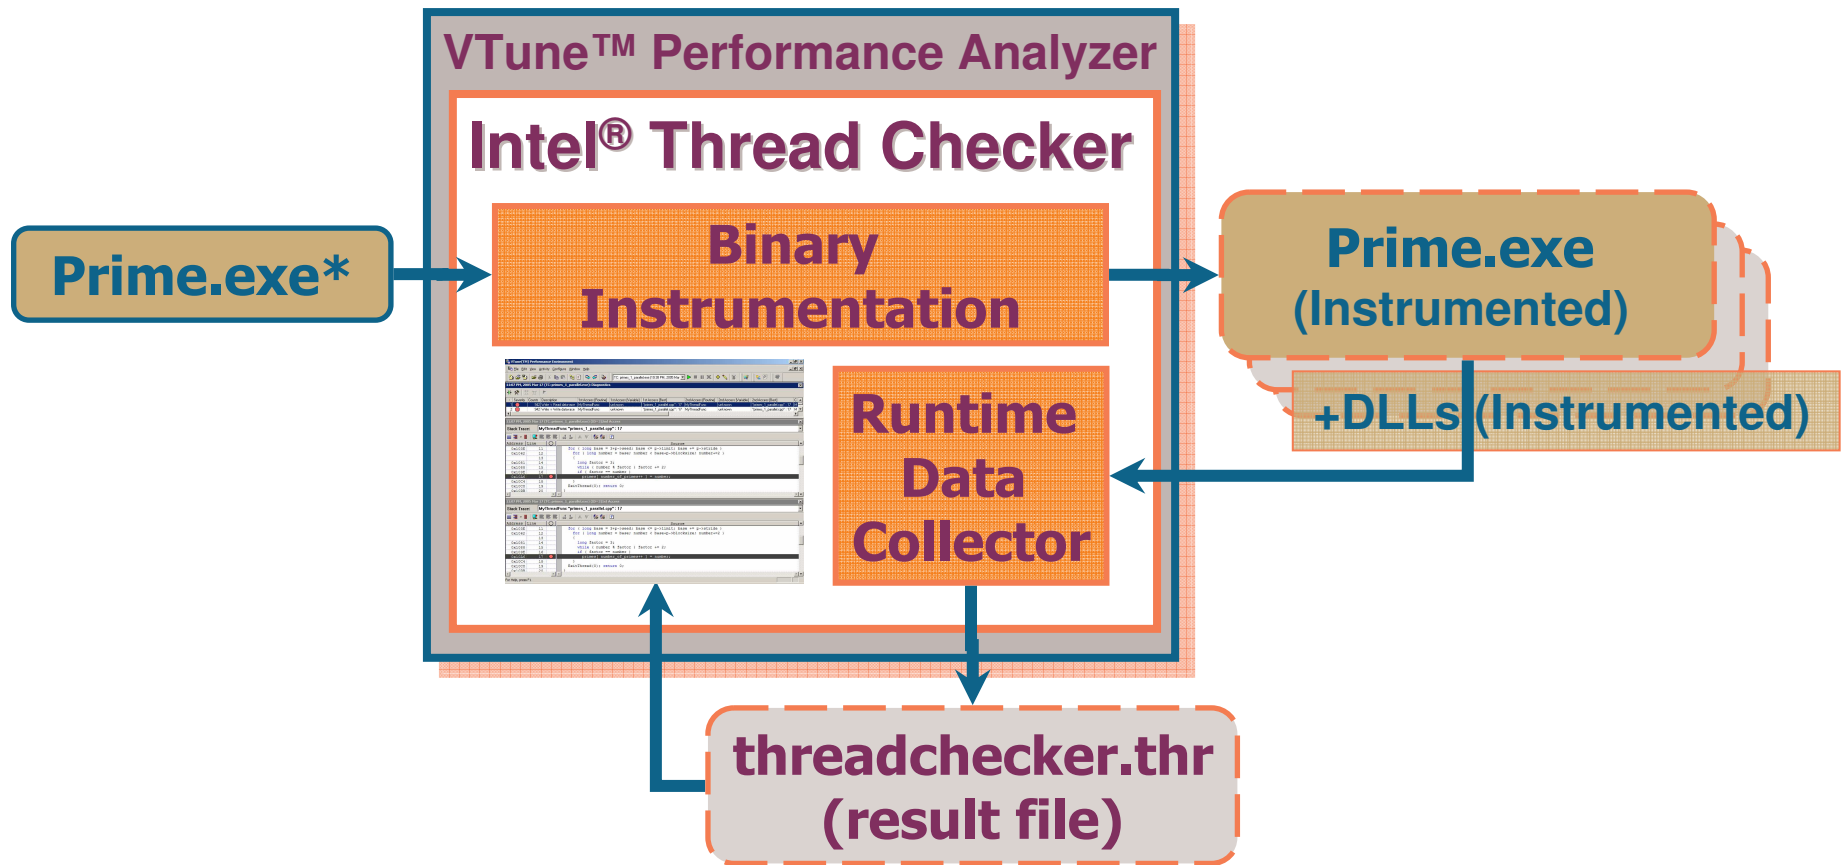
Supported environments:

- OpenMP\*
- Native thread-API on Microsoft Windows\* systems (Win32 Threads) and Linux\* (PThreads)

Analysis :

- Instrumented application
- Dynamic monitoring while application runs

# Intel® Thread Checker



\* Compiled: /Zi, Linked: /debug /fixed:no


## Demo #3

- Intel® Thread Checker analysis to find potential thread-specific ( OpenMP\*-specific) coding issues

```
#pragma omp parallel for private (j, limit, prime)
for(i = start; i <= end; i += 2){
    limit = (int) sqrt((float)i) + 1;
    prime = 1; /* assume number is prime */
    for ( j=3; (prime && (j <= limit); j+=2 )
        if (i%j == 0) prime = 0;
    if (prime)
        number_of_primes++;
}
```


# Pre-emptive context switching and shared memory conflicts (Data Races)

number\_of\_primes++



```
<thread 1>
mov ecx, dword ptr [number_of_primes]
add ecx, 1
mov dword ptr [number_of_primes], ecx

<thread 2>
mov ecx, dword ptr [number_of_primes]
add ecx, 1
mov dword ptr [number_of_primes], ecx
```

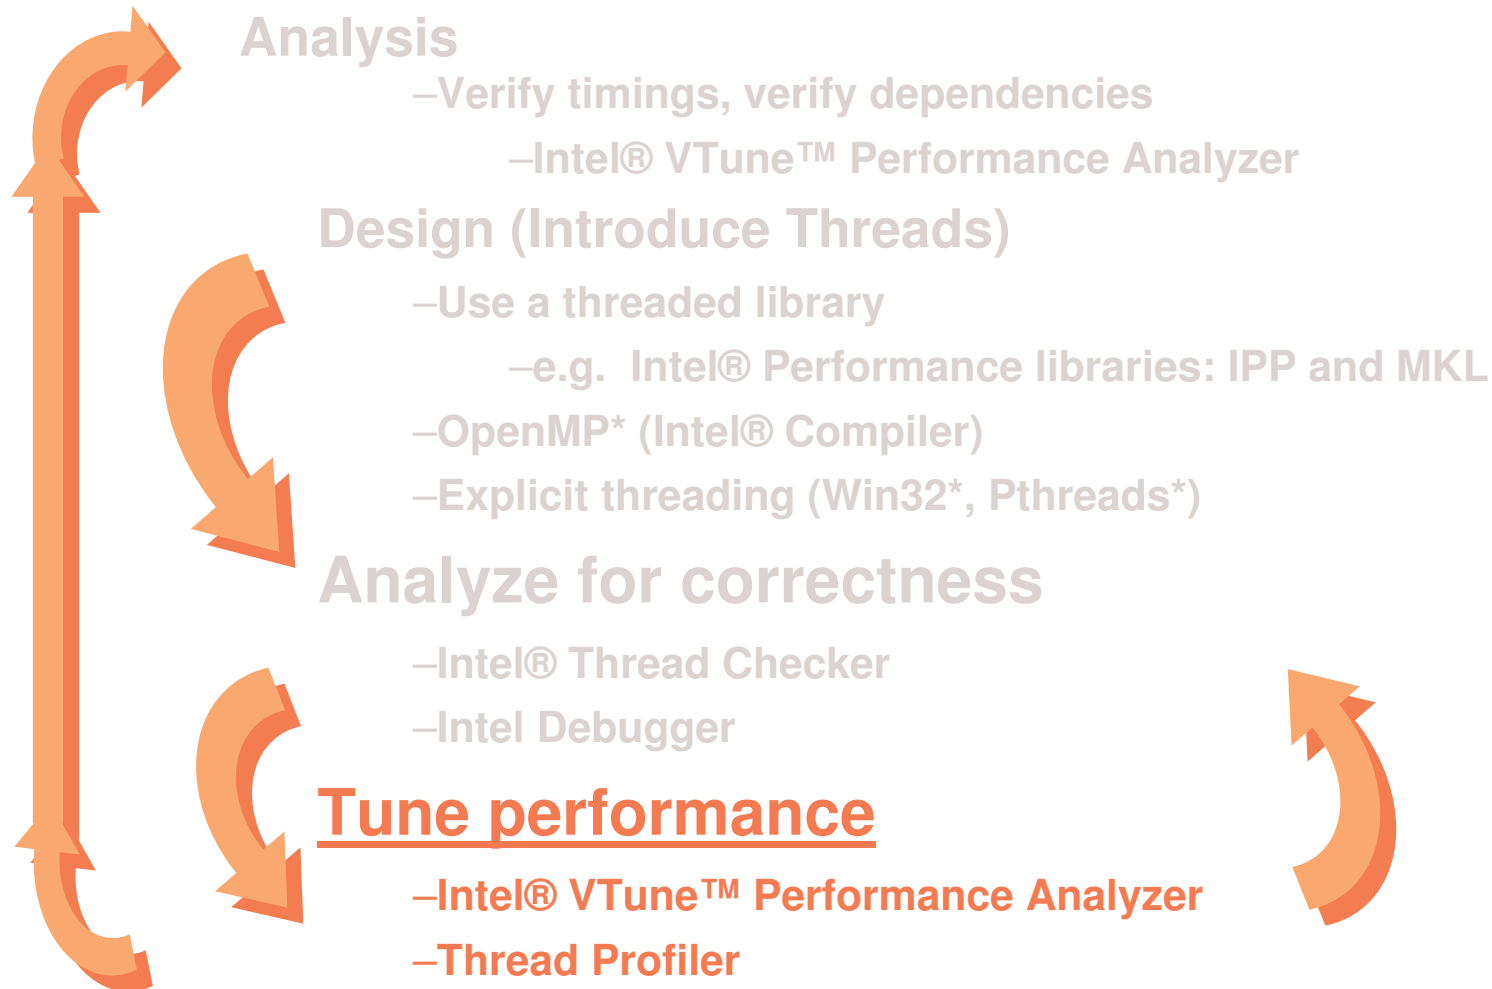


## Demo #3 continued

- Intel® Thread Checker analysis after guarding critical instruction
  - `icl /MD /Zi /Qopenmp -o prime_omp2 prime_omp2.c`

```
#pragma omp parallel for private (j, limit, prime)
for(i = start; i <= end; i += 2){
    limit = (int) sqrt((float)i) + 1;
    prime = 1; /* assume number is prime */
    for ( j=3; (prime && (j <= limit)); j+=2 )
        if (i%j == 0) prime = 0;
    if (prime)
        #pragma omp critical
            number_of_primes++;
}
```

# Development Cycle



# Intel<sup>®</sup> VTune<sup>™</sup> Performance Analyzer


*Identifies hard to find performance bottlenecks*

## Via Callgraph

- Using instrumented application
- Some impact on execution time
- All details related to function calls collected at run-time
- Provides call relationships, frequency of calls, time spent in functions
- Quickly identifies critical path in call-tree

## Via Sampling


- Using statistical sampling approach to analyze where time (CPU cycles) is spent or micro-architectural event like cache misses are occurring
- Processor is interrupted in fixed small intervals; execution context is saved for later analysis
- Almost not intrusive at all – less than 5%
- GUI provide hierarchical display of sampling data in many ways like e.g. "Sampling over execution time"





## Demo #4

- Using *Sampling over Time* feature of Intel® VTune™ Performance Analyzer for initial load balancing checking



Threads are not balanced


# Assigning Loop Iterations in OpenMP\*: Which Schedule to Use

Schedule Clause	When To Use
STATIC	Predictable and similar work per iteration
DYNAMIC	Unpredictable, variable work per iteration
GUIDED	Special case of dynamic to reduce scheduling overhead
RUNTIME	Modify schedule at run-time via environment variable

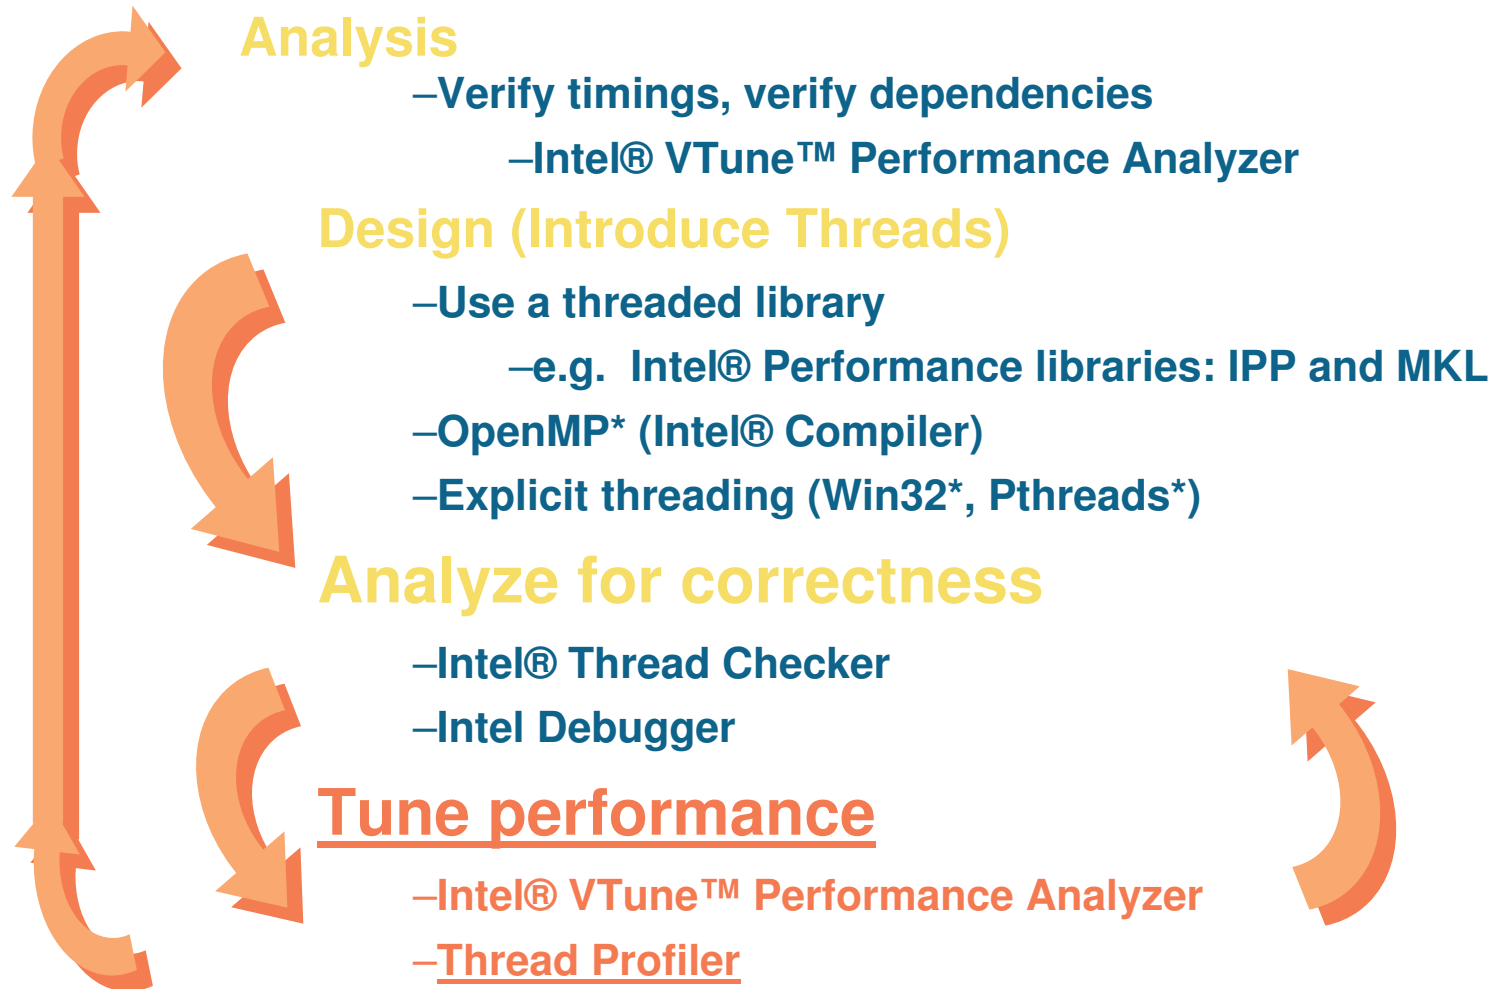
## Demo #4 continued

- Using VTune *Sampling over Time* after changing scheduling strategy to *dynamic*

```
#pragma omp parallel for schedule (dynamic)
                                private (j, limit, prime)
for(i = start; i <= end; i += 2){
    limit = (int) sqrt((float)i) + 1;
    prime = 1; /* assume number is prime */
    for ( j=3; (prime && (j <= limit)); j+=2 )
        if (i%j == 0) prime = 0;
```



# Development Cycle



# Intel® Thread Profiler



## Detection of thread-specific bottlenecks:

- Load imbalance
- Contention on synchronization objects
- Threading overhead

## Supported environments:

- OpenMP\*
- Native thread-API on Microsoft Windows\* (Win32 Threads)
- Native thread-API on Linux\* (PThreads)

## Analysis:

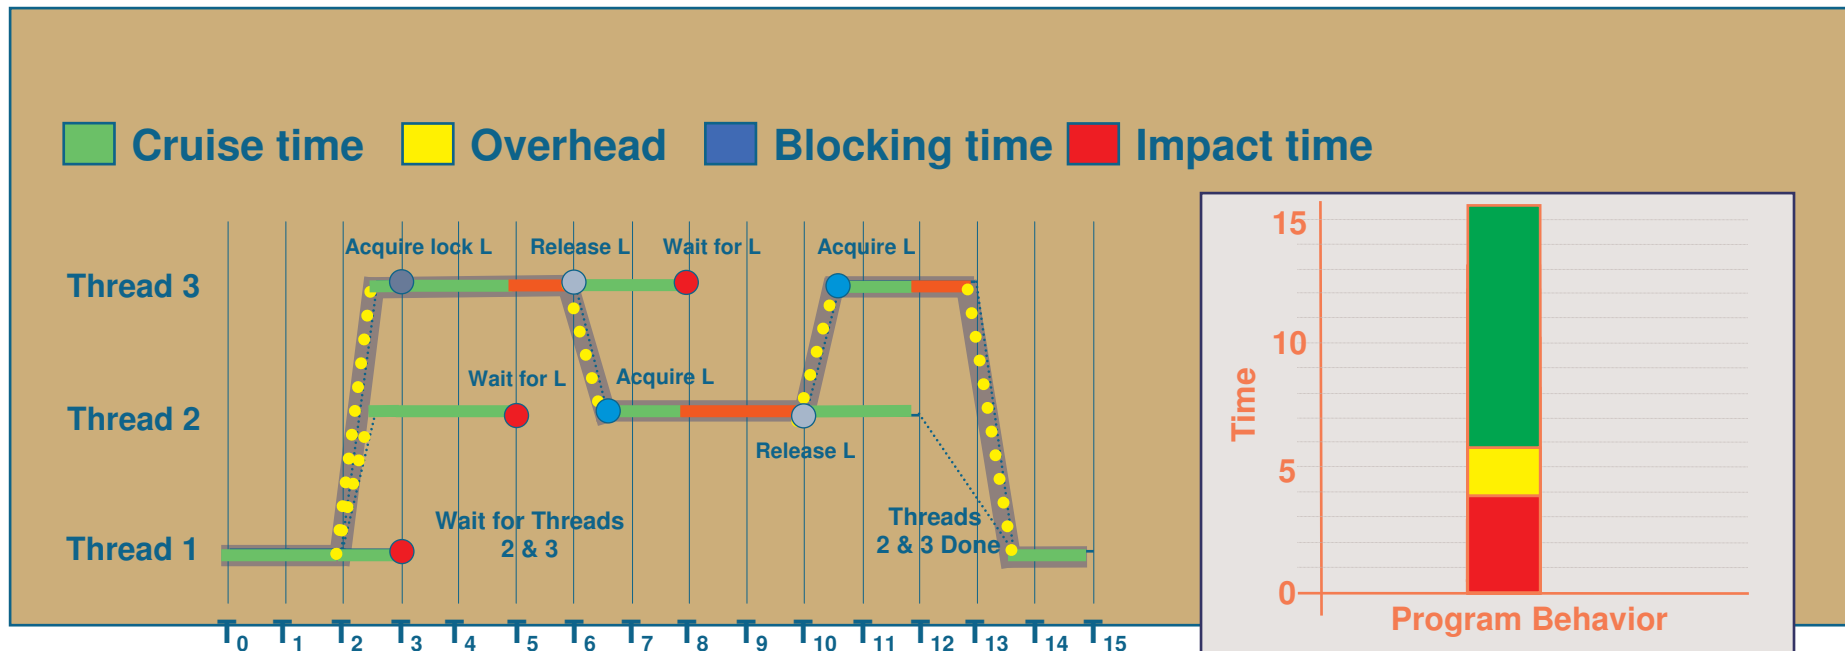
- Binary and/or source instrumentation
- Profiled run-time libraries ( `-Qopenmp_profile` )

# Execution Time Categories

Analyze program behavior along critical path execution


Record objects that cause transitions

*Categorization shown for a system configuration with 2 processors*



# Demo #5


- Using Intel® Thread Profiler to fine-tune prime application



# Demo #5 continued

- Final Intel® Thread Profiler analysis after optimizing loop scheduling and synchronization


```
#pragma omp parallel for schedule (dynamic)  
private (j, limit, prime)
```






# Demo #5 continued

- Final Intel® Thread Profiler analysis after optimizing loop scheduling and synchronization

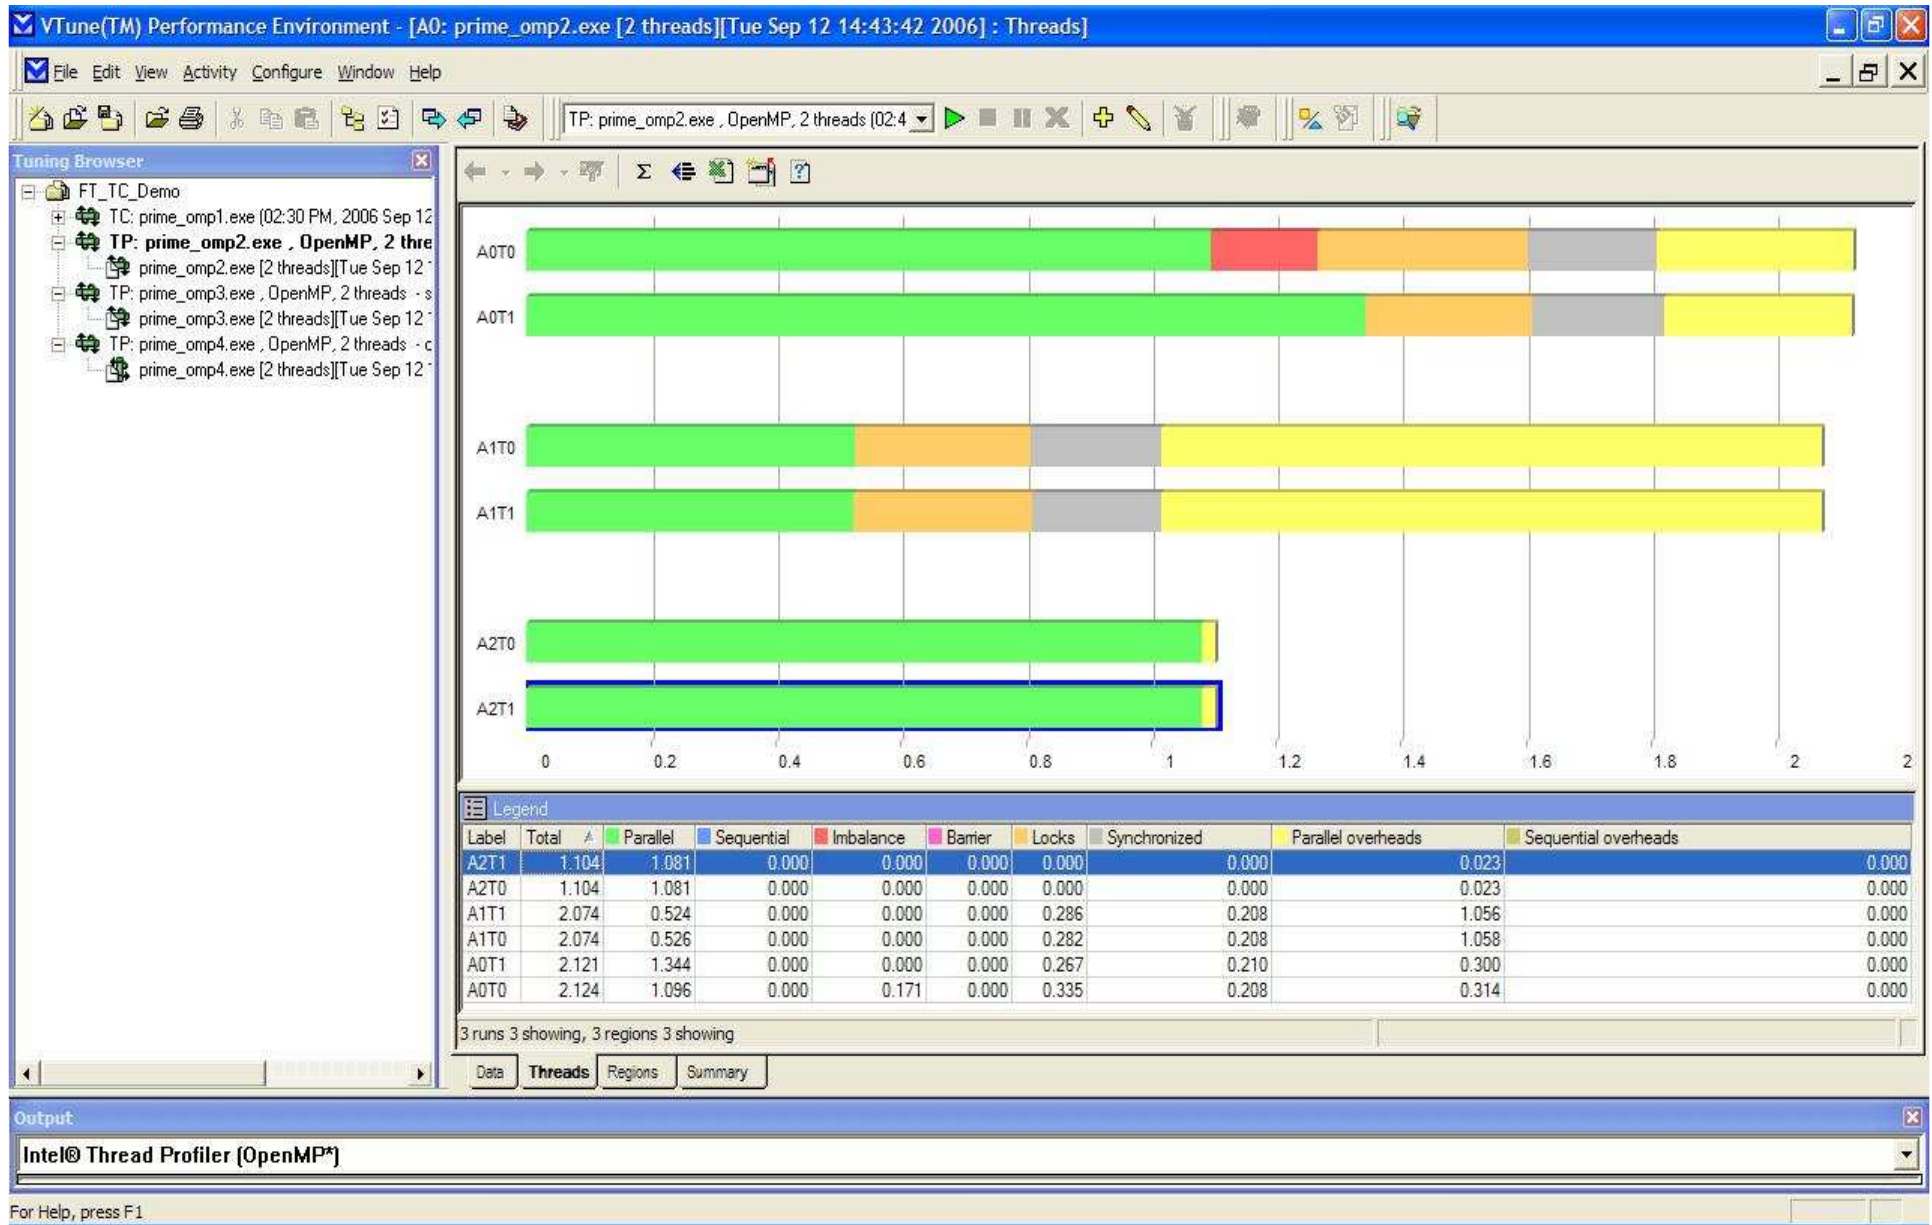


# Intel® Thread Profiler Results

Adopt better work load distribution



# Putting it together



# Agenda

Motivation for Threading

Concepts in Parallelism

Threading in Action

Common Threading Errors

Existing Code and Intel® Tools

Summary

# Summary

**Concurrent compute threads are required to take full advantage of multi-processor/multi-core systems**

**Properly threading your code is challenging**

- You don't have to do it all at once
- Focus on areas that will have the most impact to overall application performance
- Make sure to monitor scaling on systems with more processors than the current common customer configuration

**A thorough understanding of actual application behavior under typical loads is necessary before you consider threading**

**Intel continues to invest in hardware and software to ease the transition to threaded code**

## But Don't Take MY Word for It

“We are optimizing RenderMan’s core to be very scalable for future multi-core architectures. Intel’s Threading Tools have accelerated our development cycle dramatically. Intel’s Thread Checker for example, helped identify potential threading issues very quickly, in days compared to weeks if done otherwise. Thread Profiler, on the other hand, has helped us understand threading performance problems so we could fix them to improve scalability. The Intel Threading Tools are now an integral part of our development process.”

**Dana Batali**

**Director of RenderMan Development**

**Pixar**

## But Don't Take MY Word for It

“We found Intel ThreadChecker to be an **indispensable aid** for analyzing threaded code. We were impressed at how well it handled an application as large and complex as Maya. Based on this experience I plan to use this tool on future threading projects. Intel ThreadProfiler was very useful for analyzing bottlenecks in our threaded code. ThreadProfiler . . . showed us the reasons for the slowdown, so we were able to restructure the code for better threaded performance.”

**Martin Watt**

**Software Architect**

**Alias**

## But Don't Take MY Word for It

"We used Intel Threading Tools including Intel Thread Profiler to realize improved threaded application performance of Omni Page 15 running on Intel multi-core platforms. We look forward to using Intel Thread Profiler with its critical path analysis and selective magnification of important time regions on future thread optimization projects."

**Gyorgy Varszegi**  
**Scansoft**





Remember when  
the sky was the limit?



**Thank you**