

Thread Checker:

**Saving Valuable Time Otherwise Spent
Searching for Hard-to-find Threading Errors**

Koby Gottlieb

**Software Solution Group
Intel Corporation**

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

All dates provided are subject to change without notice.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing.

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

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

Copyright © 2004, Intel Corporation.

Agenda

- **Threading: Why do we need it now?**
 - The Dual core future
- **Challenges Unique to Threading**
- **Intel® Thread Checker product**
 - Features
 - Demo
 - How does it works?
- **The Thread Profiler feature**

Driving Parallelism

Moving from Chips/Computer to Computers/Chip

| | 2004 | 2005 | 2006* |
|-----------------------|---------|--------------------|-----------------------|
| Desktop (Perf) | 55% HT | Shipping Dual-Core | > 40% Dual-Core |
| Servers | 100% HT | Shipping Dual-Core | > 85% Dual/Multi-Core |
| Mobile (Perf) | | Shipping Dual-Core | > 70% Dual-Core |

“Exiting 2006, we believe that over 40 percent of the desktop product shipments will be dual core, over 80 percent of our server products will be multi or dual core, and over 70 percent of our mobile products will be dual core as well. We are dedicating all of our future product designs to multi-core environments. We have bet on this in terms of our software environment, our ecosystem development, and our Intel Capital infrastructure around it. We believe this is a key inflection point for the industry.”



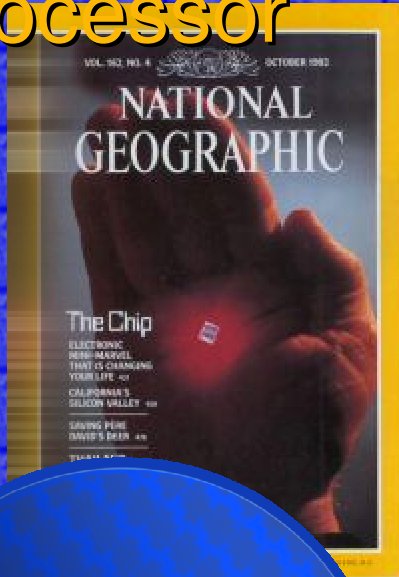
PAUL OTELLINI, IDF 7/2004

All CPU development
on Dual/Multi-Core

Unique Dual/Multi-Core
Products in all segments

* data is run rate ending 2006. CPU and memory products are not included as the property of others

Next generation Dual-Core Itanium® processor



“ Eventually **one billion** transistors, or electronic switches, may crowd a single chip, 1,000 times more than possible today. ”
National Geographic, 1982

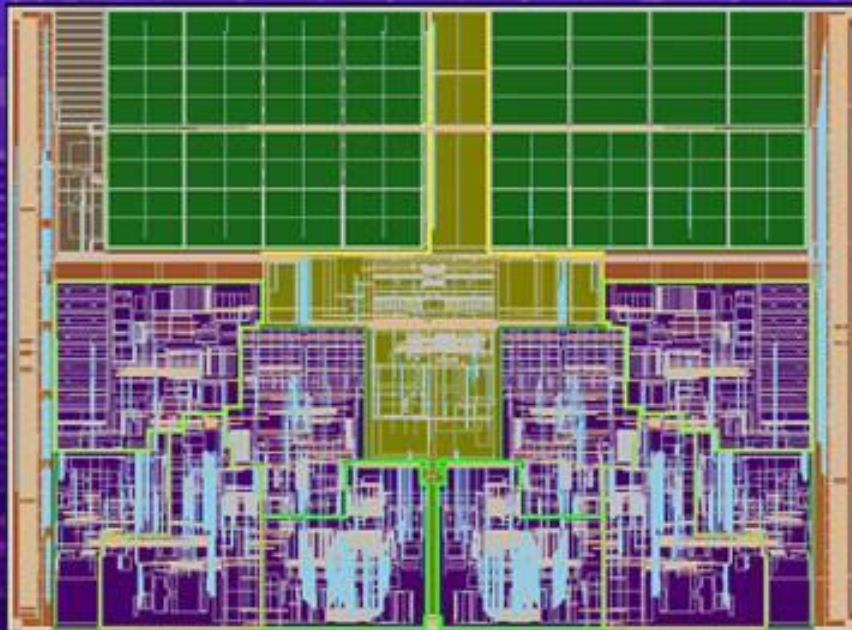
Montecito
1.7B transistors

Next Generation – Montecito

- Dual Core *and* Multithreaded
- >3X increase platform bandwidth
- Higher performance, lower power
- 24MB Cache



Future Capabilities Napa Platform



CORE 1

CORE 2

*Yonah: First mobile optimized **Dual-core** processor design from the ground up on 65nm:*

LT, VT

Calistoga chipset: Integrated graphics for superb playback

Golan: Next generation wireless solution

Efficiently Utilize Dual Cores

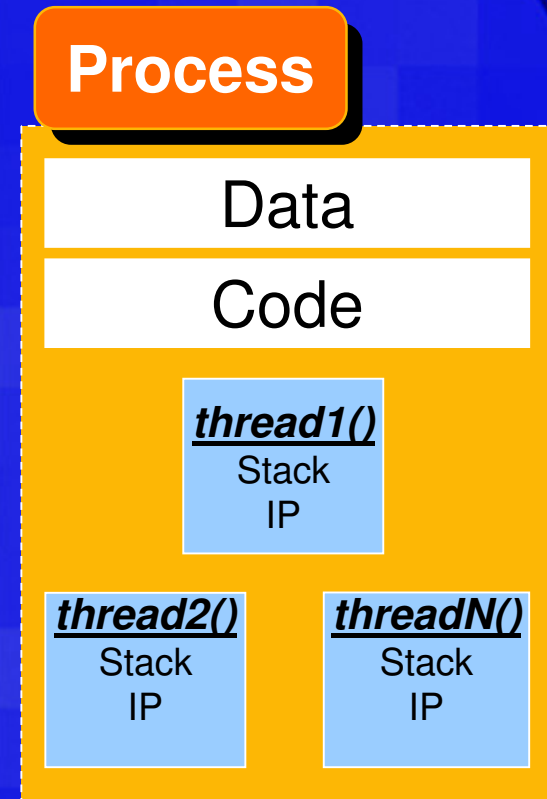
Dual-Core Systems

- One package with 2 cores
- Software impact
 - 2 Cores → 2 processors
 - 2 Cores → 2x resources

Use threads to exploit full resources of dual core processors

Threads Defined

- OS creates process for each program loaded
 - Each process executes as a separate thread
- Additional threads can be created within the process
 - Each thread has its own Stack and Instruction Pointer
 - All threads share code and data



Efficiently Utilize Dual Cores

Threading Software

- **OpenMP*** threads
 - <http://www.openmp.org/>
- **Windows*** threads
 - <http://msdn.microsoft.com/>
- **POSIX*** threads (pthreads)
 - <http://www.ieee.org/>

**If both cores fully busy, then 2x
speedup possible**

Challenges Unique to Threading

Correctness Bug: Data Races

- Suppose: $a=1$, $b=2$

Thread1

$x = a + b$

Thread2

$b = 42$

- What is value of x if:
 - Thread1 runs before Thread2? $x = 3$
 - Thread2 runs before Thread1? $x = 43$
- Data race: concurrent read, modify, write of same address

Outcome depends on thread execution order

Solving Data Races: Synchronization

Thread1

Acquire(L)

a = 1

b = 2

x = a + b

Release(L)

Thread2

Acquire(L)

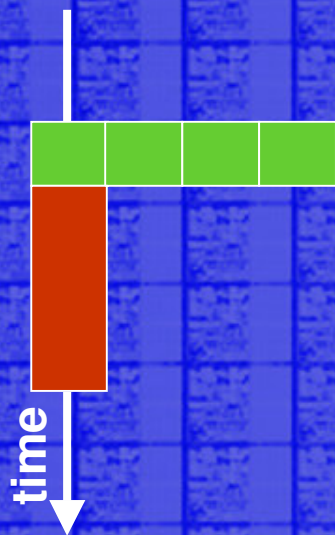
b = 42

Release(L)

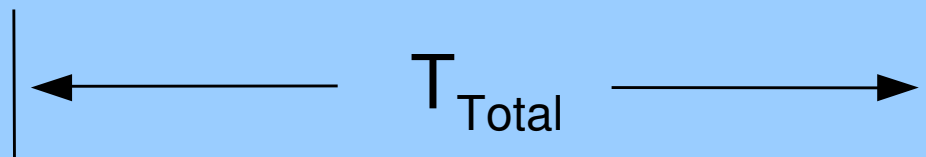
- Acquisition of mutex L ensures atomic access
 - Only one thread can hold lock at a time
- Example APIs:
 - EnterCriticalSection(), LeaveCriticalSection()
 - pthread_mutex_lock(), pthread_mutex_unlock()

Efficiently Utilize Dual Cores

Amdahl's Law



If only 1/2 of the code is parallel, 2X speedup is unlikely



$$T_{Parallel} = \{(1 - P) + \frac{P}{N} + O\}T_{Total}$$

P = parallel portion of process

N = number of processors (cores)

O = parallel overhead

Threads Intro New Class of Problems

- **Correctness bugs**

- Data races
- Deadlock
- and more...

Intel® Thread Checker
finds correctness bugs

- **Performance bottlenecks**

- Overhead
- Load balance
- and more...

Thread Profiler feature
pinpoints bottlenecks

Intel® Threading Tools can help!

Intel® Thread Checker Intro

- Identifies threading bugs in applications threaded with:
 - Microsoft* Windows* threads on Microsoft* Windows* systems
 - POSIX* pthreads on Linux* systems
 - OpenMP* on Microsoft* Windows* and Linux* systems
- Plugs into VTune™ environment
 - Microsoft* Windows* for IA-32 systems
 - Linux* for IA-32 and Itanium®-based systems

Intel® Thread Checker Analysis

- **Dynamic monitoring as software runs**
 - Data (workload) -driven execution
- **Includes monitoring of:**
 - Thread and Sync APIs used
 - Thread execution order
 - Scheduler impacts results
 - Memory accesses between threads

Only executed code path is analyzed

Intel® Thread Checker 2.0

Features

- Locates threading bugs:
 - Data races (storage conflicts)
 - Deadlocks (potential and actual)
 - Win32 threading API usage problems
 - Memory leaks and overwrites
- Isolates bugs to source code line
- Describes possible causes of errors and suggests resolutions
- Categorizes errors by severity level

Screen shot: Intel® Thread Checker

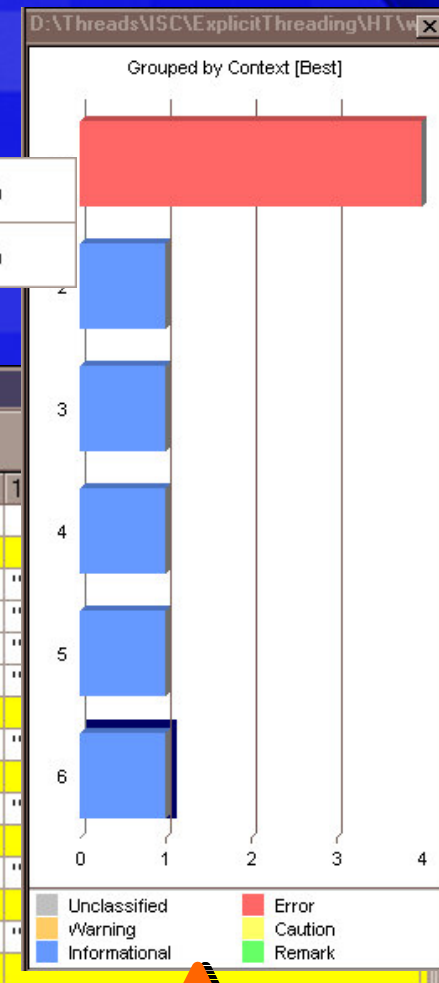
Diagnostics List

Verbose diagnostics

| | | | |
|---|------|--------|------|
| Memory write of dSum at "Pi.cpp": 23 conflicts with a prior memory read dSum at "Pi.cpp": 23 (anti dependence) | 1501 | PiFunc | dSum |
| Memory read of dSum at "Pi.cpp": 23 conflicts with a prior memory write of dSum at "Pi.cpp": 23 (flow dependence) | 1501 | PiFunc | dSum |

D:\Threads\VSCE\ExplicitThreading\HT\windows\Pi\Debug\threadchecker.thr: Diagnostics

| | Context [Best] | ID | Severity | Description | Counts | 1st Access [Routine] | 1st Access [Variable] |
|--------------------------|----------------|----|----------|--------------------------|--------|----------------------|-----------------------|
| Total | | | | | | | |
| Group 1: "Pi.cpp": 12 | | | | | | | |
| "Pi.cpp": 12 | | 0 | ● | Write -> Read data-race | 4 | main | i |
| "Pi.cpp": 12 | | 1 | ● | Read -> Write data-race | 1501 | PiFunc | dSum |
| "Pi.cpp": 12 | | 2 | ● | Write -> Read data-race | 1501 | PiFunc | dSum |
| "Pi.cpp": 12 | | 3 | ● | Write -> Write data-race | 1501 | PiFunc | dSum |
| Group 2: Whole Program 1 | | | | | | | |
| Whole Program 1 | | 4 | ● | Thread termination | 1 | main | unknown |
| Group 3: Whole Program 2 | | | | | | | |
| Whole Program 2 | | 5 | ● | Thread termination | 1 | main | unknown |
| Group 4: Whole Program 3 | | | | | | | |
| Whole Program 3 | | 6 | ● | Thread termination | 1 | main | unknown |
| Group 5: Whole Program 4 | | | | | | | |
| Whole Program 4 | | 7 | ● | Thread termination | 1 | main | unknown |
| Group 6: Whole Program 5 | | | | | | | |



**Diagnostics List
in Terse mode**

**Summary
and legend**

Screen shot: Intel® Thread Checker

Source Code View

The screenshot displays the Intel Thread Checker interface. The top window, titled "Diagnostics", shows a list of detected issues. A large orange arrow points from the "Group 1: 'Pi.cpp': 12" entry in the diagnostics list to the "Stack Trace" window below. The "Stack Trace" window shows the source code for the function "PiFunc" in "Pi.cpp" at line 23. The source code is as follows:

```
start = myThreadNum+1 ;  
for (inc i = start; i < maxIterations; i+=maxThreads)  
{  
    dx = (i-0.5) * dSteps;  
    sum = sum + 1.0 / (1.0 + dx * dx);  
}  
  
return myThreadNum ; // thread exit code  
// PiFunc
```

The diagnostics list includes the following entries:

| Context [Best] | ID | Severity | Description | Count | 1st Access [Routine] | 1st Access [Variable] | 1st Access [Best] | 2 |
|---------------------------------|----|----------|--|-------|----------------------|-----------------------|-------------------|---|
| Group 1: 'Pi.cpp': 12 | | | | | | | | |
| "Pi.cpp": 12 | 0 | Warning | Memory read of pArg at "Pi.cpp": 13 conflicts with a prior memory write of i at "Pi.cpp": 50 (flow dependence) | 4 | main | i | "Pi.cpp": 50 | P |
| "Pi.cpp": 13 | 1 | Warning | Memory read of dSum at "Pi.cpp": 23 conflicts with a prior memory read of dSum at "Pi.cpp": 23 (flow dependence) | 1501 | PiFunc | dSum | "Pi.cpp": 23 | P |
| "Pi.cpp": 12 | 2 | Warning | Memory read of dSum at "Pi.cpp": 23 conflicts with a prior memory write of dSum at "Pi.cpp": 23 (flow dependence) | 1501 | PiFunc | dSum | "Pi.cpp": 23 | P |
| "Pi.cpp": 12 | 3 | Warning | Memory write of dSum at "Pi.cpp": 23 conflicts with a prior memory write of dSum at "Pi.cpp": 23 (output dependence) | 1501 | PiFunc | dSum | "Pi.cpp": 23 | P |
| Group 2: Whole Program 1 | | | | | | | | |

Each Diagnostics in List links to its source code line(s)

Screen shot: Intel® Thread Checker

Help with Diagnostics

The screenshot shows the Intel Thread Checker diagnostics window. The main table lists memory access conflicts. A right-click context menu is open over a row, showing options like 'Open Graphical Summary', 'Group By', 'Configure Columns', 'Expand All', 'Collapse All', 'Update', 'What's this column?', and 'What's this diagnostic?'. Two orange callout boxes provide instructions: '1) Right-click here ...' and '2) More help!'.

| Context [Best] | ID | Severity | Description | Count | 1st Access [Routine] | 1st Access [Variable] | 1st Access [Best] | 2 |
|---------------------------------|----|----------|---|--------|----------------------|-----------------------|-------------------|---|
| Group 1: "Pi.cpp": 12 | | | | | | | | |
| "Pi.cpp": 12 | 0 | 1 | Memory read of pArg at "Pi.cpp": 13 conflicts with a prior memory write of i at "Pi.cpp": 50 (flow dependence) | 4 main | i | "Pi.cpp": 50 | "Pi.cpp": 12 | P |
| "Pi.cpp": 12 | 1 | 1 | Memory write of dSum at "Pi.cpp": 23 conflicts with a prior memory read of dSum at "Pi.cpp": 23 (flow dependence) | 1000 | dSum | dSum | "Pi.cpp": 23 | P |
| "Pi.cpp": 12 | 2 | 1 | Memory read of dSum at "Pi.cpp": 23 conflicts with a prior memory write of dSum at "Pi.cpp": 23 (flow dependence) | 1000 | dSum | dSum | "Pi.cpp": 23 | P |
| "Pi.cpp": 12 | 3 | 1 | Memory write of dSum at "Pi.cpp": 23 conflicts with a prior memory read of dSum at "Pi.cpp": 23 (flow dependence) | 1000 | dSum | dSum | "Pi.cpp": 23 | P |
| Group 2: Whole Program 1 | | | | | | | | |

Stack Trace: **Pi.cpp": 23**

Line 16
Line 17

24
25
26
27
28
29
30
31
32
33

```
return myThreadNum ; // thread exit code  
// PiFunc
```

Context Definition 1st Access 2nd Access Stack Trace

Intel® Threading Tools

Intel® Thread Checker Demo

Find threading bugs faster

The secrets behind the tool

- Automatic application instrumentation
 - Binary instrumentation – VTune™ analyzer technology
 - Source instrumentation – Compiler technology
- Semantic replacement of Threading functions
 - Observe thread management and synchronization
 - Observe system functions which modify arguments
- Partial order of threads drives analysis
- Execution generates annotated address trace
- Produce error diagnostics in a single execution

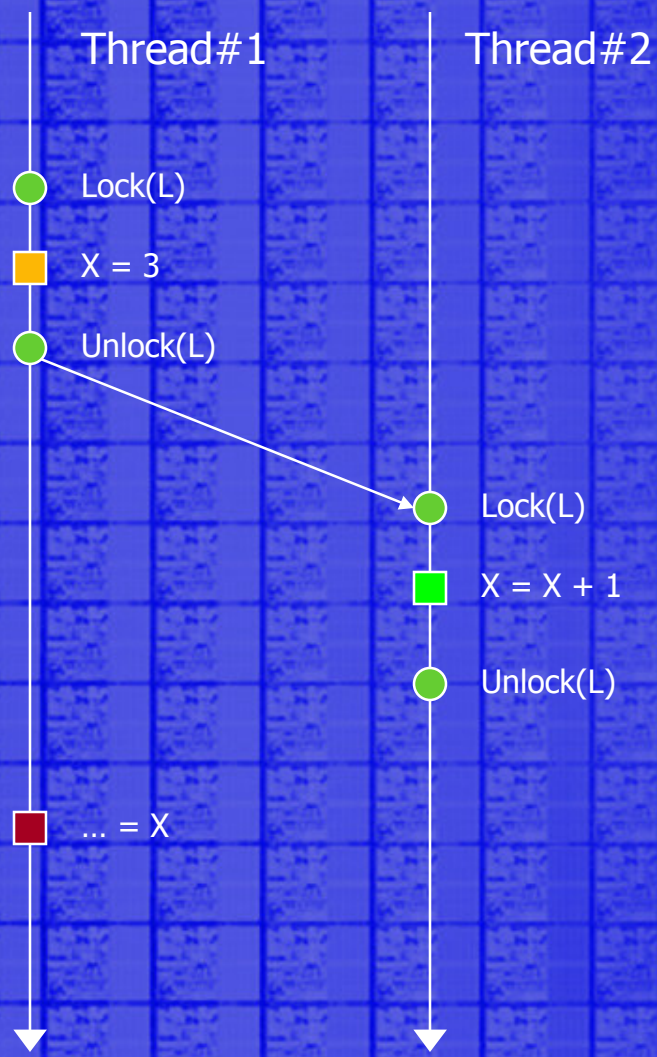
Data Structures

- Threads maintain an integer clock
- Objects (Thread & Sync) have a time vector
 - $T\#1 = [0, 0]$
 - Vector length is the number of threads
 - Hold partial order relationship between objects
- Memory has a shadow cell
 - $X = (T\#, 0)$
 - Record last accesses (thread id and time)
- The Max function is used to merge the vectors
 - $\text{Merge}([A,B], [C,D]) = [\text{Max}(A,C), \text{Max}(B,D)]$

Actions

- ACQUIRE sync action (e.g. LOCK API)
 - The lock is merged with the thread
 - The thread is advanced
- RELEASE sync action (e.g. UNLOCK API)
 - The thread is advanced
 - The thread is merged with the lock
- ACCESS memory action (e.g. READ/ WRITE)
 - Check shadow cell for conflict
 - Updated shadow cell with threads info

Example



Start: T#1 = [1,0] ; T#2 = [0,1] ; L = [0,0]; X = ()

T#1 = [2,0]

compare X=() with T#1=[2,0] : **FIRST ACCESS**
X = (T#1, 2)

T#1 = [3,0]; L = [3,0]

T#2 = [3,2]

compare X=(T#1, 2) with T#2=[3,2] : **OK**
X = (T#2, 2)

T#2 = [3,3]; L = [3,3]

compare X=(T#2 , 2) with T#1 [3,0] : **ERROR**
X = (T#1, 3)

Intel® Threading Tools

“Intel® Thread Checker helped Siemens by identifying issues in software we develop and in software we purchase from third parties. We use Intel Thread Checker to improve the quality of our software and look forward to expanding the use of the tool in more of our software development groups.”

-Andreas Dietrich, Research and Development Image Processing, Siemens Medical Solutions

“Using Intel Thread Checker we discovered two elusive bugs on the very first day, as well as numerous inconsistencies and opportunities for performance improvement. We were pleasantly surprised because our product, AcuSolve*, has been running successfully on multiple platforms for many years. We have now incorporated it in our basic development and release process.” *-Farzin Shakib, President ACUSIM Software, Inc.*

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

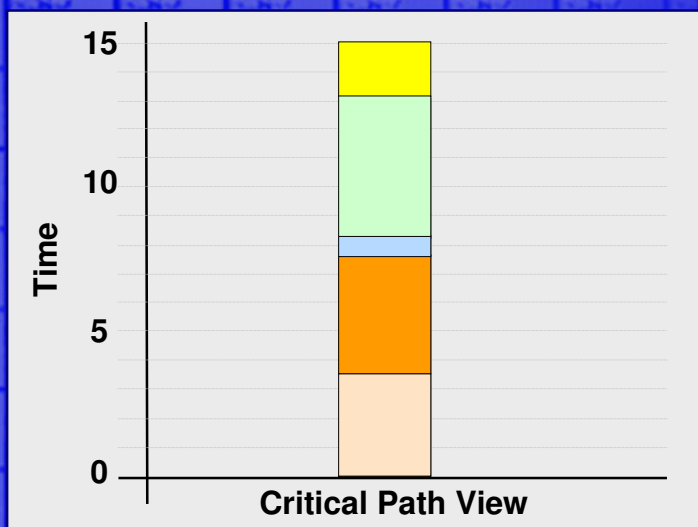
The Thread Profiler Feature

- Pinpoints threading performance bottlenecks in apps threaded with:
 - Microsoft* Windows* threads on Microsoft* Windows* systems
 - POSIX* pthreads on Linux* systems
 - OpenMP* on Microsoft* Windows* and Linux* systems
- Plugs into VTune™ environment
 - Microsoft* Windows* for IA-32 systems
 - Linux* for IA-32 systems

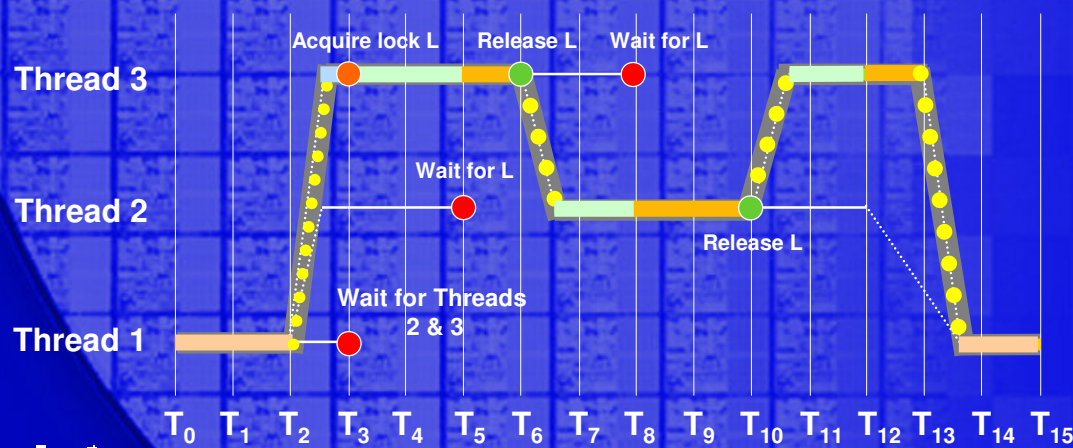
Thread Profiler Feature Analysis

- **Monitors execution flows to find Critical Path**
 - Longest execution flow is the Critical Path
- **Analyzes Critical Path**
 - System utilization
 - Over-subscribed vs. under-subscribed
 - Thread state transitions
 - Blocked -> Running
- **Captures threads timeline**
 - Visualize threading structure

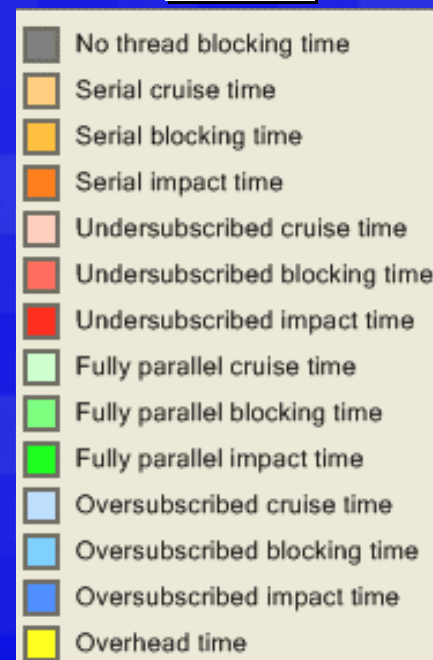
Thread Profiler Critical Path



Analysis shown for 2-way system



- Start with the critical path
- Separate according to system utilization
- Add overhead
- Further analyze by thread state



Exploiting Dual Core Systems: The Intel® Threading Tools

- Add threads to realize full performance benefits of multi cores
- Intel® Threading Tools can help:
 - Intel® Thread Checker finds threading bugs
 - Thread Profiler pinpoints threading bottlenecks

**Intel® Threading Tools shortens
development cycle for threaded apps**

Backup

Collateral

- **Intel® Threading Tools**
<http://www.intel.com/software/products/>
- **OpenMP* threads**
 - <http://www.openmp.org/>
- **Windows* threads**
 - <http://msdn.microsoft.com/>
- **POSIX* threads (pthreads)**
 - <http://www.ieee.org/>

Additional Reference Materials

- A comprehensive source of information on tools and techniques for developers of software for dual core systems
 - **“Programming with Hyper-Threading Technology”**
How to Write Multithreaded Software for Intel® IA-32 Processors
Richard Gerber and Andrew Binstock, ISBN 0-9717861-4-3
- More info at www.intel.com/intelpress

Purchase Intel Press books:

- IDF bookstore
(20% discount during the conference)
- ShopIntel.com
- Amazon.com

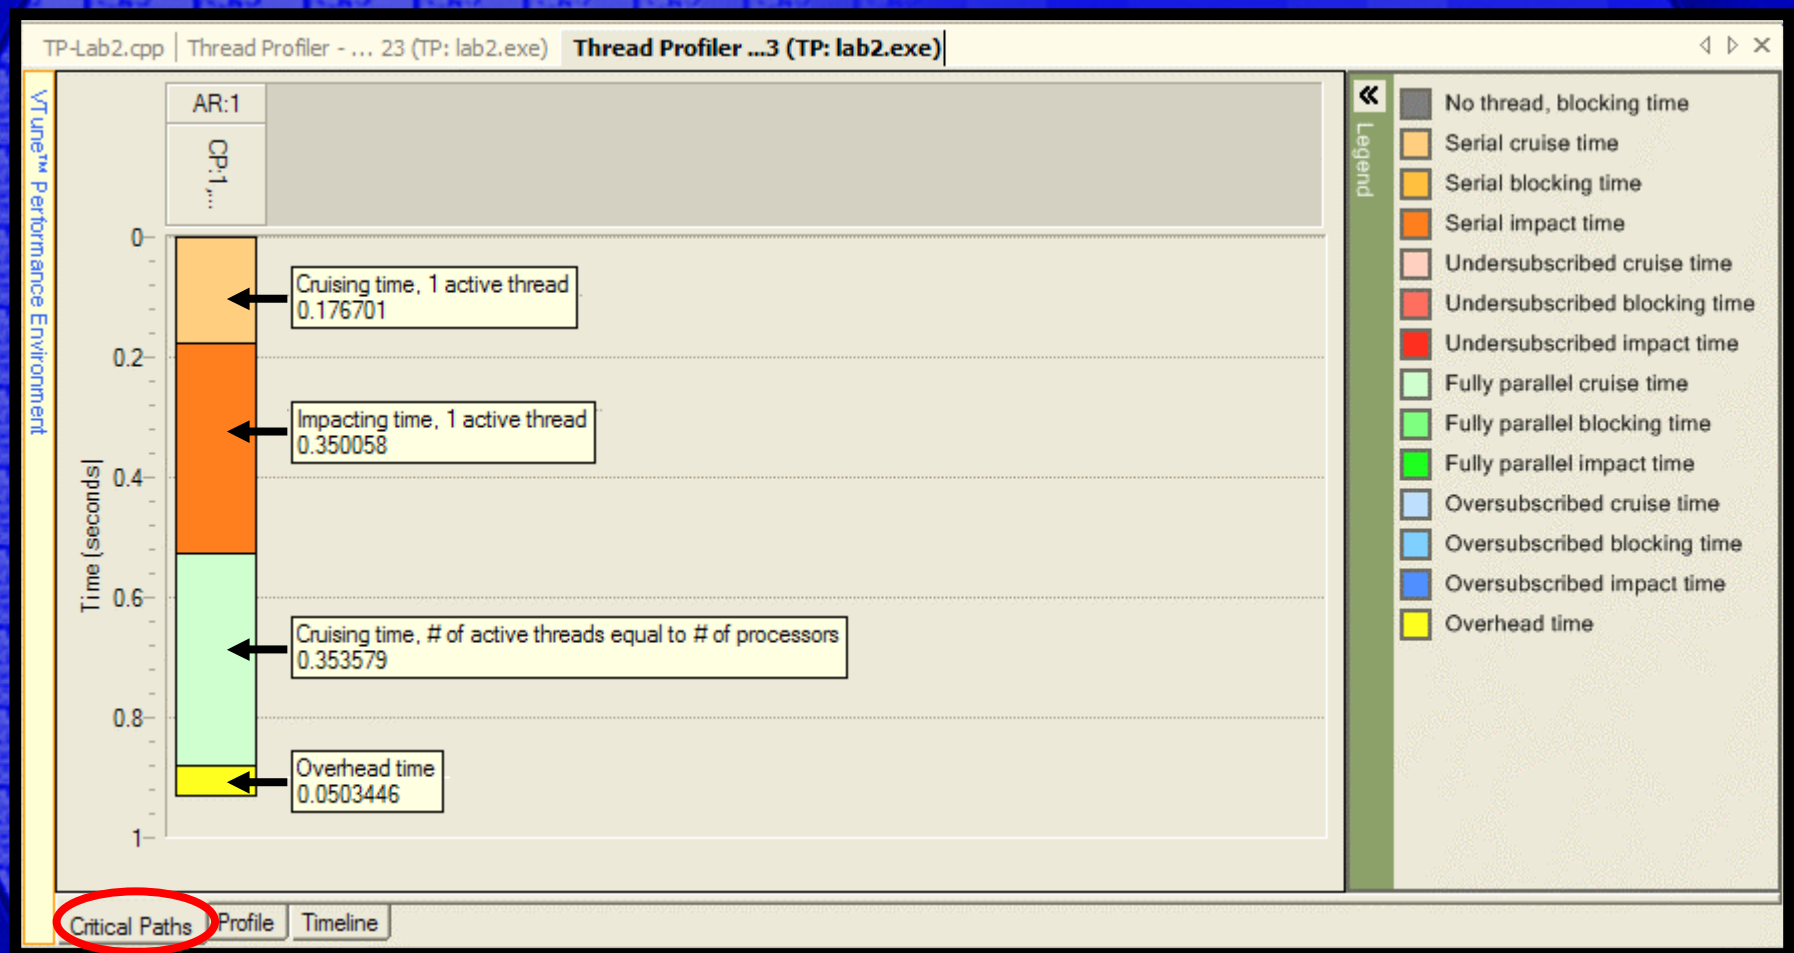


Acronyms

- Pthreads: POSIX* threads

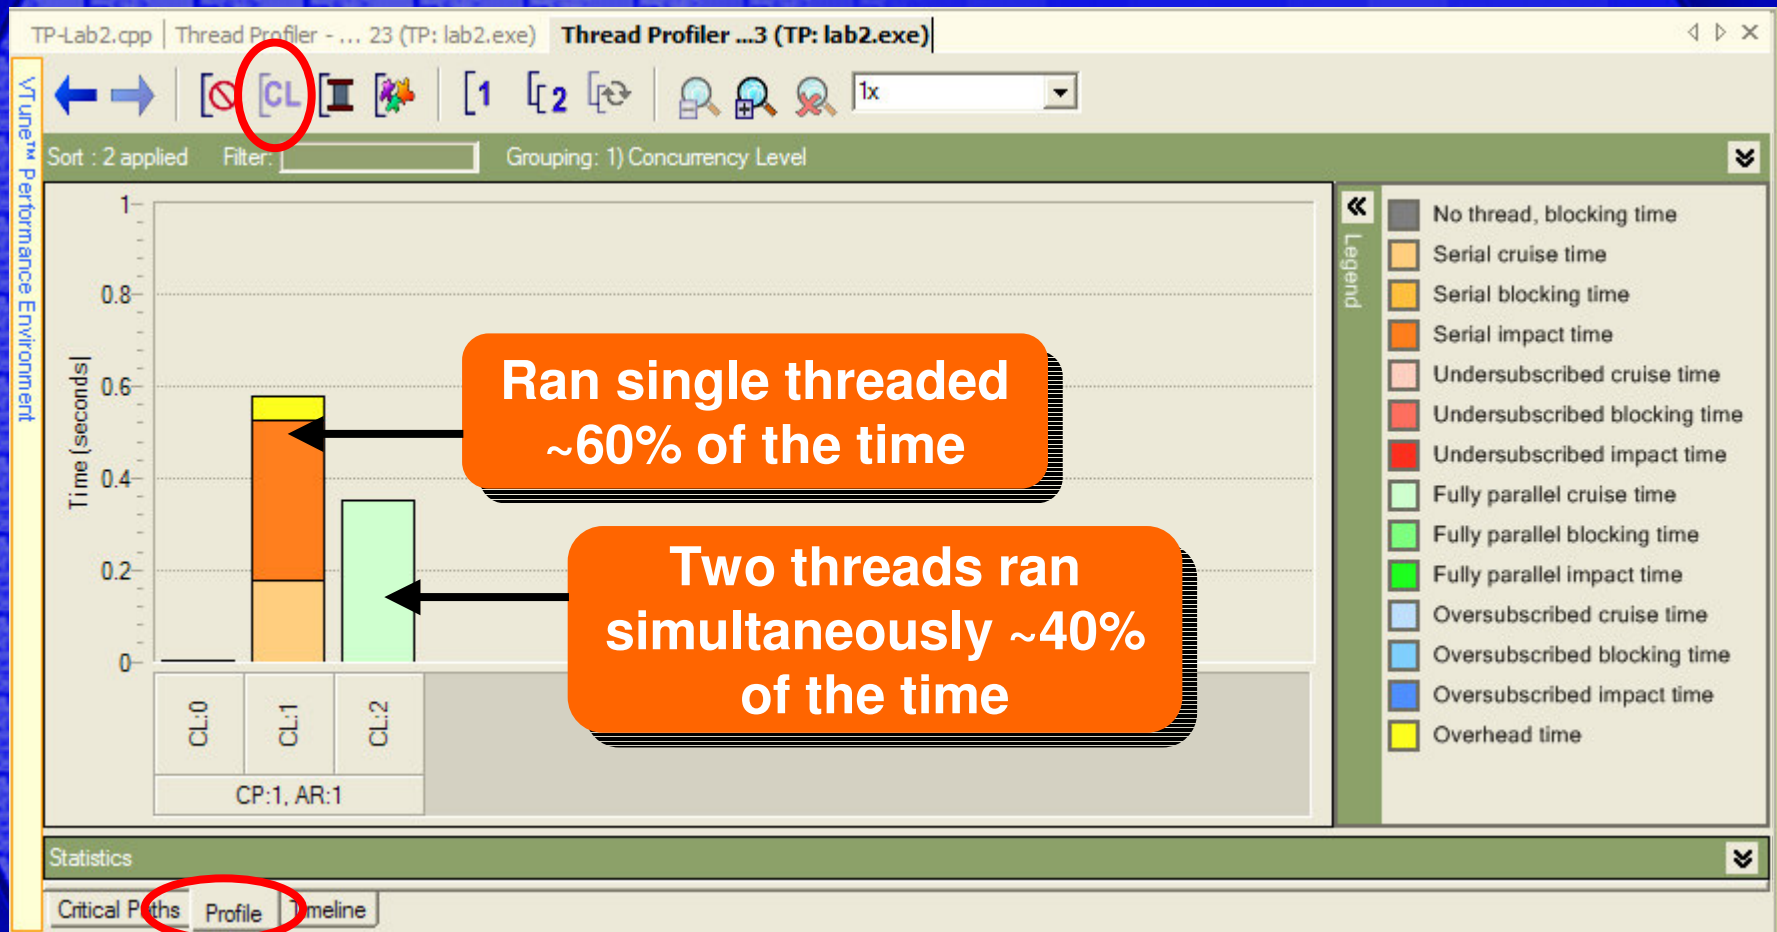
Backup screen shot: The Thread Profiler Feature

Critical Path View



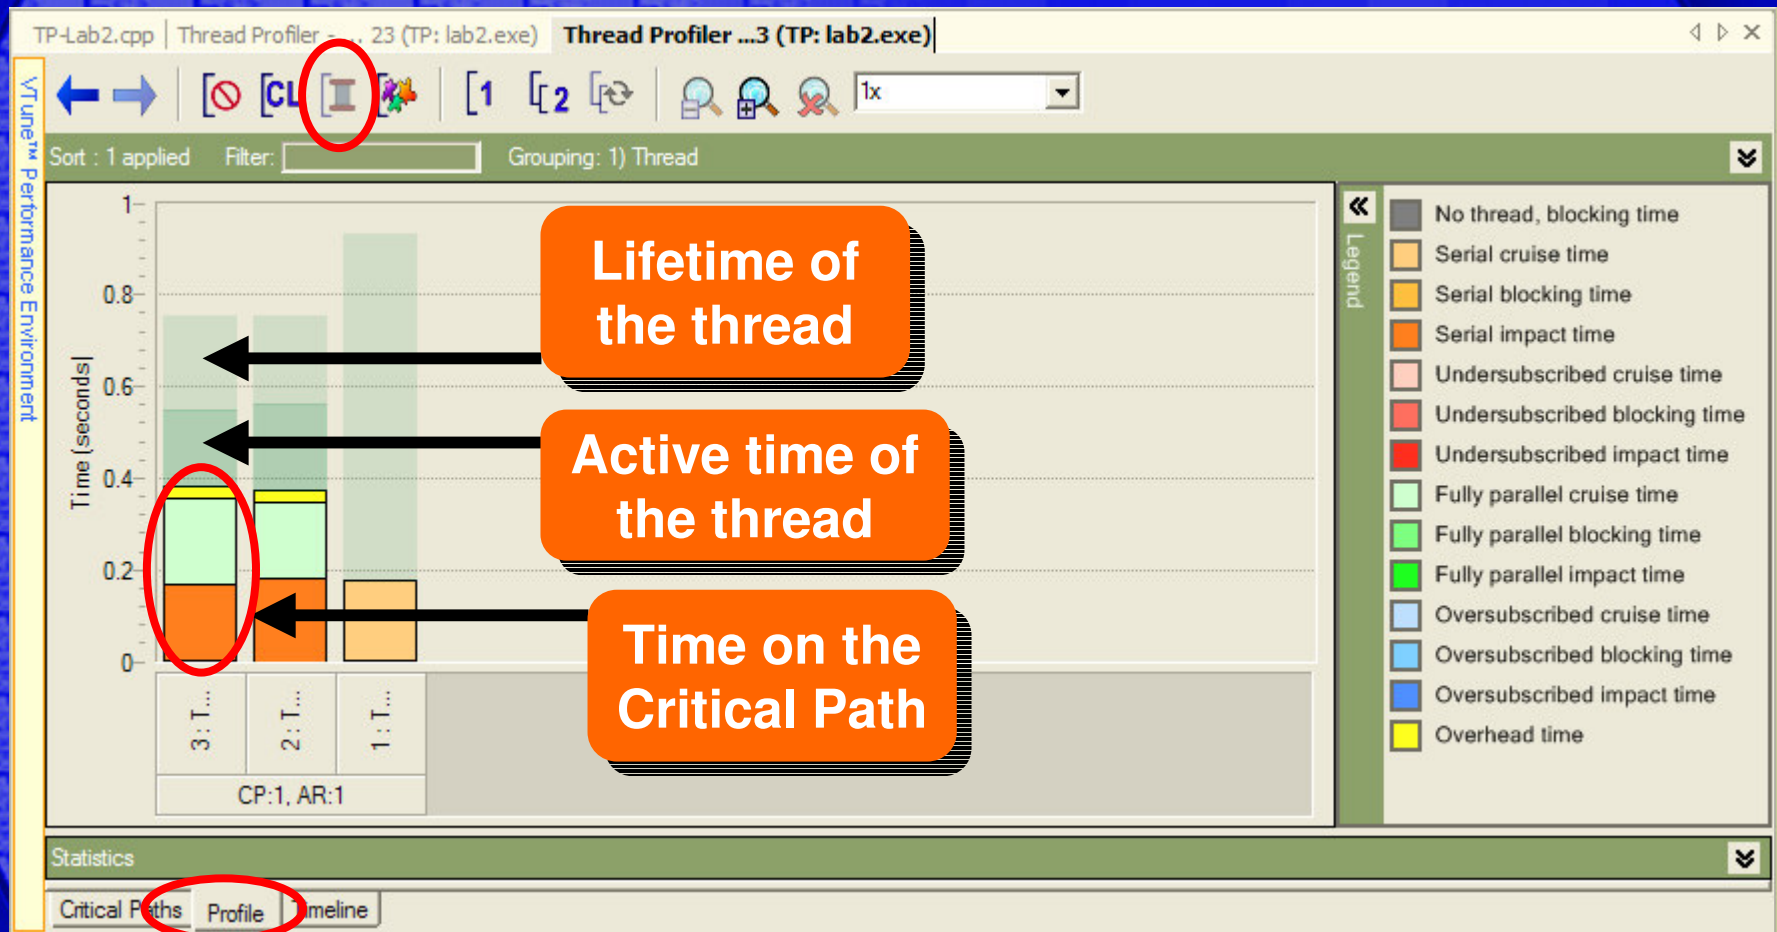
Backup screen shot: The Thread Profiler Feature

Profile Views: Concurrency Levels



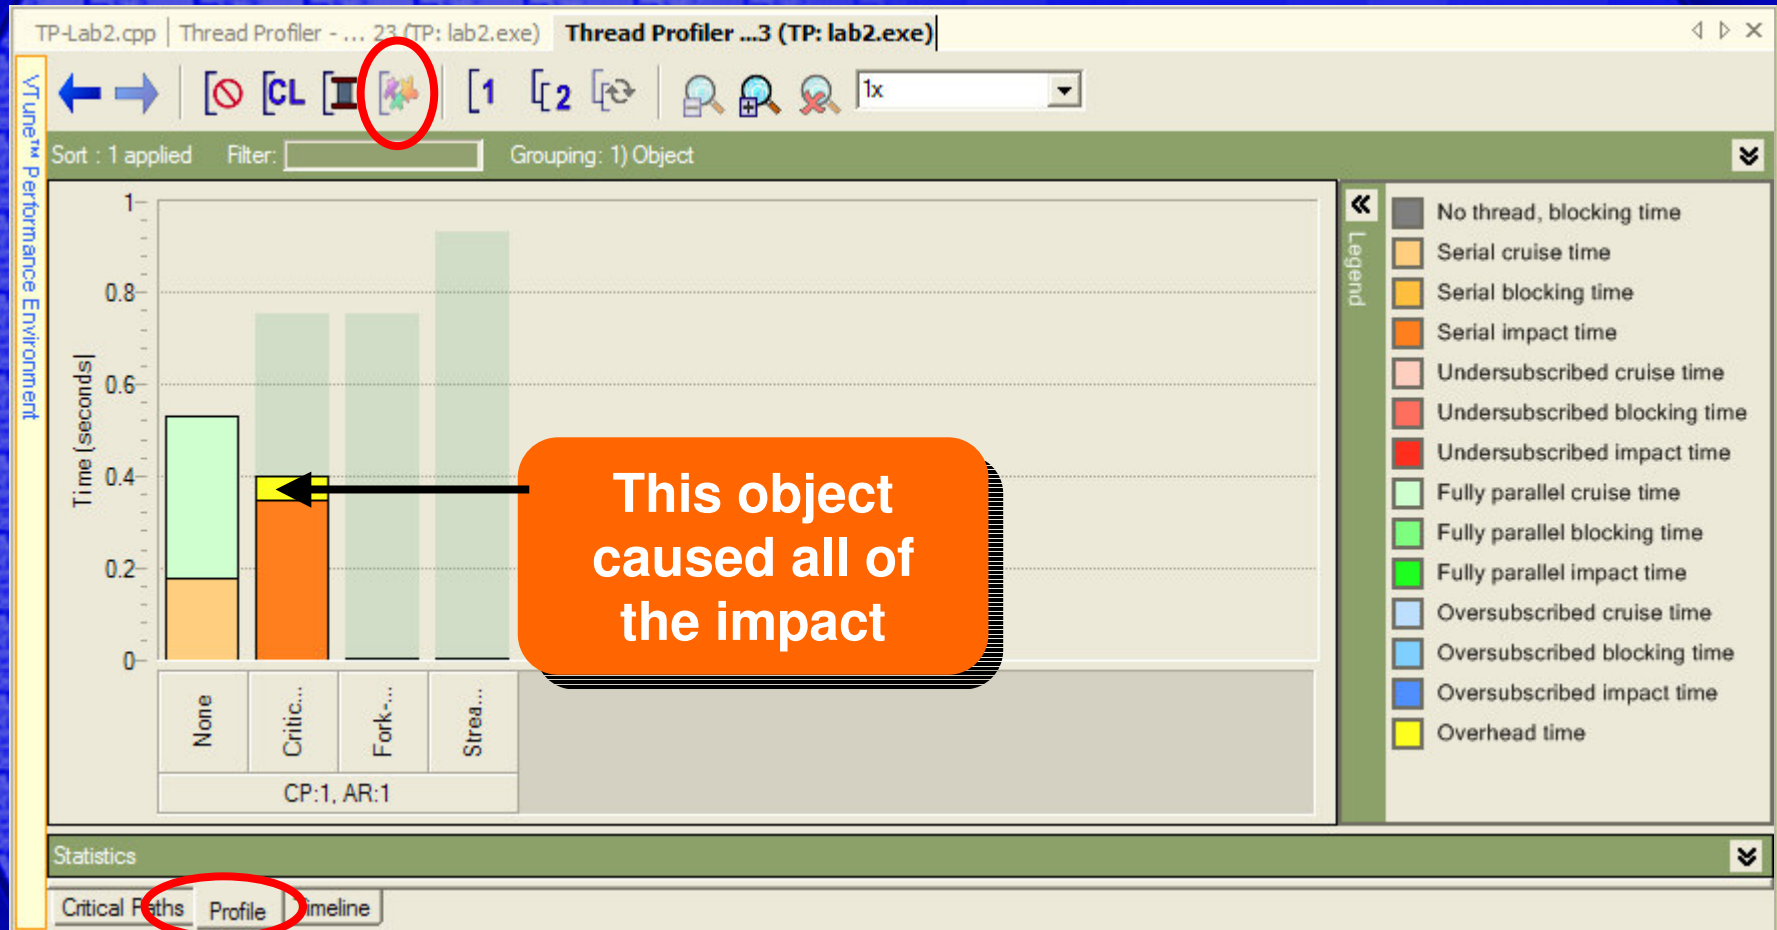
Backup screen shot: The Thread Profiler Feature

Profile Views: Threads View



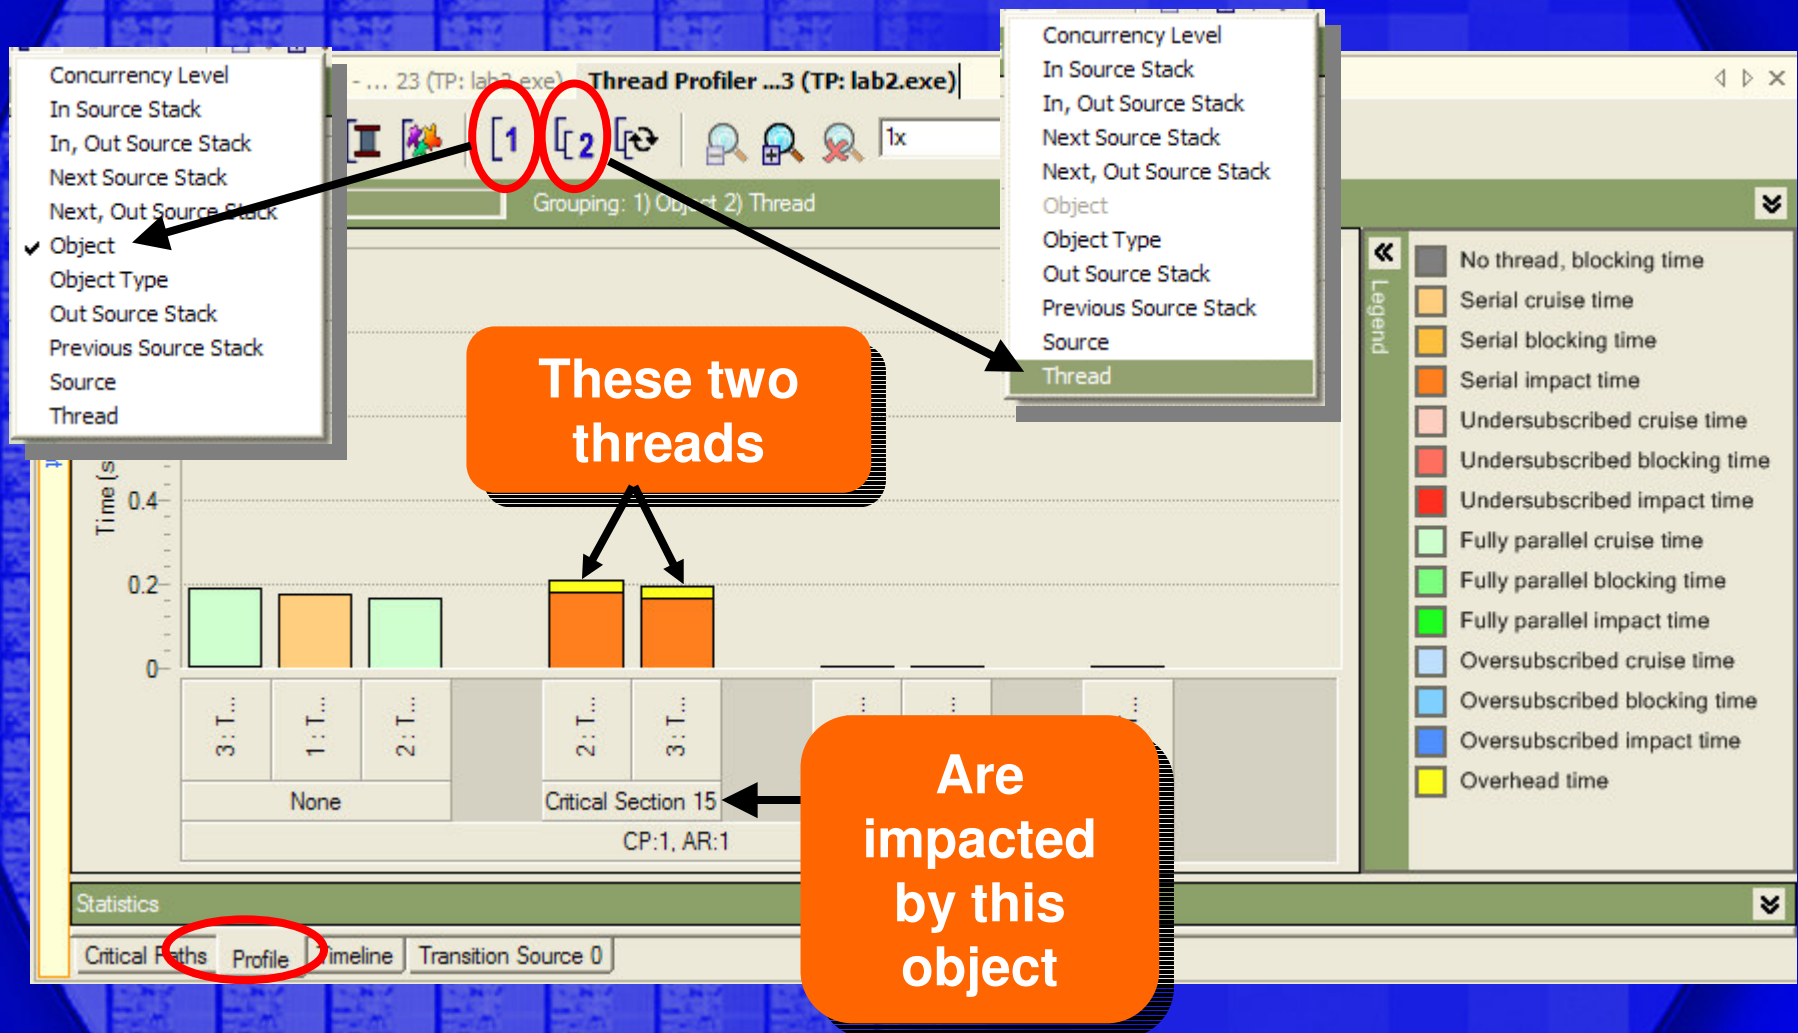
Backup screen shot: The Thread Profiler Feature

Profile Views: Objects View



Backup screen shot: The Thread Profiler Feature

Profile Views: Grouping



Backup screen shot: The Thread Profiler Feature

Timeline View



Performance Penalty: Synchronization

- **Thread blocked waiting for Mutex**
 - Thread not running, so no parallelism
- **Mutex Release, Acquire takes time**
 - Release marks mutex free
 - Acquire must check for free
 - If free, mark as in use
 - If not free, thread put to sleep
 - Costs context switch out and in of processor

Create private copies of frequently accessed data to reduce required synchronization