



Getting Started Tutorial: Finding Hotspots

Intel® VTune™ Amplifier XE 2013 for Windows* OS

Fortran Sample Application Code

Document Number: 327358-001

[Legal Information](#)

Contents

Legal Information	5
Overview	7
Chapter 1: Navigation Quick Start	
Chapter 2: Finding Hotspots	
Visual Studio* IDE: Choose Project and Build Application.....	14
Standalone GUI: Build Application and Create New Project.....	17
Run Hotspots Analysis.....	20
Interpret Hotspots Results.....	21
Resolve Issue.....	24
Run Concurrency Analysis.....	25
Interpret Concurrency Results.....	26
Run Locks and Waits Analysis.....	28
Interpret Locks and Waits Results.....	29
Remove Locks.....	32
Compare with Previous Result.....	33
Chapter 3: Summary	
Chapter 4: Key Terms	

Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. 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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to: http://www.intel.com/products/processor_number/

BlueMoon, BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Cilk, Core Inside, E-GOLD, Flexpipe, i960, Intel, the Intel logo, Intel AppUp, Intel Atom, Intel Atom Inside, Intel CoFluent, Intel Core, Intel Inside, Intel Insider, the Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel Sponsors of Tomorrow., the Intel Sponsors of Tomorrow. logo, Intel StrataFlash, Intel vPro, Intel Xeon Phi, Intel XScale, InTru, the InTru logo, the InTru Inside logo, InTru soundmark, Itanium, Itanium Inside, MCS, MMX, Pentium, Pentium Inside, Puma, skool, the skool logo, SMARTi, Sound Mark, Stay With It, The Creators Project, The Journey Inside, Thunderbolt, Ultrabook, vPro Inside, VTune, Xeon, Xeon Inside, X-GOLD, XMM, X-PMU and XPOSYS are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Copyright © 2009-2012, Intel Corporation. All rights reserved.

Overview



Discover how to use Hotspots analysis of the Intel® VTune™ Amplifier XE to understand where your application is spending time, identify *hotspots* - the most time-consuming program units, and detect how they were called.

Hotspots analysis is useful to analyze the performance of both serial and parallel applications.

About This Tutorial	This tutorial uses the sample <code>tachyon</code> and guides you through basic steps required to analyze the code for hotspots.
Estimated Duration	10-15 minutes.
Learning Objectives	After you complete this tutorial, you should be able to: <ul style="list-style-type: none">• Choose an analysis target.• Choose the Hotspots analysis type.• Run the Hotspots analysis to locate most time-consuming functions in an application.• Analyze the function call flow and threads.• Analyze the source code to locate the most time-critical code lines.• Compare results before and after optimization.
More Resources	<ul style="list-style-type: none">• Intel® Parallel Studio XE tutorials (HTML, PDF): http://software.intel.com/en-us/articles/intel-software-product-tutorials/• Intel® Parallel Studio XE support page: http://software.intel.com/en-us/articles/intel-parallel-studio-xe/

Navigation Quick Start



Intel® VTune™ Amplifier XE, an Intel® Parallel Studio XE tool, provides information on code performance for users developing serial and multithreaded applications on Windows* and Linux* operating systems. VTune Amplifier XE helps you analyze algorithm choices and identify where and how your application can benefit from available hardware resources.

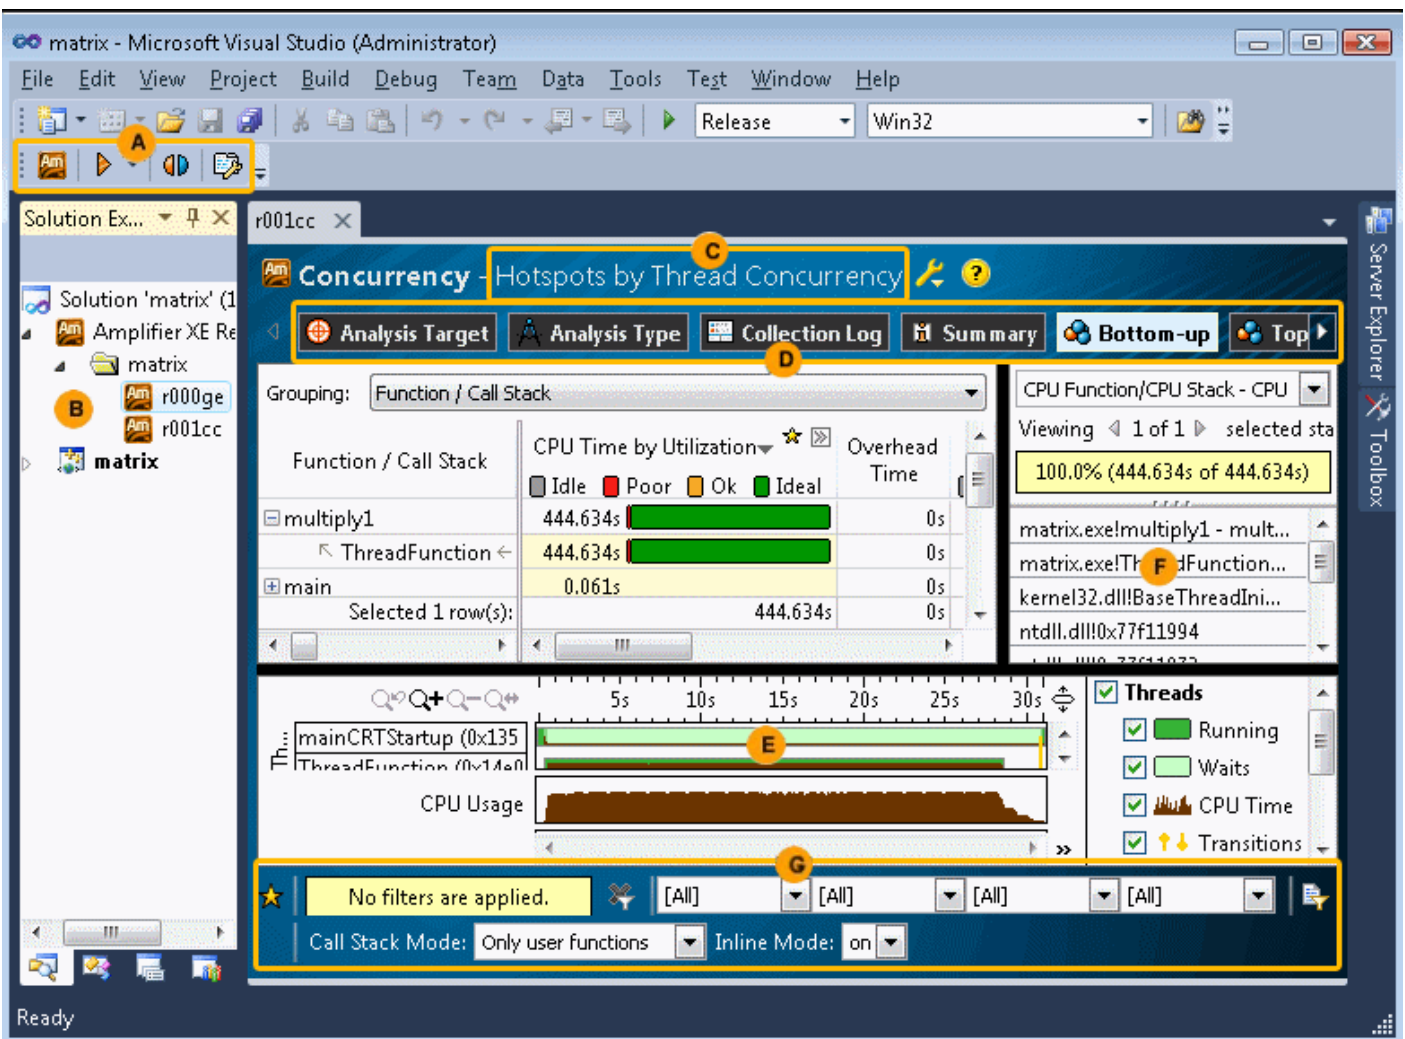
VTune Amplifier XE Access

To access the VTune Amplifier XE in the Visual Studio* IDE: From the Windows* **Start** menu, choose **Intel Parallel Studio XE 2013 > Parallel Studio XE 2013 with [VS2008 | VS2010]**.

To access the Standalone VTune Amplifier XE GUI, do one of the following:

- From the Windows* **Start** menu, choose **Intel Parallel Studio XE 2013 > Intel VTune Amplifier XE 2013**.
- From the Windows* **Start** menu, choose **Intel Parallel Studio XE 2013 > Command Prompt > Parallel Studio XE > IA-32 Visual Studio [2008 | 2010] mode** to set your environment, then type `amplxe-gui`.

VTune Amplifier XE/Visual Studio* IDE Integration



- A** Use the VTune Amplifier XE toolbar to configure and control result collection.
- B** VTune Amplifier XE results *.amplxe show up in the **Solution Explorer** under the **My Amplifier XE Results** folder. To configure and control result collection, right-click the project in the **Solution Explorer** and select the **Intel VTune Amplifier XE 2013** menu from the pop-up menu. To manage previously collected results, right-click the result (for example, r000ge.amplxe) and select the required command from the pop-up menu.
- C** From the drop-down menu, select a *viewpoint*, a preset configuration of windows/panes for an analysis result. For each analysis type, you can switch among several viewpoints to focus on particular performance metrics.
- D** Click the buttons on navigation toolbars to change window views and toggle window panes on and off.
- E** In the **Timeline** pane, analyze the thread activity and transitions presented for the user-mode sampling and tracing analysis results (for example, Hotspots, Concurrency, Locks and Waits) or analyze the distribution of the application performance per metric over time for the event-based sampling analysis results (for example, Memory Access, Bandwidth Breakdown).
- F** Use the **Call Stack** pane to view call paths for a function selected in the grid.

- G** Use the filter toolbar to filter out the result data according to the selected categories.

Standalone VTune Amplifier XE GUI

The screenshot displays the Intel VTune Amplifier XE 2013 GUI. The main window shows the 'Hotspots by CPU Usage' view. The interface is annotated with letters A through H:

- A**: File, View, Help menu.
- B**: Primary toolbar with icons for analysis and project management.
- C**: Project Navigator showing a tree view of projects (My Amplifier, r000hs, r001lw).
- D**: Results grid showing CPU Time by Utilization and Overhead Time for various functions.
- E**: Drop-down menu for selecting a viewpoint (e.g., Bottom-up, Top-down Tree).
- F**: Window tabs for Analysis Target, Analysis Type, Summary, Bottom-up, Top-down Tree, and Tasks.
- G**: Grouping drop-down menu set to 'Function / Call Stack'.
- H**: Filter toolbar at the bottom showing 'No filters are applied' and 'Call Stack Mode: Only user functions'.

Function / Call Stack	CPU Time by Utilization	Overhead Time
NQUEENS_ip_SETQUEEN	15.674ms	0ms
NQUEENS_ip_SETQUEE	15.674ms	0ms
chkstk	15.600ms	0ms
_kmp_x86_pause	15.526ms	0ms
Selected 1 row(s):		15.674ms

- A** Use the VTune Amplifier XE menu to control result collection, define and view project properties, and set various options.
- B** Configure and manage projects and results, and launch new analyses from the primary toolbar. Click the **Project Properties** button on this toolbar to manage result file locations.
- C** The **Project Navigator** provides an iconic representation of your projects and analysis results. Click the **Project Navigator** button on the toolbar to enable/disable the **Project Navigator**.
- D** Newly completed and opened analysis results along with result comparisons appear in the results tab for easy navigation.
- E** Use the drop-down menu to select a *viewpoint*, a preset configuration of windows/panes for an analysis result. For each analysis type, you can switch among several viewpoints to focus on particular performance metrics. Click the yellow question mark icon to read the viewpoint description.
- F** Switch between window tabs to explore the analysis type configuration options and collected data provided by the selected viewpoint.
- G** Use the **Grouping** drop-down menu to choose a granularity level for grouping data in the grid.

1 *Getting Started Tutorial: Finding Hotspots*

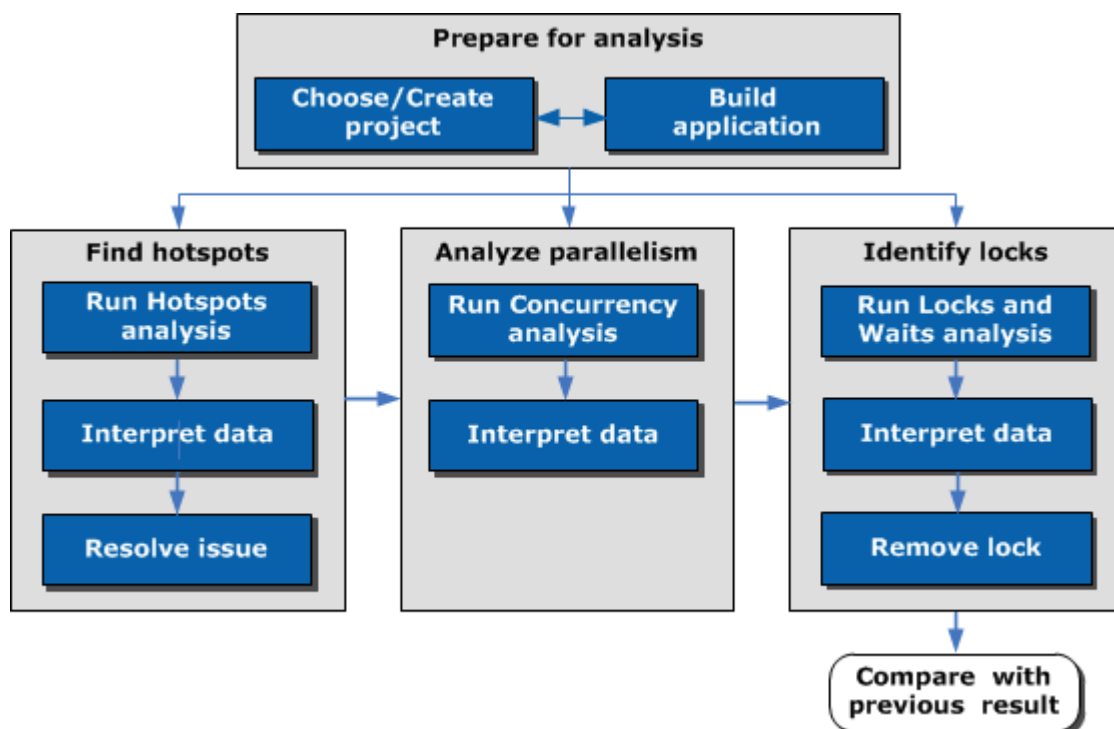


Use the filter toolbar to filter out the result data according to the selected categories.

Finding Hotspots



You can use the Intel® VTune™ Amplifier XE to identify and analyze hotspot functions in your serial or parallel application by performing a series of steps in a workflow. This tutorial guides you through these workflow steps while using a sample multithreaded application named `nqueens_parallel`.



Step 1: Prepare for analysis	<p>Do one of the following:</p> <ul style="list-style-type: none"> In the Visual Studio* IDE: Choose a project, verify settings, and build application In the Standalone Intel VTune Amplifier XE GUI: Build an application to analyze for hotspots and create a new VTune Amplifier XE project
Step 2: Find hotspots	<ul style="list-style-type: none"> Choose and run the Hotspots analysis. Interpret the result data. Resolve issue.
Step 3: Analyze parallelism	<ul style="list-style-type: none"> Choose and run the Concurrency analysis. Interpret the result data.
Step 4: Identify locks	<ul style="list-style-type: none"> Choose and run the Locks and Waits analysis. Interpret the result data. Remove lock.
Step 5: Check your work	<p>Re-build the target, re-run the Locks and Waits analysis, and compare the result data before and after optimization.</p>

Visual Studio* IDE: Choose Project and Build Application



Before you start analyzing your application target for hotspots, do the following:

1. Get software tools.
2. Choose a project.
3. Verify optimal compiler/linker options.
4. Configure target properties.
5. Build the target in the release mode.
6. Create a performance baseline.

Get Software Tools

You need the following tools to try tutorial steps yourself using the `nqueens_fortran` sample application:

- VTune Amplifier XE, including sample applications
- .zip file extraction utility
- Supported Fortran compiler (see Release Notes for more information)

Acquire Intel VTune Amplifier XE

If you do not already have access to the VTune Amplifier XE, you can download an evaluation copy from <http://software.intel.com/en-us/articles/intel-software-evaluation-center/>.

Install and Set Up VTune Amplifier XE Sample Applications

1. Copy the `nqueens_fortran.zip` file from the `<install-dir>\samples\<locale>\Fortran\` directory to a writable directory or share on your system. The default installation path is `C:\[Program Files]\Intel\VTune Amplifier XE 2013\`.
2. Extract the sample from the .zip file.



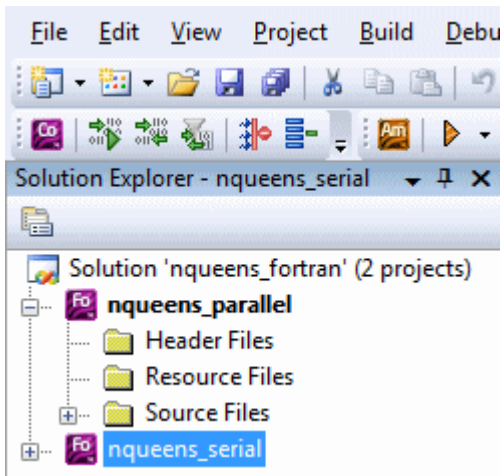
NOTE

- Samples are non-deterministic. Your screens may vary from the screen captures shown throughout this tutorial.
 - Samples are designed only to illustrate the VTune Amplifier XE features; they do not represent best practices for creating code.
-

Choose a Project

Choose a project with the analysis target in the Visual Studio IDE as follows:

1. From the Visual Studio menu, select **File > Open > Project/Solution....**
The **Open Project** dialog box opens.
2. In the **Open Project** dialog box, browse to the location you used to extract the `nqueens_fortran.zip` file and select the `nqueens_fortran.sln` file.



The solution is added to Visual Studio IDE and shows up in the Solution Explorer.

3. In the Solution Explorer, right-click the **nqueens_parallel** project and select **Project > Set as StartUp Project**.

nqueens_parallel appears in bold in the Solution Explorer.



NOTE `nqueens_serial` project provides a single-threaded version of the sample application. This project is provided for your reference only and not described in this tutorial.

When you choose a project in Visual Studio IDE, the VTune Amplifier XE automatically creates the `config.amplxproj` project file and sets the `nqueens_parallel` application as an analysis target in the project properties.

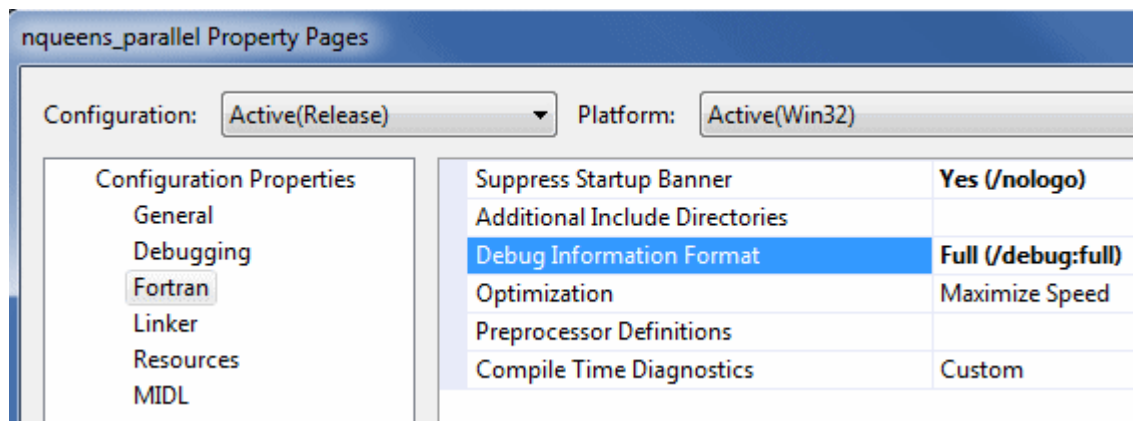
Verify Optimal Compiler/Linker Options

Configure Visual Studio project properties to generate the debug information for your application so that the VTune Amplifier XE can open the source code.

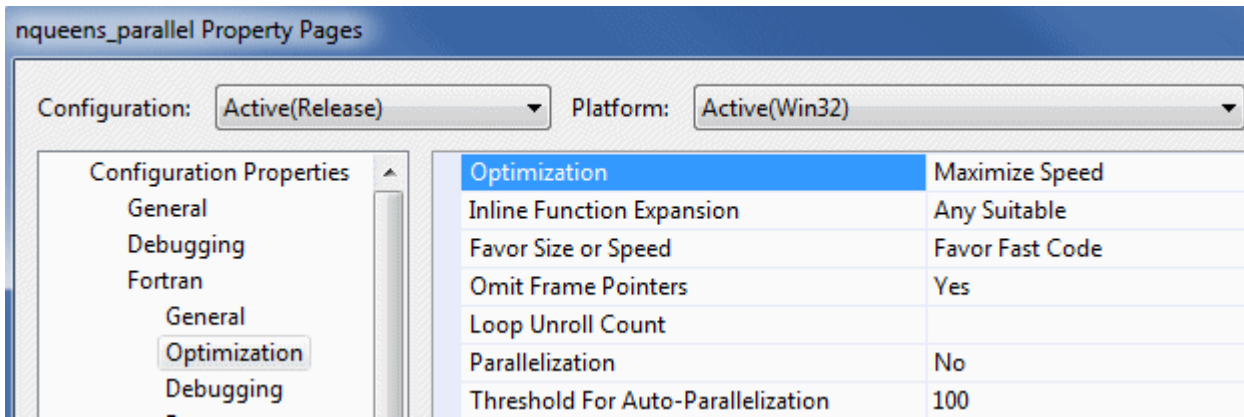
1. Select the **nqueens_parallel** project and go to **Project > Properties**.

The **nqueens_parallel Property Pages** dialog box opens.

2. Select **Configuration Properties > General** and make sure the selected **Configuration** (top of the dialog) is **Release**.
3. Select **Fortran > General** pane and specify the **Debug Information Format** as **Full (/debug:full)**.



4. Select **Fortran > Optimization** pane and set the **Optimization** to **Maximize Speed**



Configure Target Properties

1. Right-click the `nqueens_parallel` project and select **Properties** from the context menu.
The **nqueens_parallel Property Pages** dialog box opens.
2. In the **nqueens_parallel Property Pages** dialog box, select **Configuration Properties > Debugging**.
3. In the right pane, specify the **Command Arguments** as **14**.

Build the Target in the Release Mode

Build the target in the Release mode with full optimizations, which is recommended for performance analysis.

1. Go to the **Build > Configuration Manager...** dialog box and select the **Release** mode for your target project.
2. From the Visual Studio menu, select **Build > Build nqueens_parallel**.

The `nqueens_parallel.exe` application is built with your default Fortran compiler. For this tutorial, Intel® Fortran Compiler XE is used.

NOTE The build configuration for `nqueens_parallel` may initially be set to Debug, which is typically used for development. When analyzing performance issues with the VTune Amplifier XE, you are recommended to use the Release build with normal optimizations. In this way, the VTune Amplifier XE is able to analyze the realistic performance of your application.

Create a Performance Baseline

Run the application to create a performance baseline that will be used to identify optimization you achieve during performance tuning with the VTune Amplifier XE.

1. From the Visual Studio menu, select **Debug > Start Without Debugging**.

The `nqueens_parallel.exe` application starts running.

NOTE Before you start the application, minimize the amount of other software running on your computer to get more accurate results.

```
Starting nqueens solver for size 14 with 4 thread(s)
Number of solutions: 365596
Correct Result!
Calculations took 31215ms.
Press any key to continue . . . _
```


- Note the execution time displayed in the window caption. In the example above, the execution time is 31215 milliseconds.



- Run the application several times, note the execution time for each run, and use the average number. This helps to minimize skewed results due to transient system activity.
- The screenshots and execution time data provided in this tutorial are created on a system with 4 CPU cores. Your data may vary depending on the number and type of CPU cores on your system.

Recap

You chose the target for the Hotspots analysis, set up your environment to enable generating symbol information for your binary files, built the target in the Release mode, and created the performance baseline. Your application is ready for analysis.

Key Terms

- Baseline
- Hotspots analysis
- Target

Next Step

[Run Hotspots Analysis](#)

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Standalone GUI: Build Application and Create New Project



Before you start analyzing your application target for hotspots, do the following:

- Get software tools.
- Build application.
If you build the code in Visual Studio*, make sure to:
 - Verify optimal Fortran compiler options.
 - Build the target in the release mode.
- Create a performance baseline.
- Create a VTune Amplifier XE project.

Get Software Tools

You need the following tools to try tutorial steps yourself using the `nqueens_fortran` sample application:

- VTune Amplifier XE, including sample applications
- .zip file extraction utility

Configure Target Properties

1. Right-click the `nqueens_parallel` project and select **Properties** from the context menu.
The **nqueens_parallel Property Pages** dialog box opens.
2. In the **nqueens_parallel Property Pages** dialog box, select **Configuration Properties > Debugging**.
3. In the right pane, specify the **Command Arguments** as **14**.

Build the Target in the Release Mode

Build the target in the Release mode with full optimizations, which is recommended for performance analysis.

1. Go to the **Build > Configuration Manager...** dialog box and select the **Release** mode for your target project.
2. From the Visual Studio menu, select **Build > Build nqueens_parallel**.

The `nqueens_parallel.exe` application is built with your default Fortran compiler. For this tutorial, Intel® Fortran Compiler XE is used.



NOTE The build configuration for `nqueens_parallel` may initially be set to Debug, which is typically used for development. When analyzing performance issues with the VTune Amplifier XE, you are recommended to use the Release build with normal optimizations. In this way, the VTune Amplifier XE is able to analyze the realistic performance of your application.

Create a Performance Baseline

Run the application to create a performance baseline that will be used to identify optimization you achieve during performance tuning with the VTune Amplifier XE.

1. From the Visual Studio menu, select **Debug > Start Without Debugging**.

The `nqueens_parallel.exe` application starts running.



NOTE Before you start the application, minimize the amount of other software running on your computer to get more accurate results.

```
Starting nqueens solver for size 14 with 4 thread(s)
Number of solutions: 365596
Correct Result?
Calculations took 31215ms.
Press any key to continue . . . _
```

2. Note the execution time displayed in the shell window caption. In the example above, the execution time is 31215 milliseconds.



- Run the application several times, note the execution time for each run, and use the average number. This helps to minimize skewed results due to transient system activity.
- The screenshots and execution time data provided in this tutorial are created on a system with 4 CPU cores. Your data may vary depending on the number and type of CPU cores on your system.

Create a Project

1. From the **Start** menu select **Intel Parallel Studio XE 2013 > Intel VTune Amplifier XE 2013** to launch the VTune Amplifier XE standalone GUI.
2. Create a new project via **File > New > Project...**

The **Create a Project** dialog box opens.

3. Specify the project name `nqueens_parallel` that will be used as the project directory name.
VTune Amplifier XE creates the `nqueens_parallel` project directory under the `%USERPROFILE%\My Documents\My Amplifier Projects` directory and opens the **Project Properties: Target** dialog box.
4. In the **Target** tab, select the **Application to Launch** target type and specify your target as follows:
 - For the **Application** field, browse to: `<sample_code_dir>`, for example: `C:\samples\nqueens_fortran\vf9\nqueens_parallel\Release\nqueens_parallel.exe`.
5. In the **Application parameters** field, specify the task size for this target: 14.
6. Click **OK** to apply the settings and exit the **Project Properties** dialog box.

Recap

You set up your environment to enable generating symbol information for your binary files, built the target in the Release mode, created the performance baseline, and created the VTune Amplifier XE project for your analysis target. Your application is ready for analysis.

Key Terms

- Baseline
- Target

Next Step

Run Hotspots Analysis

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.


Notice revision #20110804

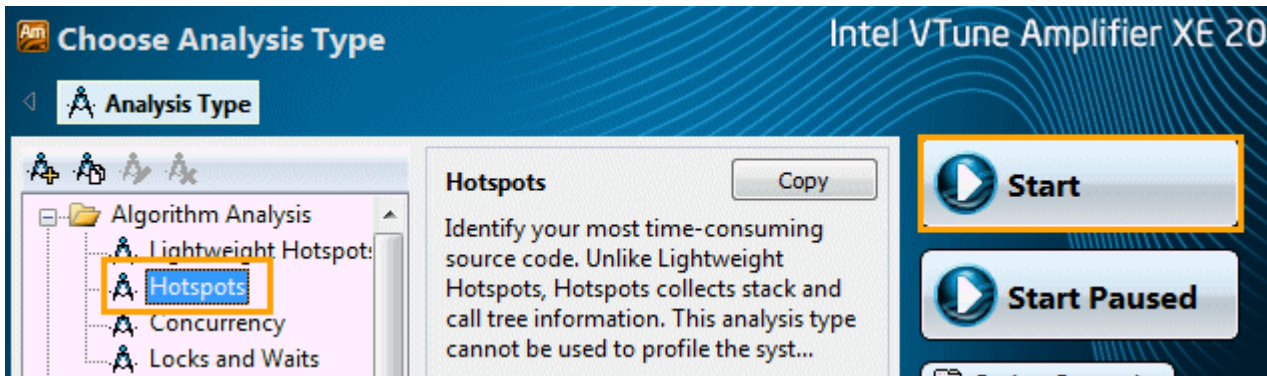
Run Hotspots Analysis



In this tutorial, you run the Hotspots analysis to identify the hotspots that took much time to execute.

To run an analysis:

1. From the VTune Amplifier XE toolbar, click the  **New Analysis** button.
VTune Amplifier XE result tab opens with the **Analysis Type** window active.
2. On the left pane of the **Analysis Type** window, locate the analysis tree and select **Algorithm Analysis > Hotspots**.
The right pane is updated with the predefined settings for the Hotspots analysis.
3. Click the **Start** button on the right command bar.



VTune Amplifier XE launches the `nqueens_parallel` application that makes calculations, displays the execution time, and exits. VTune Amplifier XE finalizes the collected results and opens the analysis results in the **Hotspots** viewpoint.

To make sure the performance of the application is repeatable, go through the entire tuning process on the same system with a minimal amount of other software executing.



NOTE This tutorial explains how to run an analysis from the VTune Amplifier XE graphical user interface (GUI). You can also use the VTune Amplifier XE command-line interface (`amplxe-cl` command) to run an analysis. For more details, check the *Command-line Interface Support* section of the VTune Amplifier XE Help.

Key Terms

- Elapsed time
- Finalization
- Hotspot
- Viewpoint

Next Step

[Interpret Hotspots Results](#)

Interpret Hotspots Results




When the sample application exits, the Intel® VTune™ Amplifier XE finalizes the results and opens the **Hotspots** viewpoint where each window or pane is configured to display code regions that consumed a lot of CPU time. To interpret the data on the sample code performance, do the following:

1. [Explore application-level performance.](#)
2. [Analyze the most time-consuming functions.](#)
3. [Identify the hotspot code region.](#)



NOTE The screenshots and execution time data provided in this tutorial are created on a system with 4 CPU cores. Your data may vary depending on the number and type of CPU cores on your system.

Explore Application-level Performance

Start analysis with the **Summary** window that opens by default when data collection completes. To interpret the data, hover over the question mark icons  to read the pop-up help and better understand what each performance metric means.

Elapsed Time: [ⓘ] 32.440s
Total Thread Count: 4
CPU Time: [ⓘ] 108.833s
Paused Time: [ⓘ] 0s

Note that **CPU Time** for the sample application is equal to 108.833 seconds. It is the sum of CPU time for all application threads. **Total Thread Count** is 4, so the sample application is multi-threaded.

The **Top Hotspots** section provides data on the most time-consuming functions (*hotspot functions*) sorted by CPU time spent on their execution.

Function	CPU Time [ⓘ]
NQUEENS_ip_SETQUEEN	105.684s
chkstk	2.697s
NQUEENS_ip_SOLVE	0.187s
_kmp_fork_barrier	0.125s
NtDelayExecution	0.064s
[Others]	0.076s

For the sample application, the `NQUEENS_ip_SETQUEEN` function, which took 105.684 seconds to execute, shows up at the top of the list as the hottest function.

The `[Others]` entry at the bottom shows the sum of CPU time for all functions not listed in the table.

Analyze the Most Time-consuming Functions

Click the **Bottom-up** tab to explore the **Bottom-up** pane. By default, the data in the grid is sorted by **Function/Call Stack**.

Analyze the **CPU Time** column values. This column is marked with a yellow star as the Data of Interest column. It means that the VTune Amplifier XE uses this type of data for some calculations (for example, filtering, stack contribution, and others). Functions that took most CPU time to execute are listed on top.

The `NQUEENS_ip_SETQUEEN` function took 105.684 seconds to execute.

Function / Call Stack	CPU Time [ⓘ] ★	Module	Function (Full)
⊕ NQUEENS_ip_SETQUEEN	105.684s	nqueens_parallel.exe	NQUEENS_ip_SETQUEEN
⊕ chkstk	2.697s	nqueens_parallel.exe	chkstk
⊕ NQUEENS_ip_SOLVE	0.187s	nqueens_parallel.exe	NQUEENS_ip_SOLVE
⊕ _kmp_fork_barrier	0.125s	libiomp5md.dll	_kmp_fork_barrier(int, ...)
⊕ NtDelayExecution	0.064s	ntdll.dll	NtDelayExecution

Double-click the hotspot function to open the source and identify the most time-critical code lines.

Identify the Hotspot Code Region

Source Line	Source	CPU Time
116		
117	recursive subroutine setQueen (queens, row, col)	7.171s
118	implicit none	
119	integer, intent(inout) :: queens(:)	
120	integer, intent(in) :: row, col	
121	integer :: i	
122		
123	! In order to avoid a data race on the "queens" array,	
124	integer :: lcl_queens(ubound(queens,dim=1))	1.861s
125		
126	! Make copy of queens array	
127	lcl_queens = queens	54.150s
128		
129	do i=1,row-1	1.819s

The table below explains some of the features available in the **Source** window when viewing the Hotspots analysis data.

- Source** pane displaying the source code of the application if the function symbol information is available. The beginning of the hotspot function is highlighted. The source code in the **Source** pane is not editable.

If the function symbol information is not available, the **Assembly** pane opens displaying assembler instructions for the selected hotspot function. To enable the **Source** pane, make sure to build the target properly.
- Processor time attributed to a particular code line. If the hotspot is a system function, its time, by default, is attributed to the user function that called this system function.
- Source** window toolbar. Use the hotspot navigation buttons to switch between most performance-critical code lines. Hotspot navigation is based on the metric column selected as a Data of Interest. For the Hotspots analysis, this is **CPU Time**. Use the **Source/Assembly** buttons to toggle the **Source/Assembly** panes (if both of them are available) on/off.

By default, when you double-click the hotspot in the **Bottom-up** pane, the VTune Amplifier XE opens the source file related to this function. Click the hotspot navigation button to go to the code line that took the most CPU time. For the `NQUEENS_ip_SETQUEEN` function, the hottest code line is 127. This code is used to create a local copy of the `queens` array to avoid a data race. Click the **Source Editor** button on the **Source** window toolbar to open the default code editor and work on optimizing the code.

Key Terms


- CPU time
- Viewpoint

Next Step

Resolve Issue

Resolve Issue



You identified that the most time-consuming function is `NQUEENS_ip_SETQUEEN`. If you click the  **Source Editor** button from the **Source** pane, the VTune Amplifier XE opens the source `nqueens_parallel.f90` file at the hotspot line in the default code editor. You see that the OpenMP* cycle is calling the recursive `setQueen` function that initializes the `queens` array. To avoid a data race, this array is copied in each thread (see line 127):

```

115  ! Recursive routine to set a queen on the board
116
117  recursive subroutine setQueen (queens, row, col)
118    implicit none
119    integer, intent(inout) :: queens(:)
120    integer, intent(in) :: row, col
121    integer :: i
122
123    ! In order to avoid a data race on the "queens" array, create a local copy.
124    integer :: lcl_queens(ubound(queens,dim=1))
125
126    ! Make copy of queens array
127    lcl_queens = queens
128
129    do i=1,row-1
130      ! vertical attacks
131      if (lcl_queens(i) == col) return
132      ! diagonal attacks
133      if (abs(lcl_queens(i)-col) == (row-i)) return
134    end do

```

This means that the number of local copies is equal to the number of threads. Since the function is recursive, the array is also copied in every function call, which is unnecessary and creates a big overhead.

To resolve this issue, you may enable OpenMP to make a copy of the array per thread. To do this:

1. Comment out lines 124 and 127.

```

123    ! In order to avoid a data race on the "queens" array, create a local copy.
124    ! integer :: lcl_queens(ubound(queens,dim=1))
125
126    ! Make copy of queens array
127    ! lcl_queens = queens

```

2. Search and replace all `lcl_queens` entries with `queens`.
3. Edit line 159 to add the `PRIVATE(queens)` directive.

This enables the OpenMP run-time to create a private copy of the array for each thread.

```

159  !$OMP PARALLEL DO PRIVATE(queens)
160    do i=1,size
161      ! try all positions in first row
162      call SetQueen (queens, 1, i)
163    end do

```

4. Save the changes made in the source file.
5. From the Visual Studio **Build** menu, select **Rebuild nqueens_parallel**.
The project is rebuilt.
6. From Visual Studio **Debug** menu, select **Start Without Debugging** to run the application.


```
Starting nqueens solver for size 14 with 4 thread(s)
Number of solutions: 365596
Correct Result!
Calculations took 16035ms.
Press any key to continue . . .
```

Visual Studio runs `nqueens_parallel.exe`. Note that the execution time has reduced from 31215 ms to 16035 ms. This means that the proposed solution gives 15180 ms of CPU time reduction.

To identify other possible performance issues, you may run the Concurrency analysis and see how effectively your application is parallelized.

Next Step


Run Concurrency Analysis

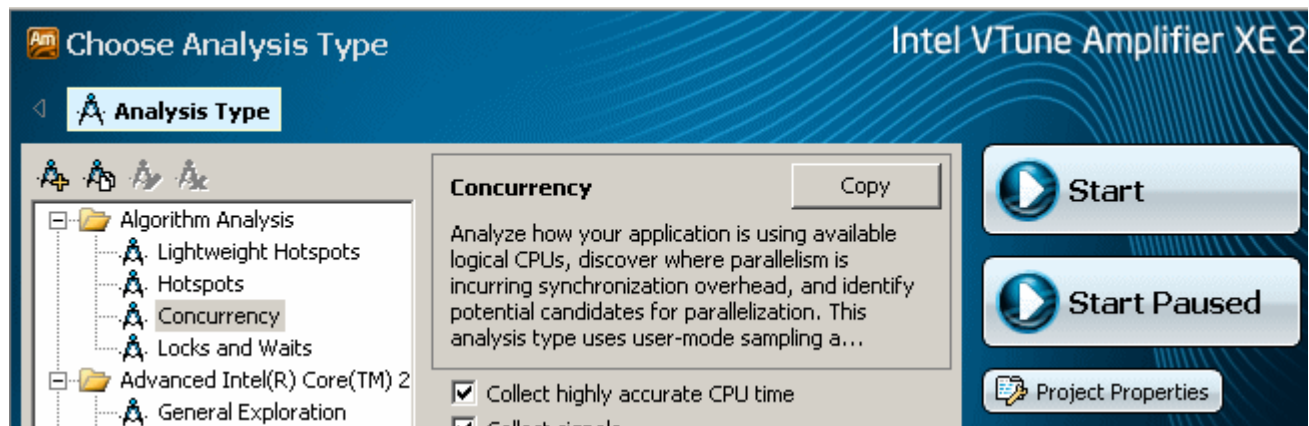
Run Concurrency Analysis



Run the Concurrency analysis to understand how effectively your application is parallelized.

To run an analysis:

1. From the VTune Amplifier XE toolbar, click the  **New Analysis** button.
VTune Amplifier XE result tab opens with the **Analysis Type** window active.
2. On the left pane of the **Analysis Type** window, locate the analysis tree and select **Algorithm Analysis > Concurrency**.
The right pane is updated with the predefined settings for the Concurrency analysis.
3. Click the **Start** button on the right command bar.



VTune Amplifier XE launches the `nqueens_parallel` application that makes calculations, displays the execution time, and exits. VTune Amplifier XE finalizes the collected results and opens the analysis results in the **Hotspots by Thread Concurrency** viewpoint.

To make sure the performance of the application is repeatable, go through the entire tuning process on the same system with a minimal amount of other software executing.



NOTE This tutorial explains how to run an analysis from the VTune Amplifier XE graphical user interface (GUI). You can also use the VTune Amplifier XE command-line interface (`amplxe-cl` command) to run an analysis. For more details, check the *Command-line Interface Support* section of the VTune Amplifier XE Help.

Key Terms

- Finalization
- Viewpoint

Next Step

[Interpret Concurrency Results](#)

Interpret Concurrency Results



When the sample application exits, the Intel® VTune™ Amplifier E finalizes the results and opens the **Hotspots by Thread Concurrency** viewpoint where each window or pane is configured to display data on application parallelism and usage of processor cores. To interpret the data on the sample code performance, do the following:

1. [Explore application-level concurrency](#)
2. [Identify the most time-consuming function.](#)



NOTE The screenshots and execution time data provided in this tutorial are created on a system with 4 CPU cores. Your data may vary depending on the number and type of CPU cores on your system.

Explore Application-level Concurrency

Start analysis with the **Summary** window that opens by default when data collection completes. To interpret the data, hover over the question mark icons ⓘ to read the pop-up help and better understand what each performance metric means.

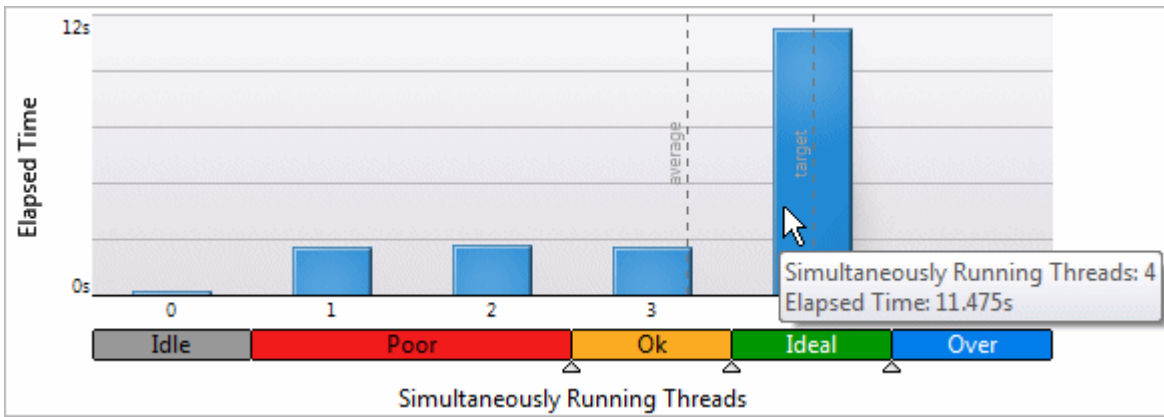
You see that after optimization the Elapsed time has reduced from 32.440 seconds to 18.351 seconds.

Elapsed Time: ⓘ	18.351s ⓘ
Overhead Time: ⓘ	0.497s
CPU Time: ⓘ	59.385s
Total Thread Count:	5
Paused Time: ⓘ	0s



NOTE The Concurrency analysis adds an overhead to the application execution. The overhead often depends on the number of threads and synchronization objects used in the application. This is the reason why Elapsed time data provided in the **Summary** window may differ from the data reported after the application launch outside of the VTune Amplifier XE.

The **Thread Concurrency Histogram** shows that the average concurrency level of the sample application is about 3.2 while the target concurrency level for this application on the 4-core system is 4. If you hover over the highest bar, you see that this application has run 4 threads for more than 11 seconds, which is categorized by the VTune Amplifier XE as Ideal processor utilization. The application has run one and two threads simultaneously for more than 4 seconds, which is classified as poor parallelization.



Identify the Most Time-consuming Function

Click the **Bottom-up** tab to switch to the **Bottom-up** window and analyze application performance by function. By default, the grid is sorted by the **CPU Time by Utilization** metric in the descending order. Select the **Function/Thread/Call Stack** grouping level from the **Grouping** menu. This granularity enables you to visualize threads where the hotspots functions were executed.

After initial optimization, the `NQUEENS_ip_SETQUEEN` function is still a bottleneck. Click the plus sign \oplus at the `NQUEENS_ip_SETQUEEN` function. You see that this function's execution was parallelized among four threads.


Concurrency - Hotspots by Thread Concurrency

Analysis Target | Analysis Type | Summary | **Bottom-up** | Top-down Tree | Tasks

Grouping: Function / Thread / Call Stack

Function / Thread / Call Stack	CPU Time by Utilization					Overhead Time	Wait Time by Utilization			
	Idle	Poor	Ok	Ideal	Over		Idle	Poor	Ok	Ideal
\oplus NQUEENS_ip_SETQUEEN	59.081s					0.422s				2.875s
\oplus OMP Worker Thread #1 (0x1c40)	17.138s					0.150s				0.747s
\oplus mainCRTStartup (0x1c80)	15.544s					0.111s				0.639s
\oplus OMP Worker Thread #2 (0x1f1c)	14.174s					0.040s				0.804s
\oplus OMP Worker Thread #3 (0x1d20)	12.226s					0.121s				0.684s
\oplus NtDelayExecution	0.097s					0s				
\oplus NQUEENS_ip_SOLVE	0.040s					0.040s				1.631s
\oplus RtlEnterCriticalSection	0.036s					0.036s				

Select these threads in the grid, right-click and choose the **Filter In by Selection** context menu option. The **Timeline** pane below is updated to display data for the selected threads only.



- Timeline area.** When you hover over the graph element, the timeline tooltip displays the time passed since the application has been launched.

- 2 Threads area** that shows the distribution of CPU time utilization per thread. Hover over a bar to see the CPU time utilization in percent for this thread at each moment of time. Dark green zones show the time threads are active. Light-green zones show the time threads were waiting.
- 3 Transitions.** The execution flow between threads where one thread signals to another thread waiting to receive that signal. You may zoom in to a time region to get more detailed view of the transitions. To do this, drag and drop to select the region and right-click to select the **Zoom In on Selection** option from the context menu.

The **Timeline** pane for the sample application shows a large number of transitions between threads, which means that the threads spent noticeable time transferring execution to each other. If you uncheck the **Transitions** display option on the right, you see that workload balance is also poor since three of four threads were waiting for OMP Worker Thread #1 to complete execution.

Run the Locks and Waits analysis to understand what prevents the sample code from effective thread concurrency and processor utilization.

Key Terms

- Thread concurrency
- Viewpoint

Next Step


[Run Locks and Waits Analysis](#)

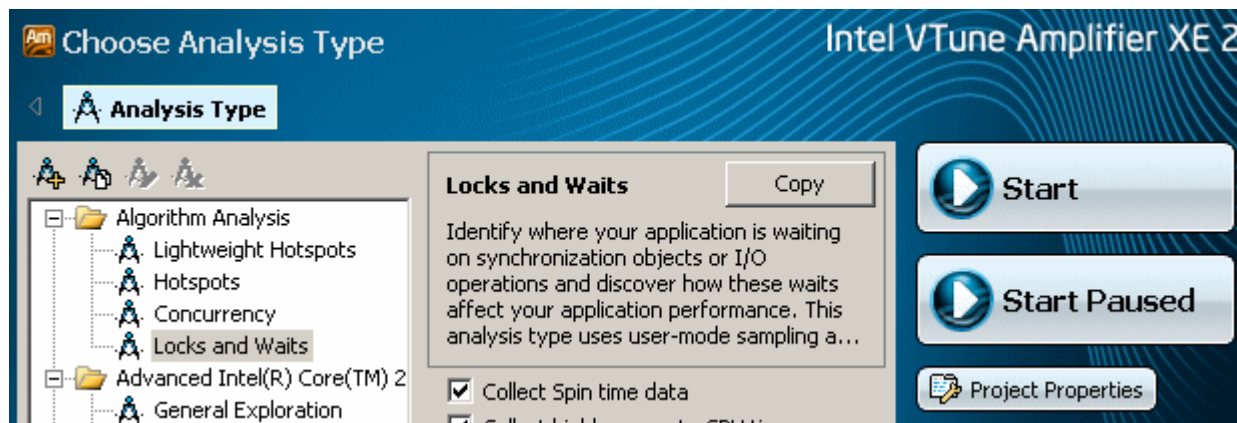
Run Locks and Waits Analysis



Run the Locks and Waits analysis to identify synchronization objects that caused contention and fix the problem in the source.

To run an analysis:

1. From the VTune Amplifier XE toolbar, click the  **New Analysis** button.
VTune Amplifier XE result tab opens with the Analysis Type window active.
2. From the analysis tree on the left, select **Algorithm Analysis > Locks and Waits**.
The right pane is updated with the default options for the Locks and Waits analysis.
3. Click the **Start** button on the right command bar.



VTune Amplifier XE launches the `nqueens_parallel` executable that makes calculations, displays the execution time, and exits. VTune Amplifier XE finalizes the collected data and opens the results in the Locks and Waits viewpoint .



- To make sure the performance of the application is repeatable, go through the entire tuning process on the same system with a minimal amount of other software executing.
- This tutorial explains how to run an analysis from the VTune Amplifier XE graphical user interface (GUI). You can also use the VTune Amplifier XE command-line interface (`amplxe-cl` command) to run an analysis. For more details, check the Command-line Interface Support section of the VTune Amplifier XE Help.

Key Terms

- Finalization
- Viewpoint

Next Step

Interpret Locks and Waits Results

Interpret Locks and Waits Results



When the sample application exits, the Intel® VTune™ Amplifier XE finalizes the results and opens the Locks and Waits viewpoint that is configured to display synchronization objects sorted by Wait time. To interpret the data on the sample code performance, do the following:

1. Identify locks.
2. Analyze source code.

Identify Locks

Click the **Bottom-up** tab to open the **Bottom-up** pane.

Sync Object / Function / Call Stack	Wait Time by Utilization	Wait Count	Spin Time
OMP Critical_NQUEENS_ip_SETQUEEN:141	11.964s [Poor]	24,551	8.147s
OMP Join Barrier_NQUEENS_ip_SOLVE:159	9.066s [Poor]	4	0.002s
Multiple Objects	0.002s [Ok]	1	0s

The table below explains the type of data provided in the **Bottom-up** pane:

- 1 Synchronization objects that control threads in the application. The hash (unique number) appended to some names of the objects identify the stack creating this synchronization object.
- 2 The utilization of the processor time when a given thread waited for some event to occur. By default, the synchronization objects are sorted by **Poor** processor utilization type. Bars showing OK or Ideal utilization (orange and green) are utilizing the processors well. You should focus your optimization efforts on functions with the longest poor CPU utilization (red bars if the bar format is selected). Next, search for the longest over-utilized time (blue bars).


This is the Data of Interest column for the Locks and Waits analysis results that is used for different types of calculations, for example: call stack contribution, percentage value on the filter toolbar.

- 3 Number of times the corresponding system wait API was called. For a lock, it is the number of times the lock was contended and caused a wait. Usually you are recommended to focus your tuning efforts on the waits with both high Wait Time and Wait Count values, especially if they have poor utilization.
- 4 Wait time, during which the CPU is busy. This often occurs when a synchronization API causes the CPU to poll while the software thread is waiting. Some Spin time may be preferable to the alternative of the increased thread context switches. However, too much Spin time can reflect lost opportunity for productive work.

In the `nqueens_parallel` sample code, there are two critical wait objects, `OMP Critical_NQUEENS_ip_SETQUEEN` and `OMP Join Barrier_NQUEENS_ip_SOLVE`, that caused redundant synchronization and took the longest Wait time and highest Wait count. The bar indicators in the **Wait Time** column indicate that most of the time for these objects processor cores were either underutilized or idle.

Analyze Source Code


Explore the source of the critical synchronization objects that caused significant Wait time and poor processor utilization. Double-click the `OMP Critical_NQUEENS_ip_SETQUEEN` object to analyze the source of the `NQUEENS_ip_SETQUEEN` wait function. Click the  button on the **Source** pane toolbar to go to the biggest hotspot code line in the function. VTune Amplifier XE highlights line 142 protected by the OpenMP* critical section.



Source Line	Source	Wait Time by Utiliza... Idle Poor Ok	Wait Count	Spin Time
139	<code>if (row == size) then</code>			
140	<code>! Change the Critical session for the Atomic directive</code>			
141	<code>!\$OMP CRITICAL</code>			
142	<code>nrOfSolutions = nrOfSolutions + 1</code>	11.964s	24,551	8.147s
143	<code>!\$OMP END CRITICAL</code>			
144	<code>else</code>			

The `NQUEENS_ip_SETQUEEN` function was waiting for almost 12 seconds while this code line was executing and most of the time the processor was either underutilized or idle. During this time, this operation was contended 24,551 times.

Hover over any transition line in the **Timeline** pane below to explore the infotip and make sure that all the transitions are caused by the `OMP Critical_NQUEENS_ip_SETQUEEN` critical section.

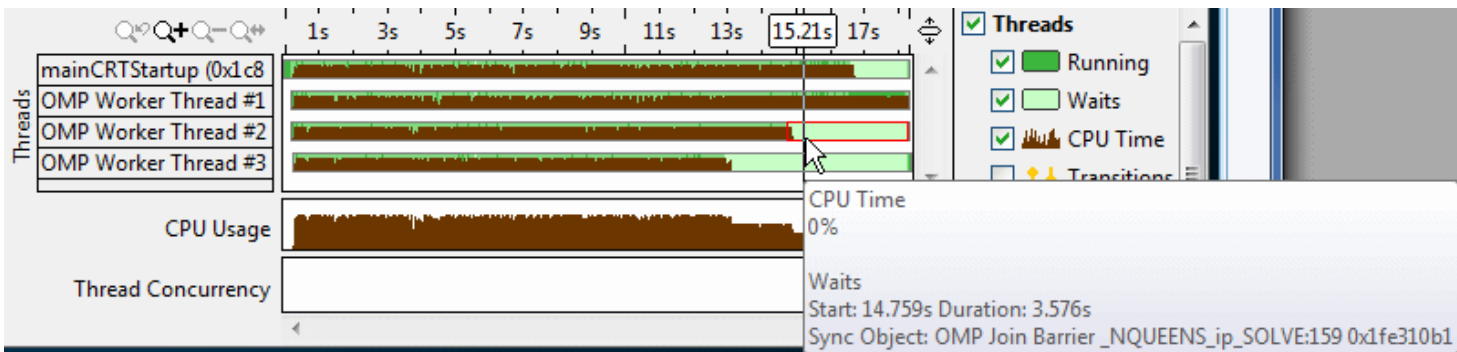



The OMP Critical_NQUEENS_ip_SETQUEEN section is the place where the application is serializing. Each thread has to wait for the critical section to be available before it can proceed. Only one thread can be in the critical section at a time.

To explore the next issue, double-click the OMP Join Barrier_NQUEENS_ip_SOLVE synchronization object to open the source function and go to the hottest line.

Sour... Line	Source	Wait Time by Utilization	Wait Count	Spin Time
152	! Main solver routine			
153	subroutine solve (queens)			
154	implicit none			
155	integer, intent(inout) :: queens(:)			
156	integer :: i			
157				
158	! Enable dynamic load scheduling			
159	!OMP PARALLEL DO PRIVATE(queens)	1.529s	1	0.030s
160	do i=1,size			

The OMP Join Barrier_NQUEENS_ip_SOLVE object creates a barrier for threads synchronization: a thread should wait until other threads complete execution. The **Timeline** pane illustrates the thread imbalance displaying light-green wait regions for each thread. If you hover over a wait region, the infotip shows that the wait happened on the OMP Join Barrier_NQUEENS_ip_SOLVE synchronization object.



You need to optimize the code to make it more concurrent. Click the  **Source Editor** button on the **Source** window toolbar to open the code editor and optimize the code.

Key Terms

- Elapsed time
- Wait time

Next Step

[Remove Lock](#)


Remove Locks



In the Source window, you located the synchronization objects that caused significant waits while the processor cores were underutilized and generated multiple wait count. To resolve the issues, do the following:

1. [Open the code editor.](#)
2. [Modify the code to remove locks.](#)
3. [Recompile the project and check the result.](#)

Open the Code Editor

Click the  **Source Editor** button to open the `nqueens_parallel.f90` file in your default editor:

```
140:      ! Change the Critical session for the Atomic directive OMP ATOMIC
141:      !$OMP CRITICAL
142:      nrOfSolutions = nrOfSolutions + 1
143:      !$OMP END CRITICAL
144:  else
145:      ! try to fill next row
146:      do i=1,size
147:          call setQueen (queens, row+1, i)
148:      end do
149:  end if
150: end subroutine SetQueen
```

Remove Locks

The critical section introduced in line 141 protects the global variable from a race condition in a multithreaded application but it spawns a redundant synchronization. To resolve this issue, you may replace the critical section with an atomic operation as follows:

1. Edit line 141 to replace the `OMP CRITICAL` with the `OMP ATOMIC` directive.
2. Comment out or remove line 143.

A threads barrier, created by OpenMP* directive in line 159, synchronizes the threads but creates a lock with long Wait time. You may resolve this by enabling dynamic load scheduling as follows:

1. Edit line 159 to add the `SCHEDULE (DYNAMIC)` directive to the OpenMP pragma:


```
141     !$OMP ATOMIC
142     nrOfSolutions = nrOfSolutions + 1
143     !$OMP END CRITICAL
144 else
145     ! try to fill next row
146     do i=1,size
147         call setQueen (queens, row+1, i)
148     end do
149 end if
150 end subroutine SetQueen
151
152 ! Main solver routine
153 subroutine solve (queens)
154     implicit none
155     integer, intent(inout) :: queens(:)
156     integer :: i
157
158 ! Enable dynamic load scheduling
159 !$OMP PARALLEL DO PRIVATE(queens) SCHEDULE(DYNAMIC)
```

2. Save your changes.

Recompile the Project and Check the Result

1. From the **Build** menu, select **Rebuild nqueens_parallel**.

The project is rebuilt.

2. From the **Debug** menu, select **Start Without Debugging** to run the application.

```
Starting nqueens solver for size 14 with 4 thread(s)
Number of solutions: 365596
Correct Result!
Calculations took 15767ms.
Press any key to continue . . .
```

Visual Studio runs the `nqueens_parallel.exe`. Note that execution time reduced from 16035 ms to 15767 ms.

Key Terms

- Wait time

Next Step

[Compare with Previous Result](#)


Compare with Previous Result

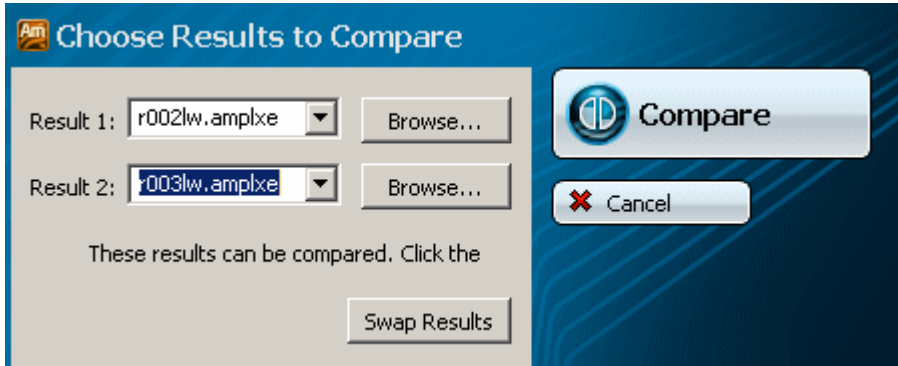


You made sure that removing the critical section gave you 268 ms of optimization in the application execution time. To understand the impact of your changes and how the CPU utilization has changed, re-run the Locks and Waits analysis on the optimized code and compare results:

1. Compare results before and after optimization.
2. Identify the performance gain by metrics.
3. Compare timeline data.

Compare Results Before and After Optimization

1. Run the Locks and Waits analysis on the modified code.
2. Click the  **Compare Results** button on the VTune Amplifier XE toolbar.
The **Compare Results** window opens.
3. Specify the Locks and Waits analysis results you want to compare:



The **Summary** window opens providing the statistics for the difference between collected results.

Identify the Performance Gain by Metrics

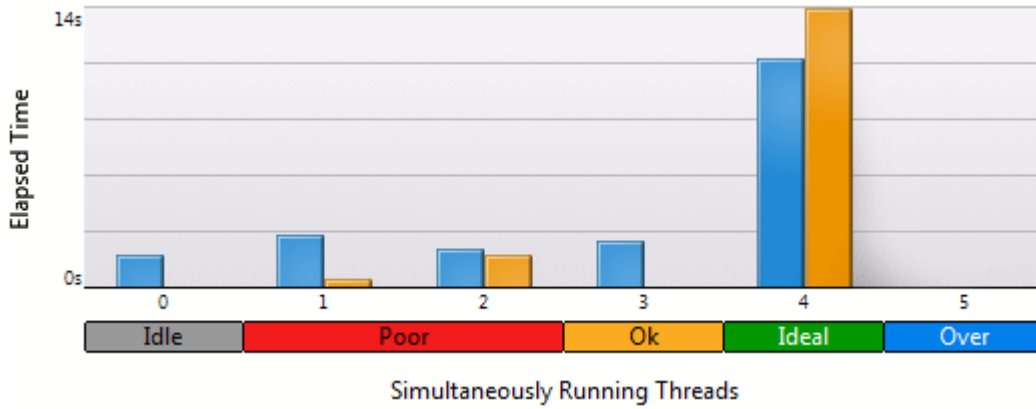
The **Result Summary** section of the **Summary** window shows that after optimization all critical metric values has reduced significantly. The **Elapsed Time** data shows the optimization of 4 seconds for the whole application. **Wait Time** decreased by 20.5 seconds, **Wait Count** - by 24,570.

Elapsed Time: [?]	20.467s - 16.446s = 4.022s
Total Thread Count:	Not changed, 5
Wait Time: [?]	41.221s - 20.705s = 20.516s
Spin Time: [?]	8.449s - 0.428s = 8.021s
Wait Count: [?]	24,662 - 92 = 24,570
CPU Time: [?]	66.448s - 59.435s = 7.013s
Paused Time: [?]	Not changed, 0s



NOTE The Locks and Waits analysis adds an overhead to the application execution. The overhead often depends on the number of threads and synchronization objects used in the application. This is the reason why Elapsed time data provided in the **Summary** window may differ from the data reported after the application launch outside of the VTune Amplifier XE.

According to the **Thread Concurrency** histogram, after optimization (an orange bar) 4 threads ran in parallel effectively utilizing CPU cores for 14 seconds, which is categorized by the VTune Amplifier XE as the Ideal processor utilization. The previous version of the application ran on 4 threads for 11.5 seconds.




In the **Bottom-up** pane, locate the OpenMP* critical section you identified as a bottleneck in your code. Since you removed it during optimization, the optimized result **r003lw** does not show any performance data for this synchronization object. If you collapse the **Wait Time:Difference by Utilization** column by clicking the button, you see that with the optimized result you got almost 12 seconds of optimization in Wait time. Using dynamic scheduling for the threads barrier gave you 4.5 seconds of optimization in Wait time.

Sync Object / Function / Call Stack	Wait Time:Difference by Utilization	Wait Time:r002lw by Utilization	Wait Time:r003lw by Utilization
OMP Critical_NQUEENS_ip_SETQ	11.964s	11.964s	
OMP Join Barrier_NQUEENS_ip_S	4.549s	9.066s	4.517s
Stream C:\Windows\Globalization	0.000s	0.000s	
Stream 0x2f221507	0.000s	0.000s	
Multiple Objects	0.000s	0.002s	0.002s

Compare Timeline Data

Open the optimized result of the Locks and Waits analysis **r003lw**, click the **Bottom-up** tab and analyze the **Timeline** pane.



The optimized result does not have transitions anymore. Though the threads are not fully balanced, the wait regions have reduced.

Compare analysis results regularly to look for regressions and to track how incremental changes to the code affect its performance. You may also want to use the VTune Amplifier XE command-line interface and run the `amplxe-cl` command to test your code for regressions. For more details, see the *Command-line Interface Support* section in the VTune Amplifier XE online help.

Key Terms

- Elapsed time


2 *Getting Started Tutorial: Finding Hotspots*



- Thread concurrency
- Wait time

Summary



You have completed the Finding Hotspots tutorial. Here are some important things to remember when using the Intel® VTune™ Amplifier E to analyze your code for hotspots:

Step	Tutorial Recap	Key Tutorial Take-aways
1. Prepare for analysis	<p>If you used the Visual Studio* IDE: You chose the target for the Hotspots analysis, configured Visual Studio to generate symbol information for your binary files, built the target in the Release mode, and created the performance baseline.</p> <p>If you used the standalone GUI: You set up your environment to enable generating symbol information for your binary files, built the target in the Release mode, created the performance baseline, and created the VTune Amplifier XE project for your analysis target.</p>	<ul style="list-style-type: none"> • Configure your project properties to get the most accurate results for user binaries and to analyze the performance of your application at the code line level. • Create a performance baseline to compare the application versions before and after optimization. Make sure to use the same workload for each application run. • Use the Project Properties: Target tab to choose and configure your analysis target. For Visual Studio* projects, the analysis target settings are inherited automatically. • Use the Analysis Type configuration window to choose, configure, and run the analysis. You can also run the analysis from command line using the <code>amplxe-cl</code> command.
2. Find hotspots	<p>You launched the Hotspots data collection that analyzed function calls and CPU time spent in each program unit of your application and identified the following hotspots:</p> <ul style="list-style-type: none"> • A function that took the most CPU time and could be a good candidate for algorithm tuning. • The code section that took the most CPU time to execute. 	<ul style="list-style-type: none"> • Start analyzing the performance of your application from the Summary window to explore the performance metrics for the whole application. Then, move to the Bottom-up window to analyze the performance per function. Focus on the <i>hotspots</i> - functions that took the most CPU time. By default, they are located at the top of the table. • Double-click the hotspot function in the Bottom-up pane or Call Stack pane to open its source code and navigate between hotspots using the Source window navigation buttons.
3. Eliminate hotspots	<p>You optimized the algorithm by enabling the OpenMP* library create a private copy of the array. You rebuilt the application and got performance gain of 15180 ms.</p>	<p>Click the  Source Editor button to open your default source editor directly from the VTune Amplifier XE Source window.</p>

Step	Tutorial Recap	Key Tutorial Take-aways
4. Analyze concurrency	You launched the Concurrency analysis and identified poor thread concurrency for the whole application execution. You analyzed the timeline and identified poor thread balance: all OpenMP threads were constantly transferring execution to each other and were waiting for all threads to complete execution.	<ul style="list-style-type: none"> Start your analysis with the Summary window. Consider the Target concurrency metric specified in the Thread Concurrency Histogram as your optimization goal. The Average metric is calculated as CPU time / Elapsed time. Use this number as another baseline for your performance measurements. The closer this number to the number of cores, the better. In the Bottom-up window, use the Filter In by Selection context menu option to focus on the performance-critical functions in the grid and analyze their performance over time in the Timeline pane.
5. Find lock	<p>You ran the Locks and Waits analysis and identified the following hotspots:</p> <ul style="list-style-type: none"> Two synchronization objects with the high Wait Time and Wait Count values and poor CPU utilization that could be locks affecting application parallelism. Your next step is to analyze the code of their wait functions. The code sections that caused significant waits and numerous transitions between threads. 	<ul style="list-style-type: none"> Use the Analysis Type configuration window to choose, configure, and run the analysis. For recently used analysis types, you may use the shortcuts to run a recent analysis: <ul style="list-style-type: none"> In the standalone interface: From the File menu, select New > [recent_analysis_type]. In Visual Studio: Click the down arrow next to the  New Analysis button on the VTune Amplifier XE toolbar and select the required analysis type from the drop-down list. In the Bottom-up window, focus on the synchronization objects that under- or over-utilized the available logical CPUs and have the highest Wait time and Wait Count values. By default, the objects with the highest Wait time values show up at the top of the window.
6. Remove lock	You optimized the application execution time by removing the unnecessary critical section that caused redundant synchronization and by adding the dynamic load scheduling.	Double-click the most time-critical synchronization object in the Bottom-up pane. This opens the source code for the wait function it belongs to. Use the hotspot navigation buttons to identify the most time-critical code lines.
7. Check your work	You ran the Locks and Waits analysis on the optimized code and compared the results before and after optimization using the Compare mode of the VTune Amplifier XE.	<p>Perform regular regression testing by comparing analysis results before and after optimization.</p> <p>From GUI, click the  Compare Results button on the VTune Amplifier XE toolbar. From command line, use the <code>amplxe-cl</code> command.</p>

Next step: Prepare your own application(s) for analysis. Then use the VTune Amplifier XE to find and eliminate hotspots.

Key Terms



baseline: A performance metric used as a basis for comparison of the application versions before and after optimization. Baseline should be measurable and reproducible.

CPU time: The amount of time a thread spends executing on a logical processor. For multiple threads, the CPU time of the threads is summed. The application CPU time is the sum of the CPU time of all the threads that run the application.






Elapsed time: The total time your target ran, calculated as follows: **Wall clock time at end of application - Wall clock time at start of application.**

finalization: A process during which the VTune Amplifier XE converts the collected data to a database, resolves symbol information, and pre-computes data to make further analysis more efficient and responsive.


hotspot: A section of code that took a long time to execute. Some hotspots may indicate bottlenecks and can be removed, while other hotspots inevitably take a long time to execute due to their nature.

target: A target is an executable file you analyze using the VTune Amplifier XE.

thread concurrency: A performance metric that helps identify how an application utilizes the processors in the system by comparing the application concurrency level (the number of active threads) and target concurrency level (by default, equal to the number of physical cores). Thread concurrency may be higher than CPU usage if threads are in the runnable state and not consuming CPU time.

Utilization Type	Default color	Description
Idle		All threads in the program are waiting - no threads are running. There can be only one node in the Summary chart indicating idle utilization.
Poor		Poor utilization. By default, poor utilization is when the number of threads is up to 50% of the target concurrency.
OK		Acceptable (OK) utilization. By default, OK utilization is when the number of threads is between 51-85% of the target concurrency.
Ideal		Ideal utilization. By default, ideal utilization is when the number of threads is between 86-115% of the target concurrency.
Over		Over-utilization. By default, over-utilization is when the number of threads is more than 115% of the target concurrency.

viewpoint: A preset result tab configuration that filters out the data collected during a performance analysis and enables you to focus on specific performance problems. When you select a viewpoint, you select a set of performance metrics the VTune Amplifier XE shows in the windows/panes of the result tab. To select the

required viewpoint, click the  button and use the drop-down menu at the top of the result tab.

Wait time: The amount of time that a given thread waited for some event to occur, such as: synchronization waits and I/O waits.

