



parallel tools platform

<http://eclipse.org/ptp>

A New and Improved Eclipse Parallel Tools Platform: Advancing the Development of Scientific Applications

Greg Watson, IBM
g.watson@computer.org

Jay Alameda, NCSA
jalameda@ncsa.uiuc.edu

Wyatt Spear, U. Oregon
wspear@cs.uoregon.edu

Beth Tibbitts, IBM
tibbitts@us.ibm.com

Jeff Overbey, UIUC
jeffreyoverbey@acm.org

Alan Humphrey, SCI
ahumphrey@sci.utah.edu

November 13, 2011

Portions of this material are supported by or based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under its Agreement No. HR0011-07-9-0002, the United States Department of Energy under Contract No. DE-FG02-06ER25752 and the SI2-SSI Productive and Accessible Development Workbench for HPC Applications, which is supported by the National Science Foundation under award number OCI 1047956





Tutorial Outline

Time (Tentative!)	Module	Topics	Presenter
8:30-9:00	1. Eclipse and PTP Installation	✦ Installation of Eclipse and PTP (can start early as people arrive)	Greg/Beth
9:00-9:30	2. Introduction & Overview	✦ Eclipse architecture & organization overview	Greg
9:30-10:00	3. Developing with Eclipse	✦ Eclipse basics; Creating a new project from CVS; Local, remote, and synchronized projects	Beth
10:00-10:30	BREAK		
10:30-12:00 Continue at 1:00-1:45	3. Developing with Eclipse (continued)	Continue from before the break... ✦ Beth: Editing C files; MPI Features; Building w/ Makefile ✦ Jay: Resource Managers and launching a parallel app ✦ Jeff: Fortran, Refactoring, other Advanced Features	Beth, Jay, Jeff
12:00 – 1:00	Lunch		
1:45-2:30	4. Debugging (45 min)	✦ Debugging an MPI program	Greg
2:30-3:00	BREAK		
3:00-4:30	5. Performance Tuning & Analysis Tools	✦ TAU, ETFw including hands-on exercise ✦ GEM, including hands-on exercise	Wyatt Spear Alan Humphrey
4:30-5:00	6. Other Tools, Wrapup	✦ NCSA HPC Workbench, Other Tools, website, mailing lists, future features	Jay/Beth

Final Slides, Installation Instructions

- ★ Please go to <http://wiki.eclipse.org/PTP/tutorials/SC11> for slides and installation instructions

Module 1: Installation

★ Objective

- ★ To learn how to install Eclipse and PTP

★ Contents

- ★ System Prerequisites
- ★ Eclipse Download and Installation of “Eclipse IDE for Parallel Application Developers” – parallel package
- ★ Installation Confirmation
- ★ Updating the PTP within your Eclipse to the latest release

About the Tutorial Installation

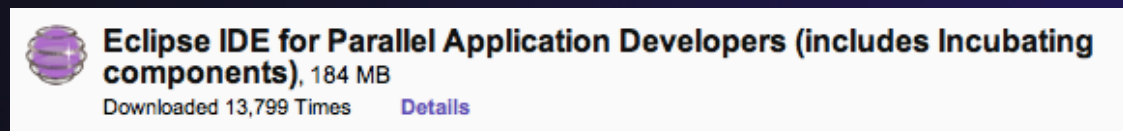
- ★ This tutorial assumes you have Eclipse and PTP pre-installed on your laptop
- ★ If you already have Eclipse installed, go directly to “Starting Eclipse”, slide 5
- ★ If you don’t have Eclipse installed, you will need to follow the handouts so that you can catch up with the rest of the class
- ★ Note: up-to-date info on installing PTP and its pre-reqs is available from the release notes:
 - ★ http://wiki.eclipse.org/PTP/release_notes/5.0
 - ★ This information may supersede these slides

System Prerequisites

- ★ Local system (running Eclipse)
 - ★ Linux (just about any version)
 - ★ MacOSX (10.5 Leopard or 10.6 Snow Leopard)
 - ★ Windows (XP on)
- ★ Java: Eclipse requires Sun or IBM Java
 - ★ Only need Java runtime environment (JRE)
 - ★ Java 1.5 or higher
 - ★ Java 1.5 is the same as JRE 5.0
 - ★ The GNU Java Compiler (GCJ), which comes standard on Linux, will not work!
 - ★ OpenJDK, distributed with some Linux distributions, has not been tested by us but should work.
 - ★ See <http://wiki.eclipse.org/PTP/installjava>

Eclipse Packages

- ★ The current version of Eclipse (3.7) is also known as “Indigo”
- ★ Eclipse is available in a number of different packages for different kinds of development
 - ★ <http://eclipse.org/downloads>
- ★ With Indigo, there is a new package directly relevant for HPC:
 - ★ Eclipse IDE for Parallel Application Developers
 - ★ **This is recommended for all new installs**



“Parallel Package”

- ★ Can also add PTP to an existing Eclipse installation



Eclipse Installation

- ★ Download the “Eclipse IDE for Parallel Application Developers” package
 - ★ <http://download.eclipse.org>
- ★ Make sure you match the architecture with that of your laptop
- ★ If your machine is Linux or Mac OS X, untar the file
 - ★ On Mac OS X you can just double-click in the Finder
- ★ If your machine is Windows, unzip the file
- ★ This creates an **eclipse** folder containing the executable as well as other support files and folders



Starting Eclipse

✦ **Linux**

- ✦ From a terminal window, enter
“<eclipse_installation_path>/eclipse/eclipse &”

✦ **Mac OS X**

- ✦ From finder, open the **eclipse** folder where you installed
- ✦ Double-click on the **Eclipse** application
- ✦ Or from a terminal window

✦ **Windows**

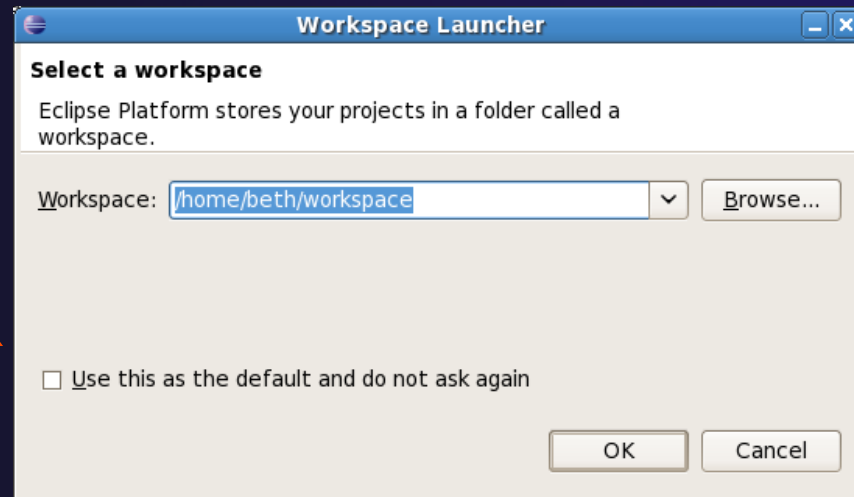
- ✦ Open the **eclipse** folder
- ✦ Double-click on the **eclipse** executable



Specifying A Workspace

- ✦ Eclipse prompts for a workspace location at startup time
- ✦ The workspace contains all user-defined data
 - ✦ Projects and resources such as folders and files
 - ✦ The default workspace location is fine for this tutorial

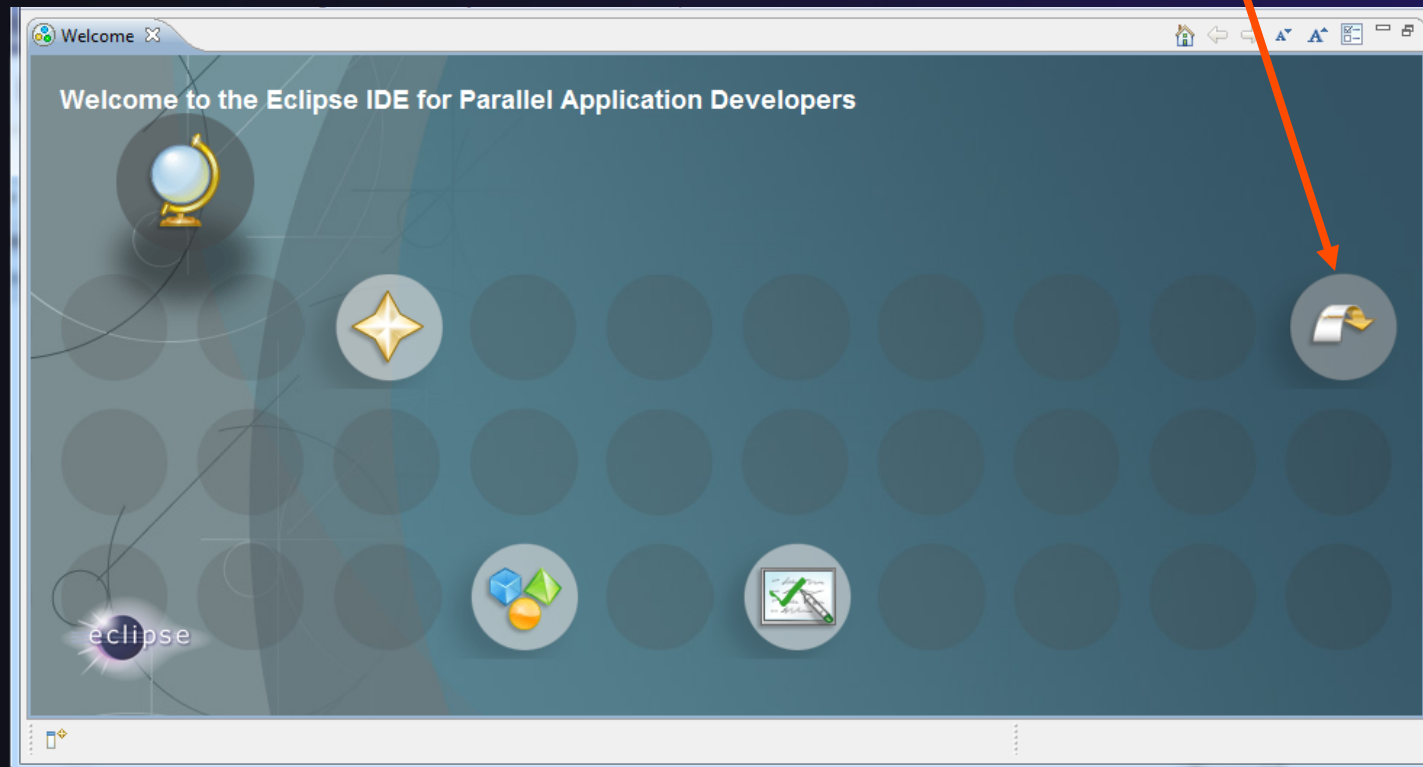
The prompt can be turned off



Eclipse Welcome Page



- ★ Displayed when Eclipse is run for the first time
Select "Go to the workbench"

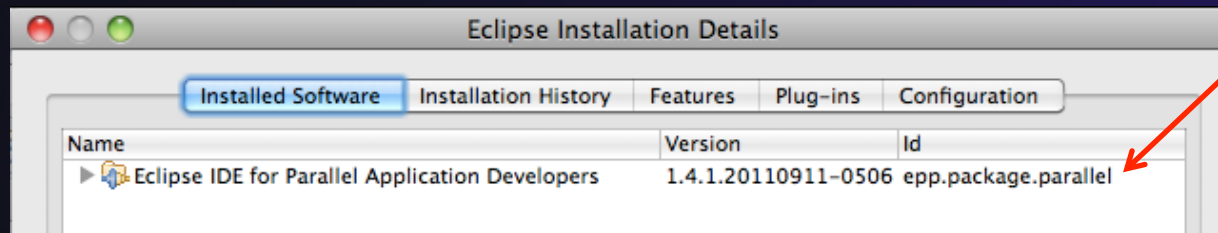




Check Installation Details

- ✦ To confirm you have installed OK
 - ✦ Mac: **Eclipse>About Eclipse**
 - ✦ Others: **Help>About**
- ✦ Choose **Installation Details**
- ✦ Confirm you have the following installed software

Differs depending on base download



- ✦ Close the dialog: **Close, OK**

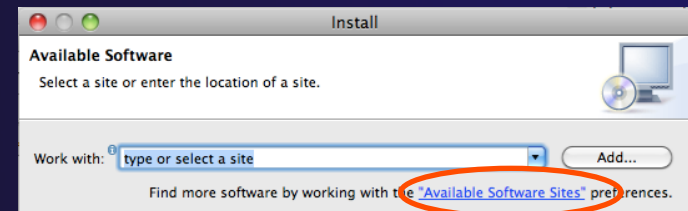
Checking for PTP Updates

- ✦ From time-to-time there may be newer PTP releases than the Indigo release
 - ✦ Indigo and "Parallel package" updates are released only in Sept and February
- ✦ PTP maintains its own update site with the most recent release
 - ✦ Bug fix releases can be more frequent than Indigo's and what is within the parallel package
- ✦ **You must enable the PTP-specific update site before the updates will be found**



Updating PTP

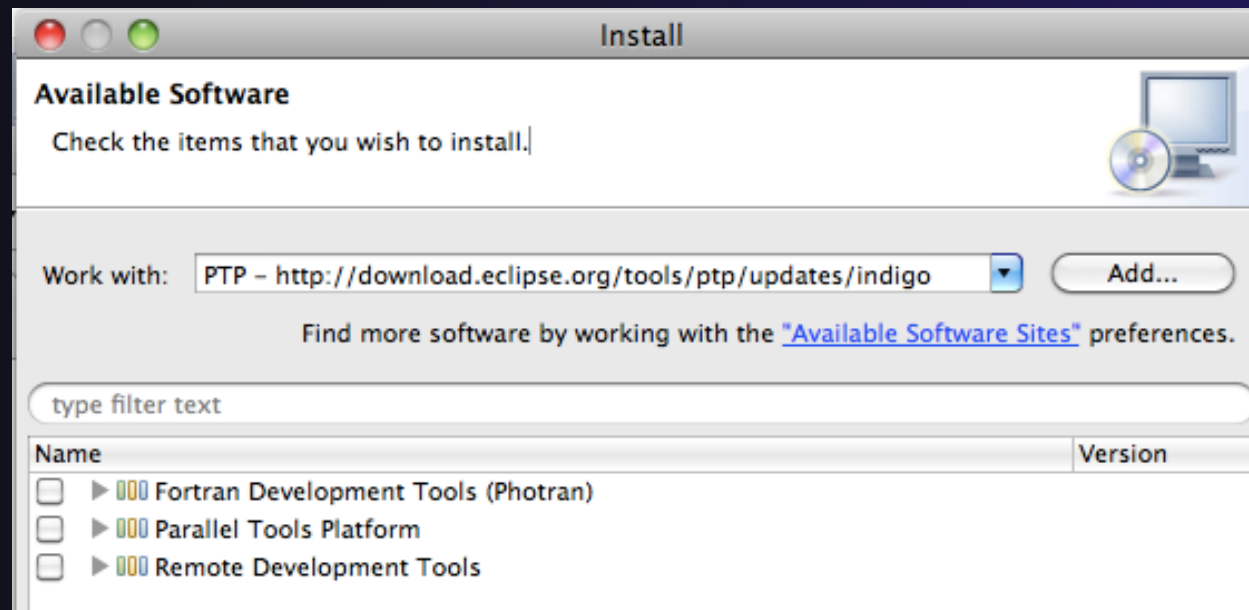
- ★ Enable PTP-specific update site
 - ★ **Help>Install New Software...**
 - ★ Click **Available Software Sites** link
 - ★ Ensure this checkbox is selected for the PTP site:
<http://download.eclipse.org/tools/ptp/updates/indigo>
 - ★ Choose **OK**
 - ★ Choose **Cancel** (to return to Eclipse workbench)
- ★ Now select **Help>Check for updates**
 - ★ If you see "No updates were found"...
 - ★ It's only because there are no updates in the "Eclipse IDE for Parallel Application Developers"
 - ★ We will update the PTP within it





Updating PTP (2)

- ★ We will get the PTP release that is more recent than what is currently (Nov. 2011) within the parallel package
- ★ Now select **Help>Install New Software...**
 - ★ In the **Work With:** dropdown box, select the PTP update site you confirmed already:

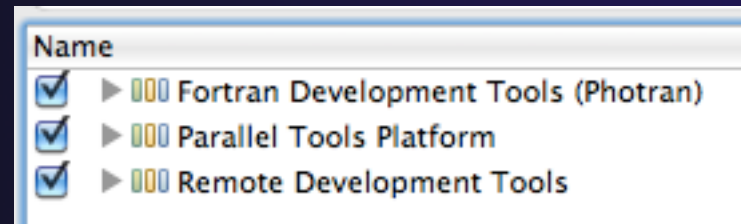




Updating PTP (3)

★ Quick and dirty:

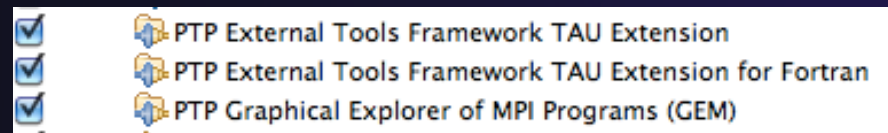
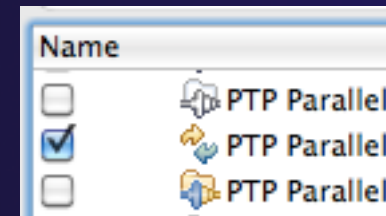
- ★ Check everything - which updates existing features and adds a few more



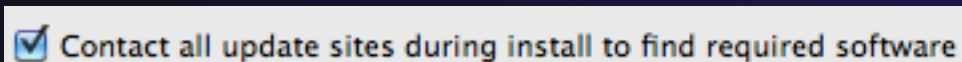
★ Detailed:

- ★ Open each feature and check the ones you want to update
- ★ Icons indicate: Grey plug: already installed and up to date
Double arrow: can be updated
Color plug: Not installed yet

Note: For this tutorial, install GEM and TAU



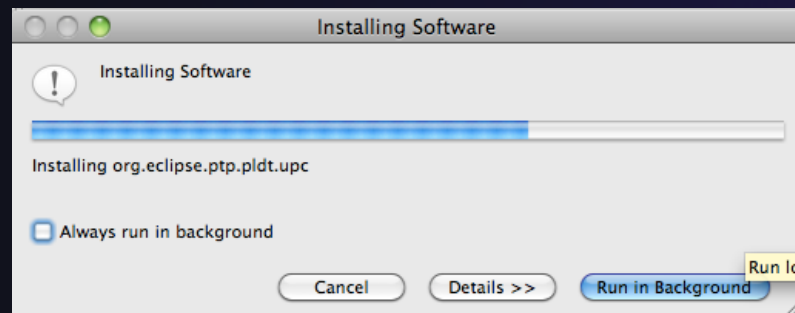
Note: if conference network is slow, consider unchecking:





Updating PTP (4)

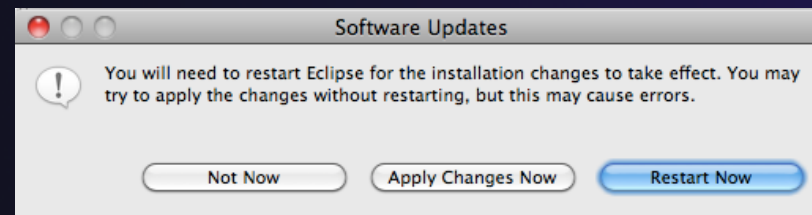
- ★ Select **Next** to continue updating PTP
- ★ Select **Next** to confirm features to install
 - ★ Note: if you are NOT on Linux: You may have to go back and un-select, under Fortran Development Tools (Photran), the “Linux Intel Fortran Compiler Support” – if it complains
- ★ Accept the License agreement and select **Finish**



Wait for installation to finish

If conference network is too slow, we have this cached on USB/CD/DVD

- ★ Select **Restart Now** when prompted

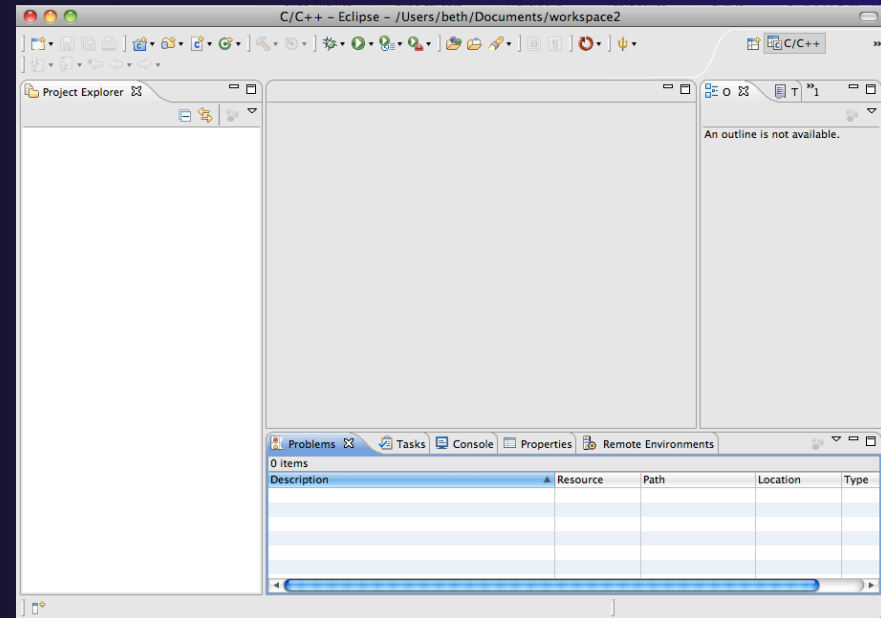




Restart after Install

- ★ If any top-level features are installed... Welcome page informs you of new features installed
- ★ We only updated PTP, so we land back at C/C++ Perspective

... Ready to go!



- ★ **Help>About** or **Eclipse > About Eclipse ...** will indicate the release of PTP installed
- ★ Further **Help>Check for Updates** will find future updates on the PTP Update site

Module 2: Introduction

- ✦ Objective
 - ✦ To introduce the Eclipse platform and PTP
- ✦ Contents
 - ✦ New and Improved Features
 - ✦ What is Eclipse?
 - ✦ What is PTP?

New and Improved Features

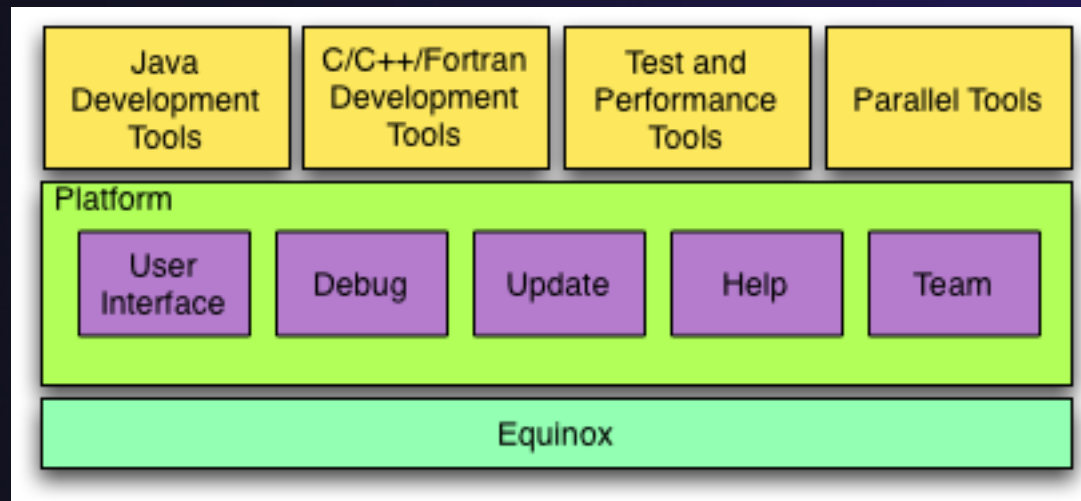
- ★ More flexible projects
 - ★ Synchronized projects overcome many problems of remote projects
 - ★ Allows development when "off-line"
 - ★ Works with non-C/C++ projects
- ★ More customizable resource managers
 - ★ Resource managers can now be added by users
 - ★ Able to have site-specific configurations
 - ★ Interactive launch using job schedulers now supported

New and Improved Features (2)

- ★ Scalable system/job monitoring
 - ★ New perspective allows monitoring of systems of virtually any size
 - ★ View shows location of jobs on cluster
 - ★ Active and inactive jobs views
- ★ Remote support for performance tools
 - ★ External Tools Framework has been extended to support remote systems
 - ★ Performance tools such as TAU can now launch and collect data from remote systems

What is Eclipse?

- ✦ A vendor-neutral open-source workbench for multi-language development
- ✦ A extensible platform for tool integration
- ✦ Plug-in based framework to create, integrate and utilize software tools

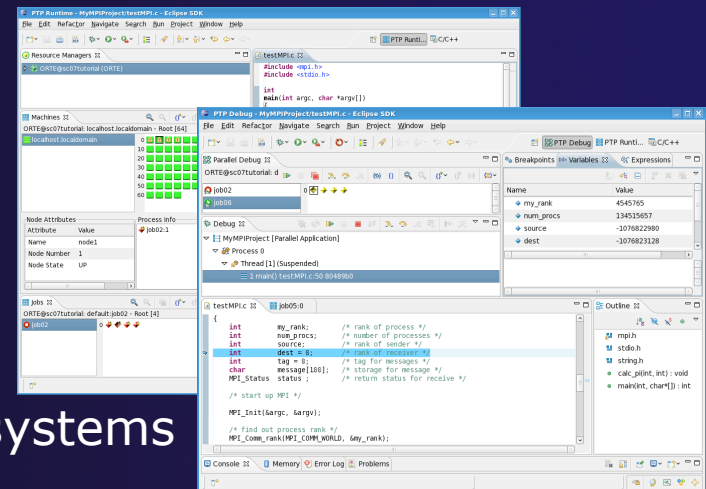


Eclipse Features

- ✦ Full development lifecycle support
- ✦ Revision control integration (CVS, SVN, Git)
- ✦ Project dependency management
- ✦ Incremental building
- ✦ Content assistance
- ✦ Context sensitive help
- ✦ Language sensitive searching
- ✦ Multi-language support
- ✦ Debugging

Parallel Tools Platform (PTP)

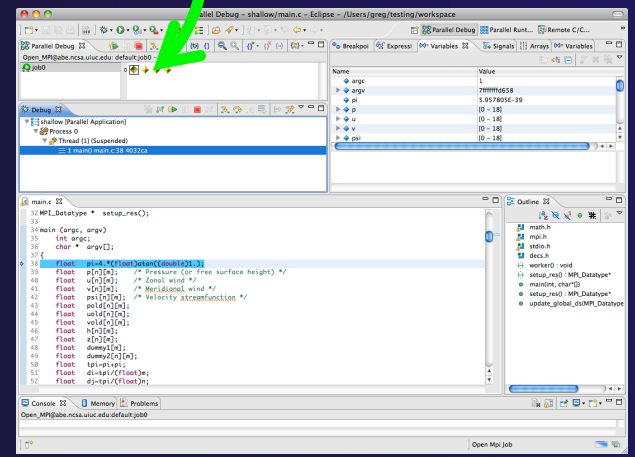
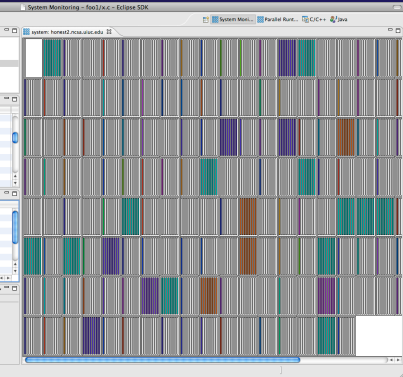
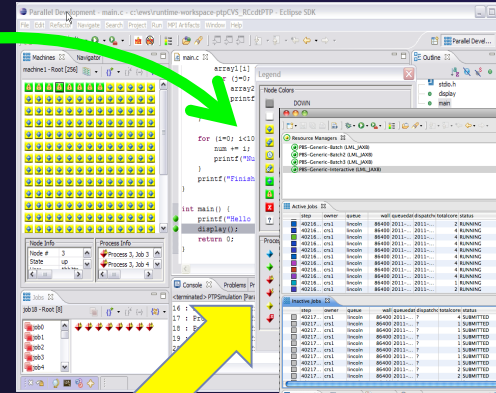
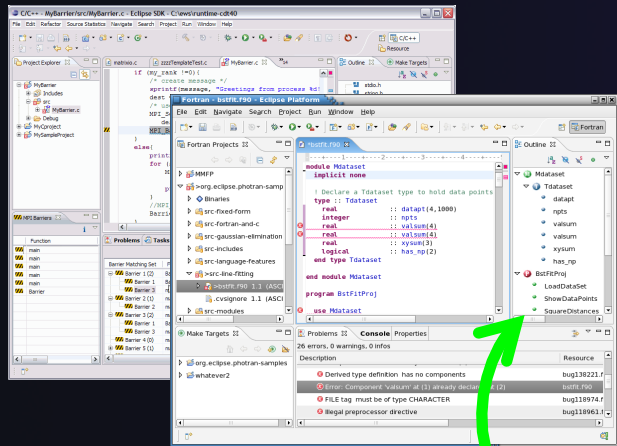
- ★ The Parallel Tools Platform aims to provide a highly integrated environment specifically designed for parallel application development
- ★ Features include:
 - ★ An integrated development environment (IDE) that supports a wide range of parallel architectures and runtime systems
 - ★ A scalable parallel debugger
 - ★ Parallel programming tools (MPI, OpenMP, UPC, etc.)
 - ★ Support for the integration of parallel tools
 - ★ An environment that simplifies the end-user interaction with parallel systems
- ★ <http://www.eclipse.org/ptp>



Eclipse PTP Family of Tools

Coding & Analysis
(C, C++, Fortran)

Launching & Monitoring



Performance Tuning
(TAU, PPW, ...)

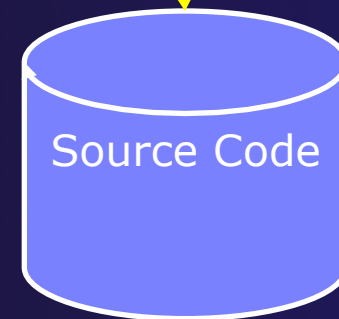
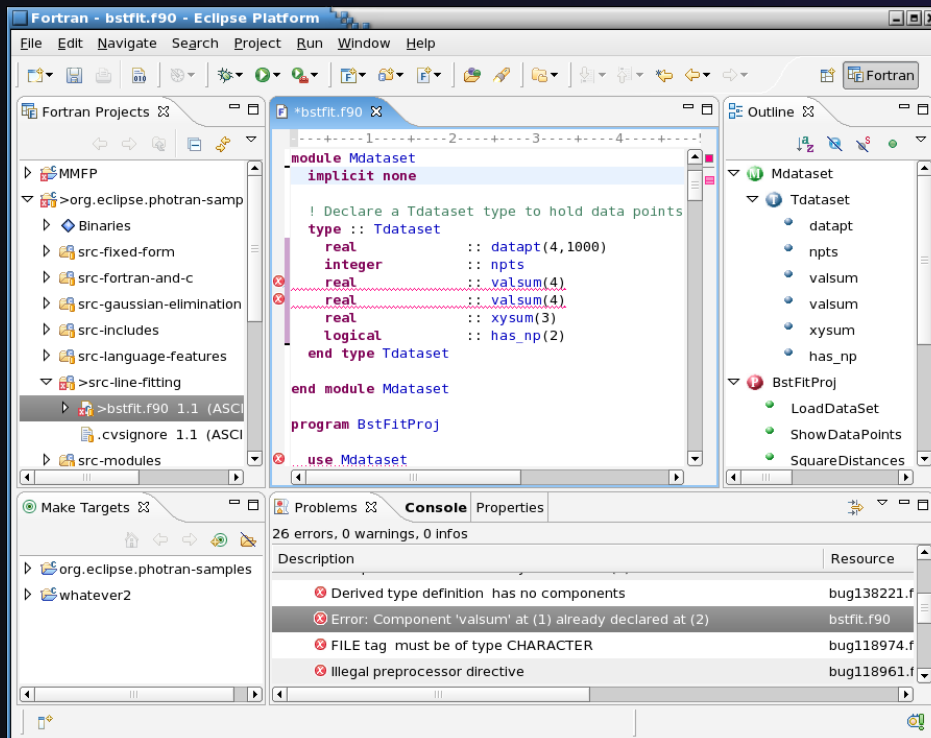
Parallel Debugging

Module 2

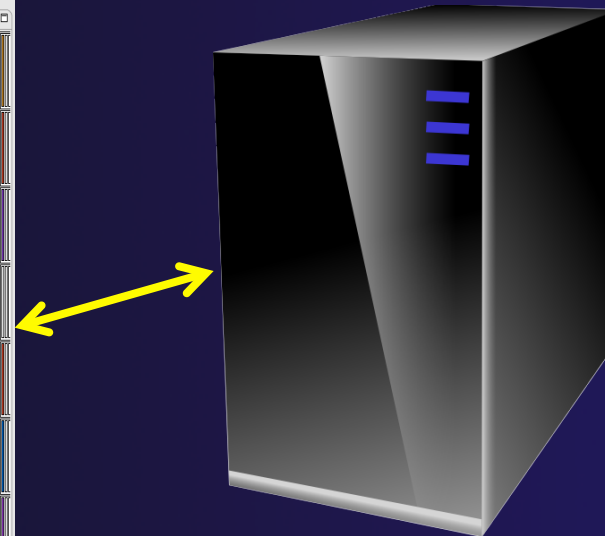
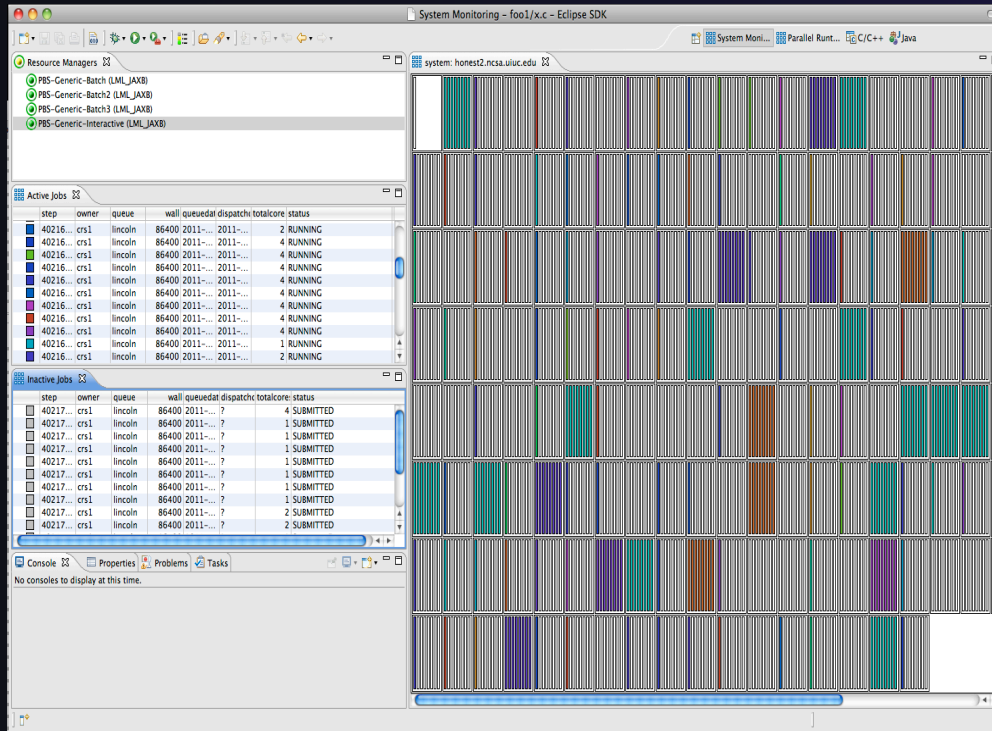
2-6

How Eclipse is Used

Editing/Compiling

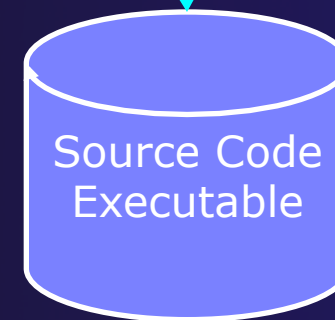
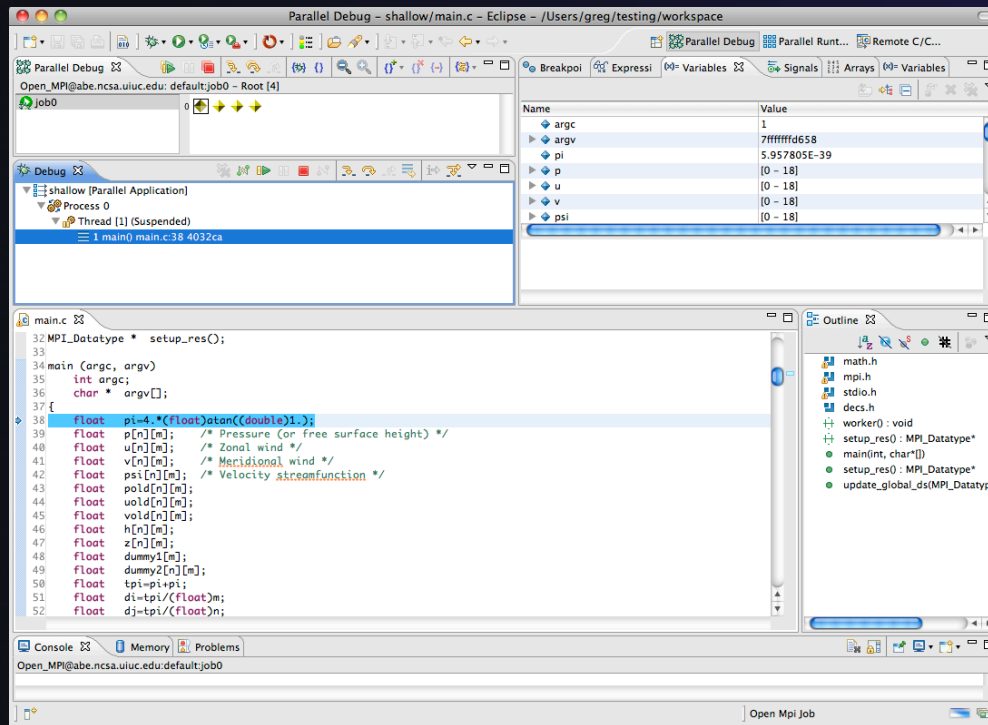


How Eclipse is Used Launching/Monitoring



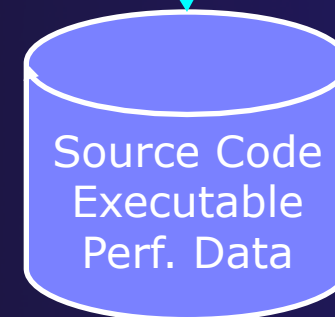
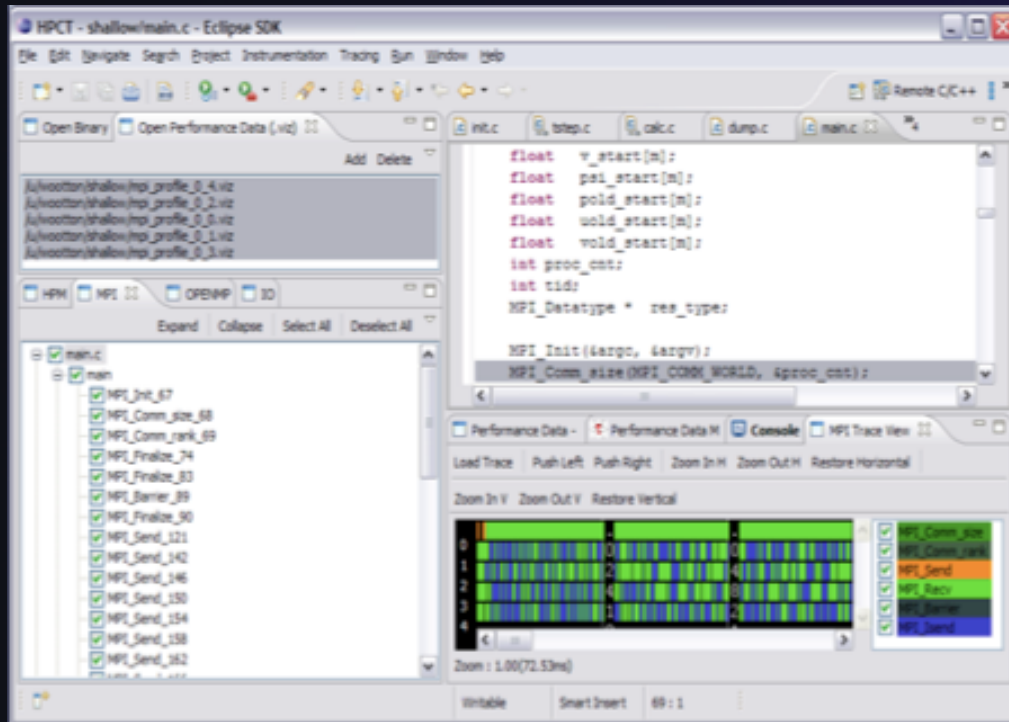
Source Code
Executable

How Eclipse is Used Debugging



How Eclipse is Used

Performance Tuning



Module 3: Developing with Eclipse

★ Objective

- ★ Learn basic Eclipse concepts: Perspectives, Views, ...
- ★ Learn about local, remote, and synchronized projects
- ★ Learn how to create and manage a C project
- ★ Learn about Eclipse editing features
- ★ Learn about Eclipse Team features
- ★ Learn about MPI features
- ★ Learn how to build and launch an MPI program on a remote system
- ★ Learn about Fortran projects
- ★ Learn about searching, refactoring, etc.

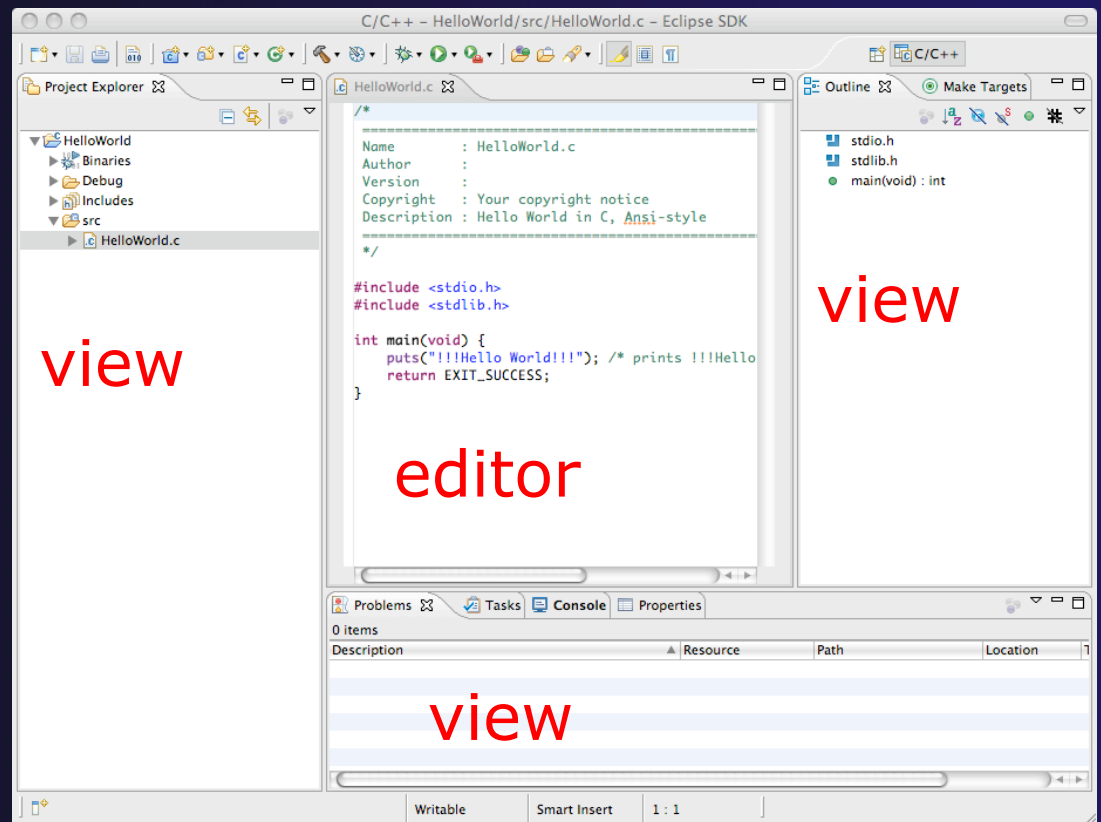
Contents

- ★ Basic Eclipse Features (3-2)
- ★ Projects In Eclipse (3-13)
- ★ Editor Features (3-24)
- ★ Team Features (3-34)
- ★ MPI Features (3-40)
- ★ Synchronizing the Project (3-56)
- ★ Building the Project (3-62)
- ★ Running: Resource Manager Configuration (3-69)
- ★ Running: Launching a Job (3-82)
- ★ Advanced Features: Searching (3-90)
- ★ Fortran Specifics (3-99)
- ★ Advanced editing: Code Templates (3-108)
- ★ Refactoring and Transformation (3-113)

Basic Eclipse Features

Eclipse Basics

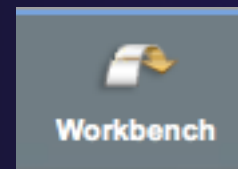
- ★ A *workbench* contains the menus, toolbars, editors and views that make up the main Eclipse window
- ★ The workbench represents the desktop development environment
 - ★ Contains a set of tools for resource mgmt
 - ★ Provides a common way of navigating through the resources
- ★ Multiple workbenches can be opened at the same time
- ★ Only one workbench can be open on a *workspace* at a time



perspective

Perspectives

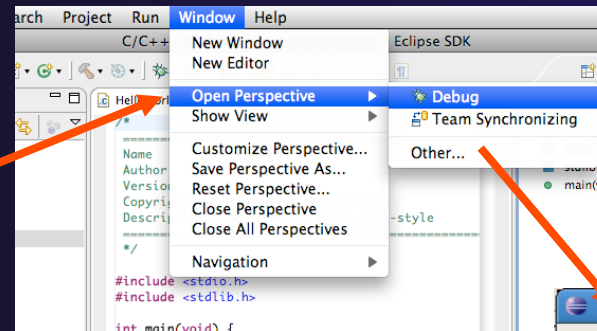
- ✦ Perspectives define the layout of views and editors in the workbench
- ✦ They are *task oriented*, i.e. they contain specific views for doing certain tasks:
 - ✦ There is a **Resource Perspective** for manipulating resources
 - ✦ **C/C++ Perspective** for manipulating compiled code
 - ✦ **Debug Perspective** for debugging applications
- ✦ You can easily switch between perspectives
- ✦ If you are on the Welcome screen now, select “Go to Workbench” now



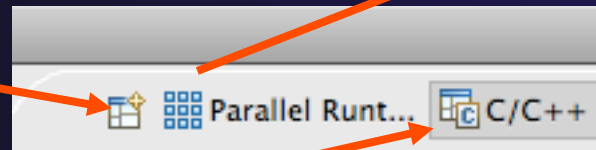
Switching Perspectives

★ Three ways of changing perspectives

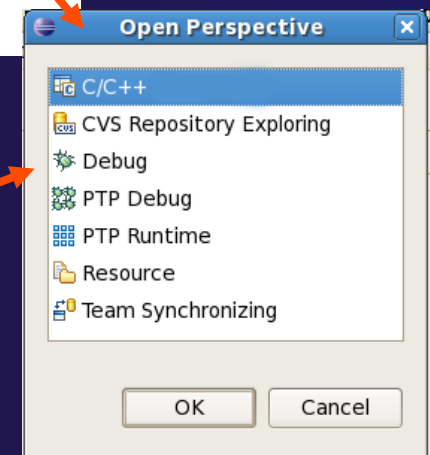
1. Choose the **Window>Open Perspective** menu option
Then choose **Other...**



2. Click on the **Open Perspective** button in the upper right corner of screen (hover over it to see names)



3. Click on a perspective shortcut button

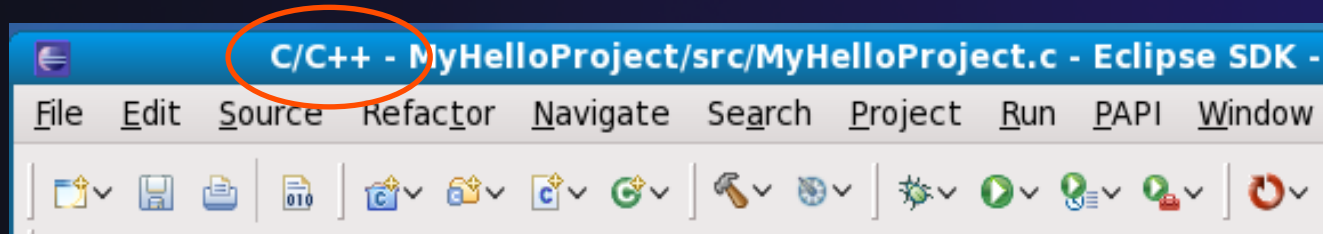


★ Switch to the C/C++ Perspective

Which Perspective?

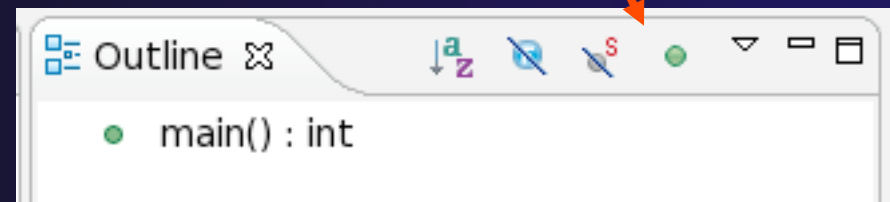
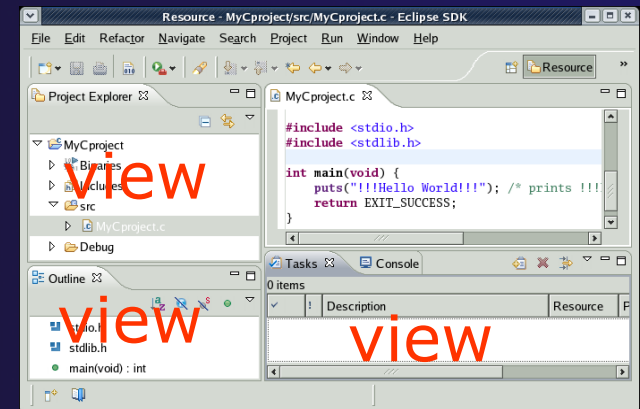


- ★ Which Perspective am in in?
See Title Bar



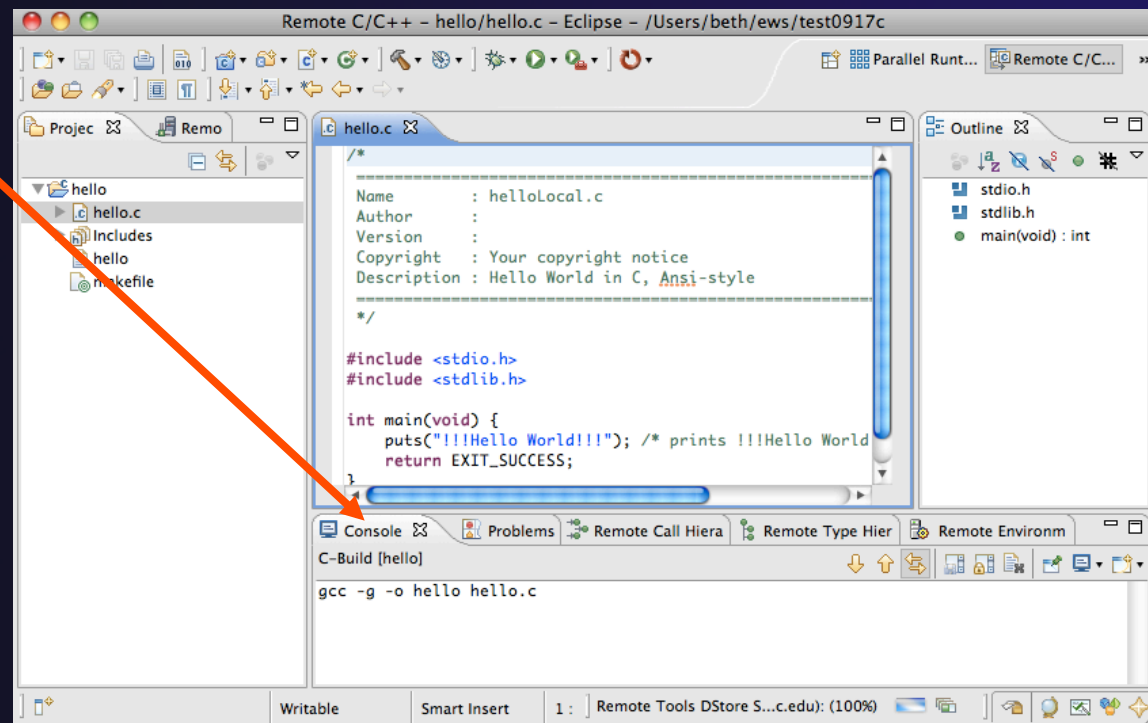
Views

- ★ The workbench window is divided up into Views
- ★ The main purpose of a view is:
 - ★ To provide alternative ways of presenting information
 - ★ For navigation
 - ★ For editing and modifying information
- ★ Views can have their own menus and toolbars
 - ★ Items available in menus and toolbars are available only in that view
 - ★ Menu actions only apply to the view
- ★ Views can be resized



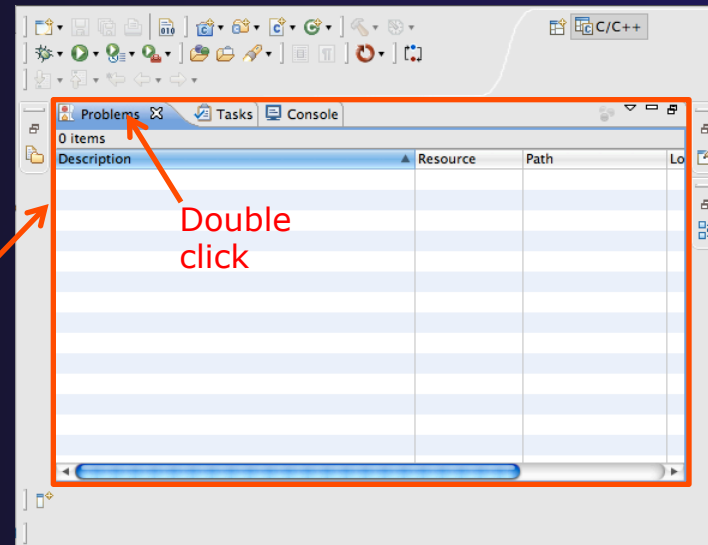
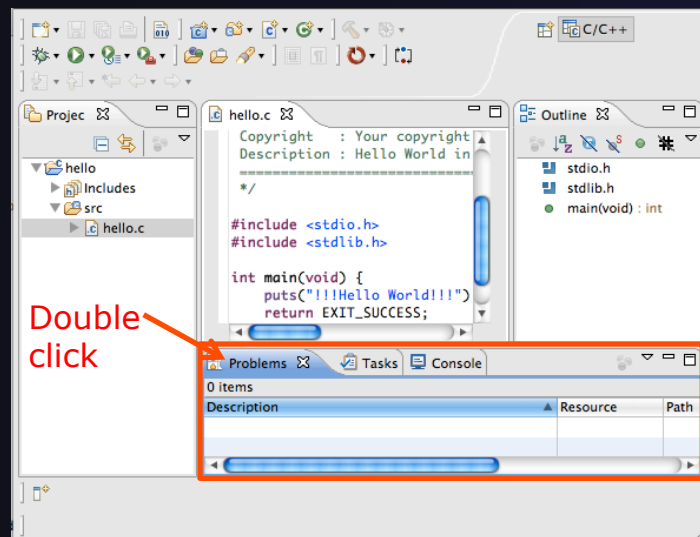
Stacked Views

- ★ Stacked views appear as tabs
- ★ Selecting a tab brings that view to the foreground



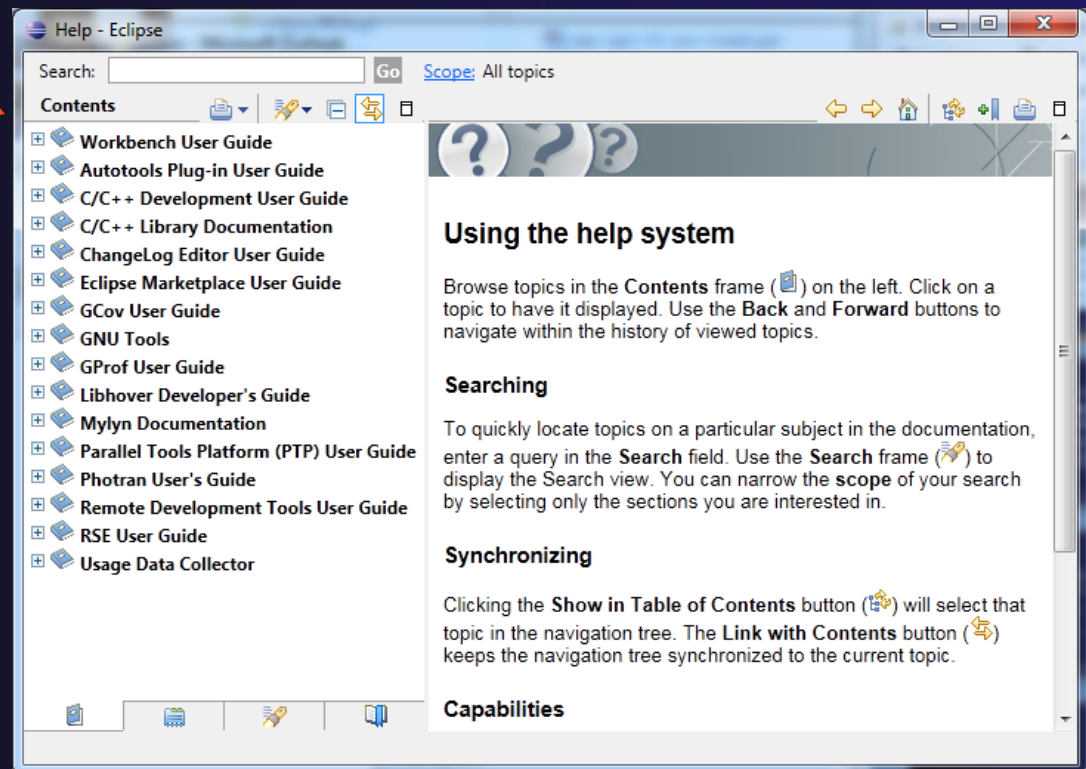
Expand a View

- ★ Double-click on a view/editor's tab to fill the workbench with its content;
- ★ Repeat to return to original size

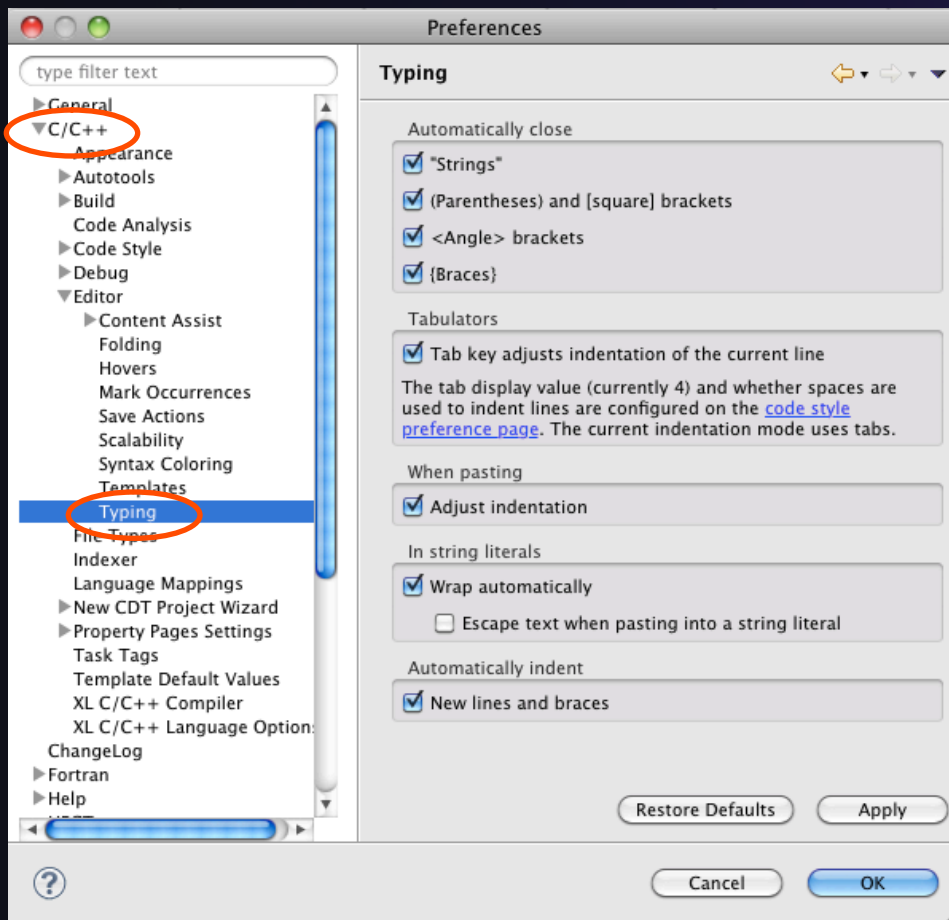


Help

- ★ To access help
 - ★ **Help>Help Contents**
 - ★ **Help>Search**
 - ★ **Help>Dynamic Help**
- ★ **Help Contents** provides detailed help on different Eclipse features *in a browser*
- ★ **Search** allows you to search for help locally, or using Google or the Eclipse web site
- ★ **Dynamic Help** shows help related to the current context (perspective, view, etc.)

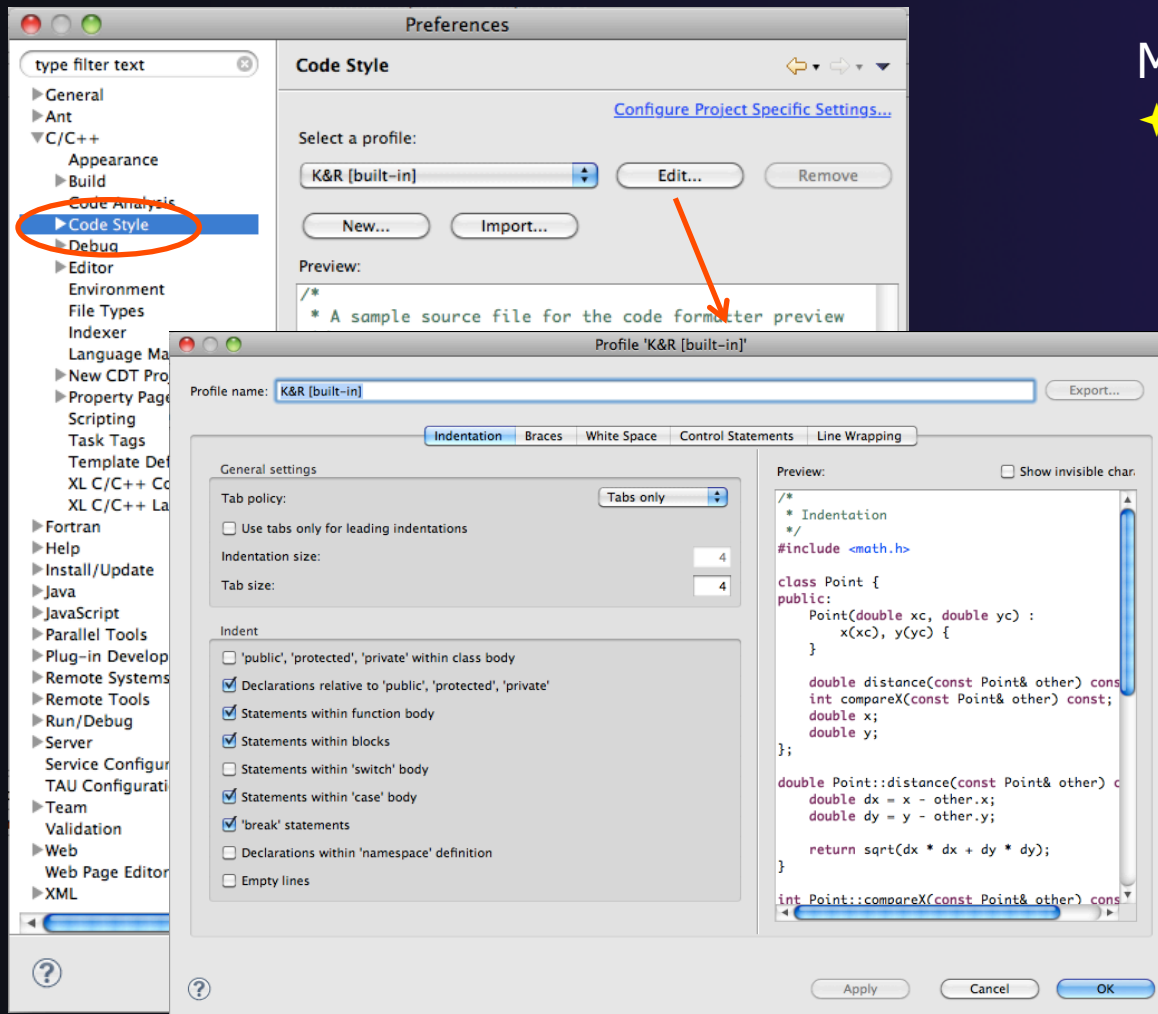


Eclipse Preferences



- ✦ Eclipse Preferences allow customization of almost everything
- ✦ To open use
 - ✦ Mac: **Eclipse>Preferences...**
 - ✦ Others: **Window>Preferences...**
- ✦ The C/C++ preferences allow many options to be altered
- ✦ In this example you can adjust what happens in the editor as you type.

Preferences Example



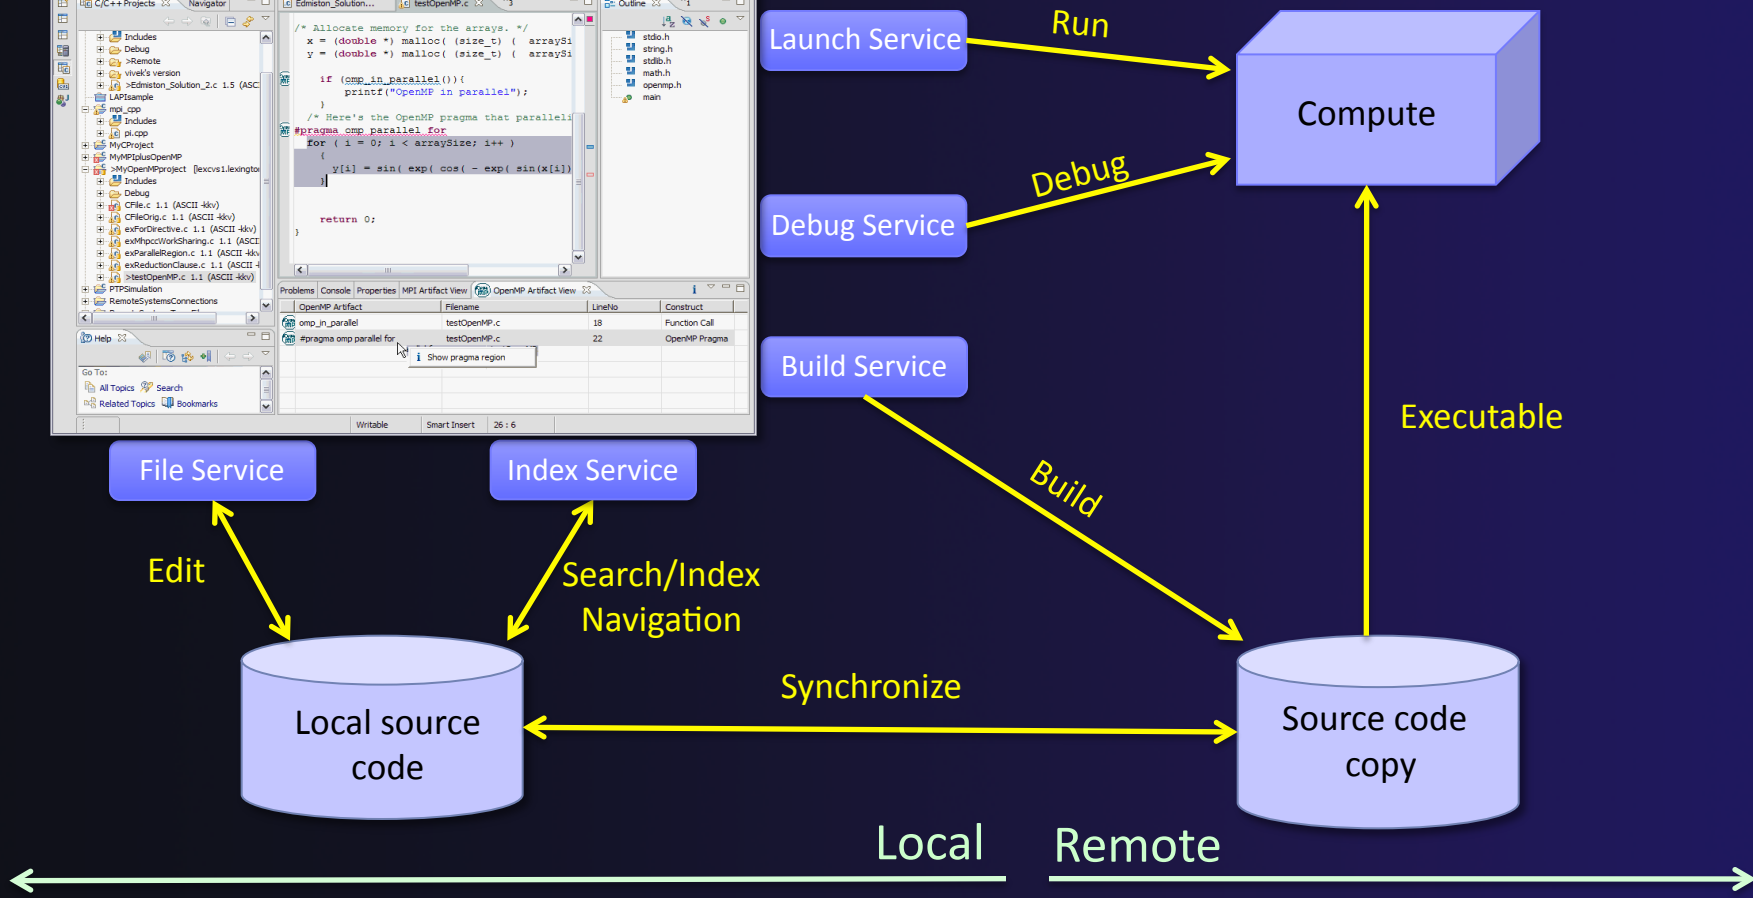
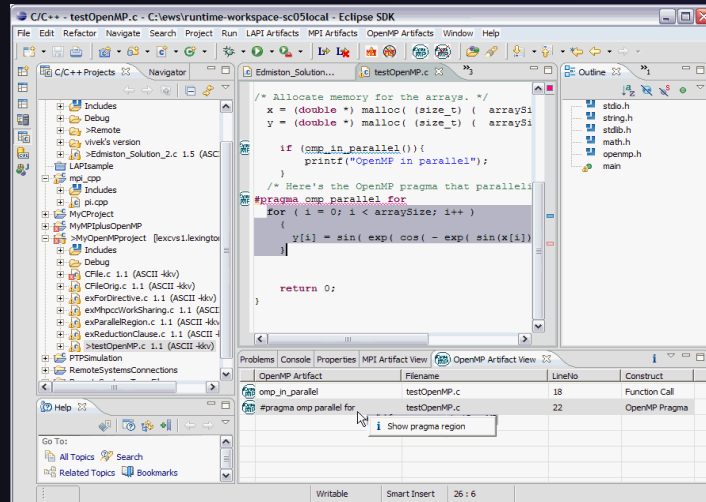
- More C/C++ preferences:
- ✦ In this example the Code Style preferences are shown
 - ✦ These allow code to be automatically formatted in different ways

Projects In Eclipse

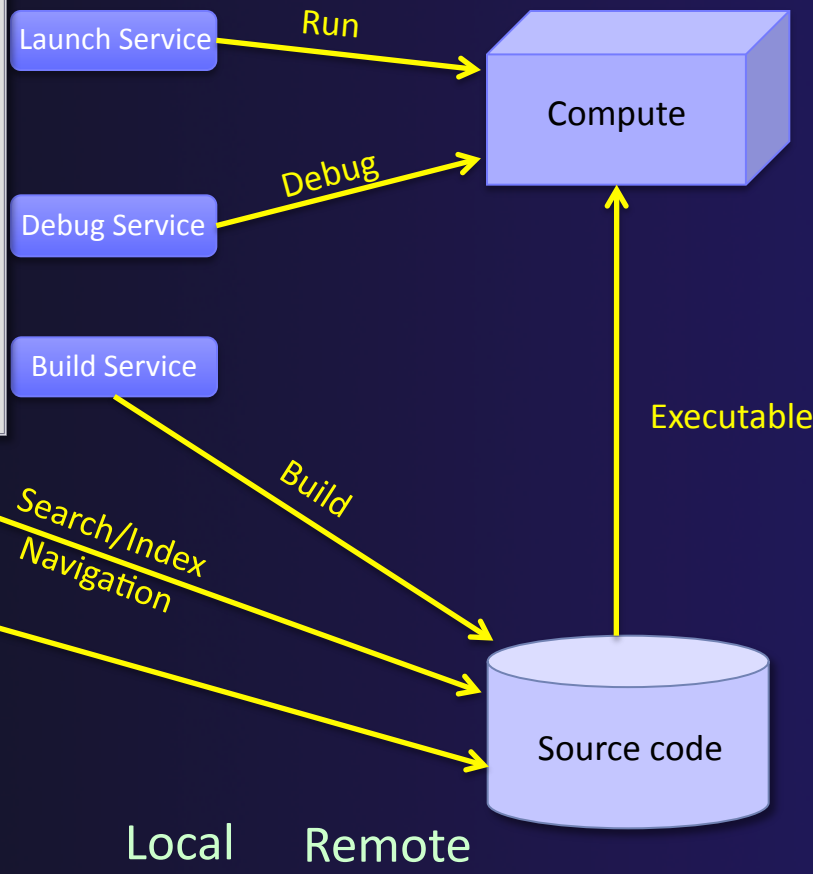
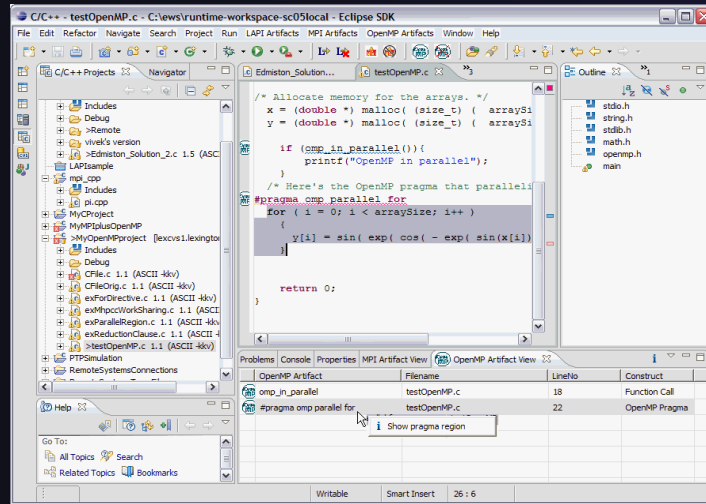
Project Types

- ✦ Local
 - ✦ Source is located on local machine, builds happen locally
- ✦ Synchronized
 - ✦ Source is local, then synchronized with remote machine(s)
 - ✦ Building and launching happens remotely (can also happen locally)
- ✦ Remote
 - ✦ Source is located on remote machine(s), build and launch takes place on remote machine(s)

Synchronized Projects



Remote Projects



C, C++, and Fortran Projects

Build types

- ★ Makefile-based
 - ★ Project contains its own makefile (or makefiles) for building the application
- ★ Managed
 - ★ Eclipse manages the build process, no makefile required

Parallel programs can be run on local machine or on a remote system

- ★ MPI (or other runtime) needs to be installed
- ★ An application built locally probably can't be run on a remote machine unless their architectures are the same

Checking out the project

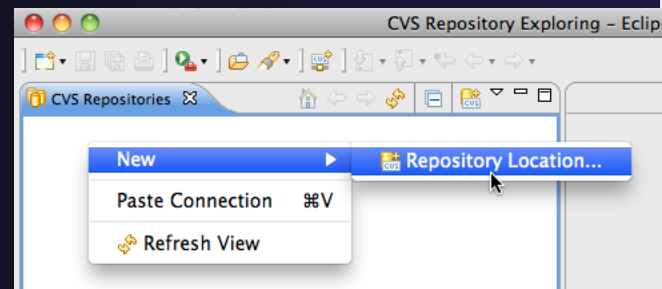
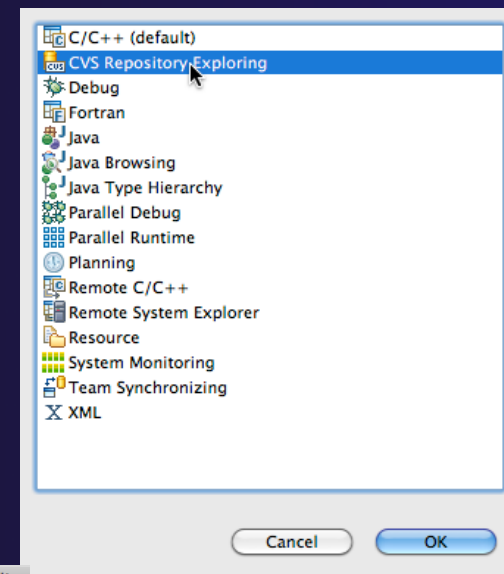
Using a Source Code Repository Introduction to Team Features



Importing a Project from CVS

- ★ Switch to **CVS Repository Exploring** perspective
 - ★ Window > Open Perspective > Other...
 - ★ Select **CVS Repository Exploring**
 - ★ Select **OK**

- ★ Right click in **CVS Repositories** view and select **New>Repository Location...**





Add CVS Repository

- ★ Enter **Host:** dev.eclipse.org
- ★ **Repository path:**
/cvsroot/tools



- ★ For anonymous access:
 - ★ **User:** anonymous
 - ★ No password is required
 - ★ **Connection type:** pserver (default)
- ★ For authorized access:
 - ★ **User:** your userid
 - ★ **Password:** your password
 - ★ **Connection type:** change to **extssh**

- ★ Select **Finish**

Add CVS Repository

Add a new CVS Repository to the CVS Repositories view

Location

Host: dev.eclipse.org
Repository path: /cvsroot/tools

Authentication

User: anonymous
Password:

Connection

Connection type: pserver
 Use default port
 Use port:

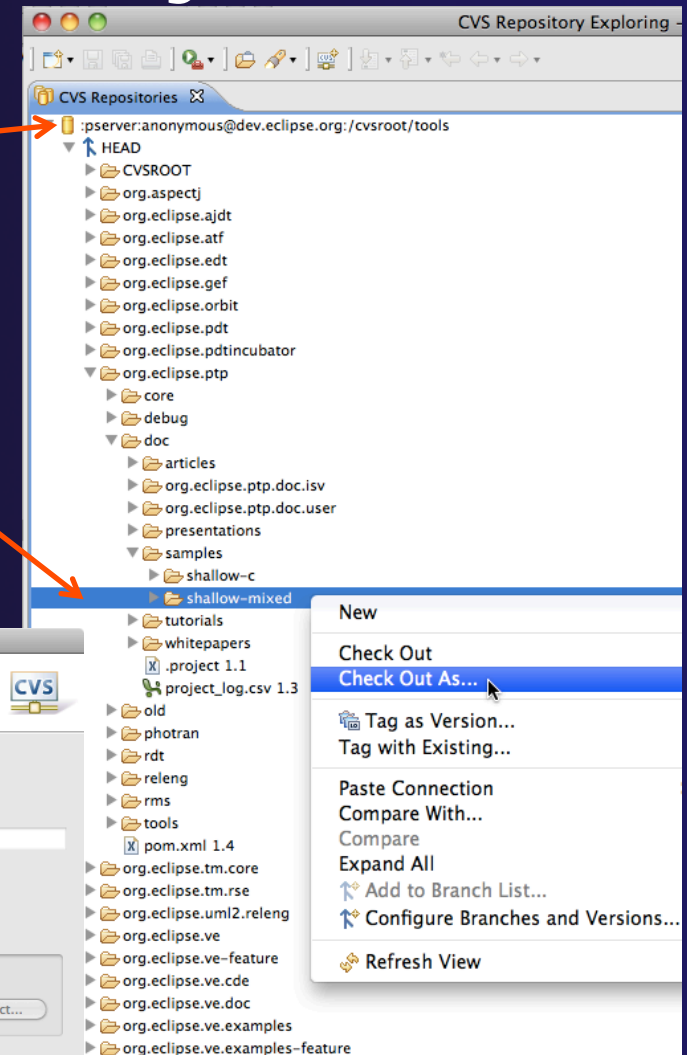
Validate connection on finish
 Save password (could trigger secure storage login)
To manage your password, please see '[Secure Storage](#)'
[Configure connection preferences...](#)

Cancel Finish

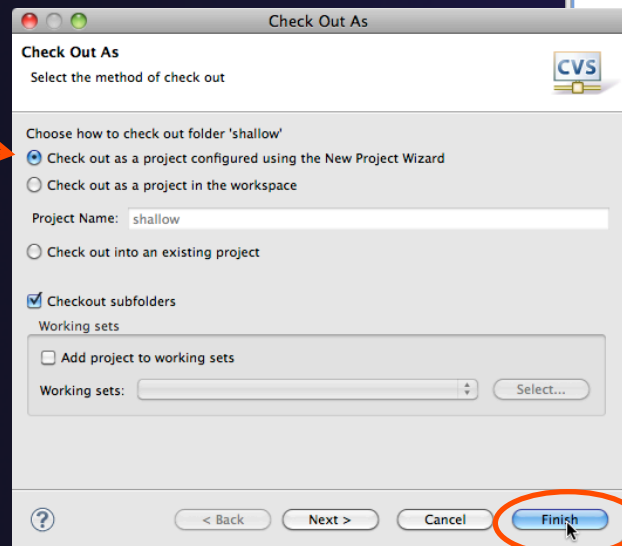
Checking out the Project



- ✦ Expand the repository location
- ✦ Expand **HEAD**
- ✦ Expand **org.eclipse.ptp, doc, and samples**
- ✦ Right click on **shallow-mixed** and select **Check Out As...**
- ✦ On **Check Out As** dialog, select **Finish**



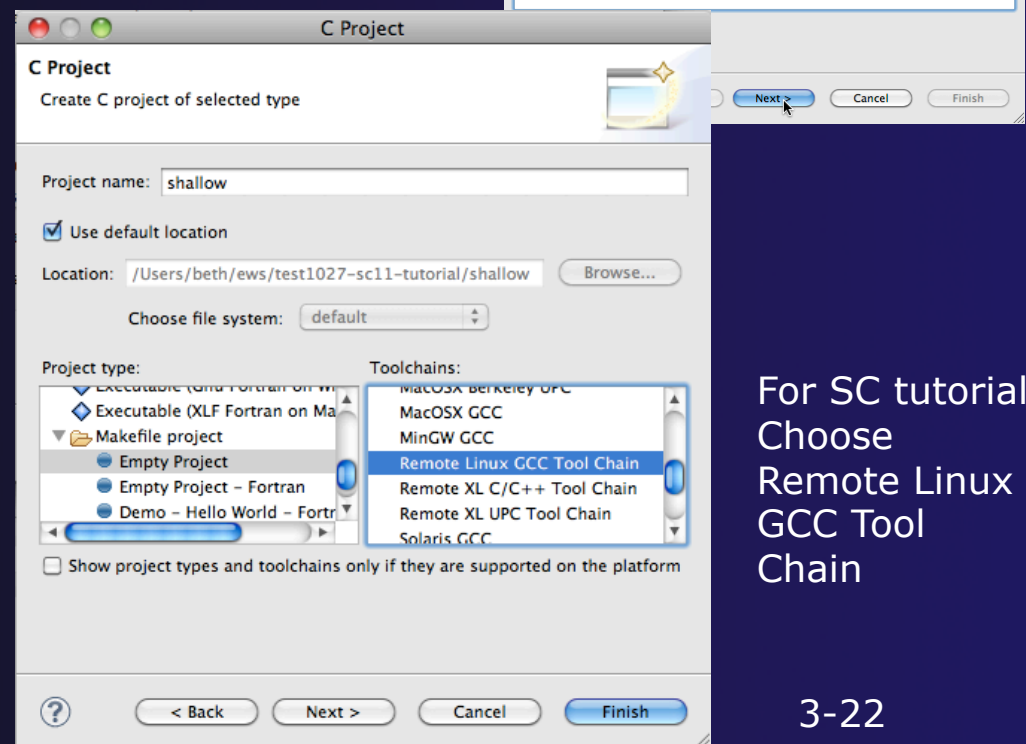
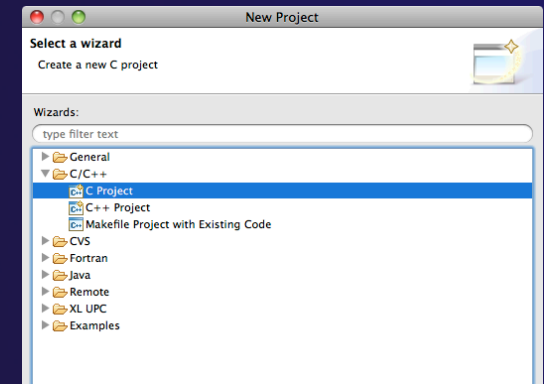
The default of "Check out as a project configured using the New Project Wizard" is what we want



New Project Wizard

As project is checked out from CVS, the **New Project Wizard** helps you configure the Eclipse information to be added to the project

- ✦ Expand **C/C++**
- ✦ Select **C Project** and click on **Next>**
- ✦ Enter 'shallow' as **Project Name**
- ✦ Under **Project type**, expand **Makefile project**
- scroll to the bottom
- ✦ Select **Empty Project**
- ✦ Select a toolchain that matches your system from **Toolchains**
 - ✦ Since we will build/run this on the remote system, choose an appropriate toolchain
 - ✦ You may need to uncheck "Show project types and toolchains only if they are supported on the platform"
- ✦ Click on **Finish**

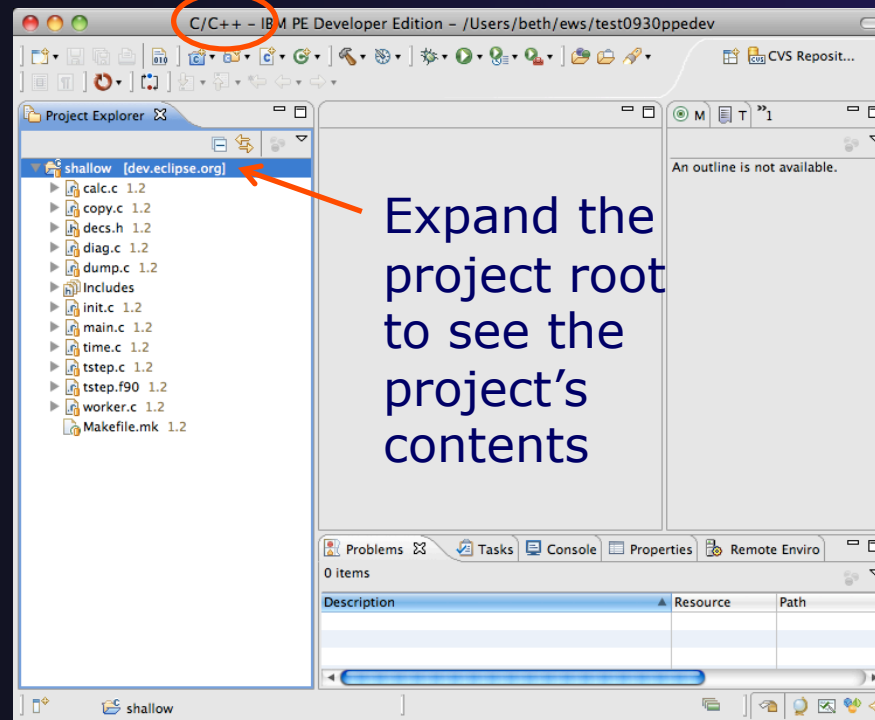
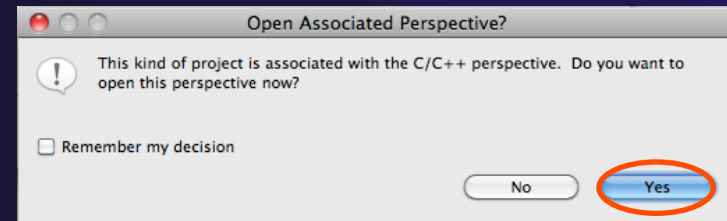


For SC tutorial
Choose
Remote Linux
GCC Tool
Chain

C/C++ Perspective



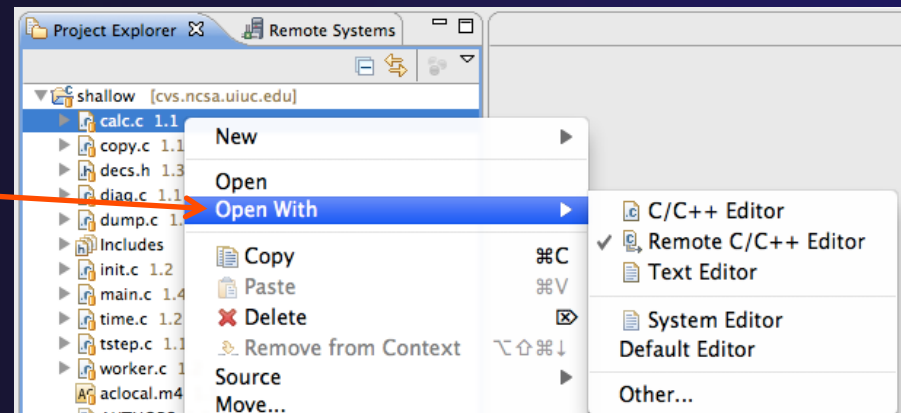
- ★ Switch to the C/C++ Perspective when Prompted
- ★ You should now see the “shallow” project in your workspace



Editor Features

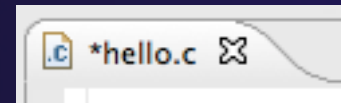
Editors

- ★ An editor for a resource (e.g. a file) opens when you double-click on a resource
- ★ The type of editor depends on the type of the resource
 - ★ .c files are opened with the C/C++ editor by default
 - ★ You can use **Open With** to use another editor
 - ★ In this case the default editor is fine (double-click)
 - ★ Some editors do not just edit raw text
- ★ When an editor opens on a resource, it stays open across different perspectives
- ★ An active editor contains menus and toolbars specific to that editor



Saving File in Editor

- ★ When you change a file in the editor, an asterisk on the editor's title bar indicates unsaved changes

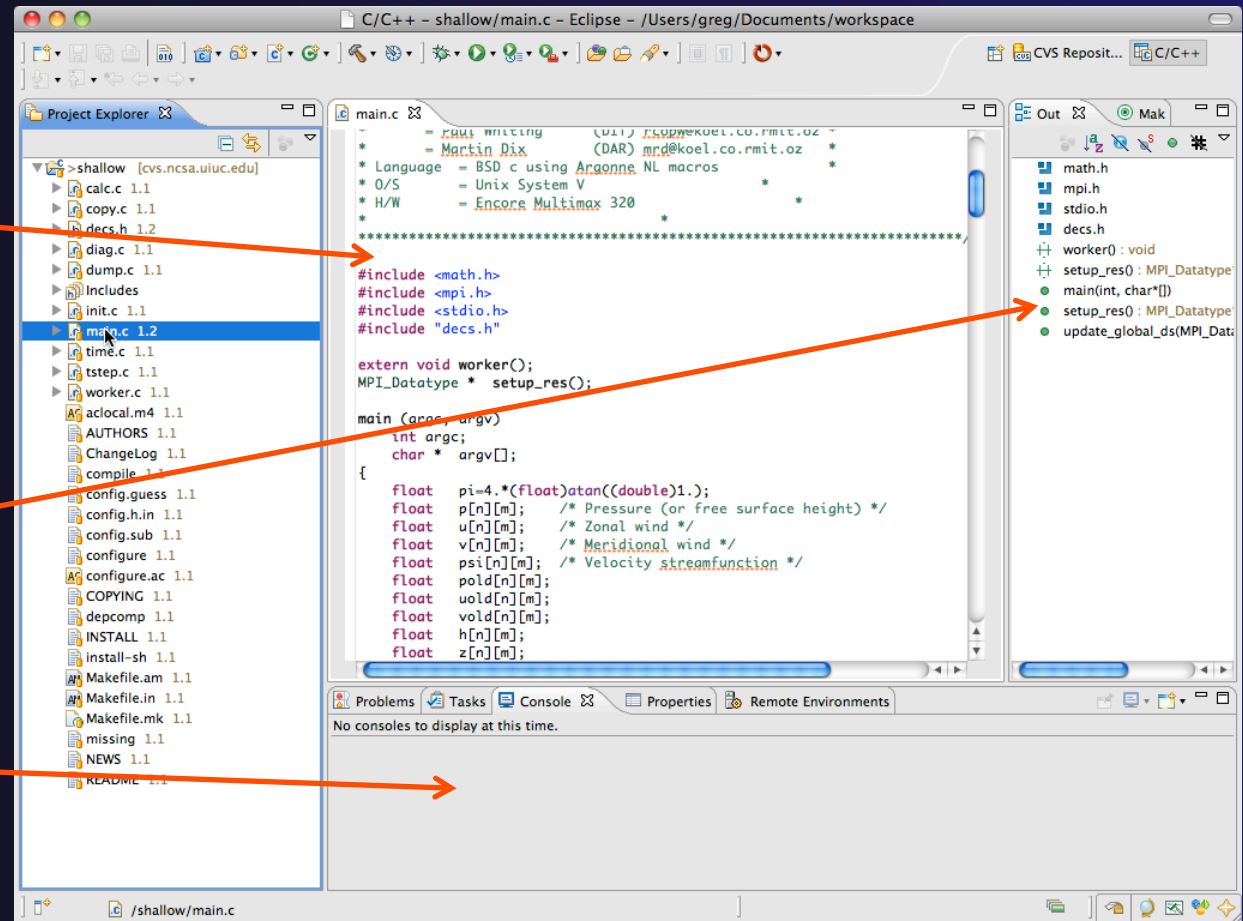


- ★ Save the changes by using Command/Ctrl-S or **File>Save**

Editor and Outline View



- ★ Double-click on source file
- ★ Editor will open in main view
- ★ Outline view is shown for file in editor
- ★ Console shows results of build, local runs, etc.

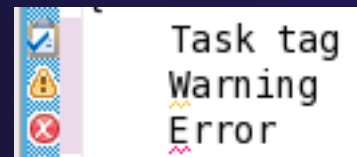


Source Code Editors & Markers

- ★ A source code editor is a special type of editor for manipulating source code
- ★ Language features are highlighted
- ★ Marker bars for showing
 - ★ Breakpoints
 - ★ Errors/warnings
 - ★ Task Tags, Bookmarks
- ★ Location bar for navigating to interesting features in the entire file

```
linear_function.c
/**
 * Returns f(x) = 3.0*x + 2.0
 */
double evaluate(double x)
{
    // TODO add semicolon to end of next line
    double y = 3.0*x + 2.0
    return y;
}
```

Icons:



Code Analysis (Codan)

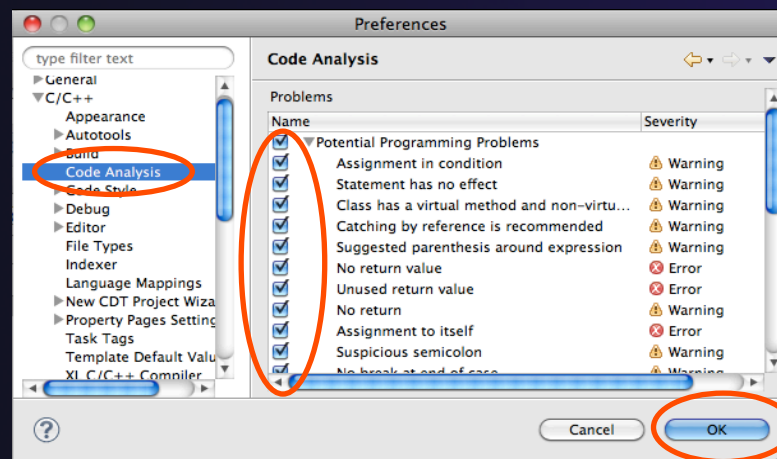
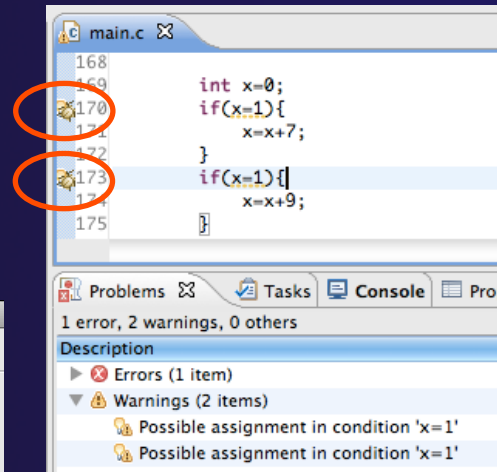
- ★ If you see bug icons in the editor marker bar, they are likely suggestions from Codan
- ★ Code checkers can flag possible errors, even if code is technically correct
- ★ To turn them off, use Preferences

Window > Preferences or Mac: Eclipse > Preferences

C/C++ > Code Analysis

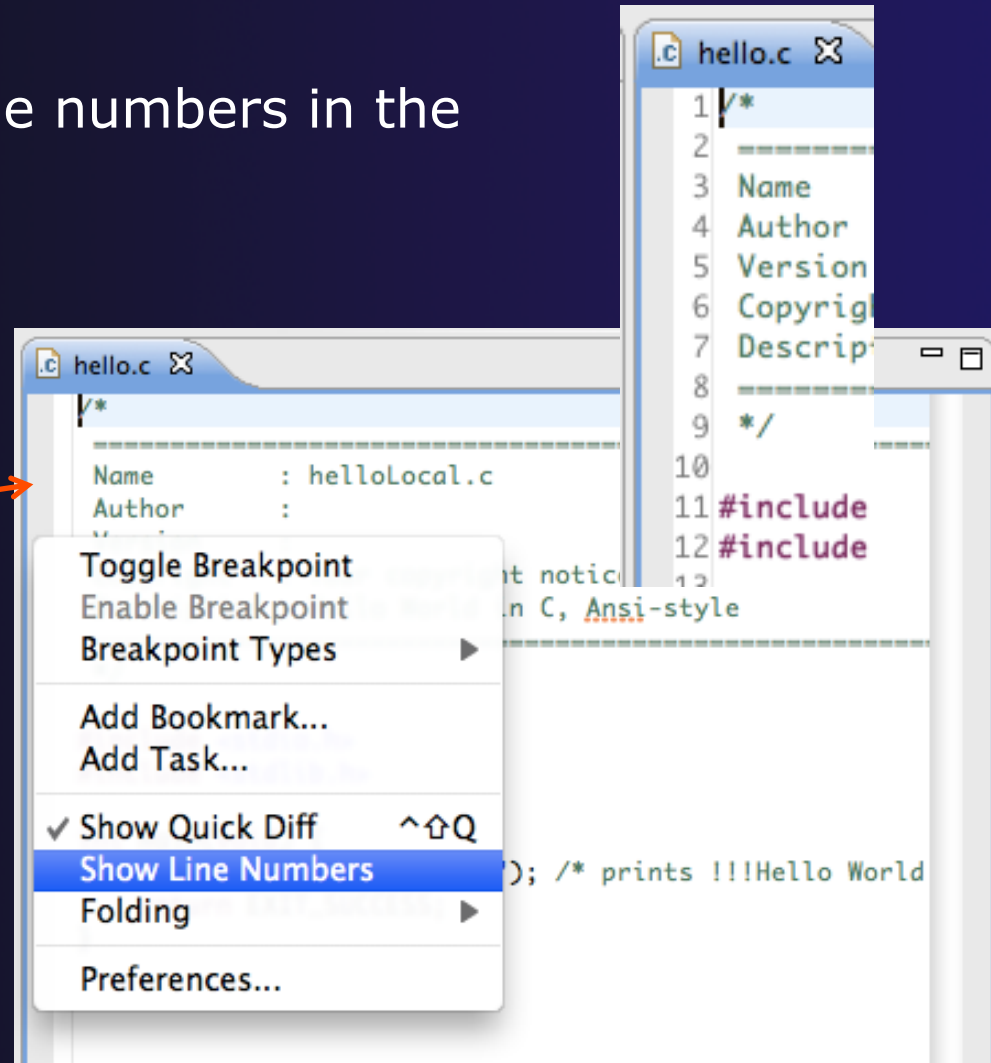
and uncheck
all problems

- ★ Select OK to close Preferences



Line Numbers

- ★ Text editors can show line numbers in the left column
- ★ To turn on line numbering:
 - ★ Right-mouse click in the editor marker bar
 - ★ Click on **Show Line Numbers**





Navigating to Other Files

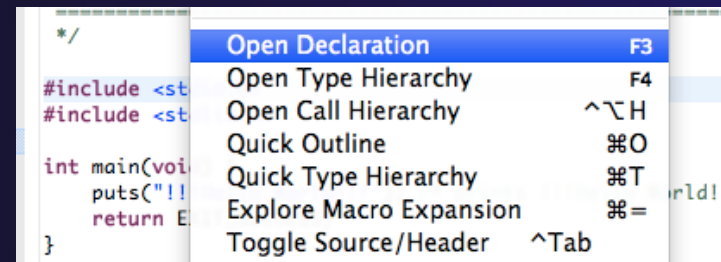
- ★ On demand hyperlink
 - ★ In main.c line 135:
 - ★ Hold down Command/Ctrl key e.g. on call to `initialise`
 - ★ Click on `initialise` to navigate to its definition in the header file (Exact key combination depends on your OS)
 - ★ E.g. Command/Ctrl and click on `initialise`
- ★ Open declaration
 - ★ Right-click and select **Open Declaration** will also open the file in which the element is declared
 - ★ E.g. in main.c line 29 right-click on `decs.h` and select **Open Declaration**

```

128 }
129
130
131 /*
132 initialise data structures and construct packets to be sent to workers
133 */
134
135 initialise(p, u, v, psi, pold, uold, vold, di, dj, z);
136 diag(1, 0, p, u, v, h, z);
137
138 for (i = 1; i < proc_cnt; i++) {
139     for (j = 0; j < n; j++) {
  
```

```

26 #include <math.h>
27 #include "decs.h"
28
29 void initialise(p, u, v, psi, pold, uold, vold, di, dj, z)
30 float p[n][m];
31 float u[n][m];
32 float v[n][m];
33 float psi[n][m];
  
```



Note: may need to left-click before right-click works



Content Assist & Templates

- ✦ Type an incomplete function name e.g. "get" into the editor, and hit **ctrl-space**
- ✦ Select desired completion value with cursor or mouse

```
13  
14 int main(void) {  
15     puts("!!!Hello World!!!"); /* prints !!!Hello World!!! */  
16     get  
17       
18     ret  
19 }  
20
```

The screenshot shows a code editor with a completion list for the word "get". The list includes:

- getchar_unlocked(void) : int
- getdelim(char ** __lineptr, * __n, int __delimit
- getenv(const char * __name) : char *
- getline(char ** __lineptr, * __n, FILE * __stream
- getloadavg(double * __loadavg, int __nelem)

An orange arrow points from the text "Select desired completion value with cursor or mouse" to the "getenv" option in the list. A blue highlight is visible under the "get" text in the code.

Press '^Space' to show Template Propos

- ✦ Code Templates: type 'for' and Ctrl-space

Hit ctrl-space again
for code templates

```
17     for  
18       
19     ret  
20 }  
21
```

The screenshot shows a code editor with a completion list for the word "for". The list includes:

- for - for loop
- for - for loop with temporary variable

The second option, "for - for loop with temporary variable", is selected and highlighted in yellow. The corresponding code template is shown to the right:

```
for (int var = 0; var < max; ++var) {  
      
}
```

More info on code templates later



Inactive code

- ★ Inactive code will appear grayed out in the CDT editor

```
260 #define VAL
261 #ifdef VAL
262     acopy_one_to_two(VAL, ds, res.indx);
263 #else
264     acopy_one_to_two(res.row, ds, res.indx);
265 #endif
```

```
260 //#define VAL
261 #ifdef VAL
262     acopy_one_to_two(VAL, ds, res.indx);
263 #else
264     acopy_one_to_two(res.row, ds, res.indx);
265 #endif
```

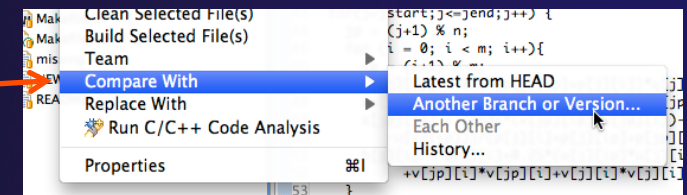
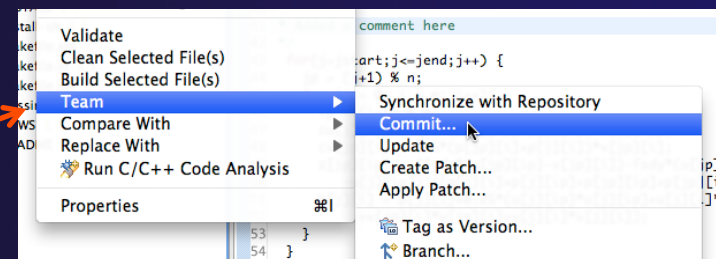
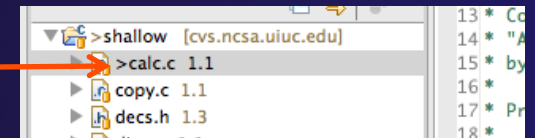
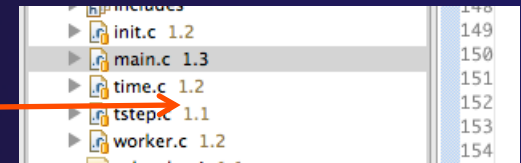
Team Features

“Team” Features

- ✦ Eclipse supports integration with multiple version control systems (VCS)
 - ✦ CVS, SVN, Git, and others
 - ✦ Collectively known as “Team” services
- ✦ Many features are common across VCS
 - ✦ Compare/merge
 - ✦ History
 - ✦ Check-in/check-out
- ✦ Some differences
 - ✦ Version numbers
 - ✦ Branching

CVS Features

- ★ Shows version numbers next to each resource
- ★ Marks resources that have changed
 - ★ Can also change color (preference option)
- ★ Context menu for Team operations
- ★ Compare to latest, another branch, or history
- ★ Synchronize whole project (or any selected resources)





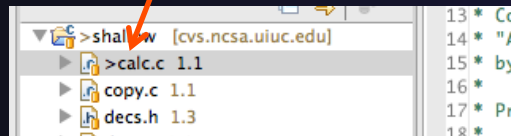
File Modification

- ✦ Open "calc.c"
- ✦ Add comment at line 40
- ✦ Save file
- ✦ File will be marked to indicate that it has been modified

```

27
28 void calcuvzh(jstart, jend, p, u, v, cu, cv, h, z, fsdx, fsdy)
29 int jstart, jend;
30 float p[n][m];
31 float u[n][m];
32 float v[n][m];
33 float cu[n][m];
34 float cv[n][m];
35 float h[n][m];
36 float z[n][m];
37 float fsdx, fsdy;
38 {
39     int i, j, ip, jp;
40     /*
41      * Added a comment here
42      */
43     for(j=jstart; j<=jend; j++) {
44         jp = (j+1) % n;
45         for (i = 0; i < m; i++){
46             ip = (i+1) % m;
47             cu[j][ip] = 0.5*(p[j][ip]+p[j][i])*u[j][ip];
48             cv[jp][i] = 0.5*(p[jp][i]+p[j][i])*v[jp][i];
49             z[jp][ip] = (fsdx*(v[jp][ip]-v[jp][i])-fsdy*(u[jp][ip]
50             -u[j][ip]))/(p[j][i]+p[j][ip]+p[jp][ip]+p[jp][i]);
51             h[j][i] = p[j][i]+0.25*(u[j][ip]*u[j][ip]+u[j][i]*u[j][i]
52             +v[jp][i]*v[jp][i]+v[j][i]*v[j][i]);
53         }
54     }

```





View Changes

- ★ Right-click on "calc.c" and select **Compare With>Latest from HEAD**
- ★ Compare editor will open showing differences between local (changed) file and the original
- ★ Buttons allow changes to be merged from right to left
- ★ Can also navigate between changes using buttons

```

C Compare
  Translation Unit
    calcuvzh
    calcuvzh
    cu
    cv
    fsdx
    fsdy

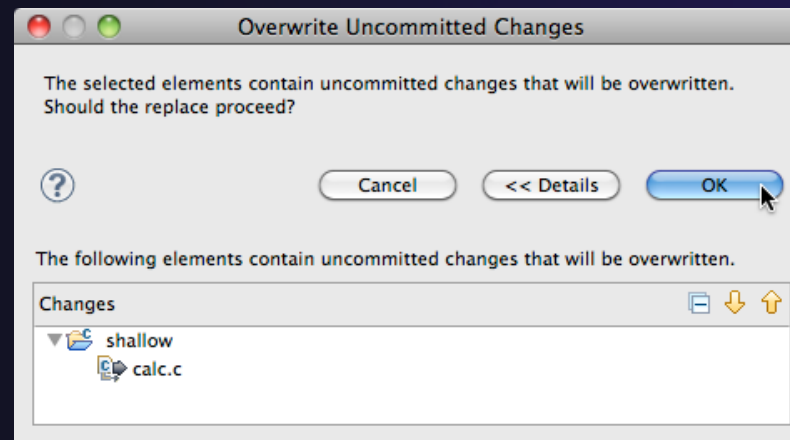
C Compare Viewer
  Local File 1.1
    32 float v[n][m];
    33 float cu[n][m];
    34 float cv[n][m];
    35 float h[n][m];
    36 float z[n][m];
    37 float fsdx, fsdy;
    38 {
    39   int i,j,ip,jp;
    40   /*
    41   * Added a comment here
    42   */
    43   for(j=jstart;j<=jend;j++) {
    44     jp = (j+1) % n;
    45     for (i = 0; i < m; i++){
    46       ip = (i+1) % m;
    47       cu[j][ip] = 0.5*(p[j][ip]+p[j][i])*u[j][i];
    48       cv[jp][i] = 0.5*(p[jp][i]+p[j][i])*v[jp][i];
    49       z[jp][ip] = (fsdx*(v[jp][ip]-v[jp][i])-f
    50       u[j][ip]) / (p[j][ip]+p[j][i]);

  Remote File 1.1
    30 float p[n][m];
    31 float u[n][m];
    32 float v[n][m];
    33 float cu[n][m];
    34 float cv[n][m];
    35 float h[n][m];
    36 float z[n][m];
    37 float fsdx, fsdy;
    38 {
    39   int i,j,ip,jp;
    40
    41   for(j=jstart;j<=jend;j++) {
    42     jp = (j+1) % n;
    43     for (i = 0; i < m; i++){
    44       ip = (i+1) % m;
    45       cu[j][ip] = 0.5*(p[j][ip]+p[j][i])*u
    46       cv[jp][i] = 0.5*(p[jp][i]+p[j][i])*v
    47       z[jp][ip] = (fsdx*(v[jp][ip]-v[jp][i]
    49       u[j][ip]) / (p[j][ip]+p[j][i]);
  
```



Revert To The Latest Version

- ★ Right-click on the “shallow” project and select **Replace With>Latest from HEAD**
- ★ Review the resources that will be replaced, then click **OK**



MPI Features

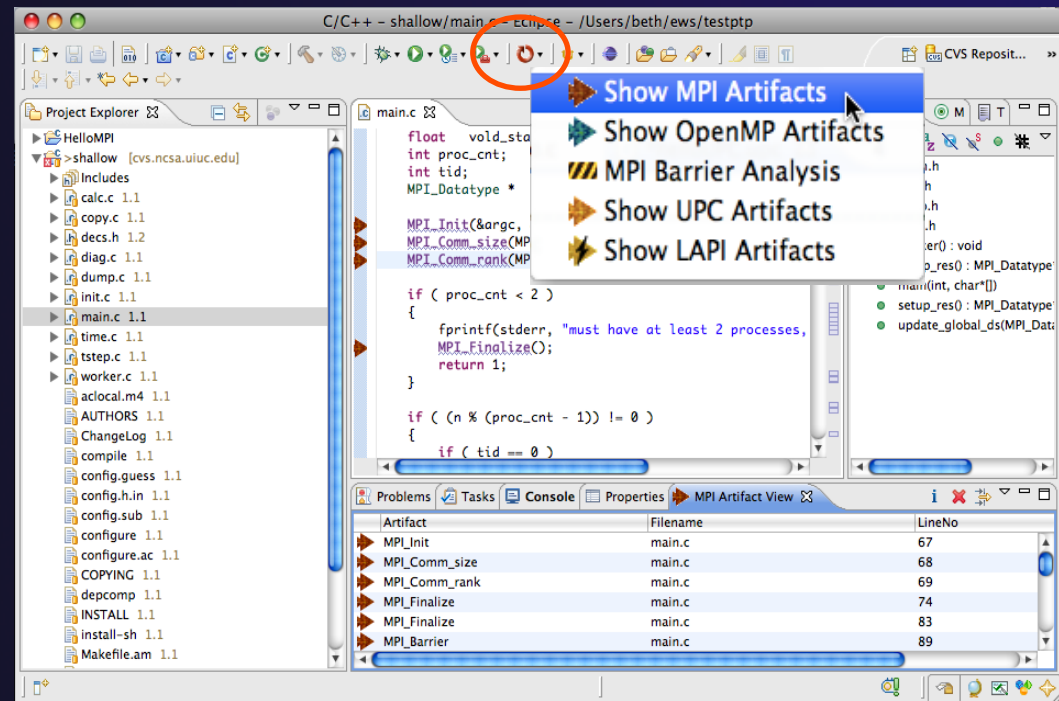
MPI-Specific Features

- ★ PTP's Parallel Language Development Tools (PLDT) has several features specifically for developing MPI code
 - ★ Show MPI Artifacts
 - ★ Code completion
 - ★ Context Sensitive Help for MPI
 - ★ Hover Help
 - ★ MPI Templates in the editor
 - ★ MPI Barrier Analysis

Show MPI Artifacts




- ★ In Project Explorer, select a project, folder, or a single source file
 - ★ The analysis will be run on the selected resources
- ★ Select **Show MPI Artifacts**
- ★ Run the analysis by clicking on drop-down menu next to the analysis button
- ★ Works on local and remote files



MPI Artifact View



- ★ Markers indicate the location of artifacts in editor
- ★ The **MPI Artifact View** lists the type and location of each artifact
- ★ Navigate to source code line by double-clicking on the artifact
- ★ Run the analysis on another file (or entire project!) and its markers will be added to the view
- ★ Click on column headings to sort
- ★ Remove markers via 

Artifact	Filename	LineNo
MPI_Init	main.c	67
MPI_Comm_size	main.c	68
MPI_Comm_rank	main.c	69
MPI_Finalize	main.c	74
MPI_Finalize	main.c	83
MPI_Barrier	main.c	89



MPI Editor Features

```

float vold_start[m];
int proc_cnt;
int tid;
MPI_Datatype * res_type;

MPI_Init(&argc, &argv);
MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt);
MPI_Comm_rank(MPI_COMM_WORLD, &tid);

MPI_B

```

- MPI_Barrier(MPI_Comm) int
- MPI_Bcast(void*, int, MPI_Datatype, int, MPI_
- MPI_Bsend(void*, int, MPI_Datatype, int, int,
- MPI_Bsend_init(void*, int, MPI_Datatype, int,
- MPI_Buffer_attach(void*, int) int
- MPI_Buffer_detach(void*, int*) int
- MPI_Barrier(MPI_Comm comm) : int
- MPI_Bcast(void * buffer, int count, MPI_Dataty
- MPI_Bsend(void * buf, int count, MPI_Datatype
- MPI_Bsend_init(void * buf, int count, MPI Data

- ★ Code completion will show all the possible MPI keyword completions
- ★ Enter the start of a keyword then press <ctrl-space>

- ★ Hover over MPI API
- ★ Displays the function prototype and a description

```

float vold_start[m];
int proc_cnt;
int tid;
MPI_Datatype * res_type;

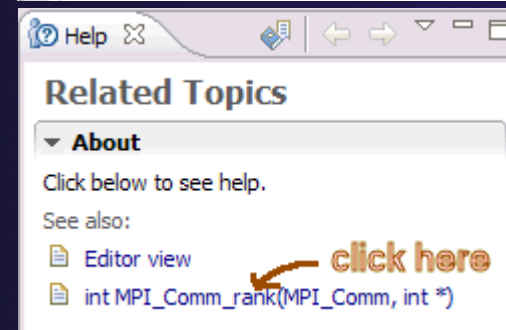
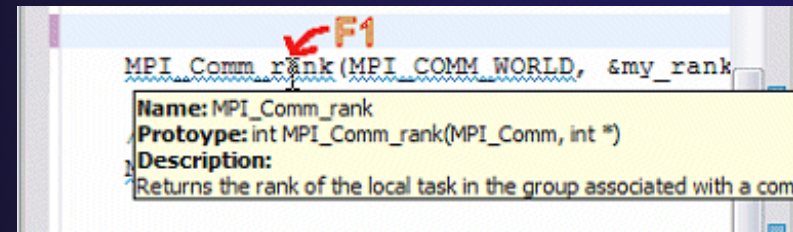
MPI_Init(&argc, &argv);
MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt);
MPI_Comm_rank(MPI_COMM_WORLD, &tid);

```

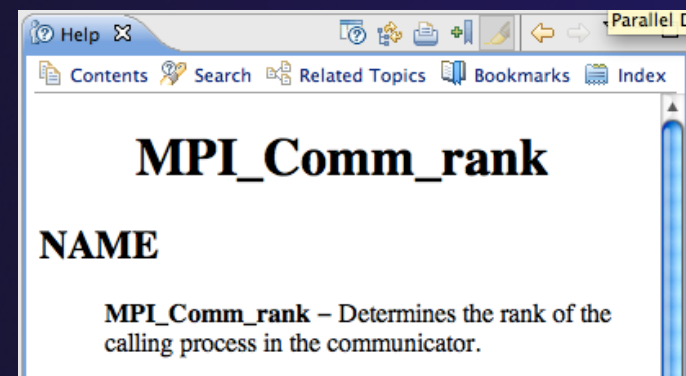
Name: MPI_Comm_rank
 Prototype: int MPI_Comm_rank(MPI_Comm, int *)
 Description:
 Returns the rank of the local task in the group associated with a communicator.

Context Sensitive Help

- ★ Click mouse, then press help key when the cursor is within a function name
 - ★ Windows: **F1** key
 - ★ Linux: **ctrl-F1** key
 - ★ MacOS X: **Help** key or **Help**►**Dynamic Help**
- ★ A help view appears (**Related Topics**) which shows additional information (You may need to click on MPI API in editor again, to populate)
- ★ Click on the function name to see more information
- ★ Move the help view within your Eclipse workbench, if you like, by dragging its title tab



Some special info has been added for MPI APIs



MPI Templates

★ Allows quick entry of common patterns in MPI programming

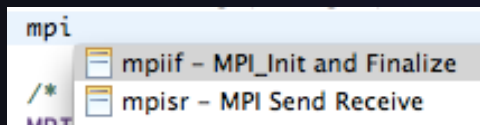
★ Example:
MPI send-receive

★ Enter:
mpisr <ctrl-space>

★ Expands to a send-receive pattern

★ Highlighted variable names can all be changed at once

★ Type mpi <ctrl-space> <ctrl-space> to see all templates



```

MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &p);
if (rank == 0){ //master task
    printf("Hello From process 0: Num processes: %d\n",p);
    for (source = 1; source < p; source++) {
        MPI_Recv(message, 100, MPI_CHAR, source, tag,
                MPI_COMM_WORLD, &status);
        printf("%s\n",message);
    }
}
else{ // worker tasks
    /* create message */
    sprintf(message, "Hello from process %d!", my_rank);
    dest = 0;
    /* use strlen+1 so that '\0' get transmitted */
    MPI_Send(message, strlen(message)+1, MPI_CHAR,
             dest, tag, MPI_COMM_WORLD);
}

```

Add more templates using Eclipse preferences!
C/C++>Editor>Templates
 Extend to other common patterns

MPI Barrier Analysis

*Local
files only*

The screenshot shows the Eclipse IDE interface for a C++ project named 'MyBarrier'. The main editor displays the source code of 'MyBarrier.c', which includes MPI_Send, MPI_Recv, and MPI_Barrier calls. Below the code, the 'Barrier Matches' and 'Barrier Errors' views show the analysis results.

Barrier Matches View:

Barrier Matching Set	Function	Filename	LineNo
Barrier 1 (2)	Barrier	MyBarrier.c	8
Barrier 1	Barrier	MyBarrier.c	8
Barrier 3	main	MyBarrier.c	41
Barrier 2 (1)	main	MyBarrier.c	31
Barrier 2	main	MyBarrier.c	31
Barrier 3 (2)	main	MyBarrier.c	41
Barrier 1	Barrier	MyBarrier.c	8
Barrier 3	main	MyBarrier.c	41
Barrier 4 (0)	main	MyBarrier.c	57
Barrier 5 (1)	main	MyBarrier.c	62

Barrier Errors View:

Barrier Matching Set	Function
Error	main
Path 1 (1 barrier(s))	
Path 2 (0 barrier(s))	
Error	main
Loop (dynamic number of barriers)	

Verify barrier synchronization in C/ MPI programs

Interprocedural static analysis outputs:

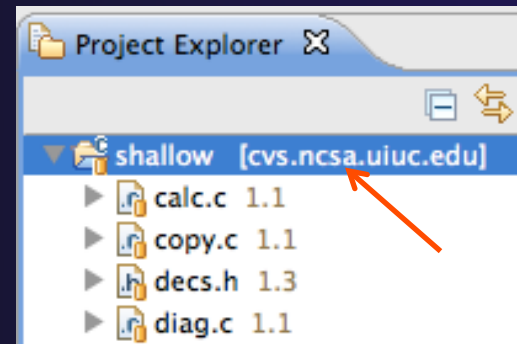
- ✦ For verified programs, lists barrier statements that synchronize together (match)
- ✦ For synchronization errors, reports counter example that illustrates and explains the error



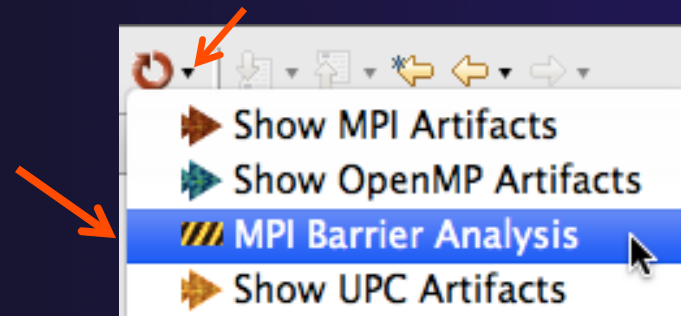
MPI Barrier Analysis – Try it

Run the Analysis:

- ★ In the Project Explorer, select the project (or directory, or file) to analyze



- ★ Select the MPI Barrier Analysis action in the pull-down menu



MPI Barrier Analysis – Try It (2)



- ★ No Barrier Errors are found (no pop-up indicating error); Two barriers are found

```
83     MPI_Finalize();
84     return 1;
85 }
86
87 if (tid != 0) {
88     worker();
89     MPI_Barrier(MPI_COMM_WORLD);
90     MPI_Finalize();
91 } else {
92
93     /* master process */
94
95     duplication = 1 / Comm.rank;
96     MPI_Finalize();
97     return 1;
98 }
```

Function	Filename	LineNo	IndexNo
main	main.c	89	1
main	main.c	206	2

MPI Barrier Analysis - views



The screenshot displays three panels from the MPI Barrier Analysis tool:

- MPI Barriers view:** A table listing barriers found in the code.

Function
main
main
main
main
main
Barrier
- Barrier Matches view:** A table showing groups of barriers that match together.

Barrier Matching Set	Function	Filename	LineNo
Barrier 1 (2)	Barrier	MyBarrier.c	8
Barrier 1	Barrier	MyBarrier.c	8
Barrier 3	main	MyBarrier.c	41
Barrier 2 (1)	main	MyBarrier.c	31
Barrier 2	main	MyBarrier.c	31
Barrier 3 (2)	main	MyBarrier.c	41
Barrier 1	Barrier	MyBarrier.c	8
Barrier 3	main	MyBarrier.c	41
Barrier 4 (0)	main	MyBarrier.c	57
Barrier 5 (1)	main	MyBarrier.c	62
- Barrier Errors view:** A table showing error paths.

Barrier Matching Set	Function
Error	main
Path 1 (1 barrier(s))	
Path 2 (0 barrier(s))	
Error	main
Loop (dynamic number of barriers)	

MPI Barriers view

Simply lists the barriers

Like MPI Artifacts view, double-click to navigate to source code line (all 3 views)

Barrier Matches view

Groups barriers that match together in a barrier set – all processes must go through a barrier in the set to prevent a deadlock

Barrier Errors view

If there are errors, a counter-example shows paths with mismatched number of barriers



Barrier Errors

- ✦ Let's cause a barrier mismatch error
- ✦ Open worker.c in the editor by double-clicking on it in Project Explorer
- ✦ At about line 125, enter a barrier:
 - ✦ Type MPI_B
 - ✦ Hit Ctl-space
 - ✦ Select MPI_Barrier
 - ✦ Add communicator arg MPI_COMM_WORLD and closing semicolon

```
120 prv = worker[PREV];
121 nxt = worker[NEXT];
122 jstart = worker[JSTART];
123 jend = worker[JEND];
124
125 MPI_B
126 /*
127 MPI_Barrier(MPI_Comm) int Blocks each task until
128 MPI_Bcast(void*, int, MPI_Datatype, int, MPI_
129 MPI_Bsend(void*, int, MPI_Datatype, int, int,
130 MPI_Bsend_init(void*, int, MPI_Datatype, int,
131 MPI_Buffer_attach(void*, int) int
132 MPI_Buffer_detach(void*, int*) int
```

```
124
125 MPI_Barrier(MPI_COMM_WORLD);
126
```

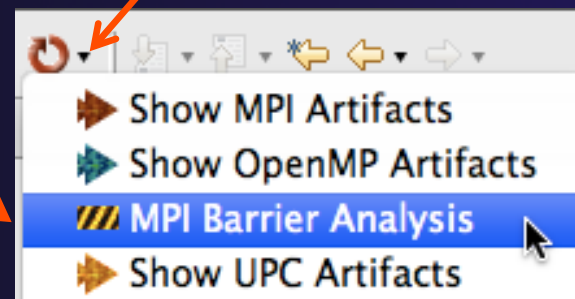


Barrier Errors (2)

- ★ Save the file
 - ★ Ctl-S (Mac Command-S) or File > Save
 - ★ Tab should lose asterisk indicating file saved



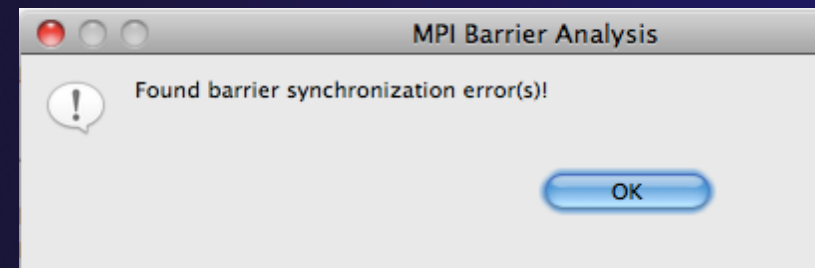
- ★ Run barrier analysis on shallow project again
 - ★ Select shallow project in Project Explorer first





Barrier Errors (3)

- ★ Barrier Error is found
- ★ Hit OK to dismiss dialog
- ★ Code diverges on line 87
 - ★ One path has 2 barriers, other has 1



Barrier Matching Set	Function	Filename	LineNo	IndexNo
▼ Error	main	main.c	87	0
▼ Path 1 (2 barrier(s))			0	0
Barrier 1	main	main.c	89	1
Barrier 3	worker	worker.c	125	3
▼ Path 2 (1 barrier(s))			0	0
Barrier 2	main	main.c	206	2

Double-click on a row in Barrier Errors view to find the line it references in the code

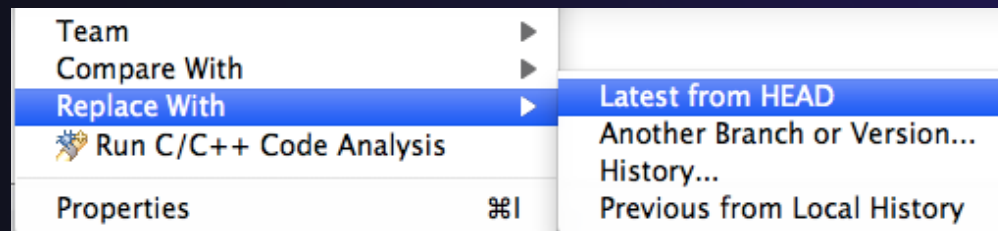


Fix Barrier Error

- ★ Fix the Barrier Error before continuing
- ★ Double-click on the barrier in worker.c to quickly navigate to it

Barrier Matching Set	Function	Filename	LineNo	IndexNo
▼ Error	main	main.c	87	0
▼ Path 1 (2 barrier(s))			0	0
Barrier 1	main	main.c	89	1
Barrier 3	worker	worker.c	104	3
▼ Path 2 (1 barrier(s))			0	0
Barrier 2	main	main.c	206	2

- ★ Remove the line and save the file
-or-
Right mouse on worker.c in Project Explorer and do **Replace With > Latest from HEAD**





Remove Barrier Markers

- ★ Run Barrier Analysis again to remove the error - and/or -
- ★ Remove the Barrier Markers via the "X" in one of the MPI Barrier views

Barrier Matching Set	Function	Filename	LineNo
▶ Error	main	main.c	87

Synchronizing the Project

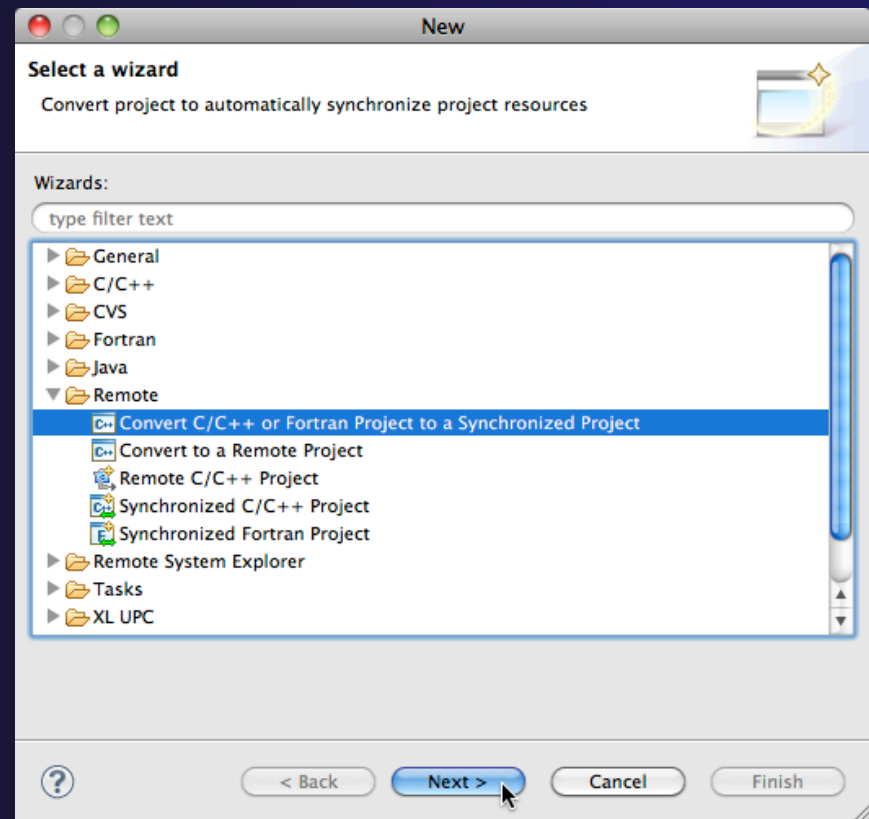
Synchronizing the Project

- ★ Because we will be running on a remote system, we must also build on that system
- ★ Source files must be available to build
- ★ We will use a synchronized project to do this
 - ★ Only needs to be done once for each project
 - ★ A synchronized project could have been created initially
- ★ Files are synchronized automatically when they are saved
- ★ A full synchronize is also performed prior to a build

Converting To Synchronized



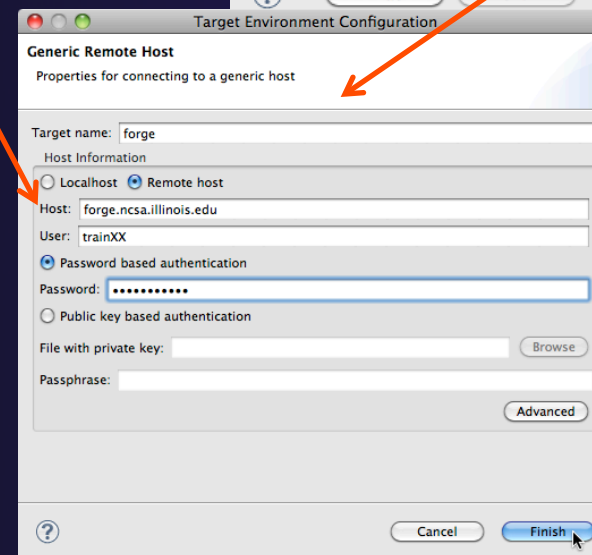
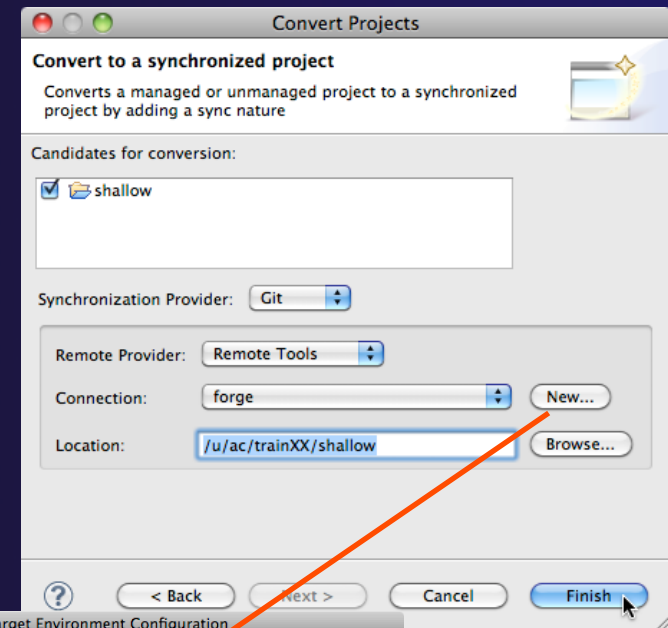
- ★ Select **File>New>Other...**
- ★ Open the Remote folder
- ★ Select **Convert C/C++ or Fortran Project to a Synchronized Project**
- ★ Click **Next>**



Convert Projects Wizard



- ✦ Select checkbox next to “shallow”
- ✦ For **Connection:**, click on **New...**
Enter as directed:
 - ✦ **Target name**
 - ✦ **Host** name of remote system
 - ✦ **User id** and **Password**
- ✦ Click **Finish** to close it
- ✦ Back in the **Convert Projects** dialog,
- ✦ Enter a directory name in the **Location** field; select **Browse...**
 - ✦ Sample: /u/ac/trainXX/shallow
 - ✦ Project files will be copied under this directory
- ✦ Click **Finish**





Synchronized Project

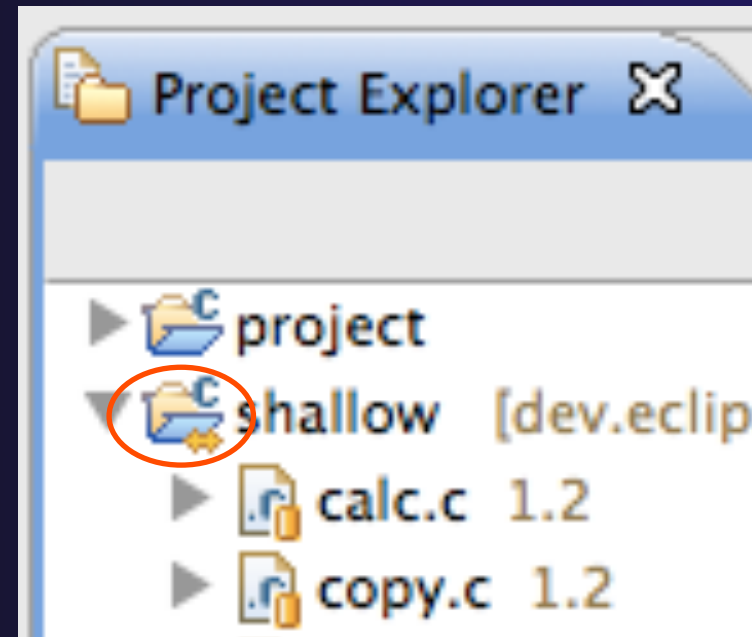
- ★ Back in the Project Explorer, decorator on project icon indicates synchronized project
- ★ Double-+ icon

- ★ Before sync

▼  shallow [dev.eclipse.org]

- ★ After sync

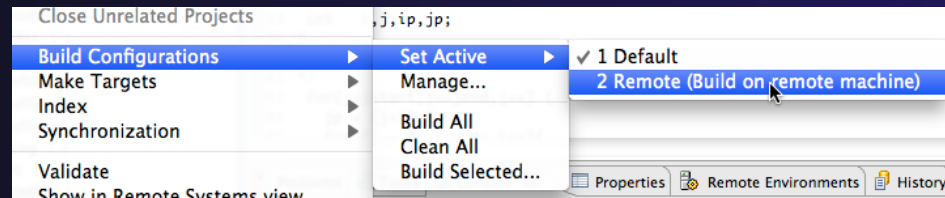
▼  shallow [dev.eclipse.org]



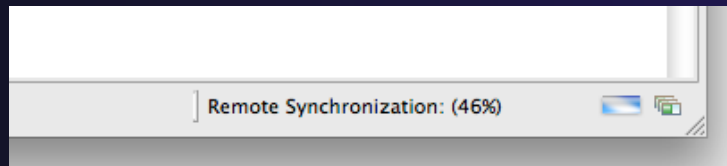


Set Active Build Configuration

- ★ The “Active” build configuration determines which system will be used for both synchronizing and building
- ★ Right-click on the project and select **Build Configurations>Set Active>Remote (Build on remote machine)**



- ★ The project should synchronize immediately

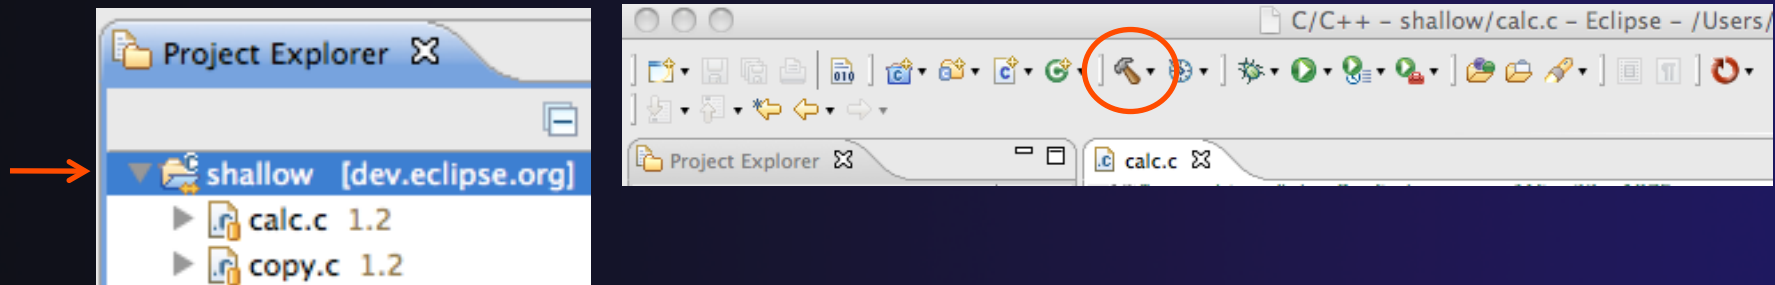


Building the Project



Building the Project

- ★ Select the project in Project Explorer, click on the  hammer button to run the build



- ★ By default, the Build Configuration assumes there is a Makefile (or makefile) for the project
- ★ In this case, there is no Makefile, so the build will fail.
See Console:

- ★ We'll see how to change it if the build command is different from 'make -f Makefile'

```

CDT Build Console [shallow]

**** Build of configuration Remote for project shallow ****

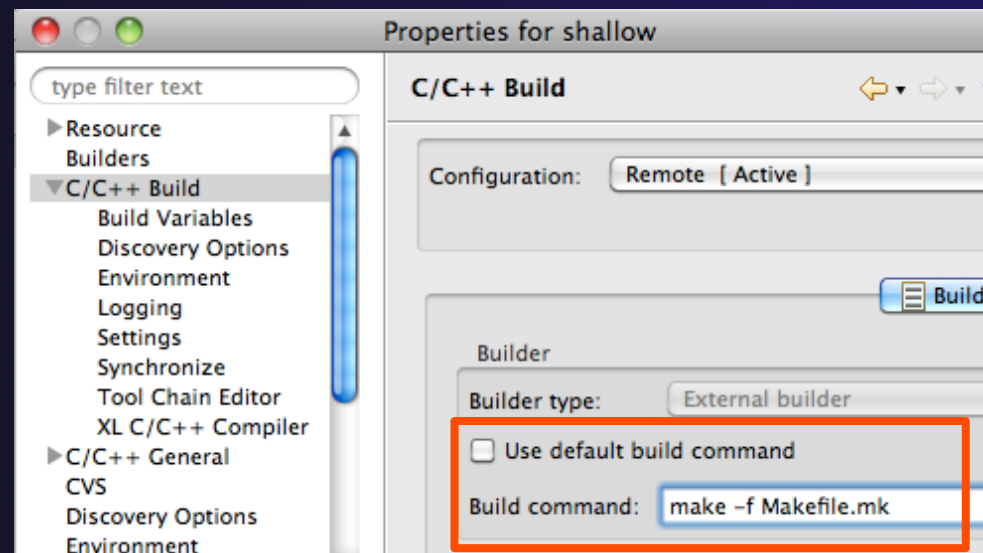
make all
make: *** No rule to make target `all'. Stop.

**** Build Finished ****
  
```

Fixing the Build Command: Editing Project Properties




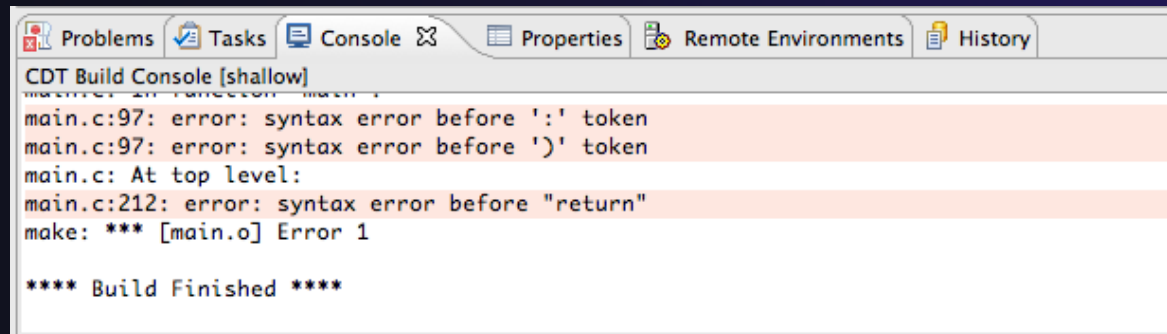
- ★ The build command is specified in the project properties
- ★ Open the properties by right-clicking on “shallow” and selecting **Properties** (bottom of the context menu list)
- ★ Click on **C/C++ Build**
- ★ Uncheck **Use default build command**
- ★ Enter “make -f Makefile.mk” in the **Build Command** field
- ★ Click **OK** to close project properties dialog





Re-Building the Project

- ★ Click on the  button again to run the build
- ★ Build output will be shown in the **Console** view

A screenshot of the CDT Build Console window. The window title is "CDT Build Console [shallow]". The console output shows several error messages: "main.c:97: error: syntax error before ':' token", "main.c:97: error: syntax error before ')' token", "main.c: At top level:", "main.c:212: error: syntax error before \"return\"", and "make: *** [main.o] Error 1". The output ends with "**** Build Finished ****". The error messages are highlighted with a light orange background.

```
CDT Build Console [shallow]
main.c:97: error: syntax error before ':' token
main.c:97: error: syntax error before ')' token
main.c: At top level:
main.c:212: error: syntax error before "return"
make: *** [main.o] Error 1

**** Build Finished ****
```

- ★ Exact output depends on your compiler

Build Problems



★ Build problems will be shown in a variety of ways

- ★ Marker on file
- ★ Marker on editor line
- ★ Line is highlighted
- ★ Marker on overview ruler
- ★ Listed in the **Problems view**

★ Double-click on line in **Problems view** to go to location of error in the editor


The screenshot shows the Eclipse IDE interface. The Project Explorer on the left shows a project named 'shallow' with various source files. The main editor window displays the code for 'main.c'. A syntax error is highlighted on line 97: 'for (i = 1; i < proc_cnt; i++) {'. The Problems view at the bottom shows three error items:

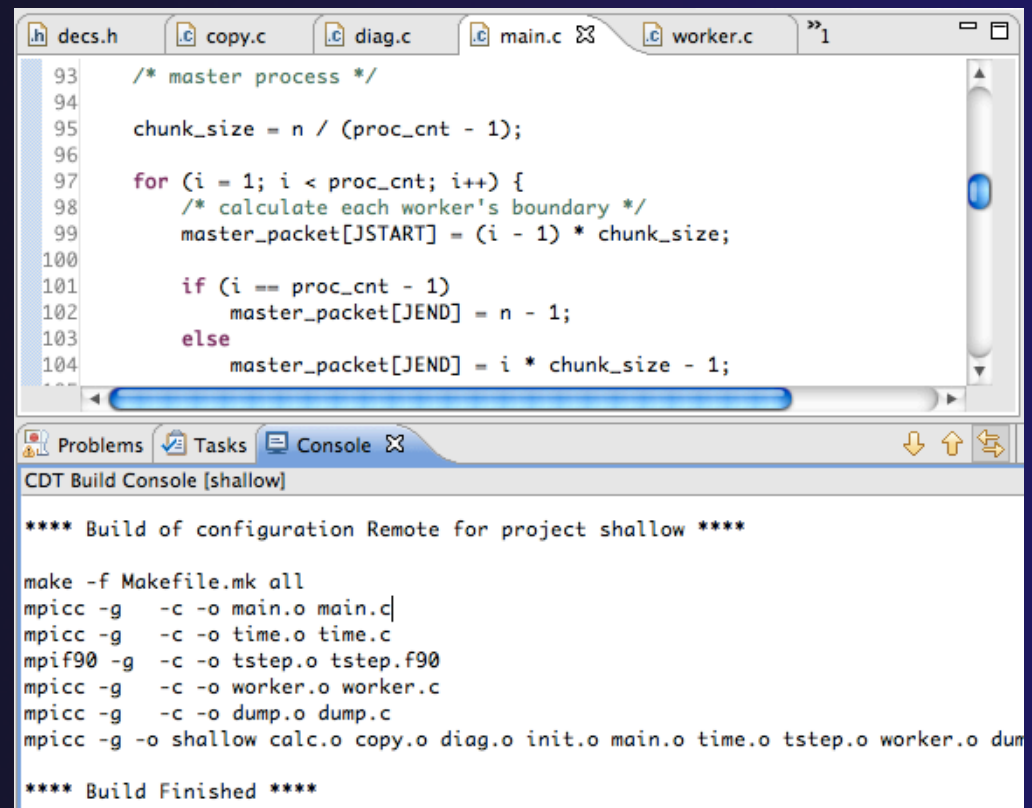
Description	Resource	Path	Location	Type
✘ syntax error before ':' token	main.c	/shallow	line 97	C/C++ Problem
✘ syntax error before ')' token	main.c	/shallow	line 97	C/C++ Problem
✘ syntax error before "return"	main.c	/shallow	line 212	C/C++ Problem

Orange arrows from the text on the left point to: the 'main.c' file in the Project Explorer, the error marker on line 97 in the editor, the highlighted line in the editor, the error marker on the overview ruler, and the error entry in the Problems view.

Fix Build Problems



- ✦ Fix errors by changing `:` to `;` on line 97
- ✦ Save the file
- ✦ Rebuild by pressing build button 
- ✦ Error markers have been removed
- ✦ Check console for correct build output



The screenshot shows an IDE window with several tabs: decs.h, copy.c, diag.c, main.c, and worker.c. The main.c tab is active, showing the following code:

```
93  /* master process */
94
95  chunk_size = n / (proc_cnt - 1);
96
97  for (i = 1; i < proc_cnt; i++) {
98      /* calculate each worker's boundary */
99      master_packet[JSTART] = (i - 1) * chunk_size;
100
101      if (i == proc_cnt - 1)
102          master_packet[JEND] = n - 1;
103      else
104          master_packet[JEND] = i * chunk_size - 1;
```

Below the code editor is a console window titled "CDT Build Console [shallow]". It displays the following output:

```
**** Build of configuration Remote for project shallow ****

make -f Makefile.mk all
mpicc -g -c -o main.o main.c
mpicc -g -c -o time.o time.c
mpif90 -g -c -o tstep.o tstep.f90
mpicc -g -c -o worker.o worker.c
mpicc -g -c -o dump.o dump.c
mpicc -g -o shallow calc.o copy.o diag.o init.o main.o time.o tstep.o worker.o dum

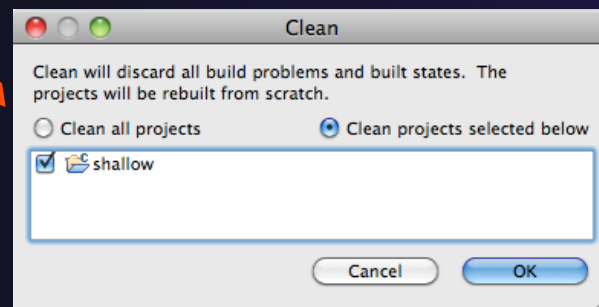
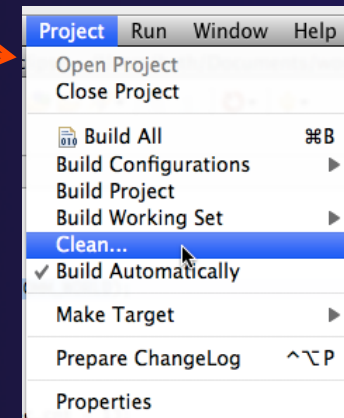
**** Build Finished ****
```



Forcing a Rebuild

- ★ If no changes have been made, make doesn't think a build is needed
- ★ **Project > Clean...**
- ★ Select your project

```
CDT Build Console [shallow]
make -f Makefile.mk all
make: Nothing to be done for `all'.
```



- ★ See console

```
CDT Build Console [shallow]
**** Clean-only build of configurat
make -f Makefile.mk clean
rm -f shallow calc.o copy.o diag.o
```

- ★ Then rebuild 

Running the Program

Resource Managers

Running the Program

- ✦ Creating a resource manager
- ✦ Starting the resource manager
- ✦ Creating a run configuration
- ✦ Running (launching) the application
- ✦ Viewing the application run



Do this
once

Much of the following setup is configuration that you only need to do once: This icon will remind you.

Resource Managers

- ★ PTP uses the term “resource manager” to refer to any subsystem that controls the resources required for launching a parallel job.
- ★ Examples:
 - ★ Batch scheduler (e.g. LoadLeveler, PBS, SLURM)
 - ★ Interactive execution (e.g. Open MPI, MPICH2, etc.)
- ★ Each resource manager controls one target system
- ★ Resource Managers can be local or remote

Monitoring/Runtime Perspectives

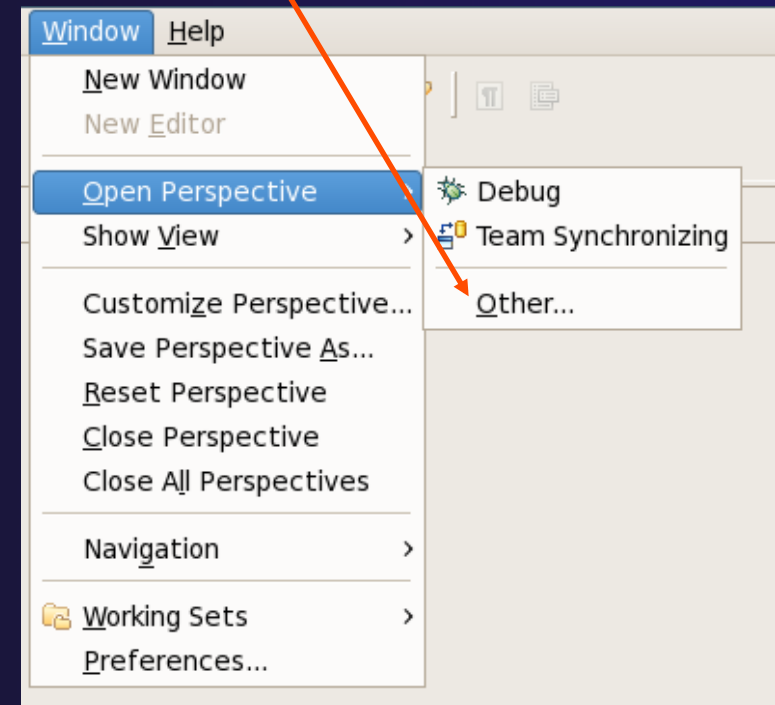
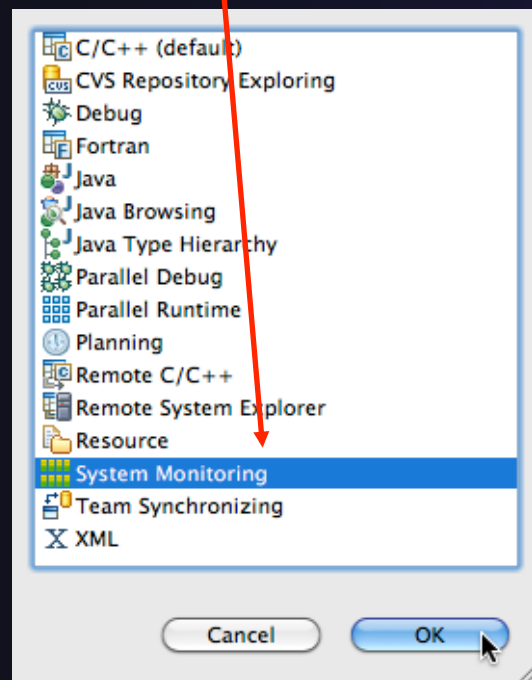
- ★ Parallel Runtime Perspective
 - ★ Used for legacy PTP Resource Managers
- ★ System Monitoring Perspective
 - ★ Used for newer Configurable Resource Managers (since PTP 5.0)
- ★ Which one is used?

Resource Manager	System Monitoring	Parallel Runtime
IBM LoadLeveler		✓
IBM Parallel Env		✓
MPICH2		✓
Open MPI		✓
PBS-Batch-Generic	✓	
PBS-Batch-Interactive	✓	
Remote Launch		✓
SLURM		✓



Preparing to Launch

- ★ Setting up a resource manager is done in the System Monitoring perspective
 - ★ (For PTP 5.0, this applies to PBS)
- ★ Select **Window>Open Perspective>Other...**
- ★ Choose **System Monitoring** and click **OK**



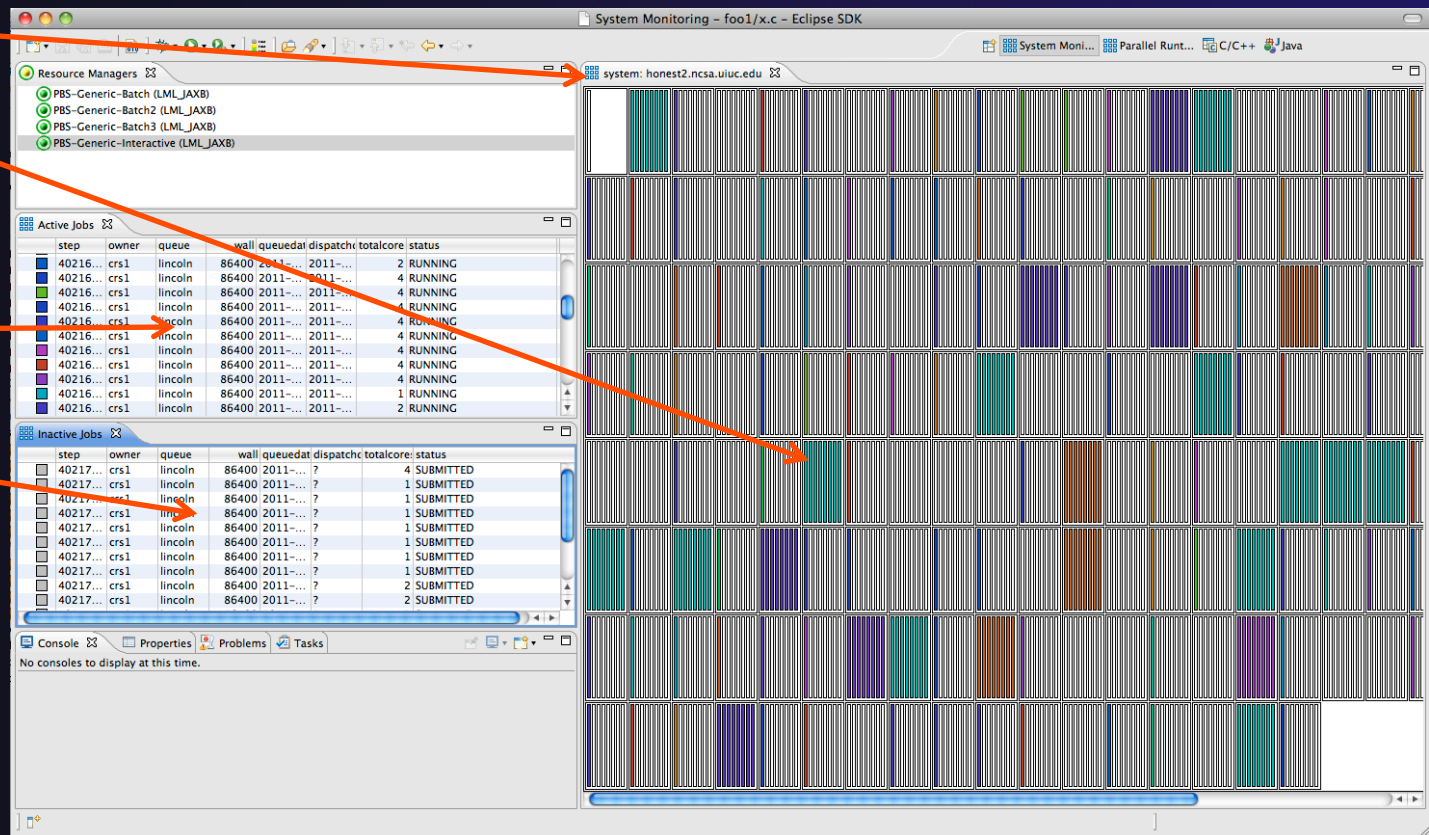
System Monitoring Perspective

✦ System view

✦ Jobs running on system

✦ Active jobs

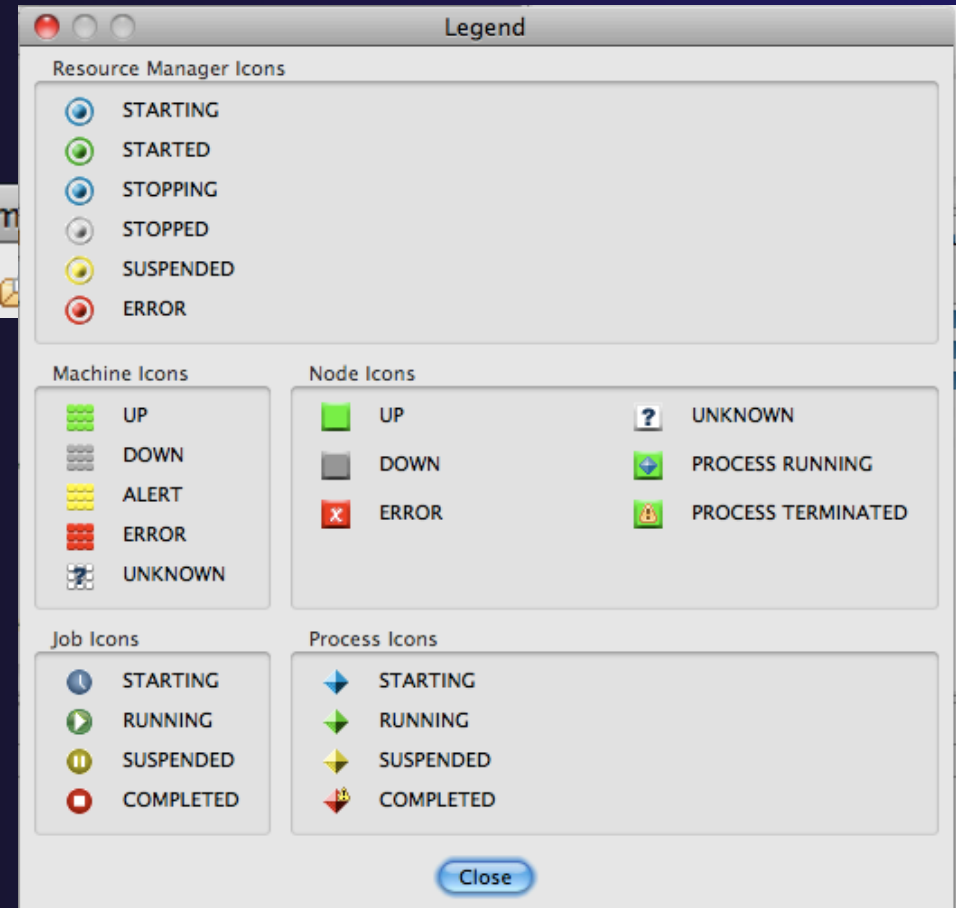
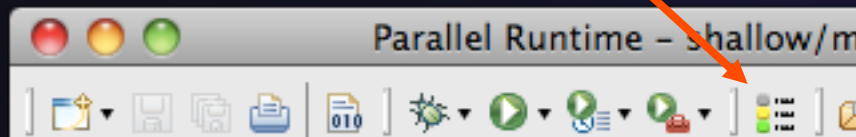
✦ Inactive jobs





About PTP Icons

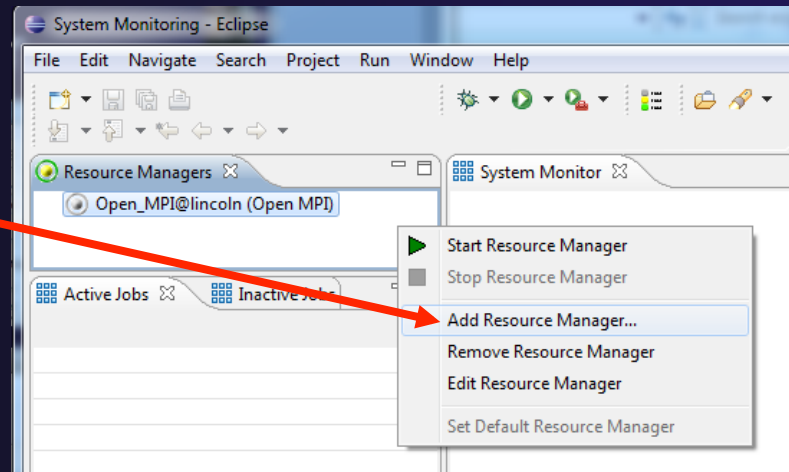
- ★ Open using legend icon in toolbar



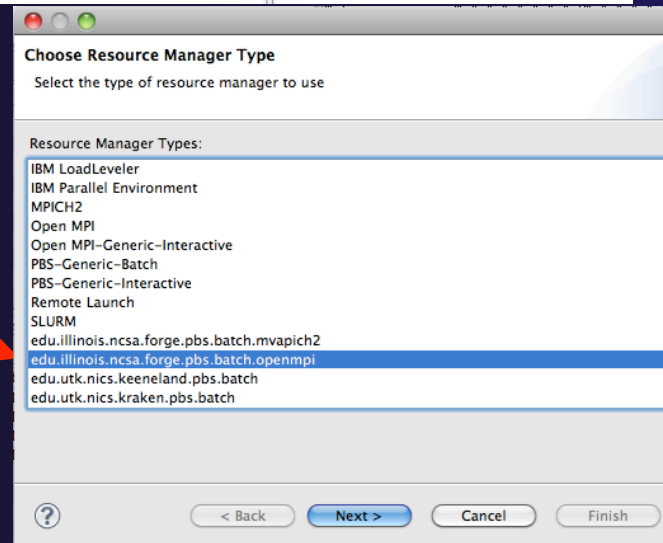
Configuring Job Scheduler



- ★ Right-click in Resource Managers view and select **Add Resource Manager**



- ★ Choose Resource Manager Type:
edu.illinois.ncsa.forge
.pbs.batch.openmpi



Do this once

- ★ Select **Next>**

Configure Control Connection



- ✦ Choose **Remote Tools** for **Remote service provider**
- ✦ Choose the remote connection you made previously
- ✦ Click **Next>**



Do this once

Control Connection configuration
Enter connection information

Remote service provider: Remote Tools

Connection name: forge

▶ Advanced Options

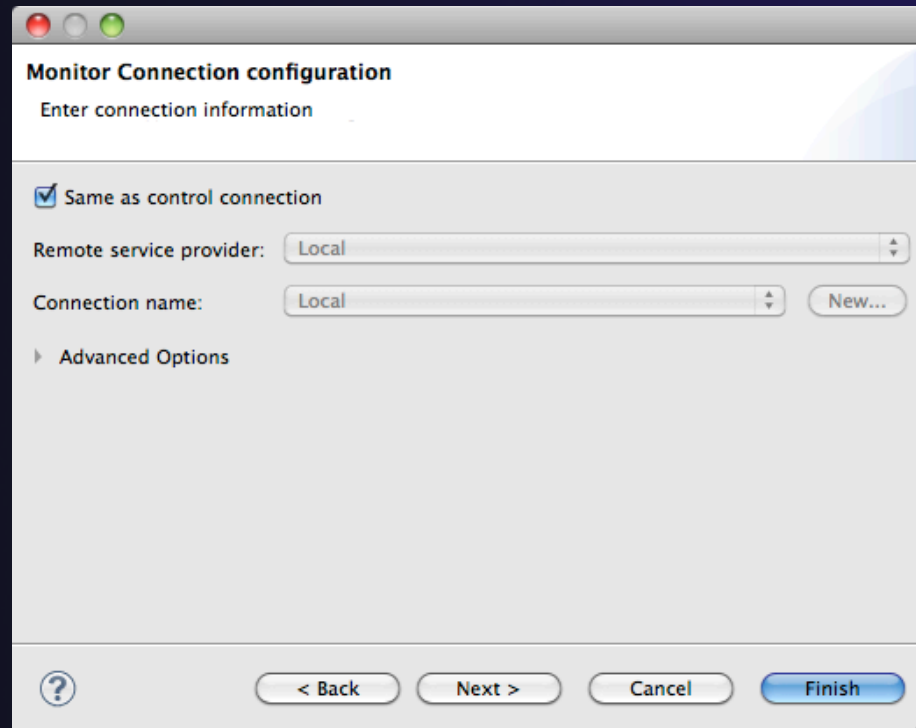
Configure Monitor Connection



- ★ Keep default Monitor Connection (same as Control Connection), click **Next**



Do this once



Monitor Connection configuration
Enter connection information

Same as control connection

Remote service provider: Local

Connection name: Local

▶ Advanced Options

Common Configuration



- ✦ Keep default name
- ✦ Can automatically start Resource Manager (leave unselected today)
- ✦ Click **Finish**



Do this once

Common Resource Manager Configuration
Change any settings for the resource manager

Name and description

Use default name and description:

Name: PBS-Generic-Batch

Description: XML Configurable Resource Manager

Startup

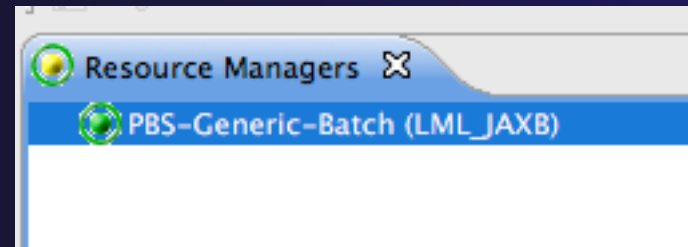
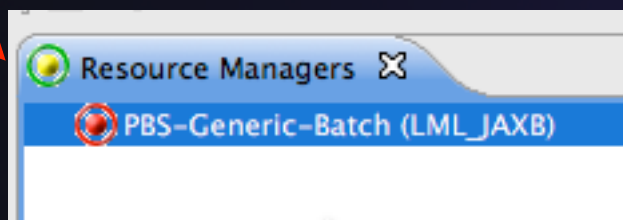
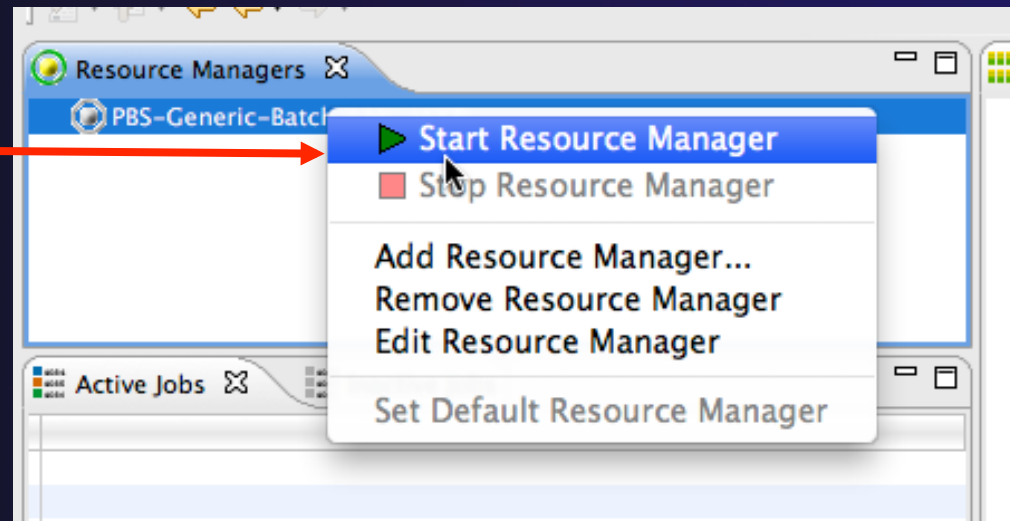
Automatically start resource manager when Eclipse starts

? < Back Next > Cancel Finish

Starting the Resource Manager



- ★ Right click on new resource manager and select **Start resource manager**
- ★ If everything is ok, you should see the resource manager change to **green**
- ★ If something goes wrong, it will change to **red**



System Monitoring



forge.ncsa.illinois.edu

- ★ **System** view, with abstraction of nodes for selected Resource Manager
- ★ Active and inactive jobs
- ★ Hover over node in **System** view to see job running on node in **Active Jobs** view
- ★ Hold mouse button down on a job in **Active Jobs** view to see where it is running in **System** view

Select Resource Manager

step	owner	queue	wall	queued	dispatch	totalcores	status
3716.fsched	alberto	normal	43200	201...	2011...	18	RUNNING
3718.fsched	dsouth	normal	14400	201...	2011...	6	RUNNING
3719.fsched	dsouth	normal	14400	201...	2011...	4	RUNNING
3720.fsched	alberto	normal	43200	201...	2011...	18	RUNNING
3722.fsched	dhguo	normal	12600	201...	2011...	1	RUNNING
3723.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3724.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3725.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3726.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3727.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3728.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3729.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3730.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING
3731.fsched	cheatham	normal	43200	201...	2011...	16	RUNNING

```

**** Build of configuration Remote for project shallow ****
make -f Makefile.mk all
mpicc -g -c -o calc.o calc.c
mpicc -g -c -o copy.o copy.c
  
```

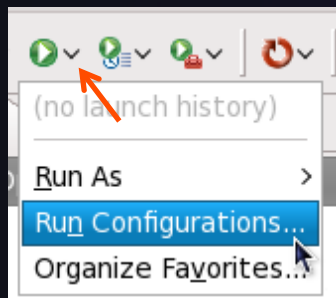
One node with
16 cores


Running the Program (Launching a Job)

Create a Run Configuration

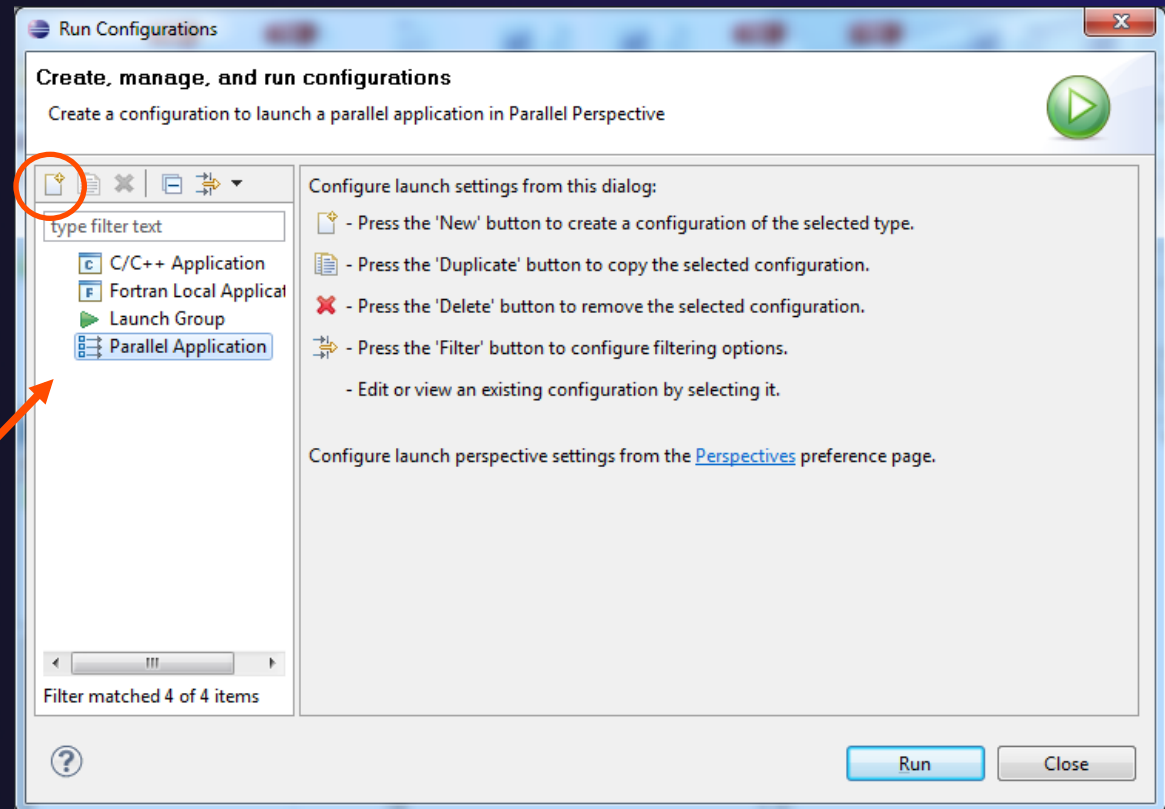


Do this once



- ★ Open the run configuration dialog **Run>Run Configurations...**
- ★ Select **Parallel Application**
- ★ Select the **New** button 

Or, just double-click on **Parallel Application** to create a new one



Depending on which flavor of Eclipse you installed, you might have more choices in Application types

Note: we sometimes interchange the terms "Run Configuration" and "Launch Configuration"

Complete the Resources Tab



Do this
once



- ✦ Enter a name for this run configuration, e.g. "shallow-pbs-batch"
- ✦ In **Resources** tab, select the PBS resource manager you just created (edu.illinois.ncsa.forge....)
- ✦ Select the destination queue – **debug**
- ✦ The **MPI Command** field allows this job to be run as an MPI job
 - ✦ Choose **mpirun**
- ✦ Enter the resources needed to run this job
 - ✦ Use 1 nodes, 4 cores (MPI tasks)

The screenshot shows the 'Run Configurations' dialog box with the 'Resources' tab selected. The configuration is for a job named 'shallow-pbs-batch' using the resource manager 'edu.illinois.ncsa.forge.pbs.batch.openmpi'. The 'Queue' is set to 'debug', the 'Number of nodes' is 1, and the 'MPI Command' is 'mpirun'. The 'MPI Number of Cores' is set to 4. The 'Export Environment' checkbox is checked.

Name	Value	Description
Job Name:	ptp_job	The name assigned to the job by the qsub or qalter command.
Account:		Account to which to charge this job.
Queue:	debug	Designation of the queue to which to submit the job.
Number of nodes:	1	Number and/or type of nodes to be reserved for exclusive use.
Total Memory Needed:		Maximum amount of memory used by all concurrent processes.
Wallclock Time:	00:05:00	Maximum amount of real time during which the job can be executed.
MPI Command:	mpirun	Which mpi command to use.
MPI Number of Cores:	4	the '-np' value
Export Environment:	<input checked="" type="checkbox"/>	All variables in the qsub command's environment are to be exported.

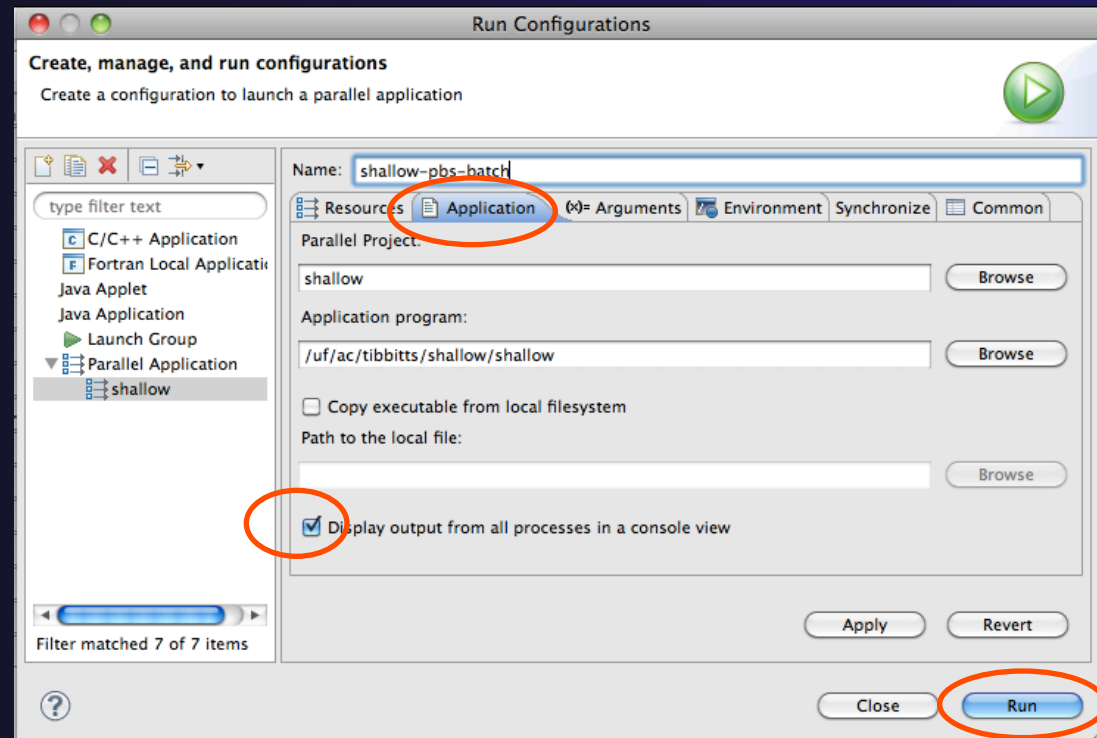
Complete the Application Tab



Do this
once



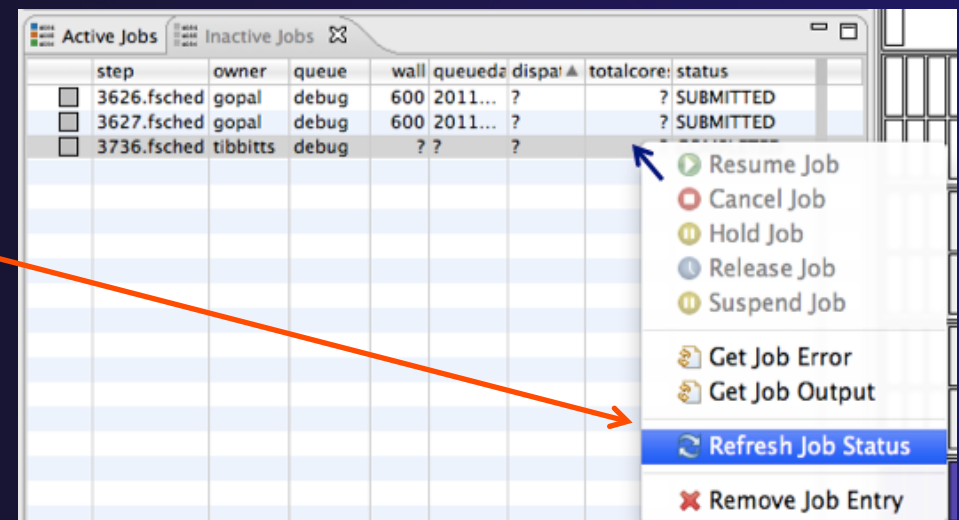
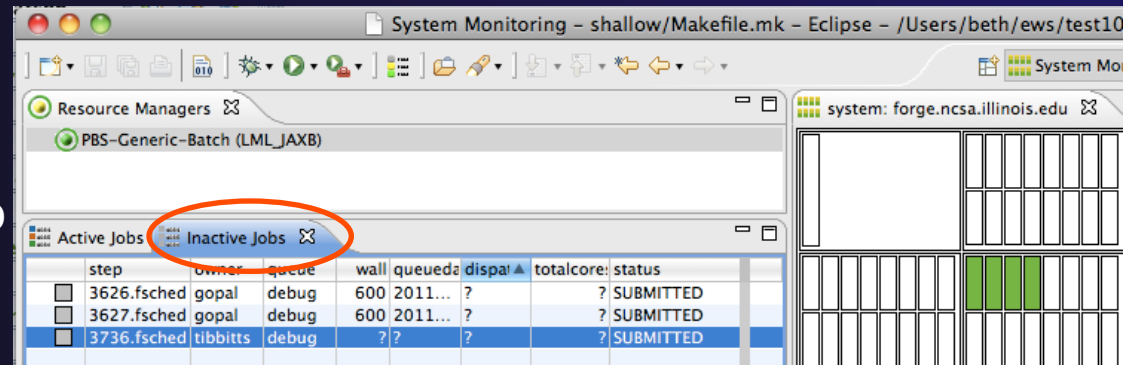
- ★ Select the **Application** tab
- ★ Choose the **Application program** by clicking the **Browse** button and locating the executable on the remote machine
 - ★ Use the same "shallow" executable
- ★ Select **Display output from all processes in a console view**
- ★ Click **Run** to submit the application to the job scheduler





Job Monitoring

- ★ Job initially appears in "Inactive Jobs", then in "Active Jobs", then returns to Inactive on completion
- ★ This short-running program may not run long enough to appear in "Active Jobs"
- ★ Status refreshes automatically every 60 sec
Or force refresh with menu
- ★ After status = COMPLETED, Can view output or error by right clicking on job, selecting appropriate output



Job Output

- ★ After status = COMPLETED, Can view output or error by right clicking on job, selecting appropriate output
- ★ Output/Error info shows in Console View

step	owner	queue	wall	queuedate	dispatch	total	status
3626.fsched	gopal	debug	600	2011-1...	?	?	SUBMITTED
3627.fsched	gopal	debug	600	2011-1...	?	?	SUBMITTED
3769.fsched	alberto	normal	43...	2011-1...	?	18	SUBMITTED
3774.fsched	dsouth	normal	14...	2011-1...	?	12	SUBMITTED
3772.fsched	tibbitts	?	??	?	?	?	COMPLETED
3773.fsched	tibbitts	debug	??	??	?	?	COMPLETED
3777.fsched	tibbitts	debug	??	??	?	?	COMPLETED
3783.fsched	tibbitts	debug	??	??	?	?	COMPLETED

Console View Output:

```

/ur/ae/tibbitts/ptp_job.o3777

Potential energy      6.505  Kinetic Energy
Total Energy         48032.016  Pot. Enstrophy

Cycle number  950  Model time in days  0.9
Potential energy      760.460  Kinetic Energy
Total Energy         48385.996  Pot. Enstrophy

Cycle number  1000  Model time in days  1.0
Potential energy     21561.033  Kinetic Energy
Total Energy         48000.496  Pot. Enstrophy
  
```

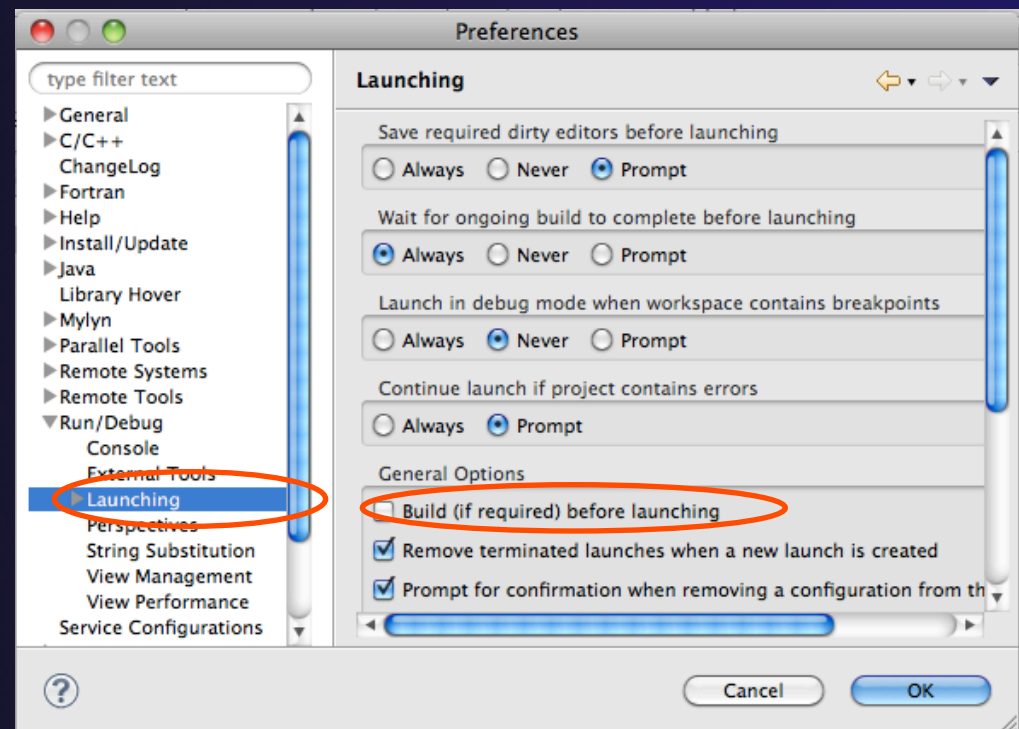
Building before Run



Do this once

- ★ If projects build prior to launch, you can turn it off.
 - ★ Go into **Preferences>Run/Debug** and click on **Launching**.
 - ★ Uncheck "**Build (if required) before launching**"

To bring up **Preferences** dialog, use Window>Preferences or Mac: Eclipse>Preferences





Exercise

- ✦ Start with your 'shallow' project
- ✦ Create and start Resource Manager
- ✦ Build; Run shallow
- ✦ See results
- ✦ Change something
 - ✦ Change m and n in decs.h
- ✦ Rebuild and re-run

Advanced Features

Searching
Fortran
Refactoring

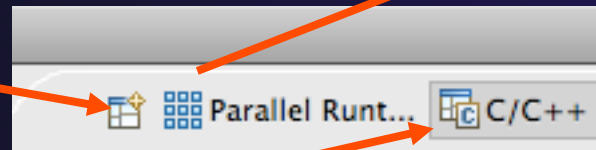
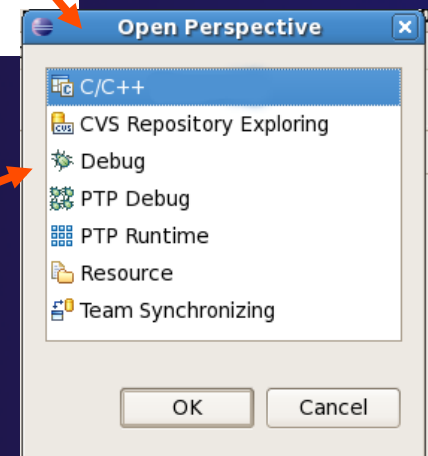
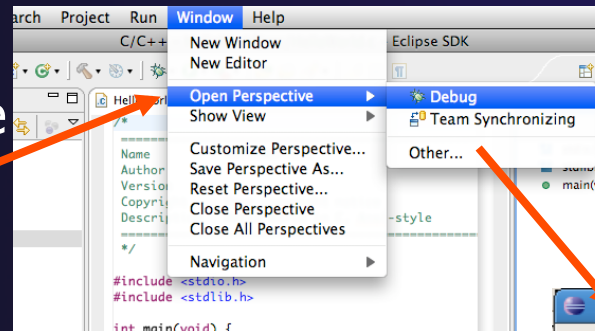
Searching

Switching Perspectives



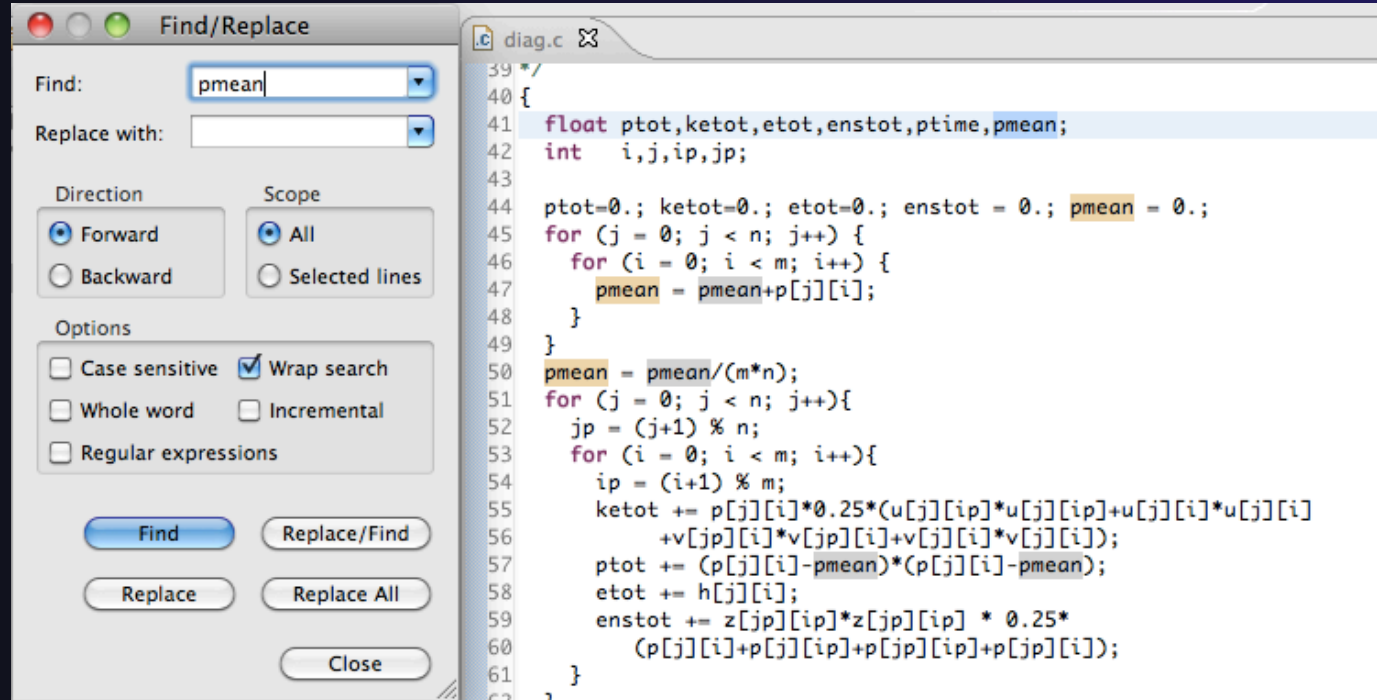
✦ Switch to C/C++ Perspective one of three ways:

1. Choose the **Window>Open Perspective** menu option
Then choose **Other...**
2. Click on the **Open Perspective** button in the upper right corner of screen (hover over it to see names)
3. Click on a perspective shortcut button



Find/Replace within Editor

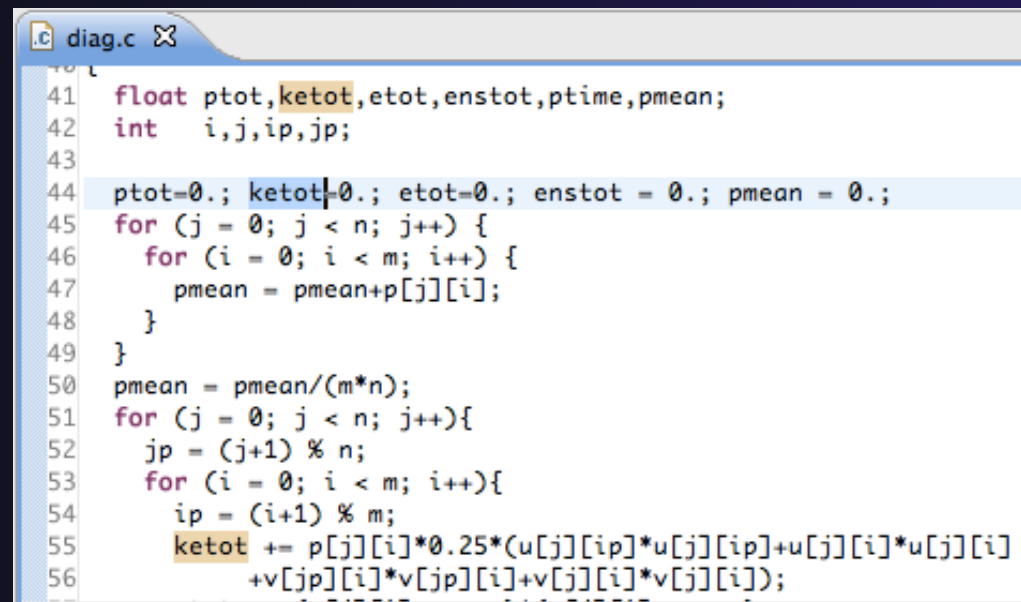
- ✦ Simple Find within editor buffer
- ✦ Ctrl-F (Mac: Command-F)



Mark Occurrences

(C/C++ Only)

- ★ Double-click on a variable in the CDT editor
- ★ All occurrences in the source file are highlighted to make locating the variable easier
- ★ Alt-shift-O to turn off (Mac: Alt-Cmd-O)



The screenshot shows a code editor window titled 'diag.c'. The code is as follows:

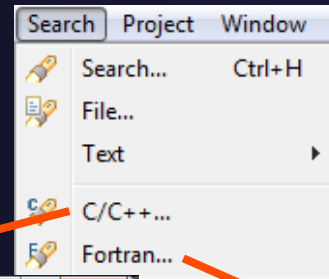
```
41 float ptot, ketot, etot, enstot, ptime, pmean;
42 int i, j, ip, jp;
43
44 ptot=0.; ketot=0.; etot=0.; enstot = 0.; pmean = 0.;
45 for (j = 0; j < n; j++) {
46     for (i = 0; i < m; i++) {
47         pmean = pmean+p[j][i];
48     }
49 }
50 pmean = pmean/(m*n);
51 for (j = 0; j < n; j++){
52     jp = (j+1) % n;
53     for (i = 0; i < m; i++){
54         ip = (i+1) % m;
55         ketot += p[j][i]*0.25*(u[j][ip]*u[j][ip]+u[j][i]*u[j][i]
56         +v[jp][i]*v[jp][i]+v[j][i]*v[j][i]);
```

In the screenshot, the variable 'ketot' is highlighted in yellow in lines 41 and 55. The line containing the first occurrence of 'ketot' (line 44) is also highlighted in light blue.

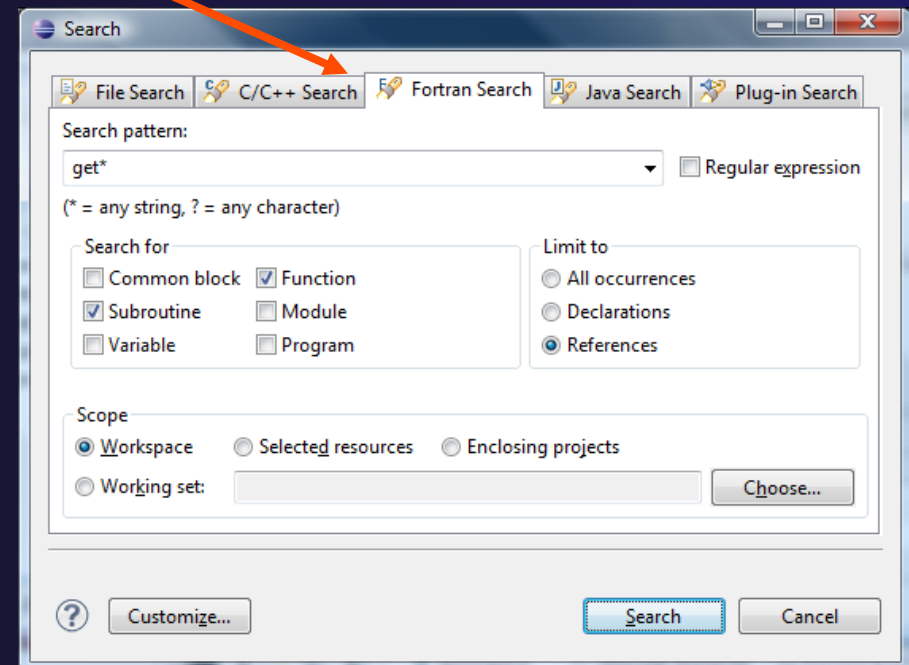
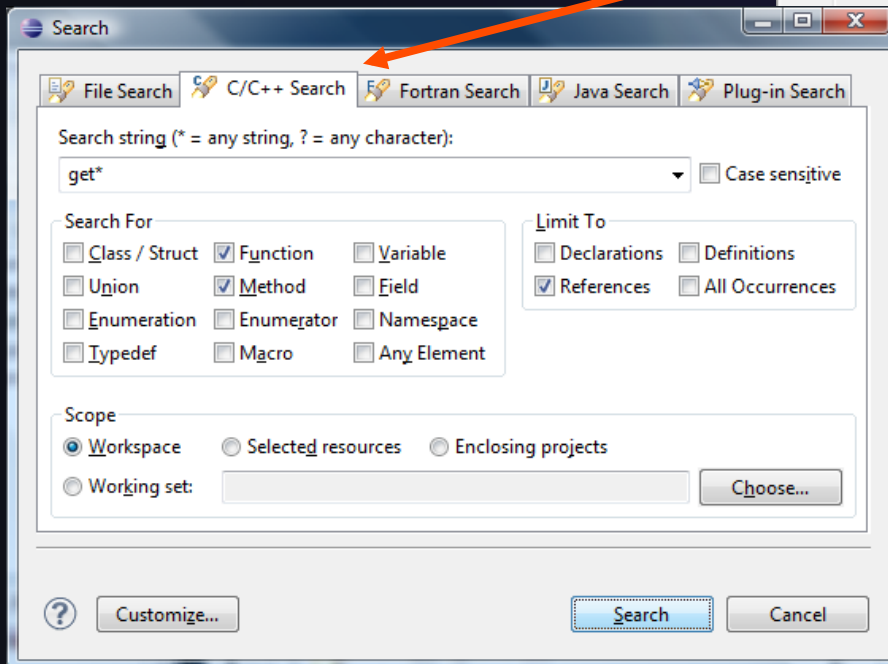
Language-Based Searching

(C/C++ and Fortran)

- ★ “Knows” what things can be declared in each language (functions, variables, classes, modules, etc.)



- ★ E.g., search for every call to a function whose name starts with “get”
- ★ Search can be project- or workspace-wide



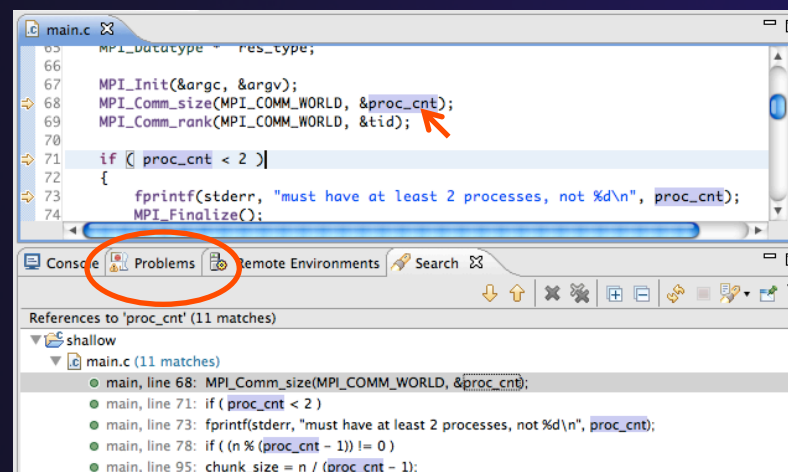
Find References

(C/C++ and Fortran)

- ★ Finds all of the places where a variable, function, etc., is used
 - ★ Right-click on an identifier in the editor
 - ★ Click **References** ▶ **Workspace** or **References** ▶ **Project**



- ★ **Search** view shows matches



Open Declaration

(C/C++ and Fortran)

- ★ Jumps to the declaration of a variable, function, etc., even if it's in a different file
- ★ Left-click to select identifier
- ★ Right-click on identifier
- ★ Click **Open Declaration**
- ★ C/C++ only:
Can also Ctrl-click (Mac: Cmd-click) on an identifier to “hyperlink” to its declaration

```

134
135 initialise(p, u, v, psi, pold, uold, vold, di, dj, z);
136 diag(1, 0., p, u, v, h, z);
137
138 for (i = 1; i < proc_cnt; i++)
139     for (j = 0; j < n; j++)
140         acopy_two_to_one(i, j, p, u, v, h, z);
141     MPI_Send(&star, 1, MPI_COMM_WORLD, j, 0, MPI_COMM_WORLD);
142
143
144     acopy_two_to_one(i, j, p, u, v, h, z);
145     MPI_Send(&star, 1, MPI_COMM_WORLD, j, 0, MPI_COMM_WORLD);
  
```

Goes to its declaration in copy.c

```

59 bcopy(src[column], dest[column], sizeof(src[column]));
60 }
61
62 acopy_two_to_one(twodim, onedim, column)
63 float twodim[n][m];
64 float onedim[m];
65 int column;
  
```



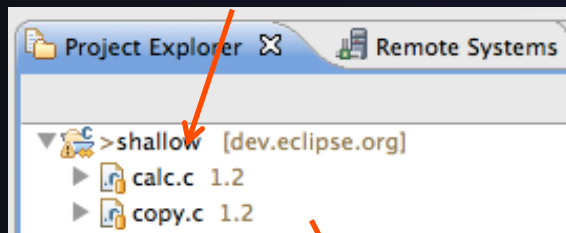
Search – Try It!

1. Find every call to `MPI_Recv` in Shallow.
2. In `worker.c`, on line 42, there is a declaration `float p[n][m]`.
 - a) What is `m` (local? global? function parameter?)
 - b) Where is `m` defined?
 - c) How many times is `m` used in the project?
3. Find every function whose name contains the word `time`

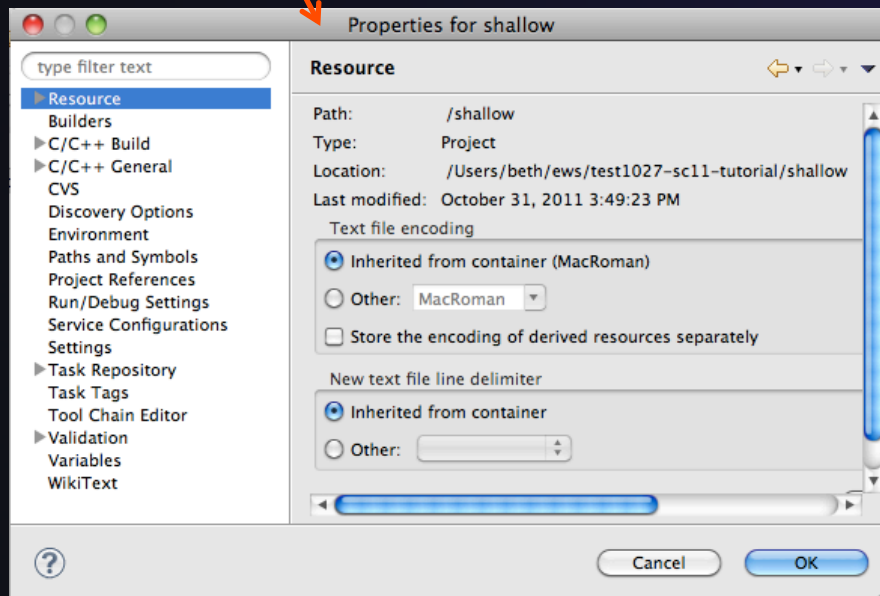
Fortran Specifics

Project Properties

- ★ Right-click Project
- ★ Select **Properties...**



- ★ *Project properties* are settings that can be changed for each project



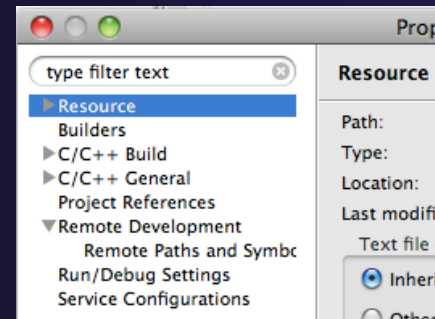
- ★ Contrast with *workspace preferences*, which are the same regardless of what project is being edited

- ★ e.g., editor colors
- ★ Set in **Window ► Preferences** (on Mac, **Eclipse ► Preferences**)
- ★ Careful! Dialog is very similar

Converting to a Fortran Project

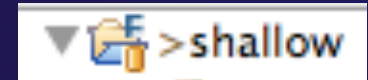
- ★ Are there categories labeled **Fortran General** and **Fortran Build** in the project properties?

No Fortran categories →



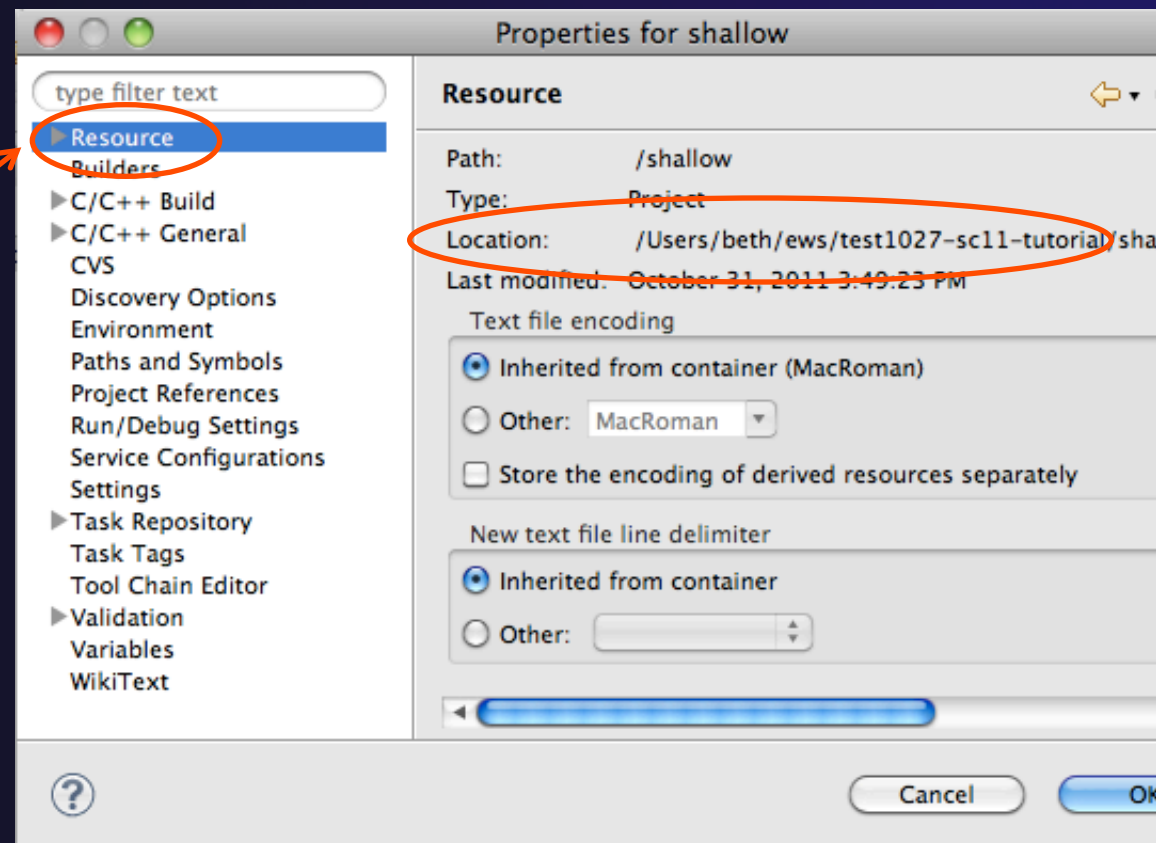
Do this once

- ★ If not, the project is not a Fortran Project
 - ★ Switch to the Fortran Perspective
 - ★ In the Fortran Projects view, right-click on the project, and click **Convert to Fortran Project**
 - ★ Don't worry; it's still a C/C++ project, too
- ★ *Every* Fortran project is also a C/C++ Project



Project Location

- ★ How to tell where a project resides?
- ★ In the project properties dialog, select the **Resource** category

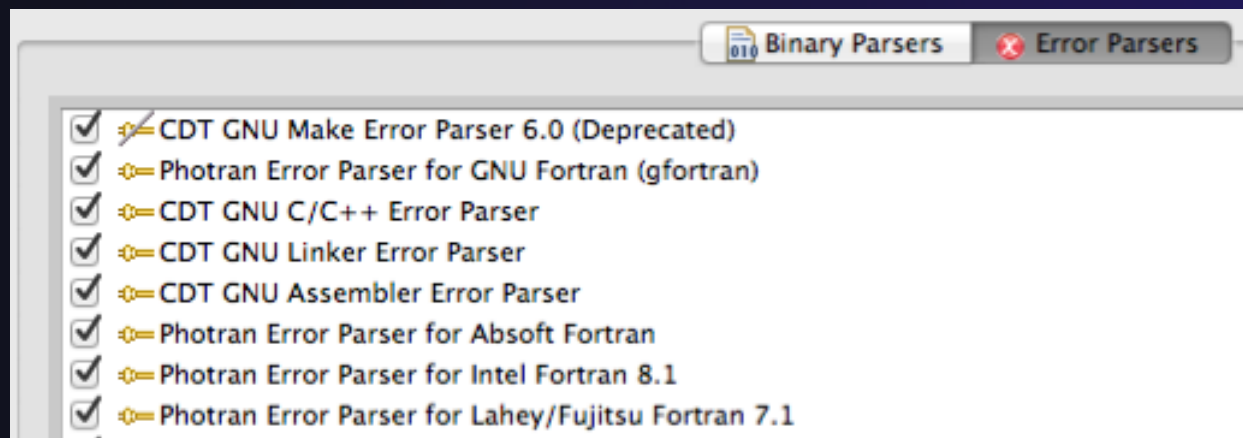


Error Parsers

- ★ Are compiler errors not appearing in the Problems view?
 - ★ Make sure the correct *error parser* is enabled
 - ★ In the project properties, navigate to **C++ Build ► Settings** or **Fortran Build ► Settings**
 - ★ Switch to the **Error Parsers** tab
 - ★ Check the error parser(s) for your compiler(s)



Do this
once



Fortran Source Form Settings

- ★ Fortran files are either *free form* or *fixed form*; some Fortran files are *preprocessed* (#define, #ifdef, etc.)

- ★ Source form determined by filename extension

- ★ Defaults are similar to most Fortran compilers:

Fixed form: .f .fix .for .fpp .ftn .f77

Free form: .f08 .f03 .f95 .f90 < unpreprocessed
 .F08 .F03 .F95 .F90 < preprocessed

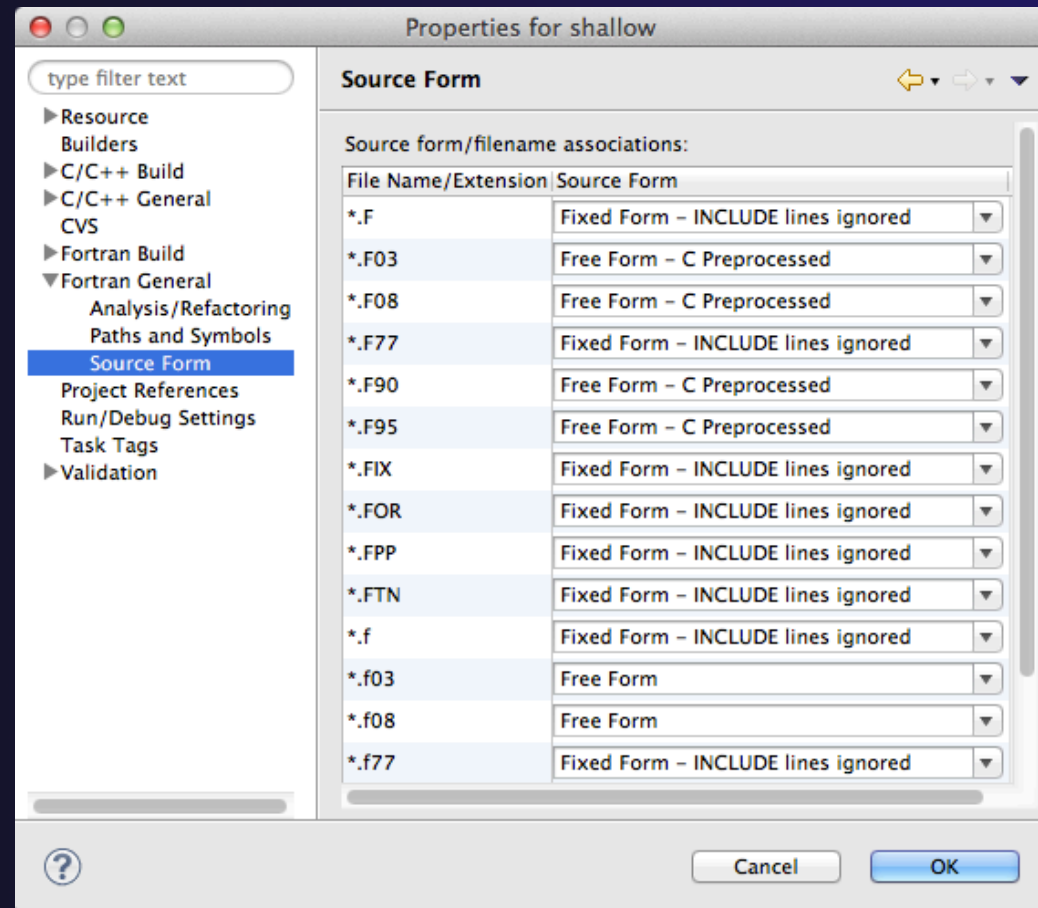
- ★ Many features *will not work* if filename extensions are associated with the wrong source form (outline view, content assist, search, refactorings, etc.)

Fortran Source Form Settings



Do this
once

- ★ In the project properties, select **Fortran General** ▶ **Source Form**
- ★ Select source form for each filename extension
- ★ Click **OK**

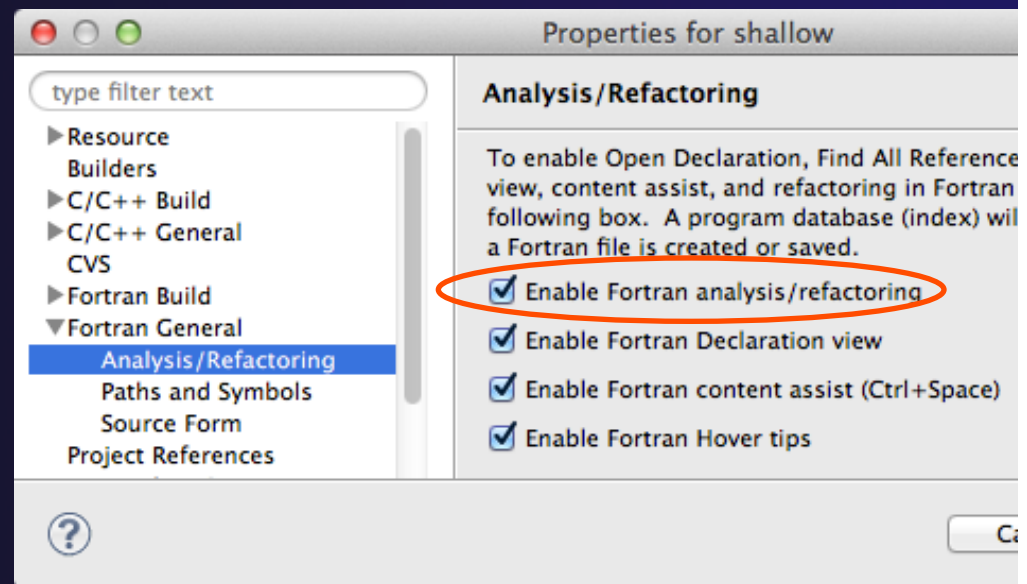


Enabling Fortran Advanced Features

- ★ Some Fortran features are *disabled* by default
- ★ Must be explicitly enabled
 - ★ In the project properties dialog, select **Fortran General ▶ Analysis/Refactoring**
 - ★ Click **Enable Analysis/Refactoring**
 - ★ Close and re-open any Fortran editors
- ★ This turns on the “Photran Indexer”
 - ★ Turn it off if it’s slow



Do this once





Project Properties – Try It!

1. Convert shallow to a Fortran project
2. Make sure errors from the GNU Fortran compiler will be recognized
3. Make sure *.f90 files are treated as “Free Form” which is unpreprocessed
4. Make sure search and refactoring will work in Fortran

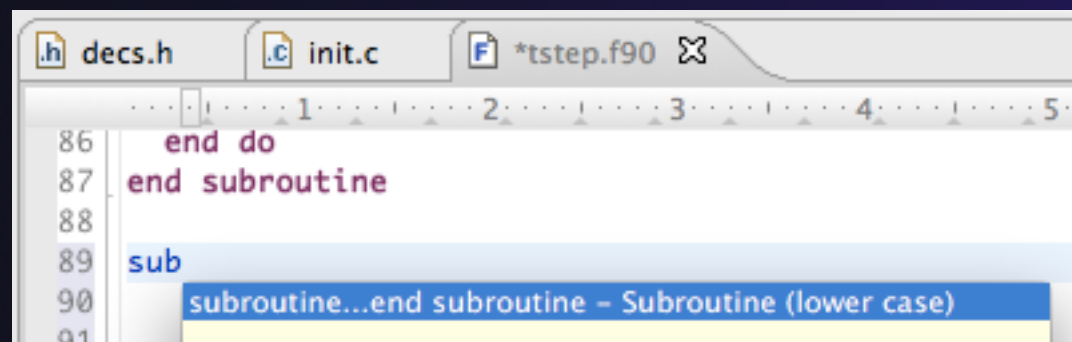
Advanced Editing

Code Templates

Code Templates

(C/C++ and Fortran)

- ★ Auto-complete common code patterns
 - ★ For loops/do loops, if constructs, etc.
 - ★ Also MPI code templates
- ★ Included with content assist proposals (when **Ctrl-Space** is pressed)
 - ★ E.g., after the last line in tstep.f90, type "sub" and press **Ctrl-Space**
 - ★ Press **Enter** to insert the template

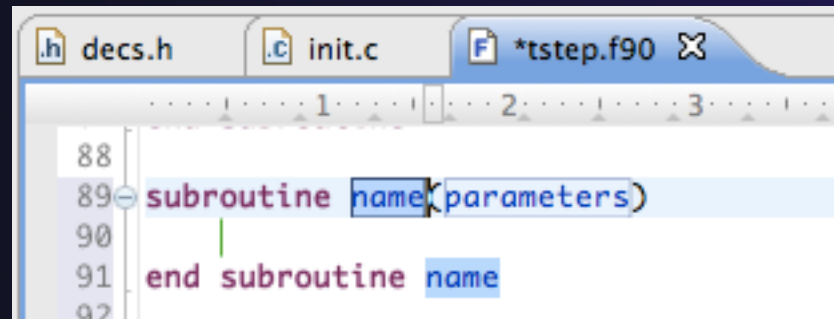


The screenshot shows an IDE window with three tabs: 'decs.h', 'init.c', and '*tstep.f90'. The active file is '*tstep.f90'. The code editor shows lines 86 to 91. Line 86 is 'end do', line 87 is 'end subroutine', and line 88 is blank. Line 89 has 'sub' typed. A dropdown menu is open below line 89, showing a suggestion: 'subroutine...end subroutine - Subroutine (lower case)'. The suggestion is highlighted in blue. Line 90 is blank, and line 91 is blank.

Code Templates (2)

(C/C++ and Fortran)

- ✦ After pressing enter to insert the code template, completion fields are highlighted



The screenshot shows a code editor window with three tabs: 'decs.h', 'init.c', and '*tstep.f90'. The active tab is '*tstep.f90'. The code is as follows:

```
88  
89 subroutine name(parameters)  
90  
91 end subroutine name  
92
```

The words 'name' and 'parameters' in line 89, and 'name' in line 91, are highlighted in blue, indicating they are completion fields.

- ✦ Press **Tab** to move between completion fields
- ✦ Changing one instance of a field changes all occurrences



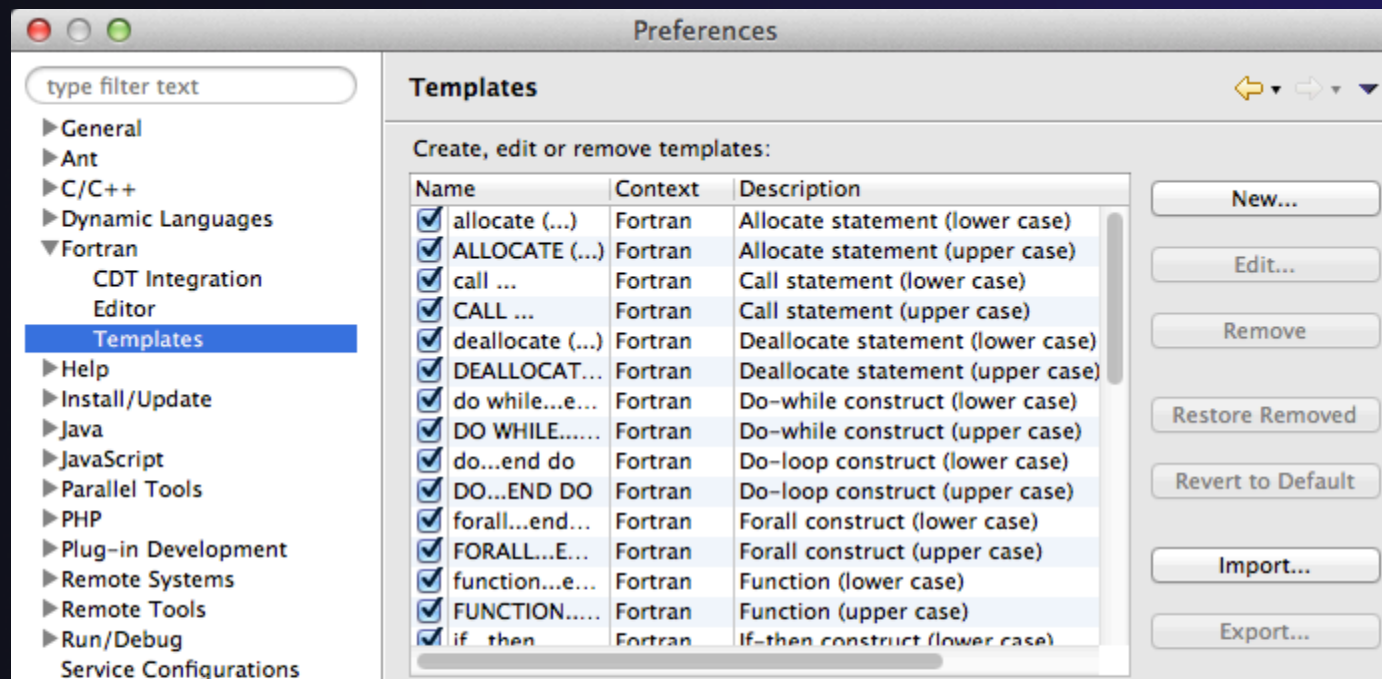
Advanced Editing – Try It!

- ✦ Open `tstep.f90` and retype the last loop nest
 - ✦ Use the code template to complete the do-loops
 - ✦ Use content assist to complete variable names

Custom Code Templates

(Fortran)

- ★ Customize code templates in **Window ▶ Preferences ▶ Fortran ▶ Templates**

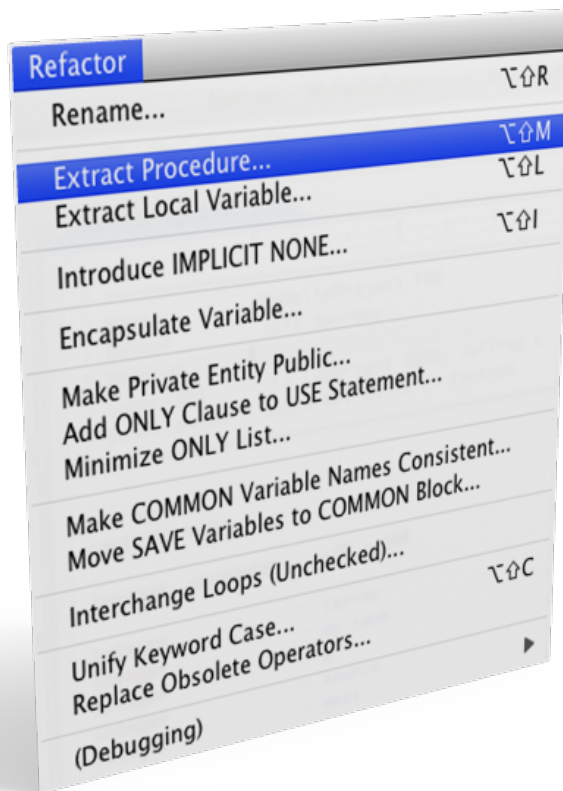


- ★ Can import/export templates to XML files

Refactoring and Transformation

Refactoring

(making changes to source code that don't affect the behavior of the program)



- ★ Refactoring is the research motivation for Photran @ Illinois
 - ★ Illinois is a leader in refactoring research
 - ★ “Refactoring” was coined in our group (Opdyke & Johnson, 1990)
 - ★ We had the first dissertation... (Opdyke, 1992)
 - ★ ...and built the first refactoring tool... (Roberts, Brant, & Johnson, 1997)
 - ★ ...and first supported the C preprocessor (Garrido, 2005)
 - ★ Photran’s agenda: refactorings for HPC, language evolution, refactoring framework
- ★ Photran 7.0: 31 refactorings

Refactoring Caveats

- ★ Photran can only refactor free form code that is *not* preprocessed

- ★ Determined by Source Form settings

(recall from earlier that these are configured in

Project Properties: Fortran General ▶ Source Form)

✓	Free Form, Unpreprocessed:	.f08	.f03	.f95	.f90		
✗	Free Form, Preprocessed:	.F08	.F03	.F95	.F90		
✗	Fixed Form:	.f	.fix	.for	.fpp	.ftn	.f77

- ★ Refactor menu will be empty if

- ★ Refactoring not enabled in project properties

(recall from earlier that it is enabled in

Project Properties: Fortran General ▶ Analysis/Refactoring)

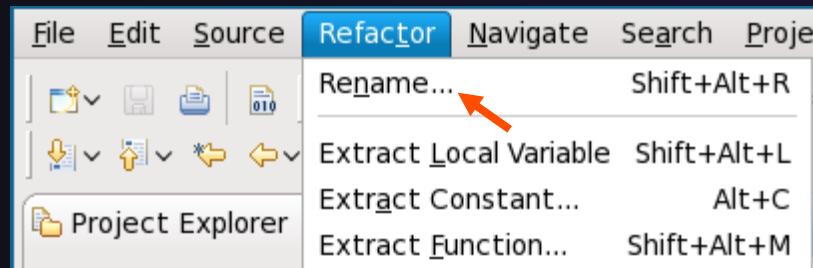
- ★ The file in the active editor is fixed form

- ★ The file in the active editor is preprocessed

Rename Refactoring

(also available in Fortran)

- ✦ Changes the name of a variable, function, etc., *including every use*
(change is semantic, not textual, and can be workspace-wide)
- ✦ Only proceeds if the new name will be legal
(aware of scoping rules, namespaces, etc.)



In Java (Murphy-Hill et al., ICSE 2008):

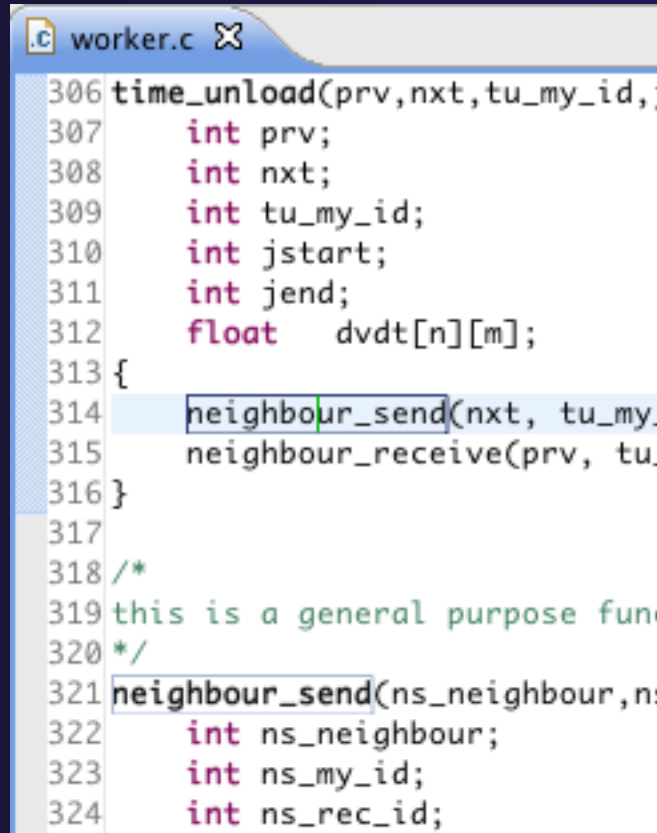
Refactoring	Uses	Percentage
Rename	179,871	74.8%
Extract Local Variable	13,523	5.6%
Move	13,208	5.5%
Extract Method	10,581	4.4%
Change Method Signature	4,764	2.0%
Inline	4,102	1.7%
Extract Constant	3,363	1.4%
(16 Other Refactorings)	10,924	4.5%

- ✦ Switch to C/C++ Perspective
- ✦ Open a source file
- ✦ In the editor, click on a variable or function name
- ✦ Select menu item **Refactor ▶ Rename**
 - ✦ Or use context menu
- ✦ Enter new name

Rename in File

(C/C++ Only)

- ★ Position the caret over an identifier.
- ★ Press **Ctrl-1** (**Command-1** on Mac).
- ★ Enter a new name. Changes are propagated within the file as you type.

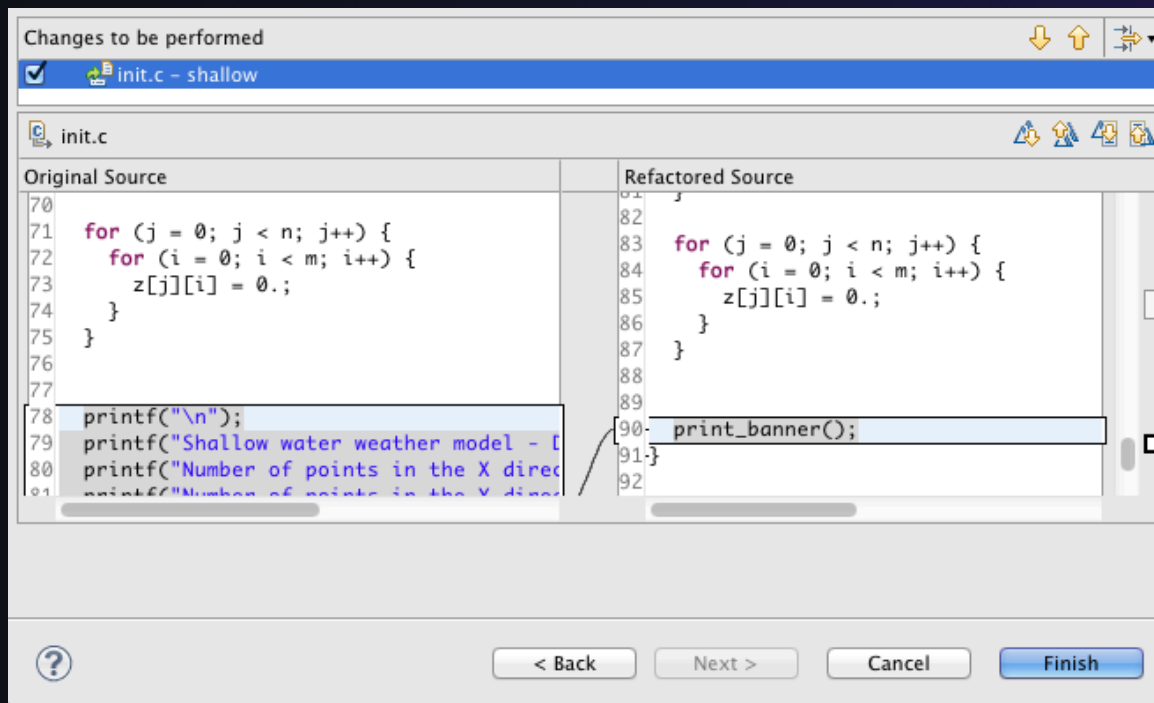


```
worker.c X
306 time_unload(prv,nxt,tu_my_id,
307     int prv;
308     int nxt;
309     int tu_my_id;
310     int jstart;
311     int jend;
312     float  dvdt[n][m];
313 {
314     neighbour_send(nxt, tu_my.
315     neighbour_receive(prv, tu.
316 }
317
318 /*
319 this is a general purpose fun
320 */
321 neighbour_send(ns_neighbour,n
322     int ns_neighbour;
323     int ns_my_id;
324     int ns_rec_id;
```

Extract Function Refactoring

(also available in Fortran - "Extract Procedure")

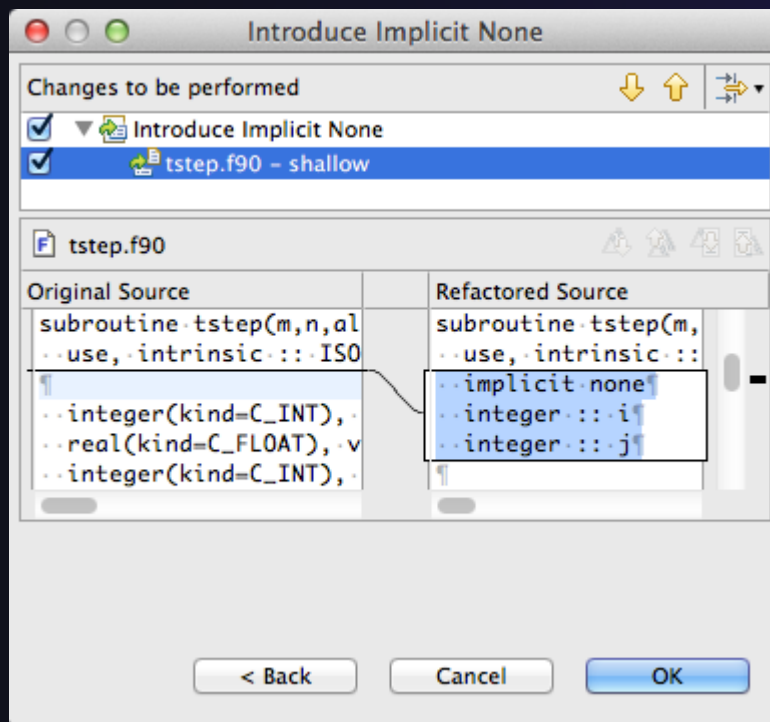
- ✦ Moves statements into a new function, replacing the statements with a call to that function
- ✦ Local variables are passed as arguments



- ✦ Select a sequence of statements
- ✦ Select menu item **Refactor** ► **Extract Function...**
- ✦ Enter new name

Introduce IMPLICIT NONE Refactoring

- ★ Fortran does not require variable declarations
(by default, names starting with I-N are integer variables; others are reals)
- ★ This adds an IMPLICIT NONE statement and adds explicit variable declarations for all implicitly declared variables

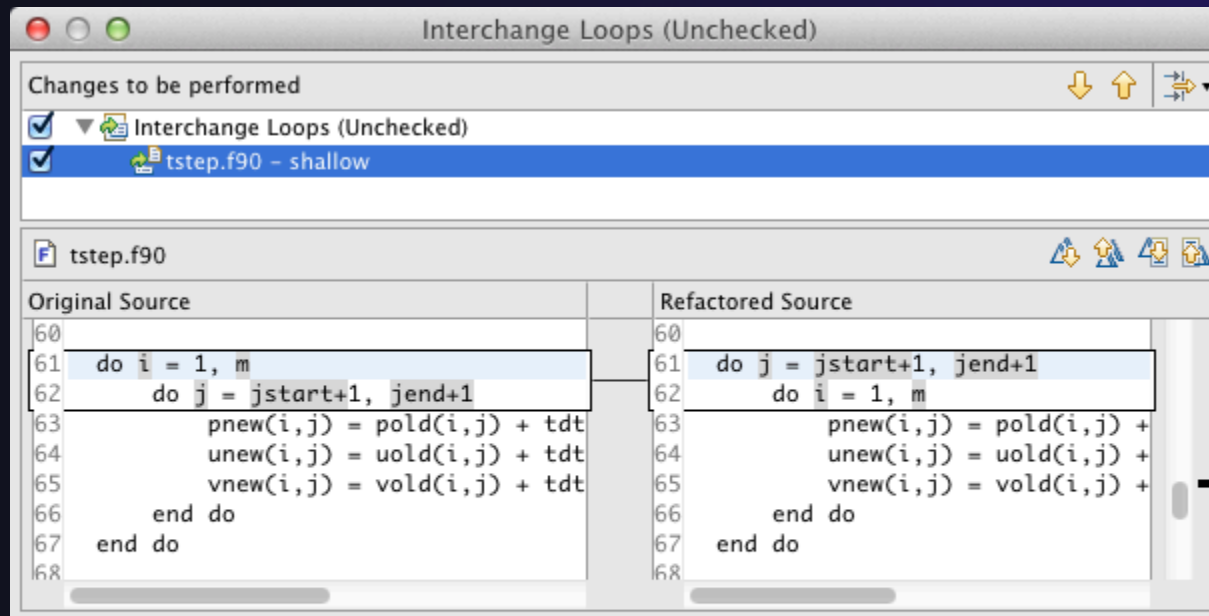


- ★ Introduce in a single file by opening the file and selecting **Refactor ▶ Coding Style ▶ Introduce IMPLICIT NONE...**
- ★ Introduce in multiple files by selecting them in the Fortran Projects view, right-clicking on the selection, and choosing **Refactor ▶ Coding Style ▶ Introduce IMPLICIT NONE...**

Loop Transformations

(Fortran only)

- ✦ **Interchange Loops** **CAUTION:** No check for behavior preservation
 - ✦ Swaps the loop headers in a two-loop nest
 - ✦ Select the loop nest, click menu item **Refactor ▶ Do Loop ▶ Interchange Loops (Unchecked)...**



Old version traverses
matrices in row-major order

New version traverses
in column-major order
(better cache performance)

Loop Transformations

(Fortran only)

★ Unroll Loop

★ Select a loop, click **Refactor** ▶ **Do Loop** ▶ **Unroll Loop...**

```
do i = 1, 10
  print *, 10*i
end do
```



Unroll 4x

```
do i = 1, 10, 4
  print *, 10*i
  print *, 10*(i+1)
  print *, 10*(i+2)
  print *, 10*(i+3)
end do
```

Original Source	Refactored Source
68	78
69 ! Don't apply time filter on first	79 end if
70 if (firststep == 0) then	80
71 do j = jstart+1, jend+1	81 do j = jstart+1, jend+1
72 do i = 1, m	82 LoopUpperBound = m
73 pold(i,j) = p(i,j)+alpha*(pne	83 do i = 1, LoopUpperBound,4
74 uold(i,j) = u(i,j)+alpha*(une	84 p(i,j) = pnew(i,j)
75 vold(i,j) = v(i,j)+alpha*(vne	85 u(i,j) = unew(i,j)
76 end do	86 v(i,j) = vnew(i,j)
77 end do	87 p((i+1),j) = pnew((i+1)
78 end if	88 u((i+1),j) = unew((i+1)
79	89 v((i+1),j) = vnew((i+1)
80 do j = jstart+1, jend+1	90 p((i+2),j) = pnew((i+2)
81 do i = 1, m	91 u((i+2),j) = unew((i+2)
82 p(i,j) = pnew(i,j)	92 v((i+2),j) = vnew((i+2)
83 u(i,j) = unew(i,j)	93 p((i+3),j) = pnew((i+3)
84 v(i,j) = vnew(i,j)	94 u((i+3),j) = unew((i+3)
85 end do	95 v((i+3),j) = vnew((i+3)
86 end do	96 end do
87 end subroutine	97 end do
88	98 end subroutine
	99
	00



Refactoring & Transformation – Try It!

In `tstep.f90`...

1. In `init.c`, extract the `printf` statements at the bottom of the file into a new function called `print_banner`
2. In `worker.c`, change the spellings of `neighbour_send` and `neighbour_receive` to American English
3. In `tstep.f90`, make the (Fortran) `tstep` subroutine `IMPLICIT NONE`

Module 4: Parallel Debugging

✦ Objective

- ✦ Learn the basics of debugging parallel programs

✦ Contents

- ✦ Launching a debug session
- ✦ The Parallel Debug Perspective
- ✦ Controlling sets of processes
- ✦ Controlling individual processes
- ✦ Parallel Breakpoints
- ✦ Terminating processes

Debugging Setup

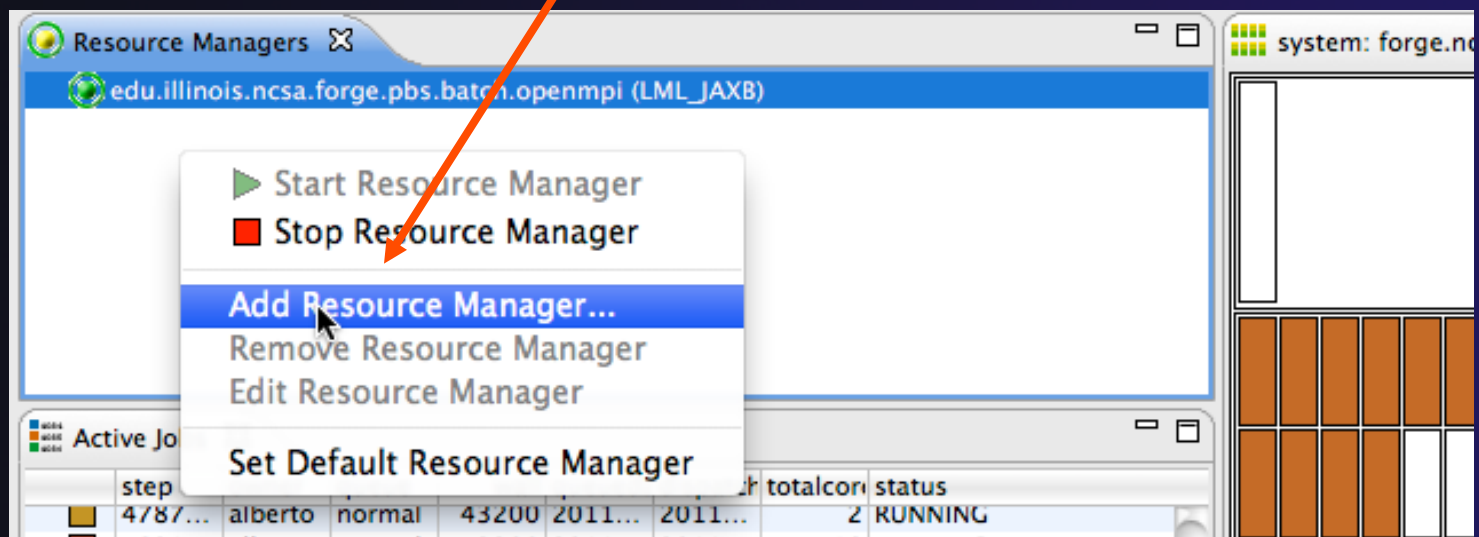
Debugging Setup

- ★ Debugging requires interactive access to the application
- ★ Since the PBS resource manager we set up previously is for batch execution, we need to set up a new one
- ★ We will use a **PBS interactive** resource manager
- ★ *This only needs to be done once*



Adding a Resource Manager

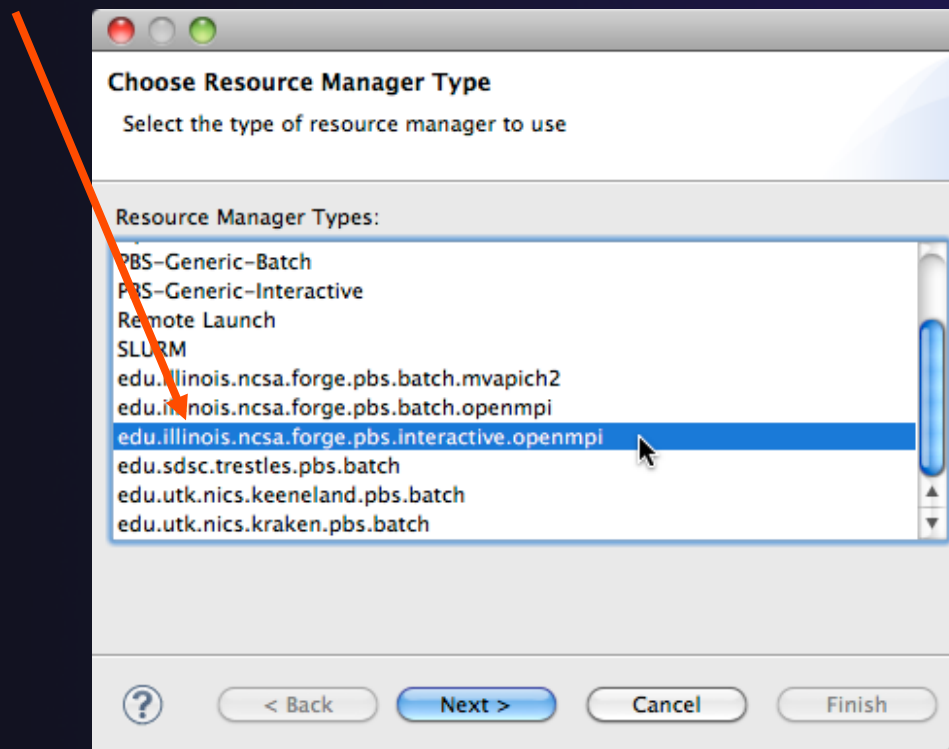
- ★ Switch to the **System Monitoring Perspective**
- ★ Right-click in **Resource Managers** view and select **Add Resource Manager...**





Selecting the Resource Manager

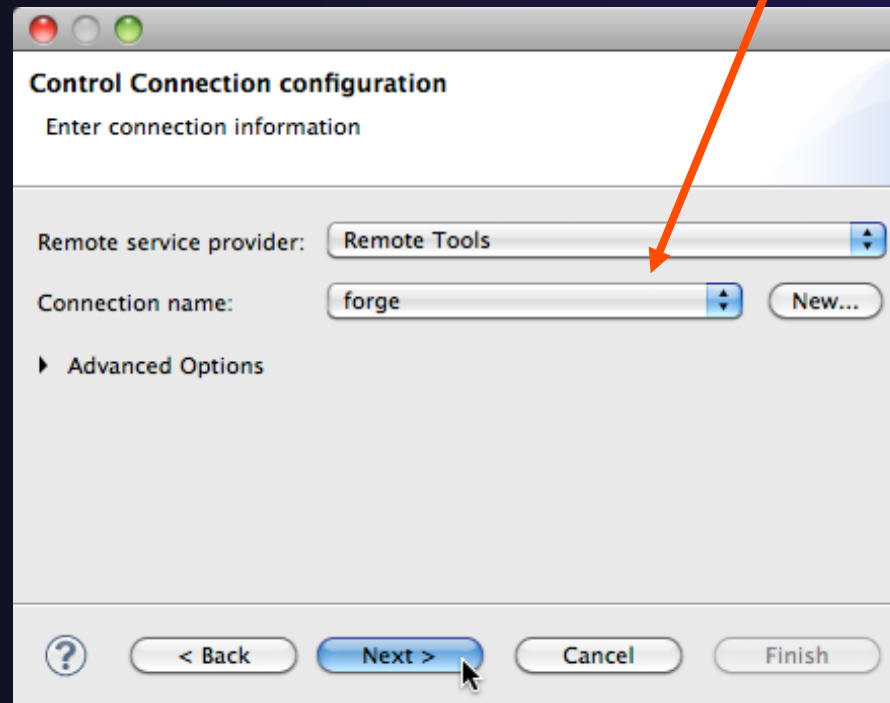
- ★ Choose the PBS interactive resource manager type that has been pre-configured for forge
- ★ **edu.illinois.ncsa.forge.pbs.interactive.openmpi**





Configuring the Resource Manager

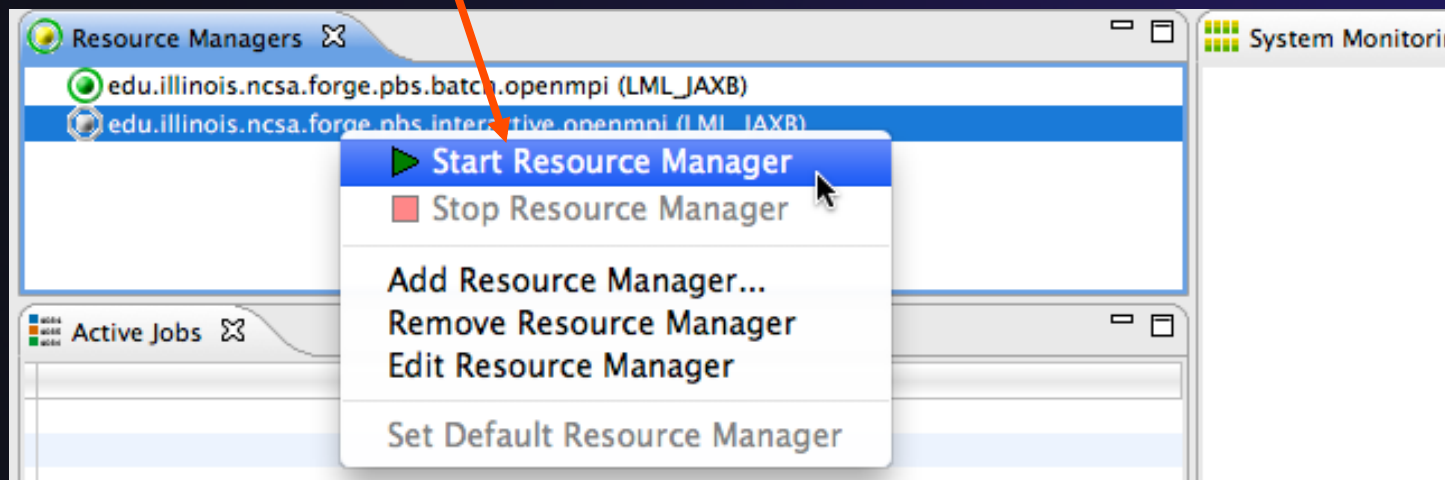
- ★ Choose **Remote Tools** for the **Remote service provider**
- ★ Choose the connection you created previously
- ★ Click **Next>** then **Finish**





Start the Resource Manager

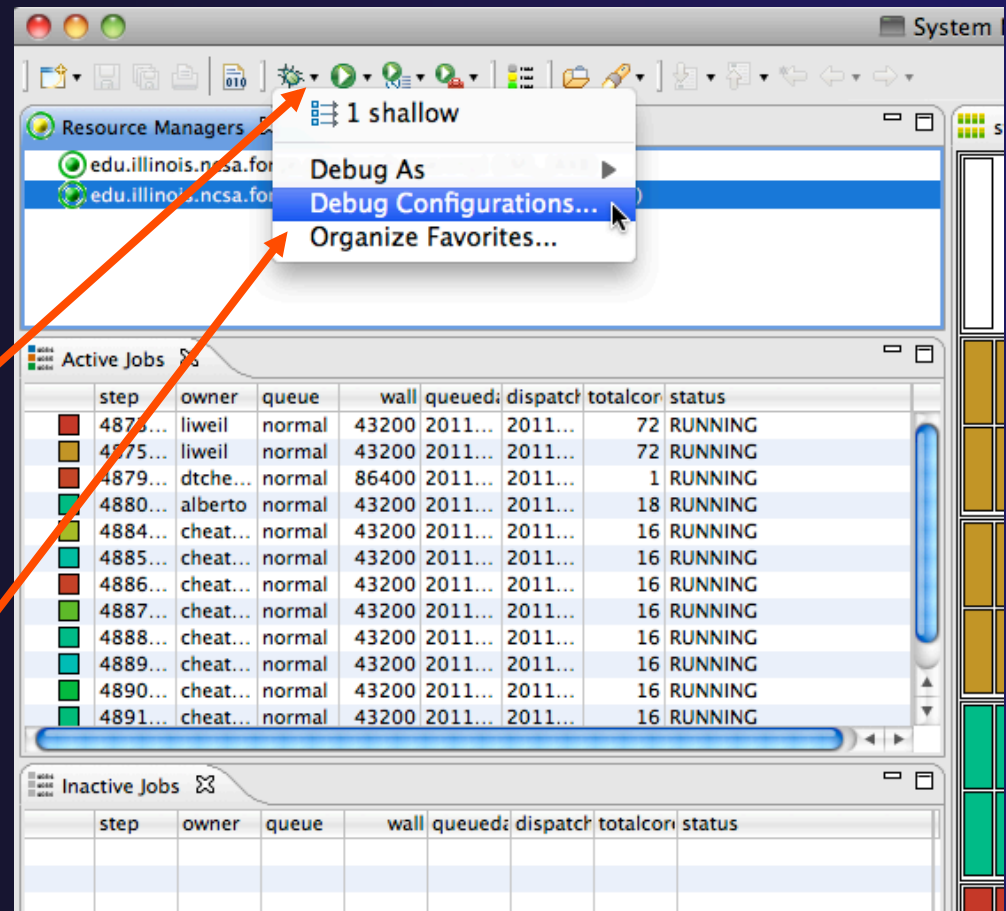
- ★ Right-click on the new resource manager and select **Start Resource Manager** from the menu



Create a Debug Configuration



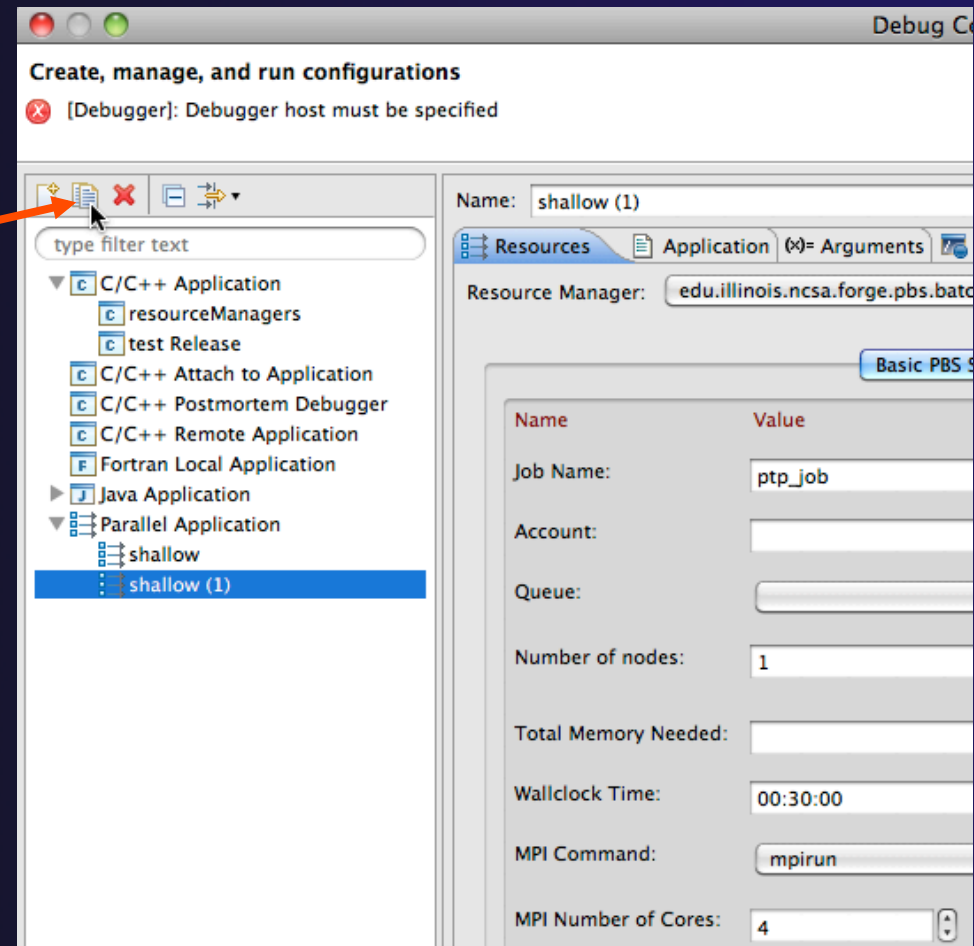
- ★ A debug configuration is essentially the same as a run configuration (like we used in module 3)
- ★ We will re-use the existing configuration and add debug information
- ★ Use the drop-down next to the debug button (bug icon) instead of run button
- ★ Select **Debug Configurations...** to open the **Debug Configurations** dialog



Copy the Existing Configuration



- ★ Select the existing configuration
- ★ Click on the **copy** button to create a duplicate configuration





Configure the Resource Tab

- ★ Select the new resource manager
- ★ Choose the **debug** queue
- ★ Choose the **mpirun** command
- ★ Select the number of cores (in this case use 4)

Debug Configurations

Create, manage, and run configurations

[Debugger]: Debugger host must be specified

Name: shallow (1)

Resources Application Arguments Environment Synchronize Debugger Source Common

Resource Manager: edu.illinois.ncsa.forge.pbs.interactive.openmpi

Basic PBS Settings

Queue: debug

Name	Value	Description
Account:		Account to which to charge this job.
Total Memory Needed:	20gb	Maximum amount of memory used by all concurrent processes in the job.
Wallclock Time:	00:30:00	Maximum amount of real time during which the job can be in the running state.
MPI Command:	mpirun	Which mpi command to use.
MPI Number of Cores:	4	the '-np' value [usually equals Nodes*ppn]

View Configuration Restore Defaults

Filter matched 12 of 20 items

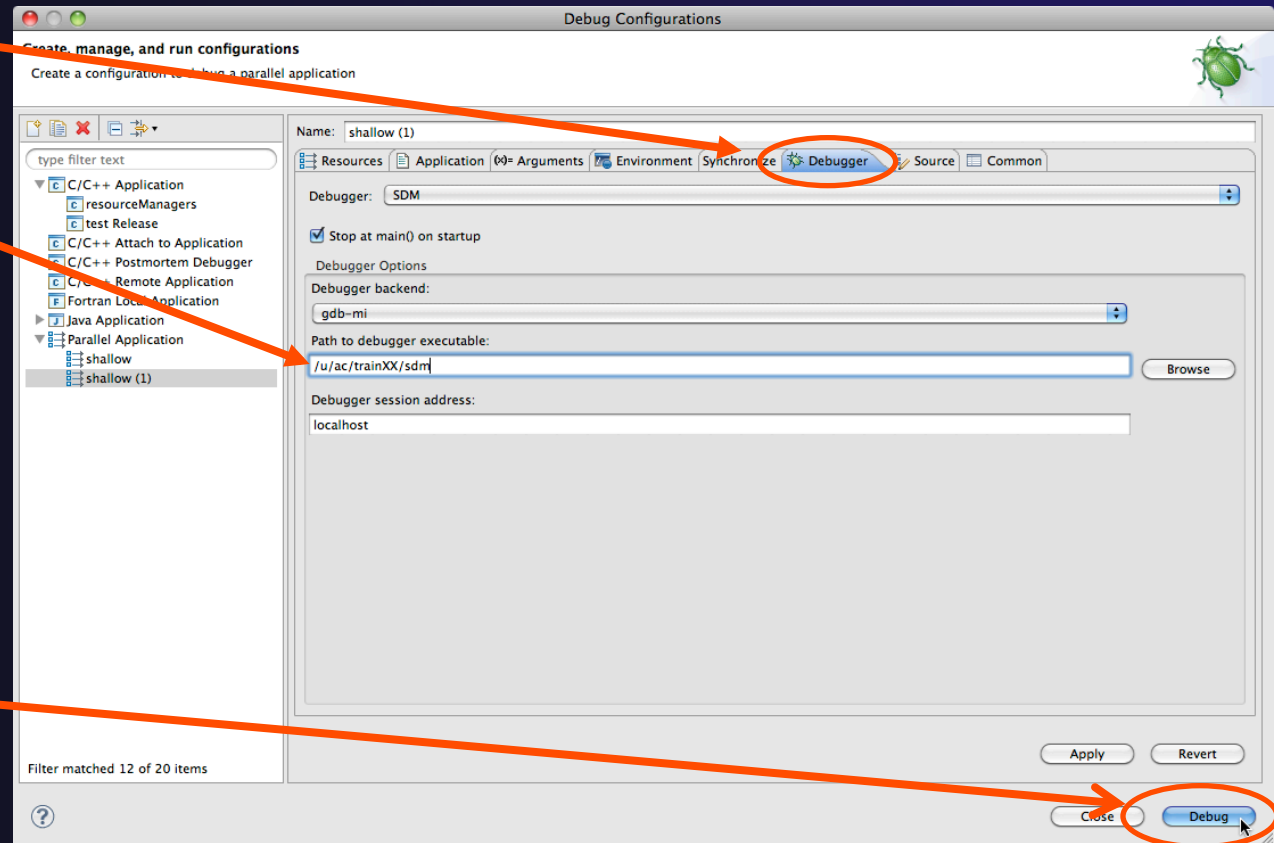
Apply Revert

Close Debug



Configure the Debug Tab

- ★ Select **Debugger** tab
- ★ Set the debugger path to the **sdm** in your home directory
- ★ Debugger session address should not need to be changed
- ★ Click on **Debug** to launch the program



Debugging the Application

The Parallel Debug Perspective (1)

- ★ **Parallel Debug view** shows job and processes being debugged
- ★ **Debug view** shows threads and call stack for individual processes
- ★ **Source view** shows a **current line marker** for all processes

Parallel Debug - shallow/main.c - Eclipse - /Users/greg/testing/workspace

Parallel Debug

Open_MPI@abe.ncsa.uiuc.edu: default:job0 - Root [4]

job0

Debug

shallow [Parallel Application]

Process 0

Thread [1] (Suspended)

1 main() main.c:38 4032ca

main.c

```

32 MPI_Datatype * setup_res();
33
34 main (argc, argv)
35     int argc;
36     char * argv[];
37 {
38     float pi=4.*((float)atan((double)1.));
39     float p[n][m]; /* Pressure (or free surface height) */
40     float u[n][m]; /* Zonal wind */
41     float v[n][m]; /* Meridional wind */
42     float psi[n][m]; /* Velocity streamfunction */
43     float pold[n][m];
44     float uold[n][m];
45     float vold[n][m];
46     float h[n][m];
47     float z[n][m];
48     float dummy1[m];
49     float dummy2[n][m];
50     float tpi=pi+pi;
51     float di=tpi/(float)m;
52     float dj=tpi/(float)n;

```

Name	Value
argc	1
argv	7fffffff658
pi	5.957805E-39
p	[0 - 18]
u	[0 - 18]
v	[0 - 18]
psi	[0 - 18]

Outline

- math.h
- mpi.h
- stdio.h
- decs.h
- worker(): void
- setup_res(): MPI_Datatype*
- main(int, char*[])
- setup_res(): MPI_Datatype*
- update_global_ds(MPI_Datatype)

Console

Memory

Problems

Open_MPI@abe.ncsa.uiuc.edu:default:job0

Open Mpi Job

The Parallel Debug Perspective (2)

- ★ **Breakpoints** view shows breakpoints that have been set (more on this later)
- ★ **Variables** view shows the current values of variables for the currently selected process in the **Debug** view
- ★ **Outline** view (from CDT) of source code

The screenshot displays the Eclipse IDE in the Parallel Debug perspective. The top toolbar shows various debugging tools. The Breakpoints view is active, showing a list of breakpoints. The Variables view displays the current values of variables for the selected process (main() main.c:38 4032ca). The Debug view shows the current state of the application, including the selected process and thread. The Outline view shows the source code structure, with the current line of code highlighted. The console at the bottom shows the output of the application.

Name	Value
argc	1
argv	7fffffff658
pi	5.957805E-39
p	[0 - 18]
u	[0 - 18]
v	[0 - 18]
psi	[0 - 18]

```

32 MPI_Datatype * setup_res();
33
34 main (argc, argv)
35     int argc;
36     char * argv[];
37 {
38     float pi=4.*(float)atan((double)1.);
39     float p[n][m]; /* Pressure (or free surface height) */
40     float u[n][m]; /* Zonal wind */
41     float v[n][m]; /* Meridional wind */
42     float psi[n][m]; /* Velocity streamfunction */
43     float pold[n][m];
44     float uold[n][m];
45     float vold[n][m];
46     float h[n][m];
47     float dummy1[n][m];
48     float dummy2[n][m];
49     float tpi=pi+pi;
50     float di=tpi/(float)m;
51     float dj=tpi/(float)n;
52

```



Stepping All Processes

- ★ The buttons in the **Parallel Debug View** control groups of processes
- ★ Click on the **Step Over** button
- ★ Observe that all process icons change to green, then back to yellow
- ★ Notice that the current line marker has moved to the next source line

The screenshot displays the Eclipse IDE interface during a parallel debug session. The top toolbar includes various debugging controls, with the 'Step Over' button (represented by a right-pointing arrow) highlighted by an orange arrow. Below the toolbar, the 'Parallel Debug' view shows a tree structure of processes, with the 'Step Over' button also highlighted by an orange arrow. The 'Debug' view shows the current process state, including 'shallow [Parallel Application]', 'Process 0 (Suspended)', and 'Thread [1] (Suspended)'. The source code editor shows the current line marker (a blue arrow) moved to line 50, which is highlighted in blue. The code defines several variables and includes a comment for 'Velocity streamfunction'.

```

38 float pi=4.*(float)atan((double)1.);
39 float p[n][m]; /* Pressure (or free surface height) */
40 float u[n][m]; /* Zonal wind */
41 float v[n][m]; /* Meridional wind */
42 float psi[n][m]; /* Velocity streamfunction */
43 float pold[n][m];
44 float uold[n][m];
45 float vold[n][m];
46 float h[n][m];
47 float z[n][m];
48 float dummy1[m];
49 float dummy2[n][m];
50 float tpi=pi+pi;
51 float di=tpi/(float)m;
52 float dj=tpi/(float)n;
53 int i, j, chunk_size, nxt, prv;
54
55 int master_packet[4];
56 float p_start[m];
57 float u_start[m];
58 float v_start[m];

```



Stepping An Individual Process

- ★ The buttons in the **Debug view** are used to control an individual process, in this case process 0
- ★ Click the **Step Over** button
- ★ You will now see two current line markers, the first shows the position of process 0, the second shows the positions of processes 1-3

Parallel Debug - shallow/main.c - Eclipse - /Users/greg/testing/works

Parallel Debug

Open_MPI@abe.ncsa.uiuc.edu: default.job0 - Root [4]

Job0 0

Debug

shallow [Parallel Application]

Process 0 (Suspended)

Thread [1] (Suspended)

1 main() main.c:51 403309

Name

- argc
- argv
- pi
- p
- u
- v
- psi

main.c

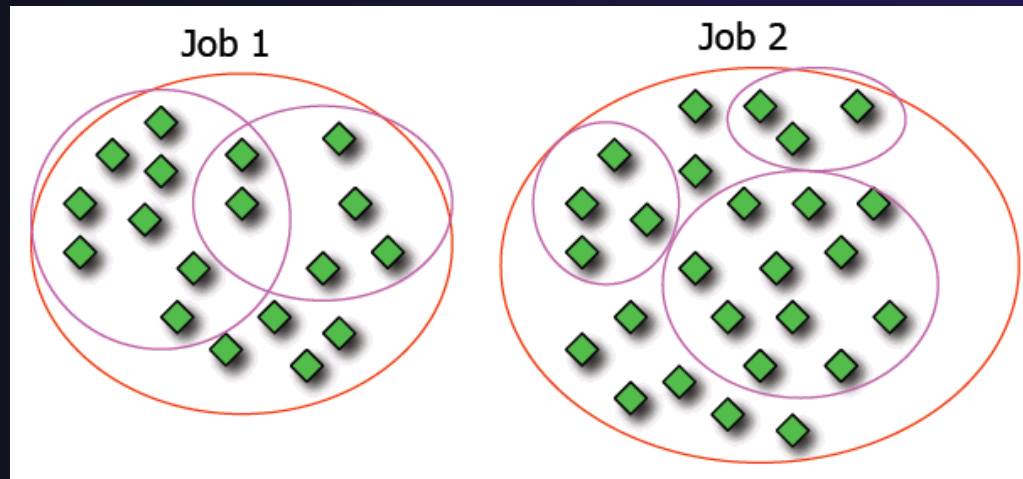
```

38 float pi=4.*(float)atan((double)1.);
39 float p[n][m]; /* Pressure (or free surface height) */
40 float u[n][m]; /* Zonal wind */
41 float v[n][m]; /* Meridional wind */
42 float psi[n][m]; /* Velocity streamfunction */
43 float pold[n][m];
44 float uold[n][m];
45 float vold[n][m];
46 float h[n][m];
47 float z[n][m];
48 float dummy1[n][m];
49 float dummy2[n][m];
50 float tpi=pi+pi;
51 float di=tpi/(float)m;
52 float dj=tpi/(float)n;
53 int i, j, chunk_size, nxt, prv;
54
55 int master_packet[4];
56 float p_start[m];
57 float u_start[m];

```

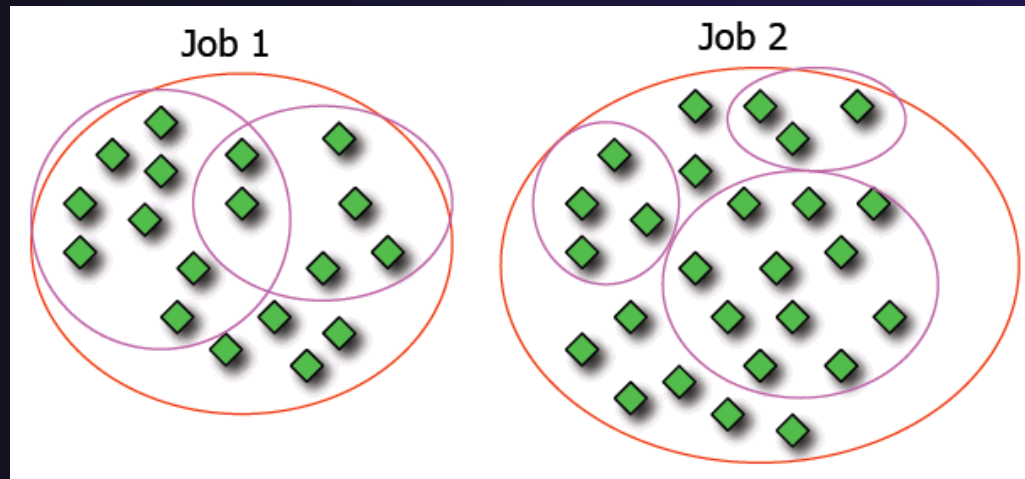

Process Sets (1)

- ★ Traditional debuggers apply operations to a single process
- ★ Parallel debugging operations apply to a single process or to arbitrary collections of processes
- ★ A process set is a means of simultaneously referring to one or more processes



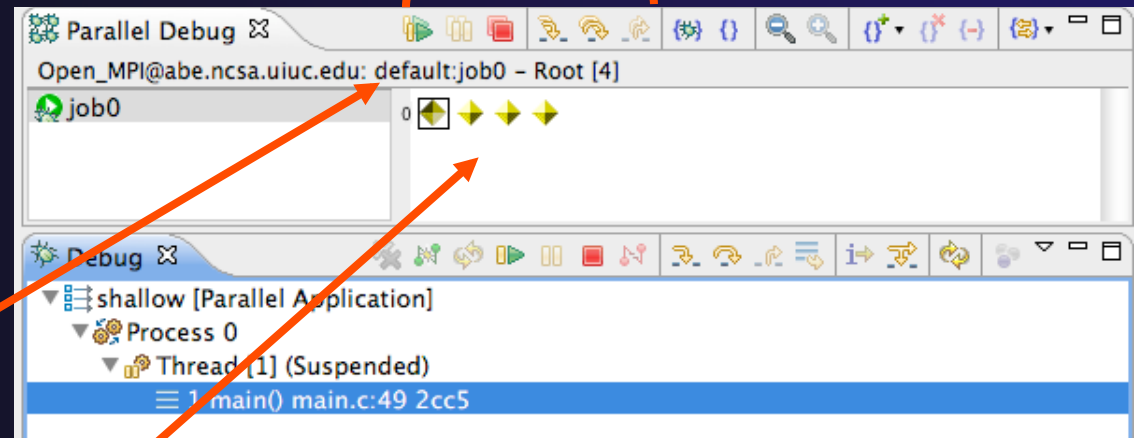
Process Sets (2)

- ★ When a parallel debug session is first started, all processes are placed in a set, called the **Root** set
- ★ Sets are always associated with a single job
- ★ A job can have any number of process sets
- ★ A set can contain from 1 to the number of processes in a job



Operations On Process Sets

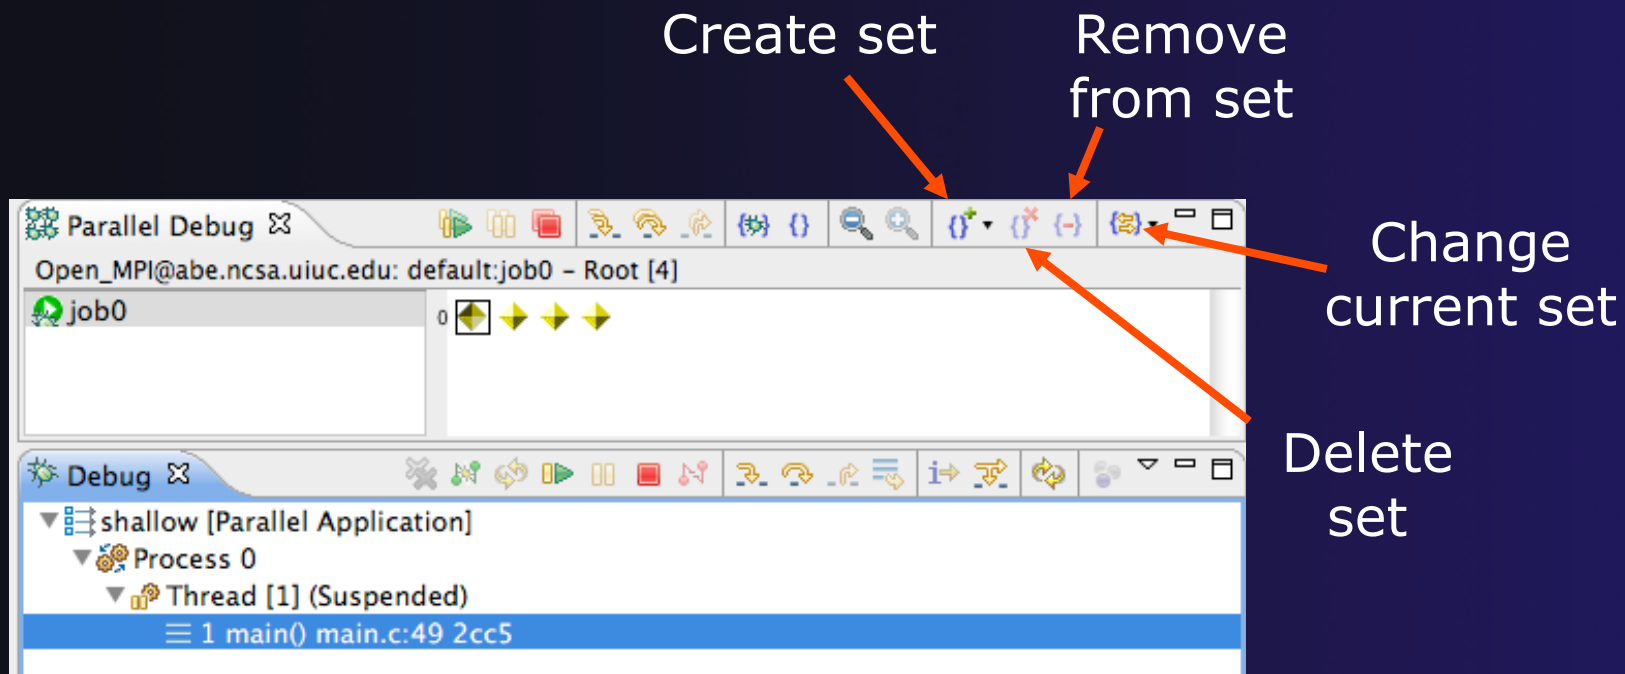
- ★ Debug operations on the **Parallel Debug view** toolbar always apply to the current set:
 - ★ Resume, suspend, stop, step into, step over, step return
- ★ The current process set is listed next to job name along with number of processes in the set
- ★ The processes in process set are visible in right hand part of the view



Root set = all processes

Managing Process Sets

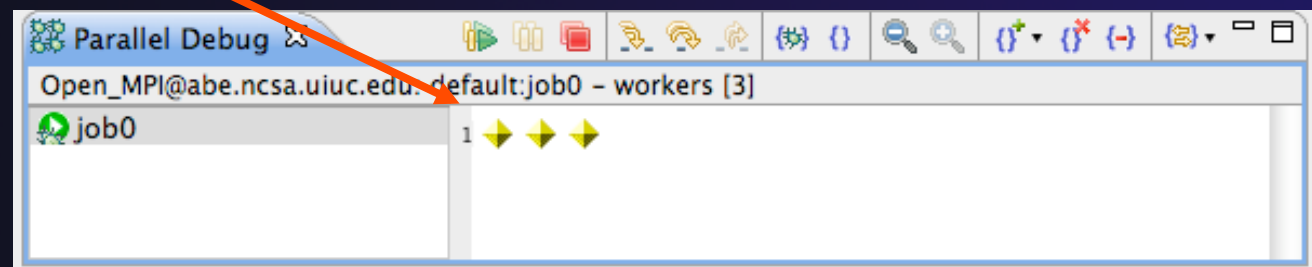
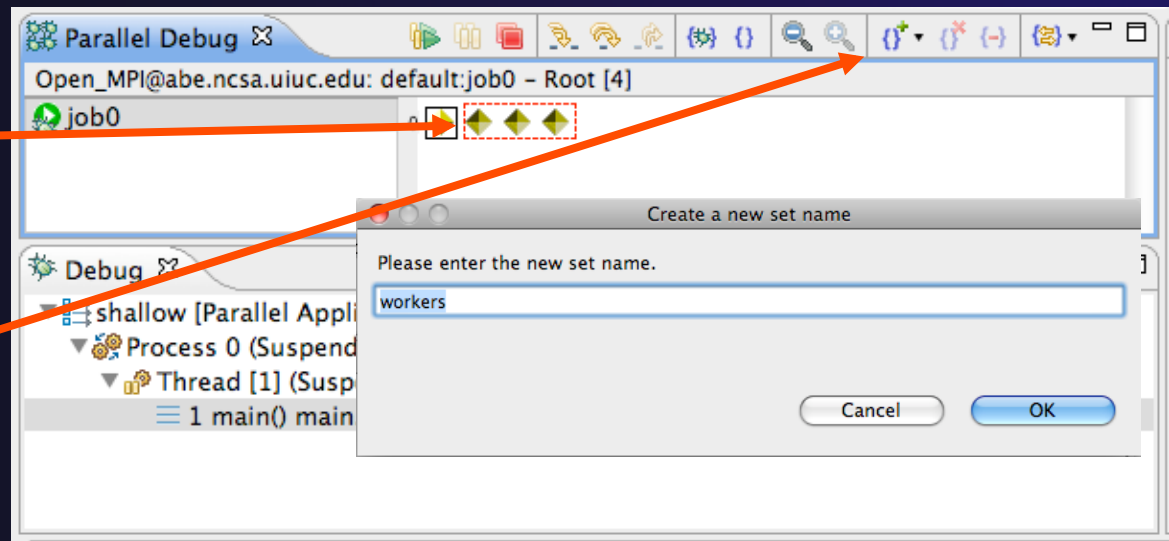
- ★ The remaining icons in the toolbar of the **Parallel Debug view** allow you to create, modify, and delete process sets, and to change the current process set





Creating A New Process Set

- ★ Select the processes you want in the set by clicking and dragging, in this case, the last three
- ★ Click on the **Create Set** button
- ★ Enter a name for the set, in this case **workers**, and click **OK**
- ★ You will see the view change to display only the selected processes



Stepping Using New Process Set



- ★ With the **workers** set active, click the **Step Over** button
- ★ You will see only the first current line marker move
- ★ Step a couple more times
- ★ You should see two line markers, one for the single master process, and one for the 3 worker processes

```

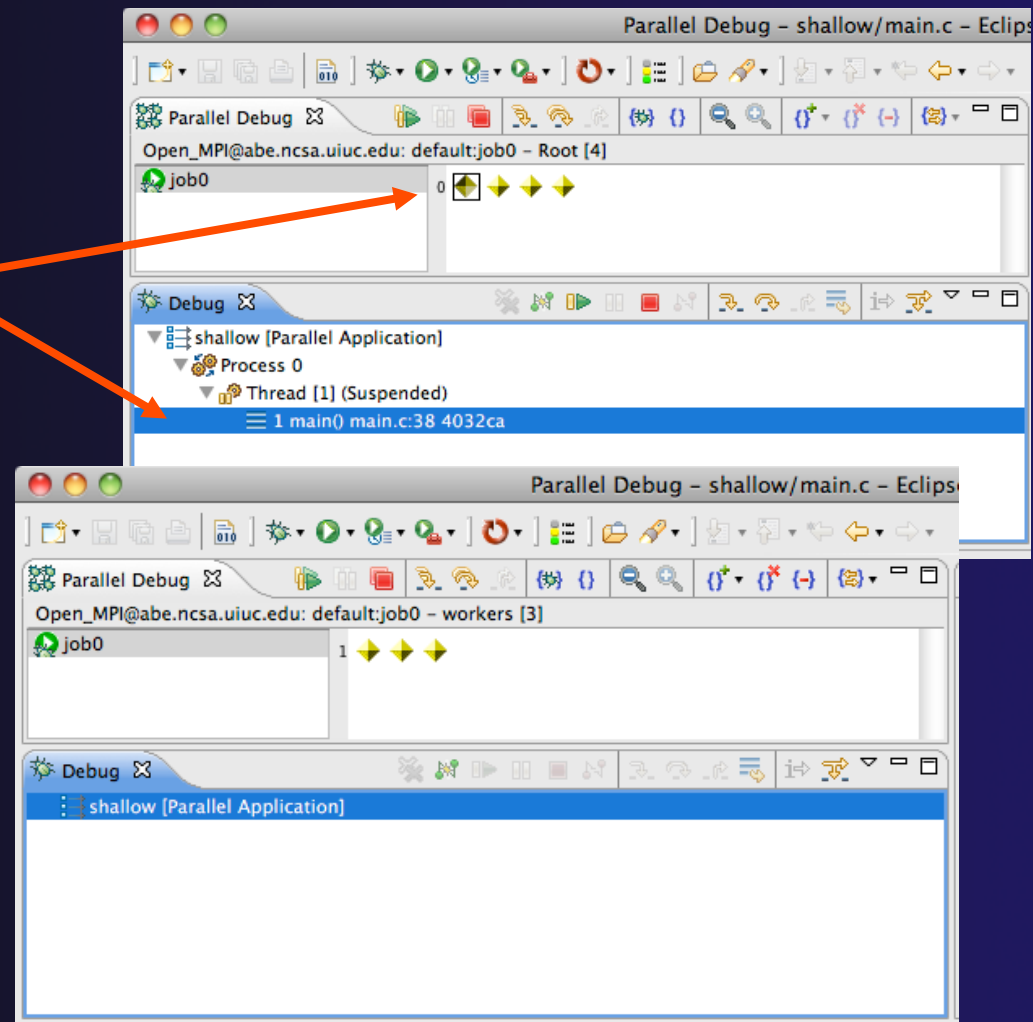
49 float dummyc[n][m];
50 float tpi=pi+pi;
51 float di=tpi/(float)m;
52 float dj=tpi/(float)n;
53 int i, j, chunk_size, nxt, prv;
54
55 int master_packet[4];
56 float p_start[m];
57 float u_start[m];
58 float v_start[m];
59 float psi_start[m];
60 float pold_start[m];
61 float uold_start[m];
62 float vold_start[m];
63 int proc_cnt;
64 int tid;
65 MPI_Datatype * res_type;
66
67 MPI_Init(&argc, &argv);
68 MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt); //hello
69 MPI_Comm_rank(MPI_COMM_WORLD, &tid);
  
```

Process Registration

- ✦ Process set commands apply to groups of processes
- ✦ For finer control and more detailed information, a process can be registered and isolated in the **Debug view**
- ✦ Registered processes, including their stack traces and threads, appear in the **Debug view**
- ✦ Any number of processes can be registered, and processes can be registered or un-registered at any time

Process Registration (2)

- ✦ By default, process 0 was registered when the debug session was launched
- ✦ Registered processes are surrounded by a box and shown in the Debug view
- ✦ The Debug view only shows registered processes in the current set
- ✦ Since the “workers” set doesn’t include process 0, it is no longer displayed in the Debug view



Registering A Process



- ★ To register a process, double-click its process icon in the **Parallel Debug view** or select a number of processes and click on the **register** button
- ★ To un-register a process, double-click on the process icon or select a number of processes and click on the **unregister** button

The screenshot displays the Parallel Debug IDE interface. The top toolbar contains various icons, including a green square with a white plus sign, which is the 'register' button. An orange arrow points from this button to the 'job0' process icon in the 'Parallel Debug' view. Another orange arrow points from the 'job0' process icon to the 'Process 3 (Suspended)' entry in the 'Debug' view. The 'Debug' view shows a tree structure with 'shallow [Parallel Application]' expanded to show 'Process 3 (Suspended)' and 'Thread [1] (Suspended)'. The 'Thread [1] (Suspended)' entry is highlighted in blue, and the text 'Individual (registered) processes' is written in red next to it. The bottom pane shows the source code for 'main.c', with line 67, 'MPI_Init(&argc, &argv);', highlighted in blue.

Parallel Debug - shallow/main.c -

Parallel Debug

Open_MPI@abe.ncsa.uiuc.edu: default:job0 - workers [3]

job0

Groups (sets) of processes

Debug

shallow [Parallel Application]

Process 3 (Suspended)

Thread [1] (Suspended)

1 main() main.c:67 403335

Individual (registered) processes

calc.c main.c test_global_bp.c

```

60 float pold_start[m];
61 float uold_start[m];
62 float vold_start[m];
63 int proc_cnt;
64 int tid;
65 MPI_Datatype * res_type;
66
67 MPI_Init(&argc, &argv);
68 MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt); //hello
69 MPI_Comm_rank(MPI_COMM_WORLD, &tid);
70
71 fprintf(stdout, "my rank is %d\n", tid);
72
73 if ( proc_cnt < 2 )
74 {

```

Current Line Marker

- ✦ The current line marker is used to show the current location of suspended processes
- ✦ In traditional programs, there is a single current line marker (the exception to this is multi-threaded programs)
- ✦ In parallel programs, there is a current line marker for every process
- ✦ The PTP debugger shows one current line marker for every group of processes at the same location

Colors And Markers

- ★ The highlight color depends on the processes suspended at that line:
 - ★ **Blue:** All registered process(es)
 - ★ **Orange:** All unregistered process(es)
 - ★ **Green:** Registered or unregistered process with no source line (e.g. suspended in a library routine)
- ★ The marker depends on the type of process stopped at that location
- ★ Hover over marker for more details about the processes suspend at that location

```

main.c
int proc_cnt;
int tid;
MPI_Datatype * res_type;

MPI_Init(&argc, &argv);
MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt);
MPI_Comm_rank(MPI_COMM_WORLD, &tid);

if ( proc_cnt < 2 )
{
    fprintf(stderr, "must have at least 2 processes, not %d\n", proc_cnt);
    MPI_Finalize();
    return 1;
}
  
```

The screenshot shows a code editor window titled 'main.c'. The code is as follows:

int proc_cnt;

int tid;

MPI_Datatype * res_type;

MPI_Init(&argc, &argv);

MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt);

MPI_Comm_rank(MPI_COMM_WORLD, &tid);

if (proc_cnt < 2)

{

 fprintf(stderr, "must have at least 2 processes, not %d\n", proc_cnt);

 MPI_Finalize();

 return 1;

}

Markers are visible on the left side of the code editor:

- A blue arrow marker points to the line `MPI_Comm_size(MPI_COMM_WORLD, &proc_cnt);`.

- An orange arrow marker points to the line `MPI_Comm_rank(MPI_COMM_WORLD, &tid);`.

- A green arrow marker points to the line `if (proc_cnt < 2)`.



Multiple processes marker



Registered process marker



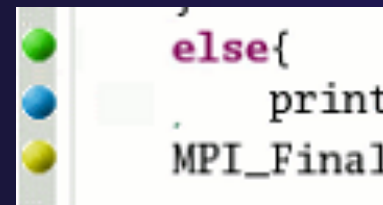
Un-registered process marker



Multiple markers at this line
 -Suspended on unregistered process: 2
 -Suspended on registered process: 1

Breakpoints

- ★ Apply only to processes in the particular set that is active in the **Parallel Debug view** when the breakpoint is created
- ★ Breakpoints are colored depending on the active process set and the set the breakpoint applies to:
 - ★ **Green** indicates the breakpoint set is the same as the active set.
 - ★ **Blue** indicates some processes in the breakpoint set are also in the active set (i.e. the process sets overlap)
 - ★ **Yellow** indicates the breakpoint set is different from the active set (i.e. the process sets are disjoint)
- ★ When the job completes, the breakpoints are automatically removed



Creating A Breakpoint



- ★ Select the process set that the breakpoint should apply to, in this case, the **workers** set
- ★ Double-click on the left edge of an editor window, at the line on which you want to set the breakpoint, or right click and use the **Parallel Breakpoint ► Toggle Breakpoint** context menu
- ★ The breakpoint is displayed on the marker bar

```
69 MPI_Comm_rank(MPI_COMM_WORLD, &tid);
70
71 fprintf(stdout, "my rank is %d\n", tid);
72
73 if ( proc_cnt < 2 )
74 {
75     fprintf(stderr, "must have at least 2 processes, not %d\n", proc_cnt);
76     MPI_Finalize();
77     return 1;
78 }
79
80 if ( (n % (proc_cnt - 1)) != 0 )
81 {
82     if ( tid == 0 )
83         fprintf(stderr, "(number of processes - 1) must be a multiple of %d\n", n);
84
85     MPI_Finalize();
86     return 1;
87 }
88
```



Hitting the Breakpoint

- ✦ Switch back to the **Root** set by clicking on the **Change Set** button
- ✦ Click on the **Resume** button in the **Parallel Debug view**
- ✦ In this example, the three worker processes have hit the breakpoint, as indicated by the yellow process icons and the current line marker
- ✦ Process 0 is still running as its icon is green
- ✦ Processes 1-3 are suspended on the breakpoint

```

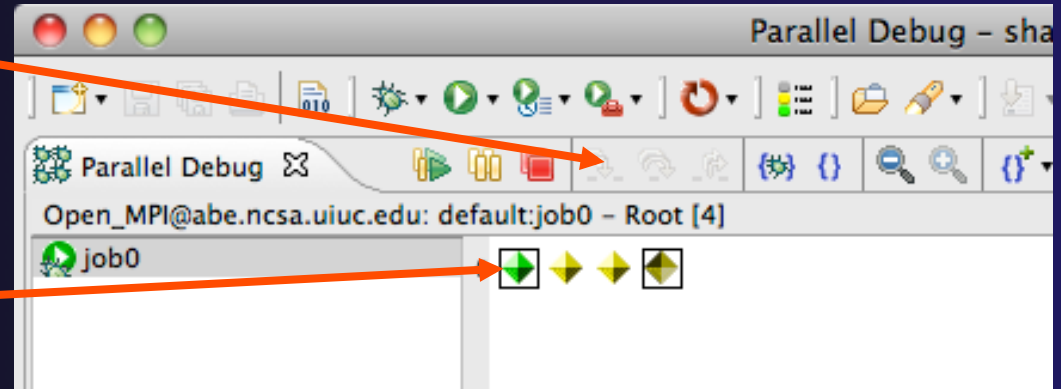
74 {
75     fprintf(stderr, "must have at least 2 processes, not %d\n", proc_cnt);
76     MPI_Finalize();
77     return 1;
78 }
79
80 if ( n % (proc_cnt - 1) != 0 )
81 {
82     if ( tid == 0 )
83         fprintf(stderr, "(number of processes - 1) must be a multiple of %d\n", n);
84
85     MPI_Finalize();
86     return 1;
87 }
88
89 if (tid != 0) {
90     worker();
91     MPI_Barrier(MPI_COMM_WORLD);
92     MPI_Finalize();
93 } else {

```

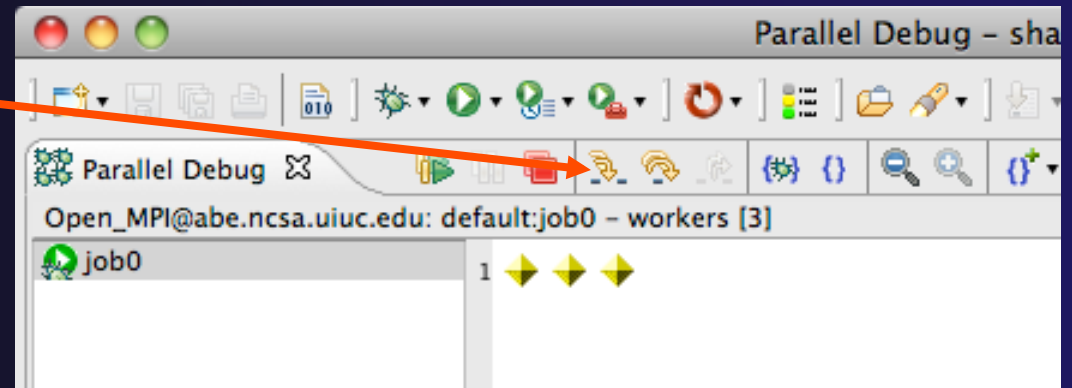


More On Stepping

- ★ The **Step** buttons are only enabled when all processes in the active set are **suspended** (yellow icon)
- ★ In this case, process 0 is still running



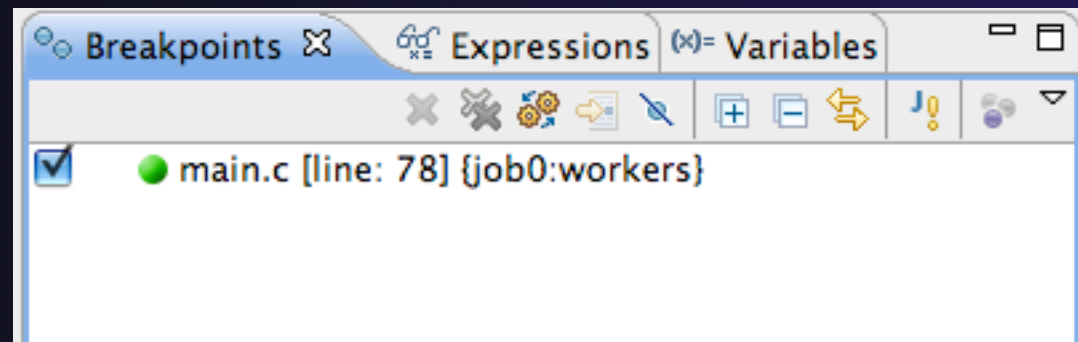
- ★ Switch to the set of suspended processes (the **workers** set)
- ★ You will now see the **Step** buttons become enabled





Breakpoint Information

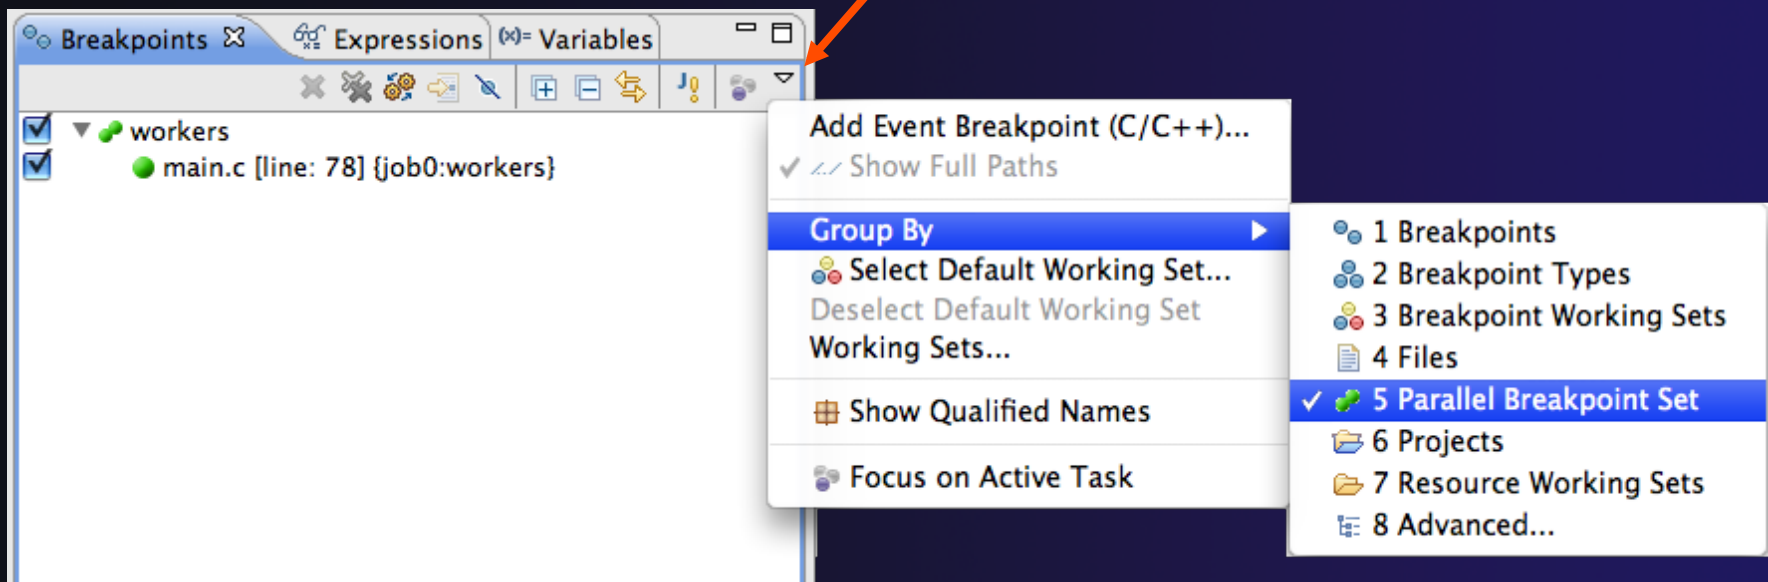
- ✦ Hover over breakpoint icon
 - ✦ Will show the sets this breakpoint applies to
- ✦ Select **Breakpoints** view
 - ✦ Will show all breakpoints in all projects





Breakpoints View

- ★ Use the menu in the breakpoints view to group breakpoints by type
- ★ Breakpoints sorted by breakpoint set (process set)



Global Breakpoints

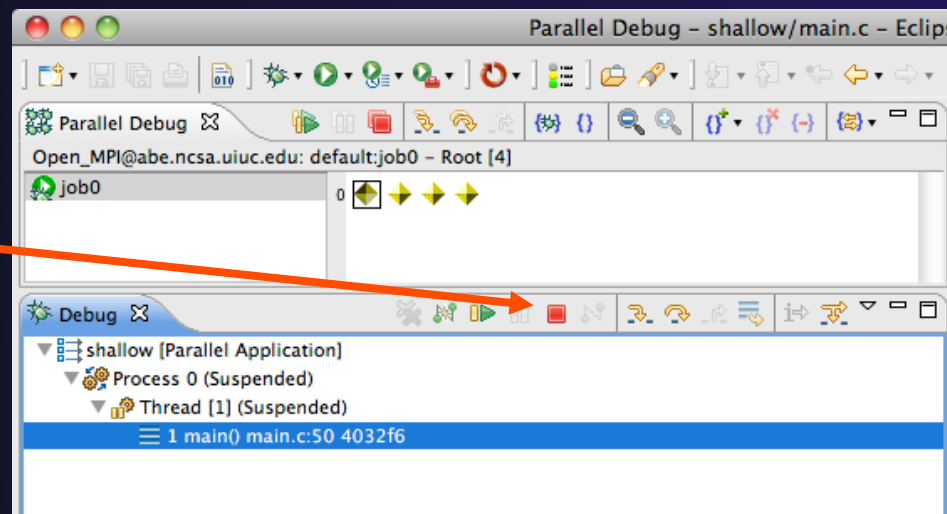
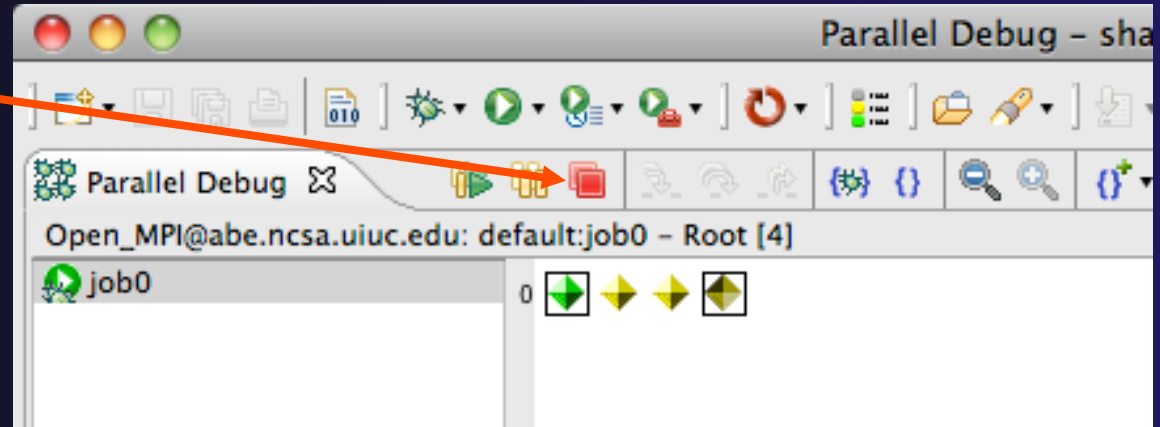
- ✦ Apply to all processes and all jobs
- ✦ Used for gaining control at debugger startup
- ✦ To create a global breakpoint
 - ✦ First make sure that no jobs are selected (click in white part of jobs view if necessary)
 - ✦ Double-click on the left edge of an editor window
 - ✦ Note that if a job is selected, the breakpoint will apply to the current set

```
if (my_rank != 0) {  
    /* create message */  
    sprintf(message, "Greetin
```



Terminating A Debug Session

- ★ Click on the **Terminate** icon in the **Parallel Debug view** to terminate all processes in the active set
- ★ Make sure the **Root** set is active if you want to terminate all processes
- ★ You can also use the terminate icon in the **Debug view** to terminate the currently selected process



Module 5: Performance Tuning and Analysis Tools

★ Objective

- ★ Become familiar with tools integrated with PTP, to help enhance performance of parallel applications

★ Contents

- ★ Performance Tuning and external tools:
 - ★ PTP External Tools Framework (ETFw), TAU
Hands-on exercise using TAU with PTP
- ★ MPI Analysis: GEM (Graphical Explorer of MPI Programs)

PTP/External Tools Framework

formerly "Performance Tools Framework"

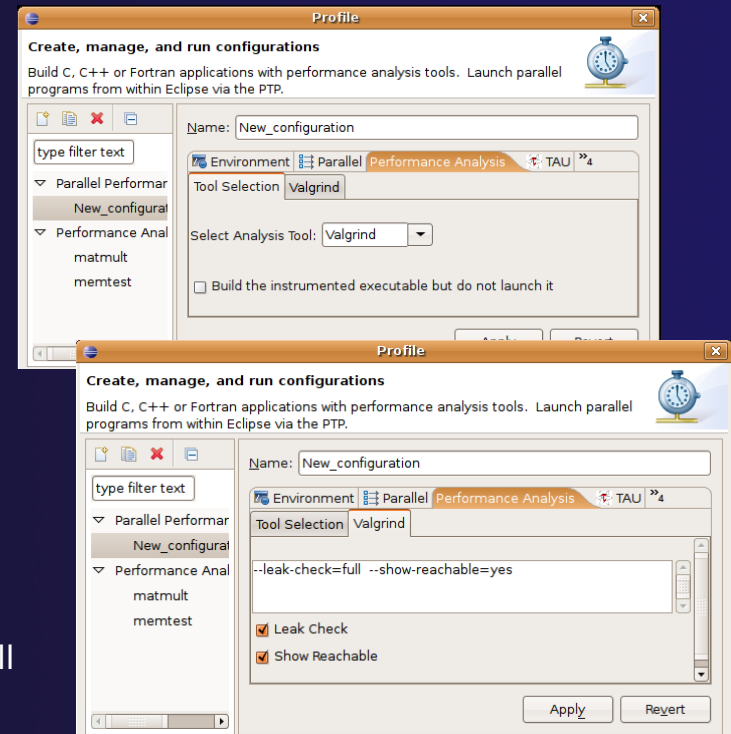
Goal:

- ★ Reduce the "eclipse plumbing" necessary to integrate tools
- ★ Provide integration for instrumentation, measurement, and analysis for a variety of performance tools
 - ★ Dynamic Tool Definitions: Workflows & UI
 - ★ Tools and tool workflows are specified in an XML file
 - ★ Tools are selected and configured in the launch configuration window
 - ★ Output is generated, managed and analyzed as specified in the workflow
 - ★ One-click 'launch' functionality
 - ★ Support for development tools such as TAU, PPW and others.
 - ★ Adding new tools is much easier than developing a full Eclipse plug-in

```

-<tool name="Valgrind">
  -<execute>
    <utility command="bash" group="inbin"/>
    -<utility command="valgrind" group="valgrind">
      -<optionpane title="Valgrind" seperatewith=" ">
        <togoption label="Leak Check" optname="--leak-check=full" tooltip="Leak Check" />
        <togoption label="Show Reachable" optname="--show-reachable=yes" tooltip="Show Reachable" />
      </optionpane>
    </utility>
  </execute>
</tool>

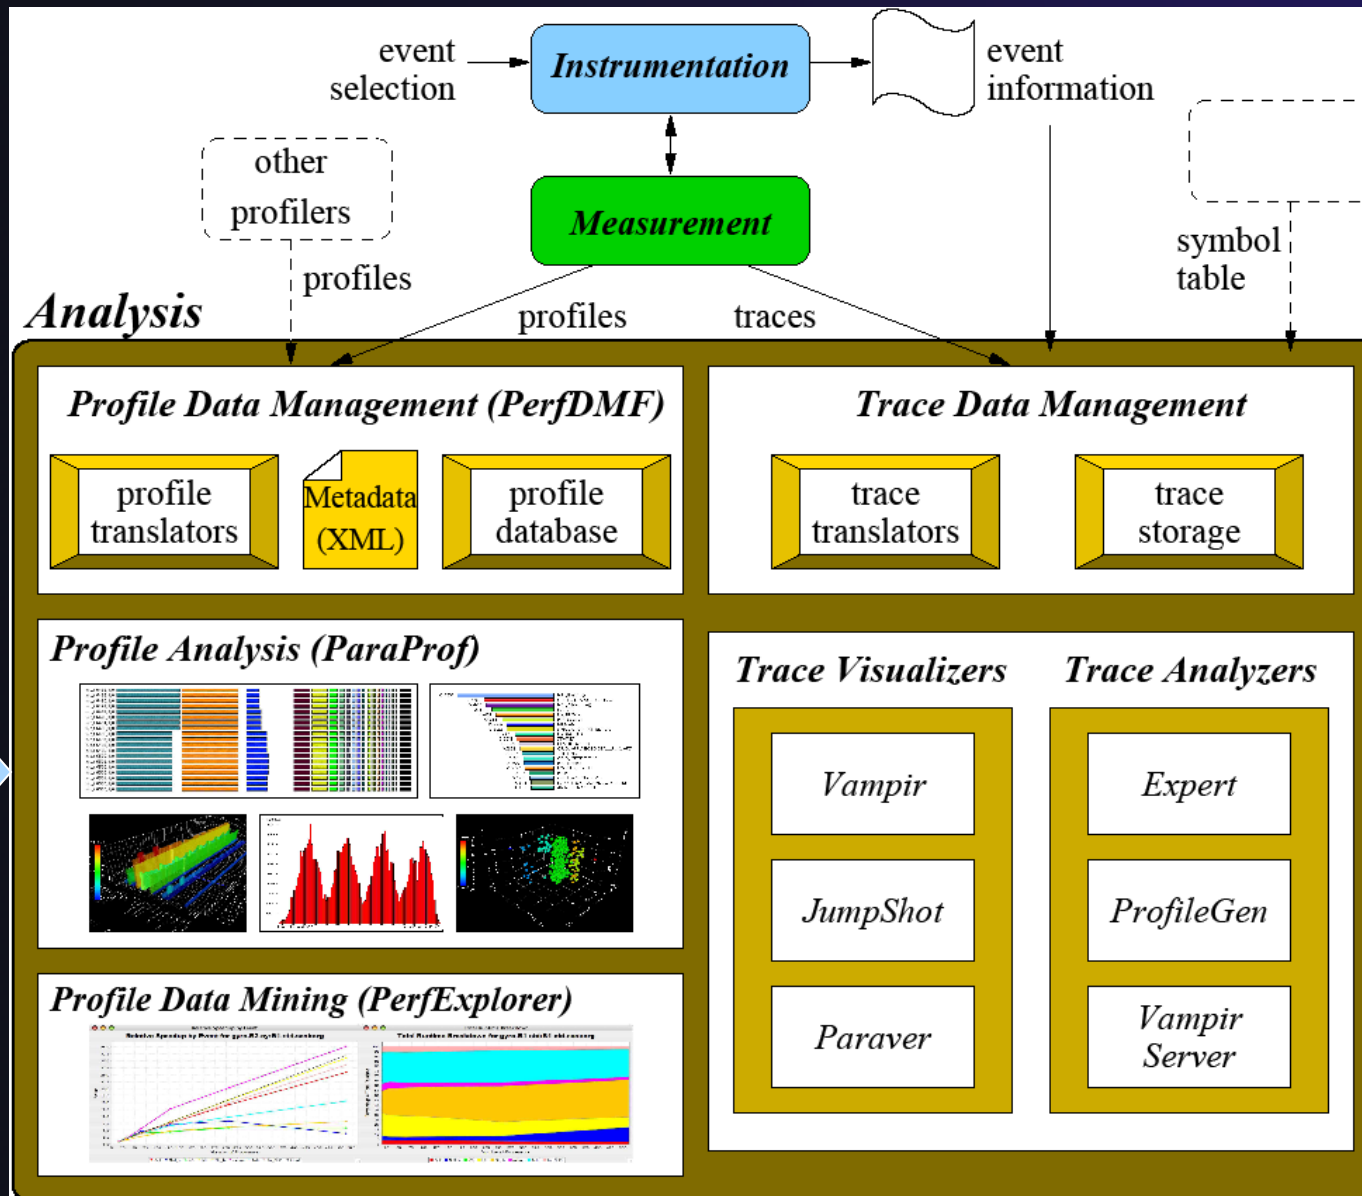
```



TAU: Tuning and Analysis Utilities

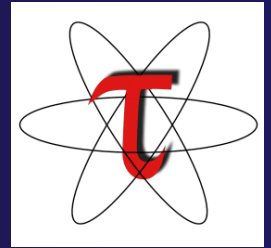
- ★ TAU is a performance evaluation tool
- ★ It supports parallel profiling and tracing
- ★ Profiling shows you how much (total) time was spent in each routine
- ★ Tracing shows you *when* the events take place in each process along a timeline
- ★ TAU uses a package called PDT for automatic instrumentation of the source code
- ★ Profiling and tracing can measure time as well as hardware performance counters from your CPU (or GPU!)
- ★ TAU can automatically instrument your source code (routines, loops, I/O, memory, phases, etc.)
- ★ TAU runs on all HPC platforms and it is free (BSD style license)
- ★ TAU has instrumentation, measurement and analysis tools
 - ★ paraprof is TAU's 3D profile browser

TAU Performance System Architecture



PTP TAU plug-ins

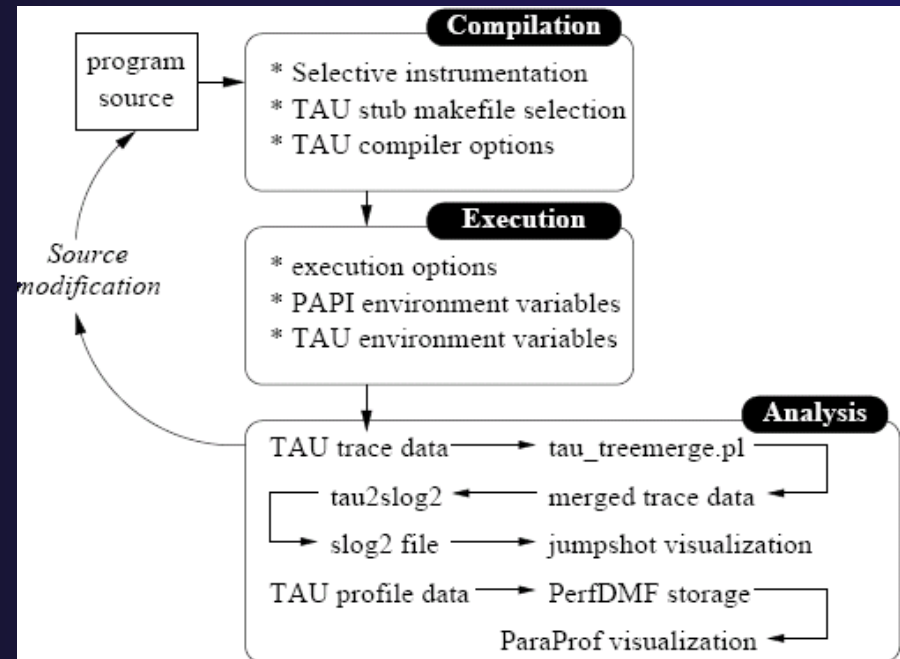
<http://www.cs.uoregon.edu/research/tau>



- ★ TAU (Tuning and Analysis Utilities)
- ★ First implementation of External Tools Framework (ETFw)
- ★ Eclipse plug-ins wrap TAU functions, make them available from Eclipse
- ★ Full GUI support for the TAU command line interface
- ★ Performance analysis integrated with development environment

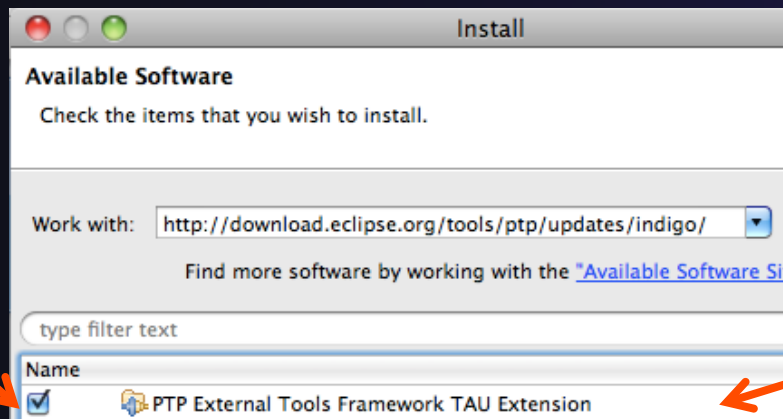
TAU Integration with PTP

- ★ TAU: Tuning and Analysis Utilities
 - ★ Performance data collection and analysis for HPC codes
 - ★ Numerous features
 - ★ Command line interface
- ★ The TAU Workflow:
 - ★ Instrumentation
 - ★ Execution
 - ★ Analysis



TAU PTP Installation

- ✦ This tutorial assumes that the TAU extensions for PTP are installed – they are not included in the “Eclipse IDE for Parallel Application Developers”
- ✦ The installation section (Module 1) shows how to install TAU and other features from the PTP update site – be sure TAU was selected



To confirm:

- ✦ Help>Install New Software...
- ✦ Select the link “What is already installed” at the bottom of the dialog
- ✦ You should see the TAU Extension

TAU/ETFW Hands-On(0)

Assumptions

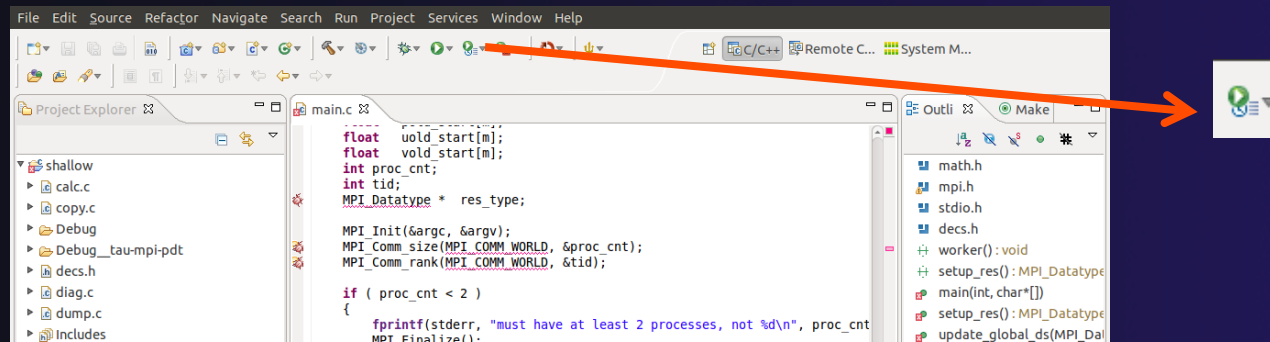
- ★ Obtain and install TAU*
 - ★ Download at tau.uoregon.edu
 - ★ The website includes setup and user guides
- ★ Set up the \$PATH on the remote machine*
 - ★ For TAU you should be able to run 'which pprof' on a remote login and see a result from your TAU bin directory
 - ★ On forge.ncsa.illinois.edu this is accomplished by placing 'module load tau' in the .modules file in the home directory
- ★ Include 'eclipse.inc' in the makefile*
 - ★ Create an empty eclipse.inc file in the same directory as the makefile
 - ★ Place 'include eclipse.inc' in the makefile after regular compiler definitions
 - ★ ETFW will modify eclipse.inc to set CC/CXX/FC variables

TAU/ETFW Hands-On(1)

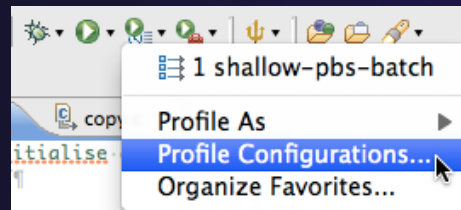
Begin Profile Configuration



- ★ The ETFW uses the same run configurations and resource managers as debugging/launching
- ★ Click on the 'Run' menu or the right side of the Profile button



- ★ From the dropdown menu select 'Profile configurations...'



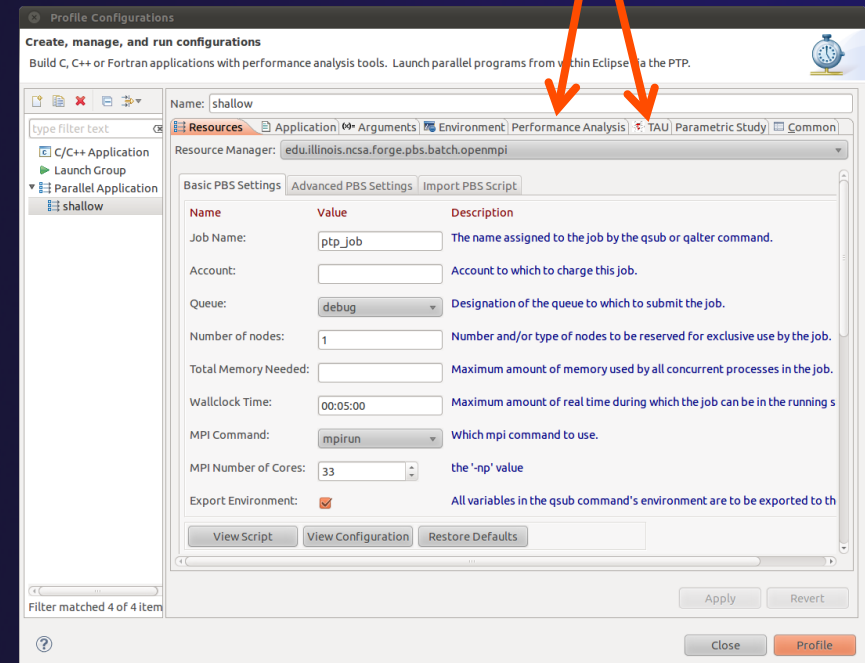
TAU/ETFW Hands-On(2)

Select Configuration



- ★ Select the shallow configuration prepared earlier
- ★ The Resource and Application configuration tabs require little or no modification
 - ★ We are using the same resource manager (edu.illinois.ncsa.forge.pbs.batch.openmpi) and PBS settings
 - ★ Since we are using a makefile project the application will be rebuilt in and run from the previously selected location

Performance Analysis and TAU tabs are present in the **Profile Configurations** dialog

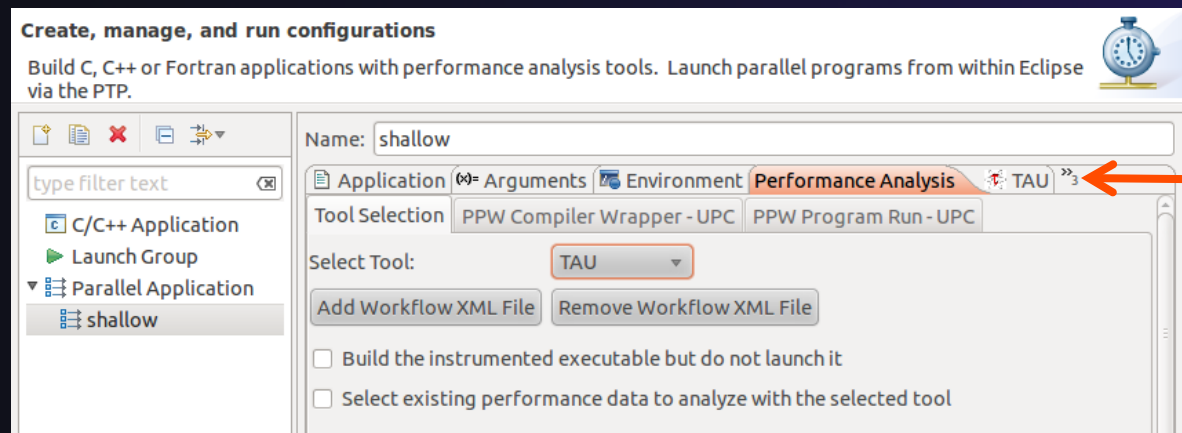


TAU/ETFW Hands-On (3)

Select Tool/Workflow



- ✦ Select the **Performance Analysis** tab and choose the TAU tool set in the 'Select Tool' dropdown box
 - ✦ Other tools may be available, either installed as plug-ins or loaded from workflow definition XML files



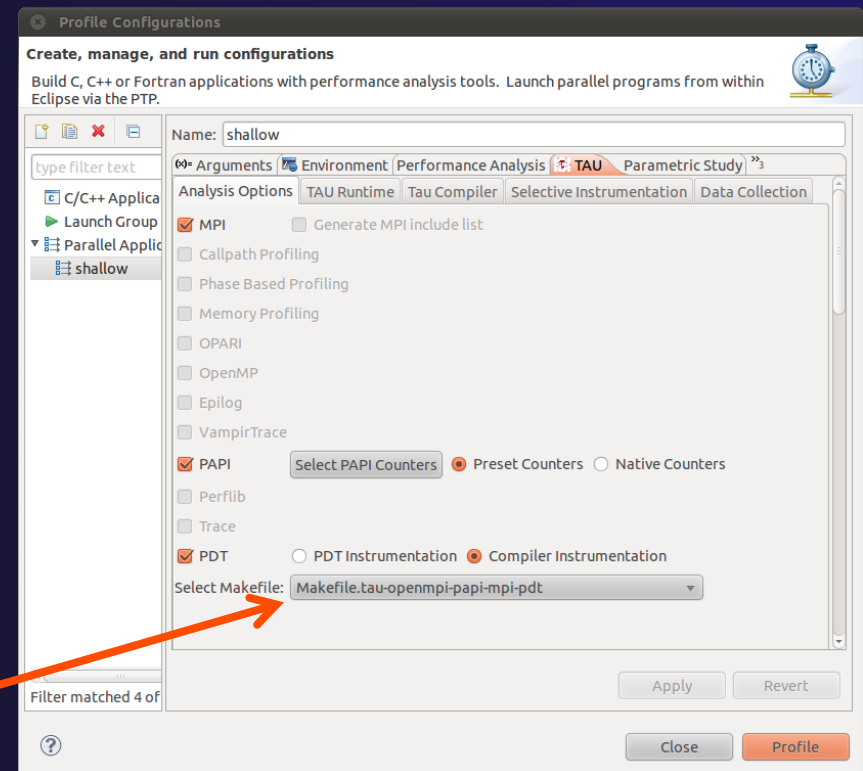
Tabs may be hidden if the window is too small

TAU/ETFW Hands-On (4)

Select TAU Configuration



- ✦ Select the **TAU** tab
- ✦ Choose the TAU stub makefile:
 - ✦ All TAU configurations in remote installation are available
 - ✦ Check MPI, PAPI and PDT checkboxes to filter listed makefiles
 - ✦ Make your selection in the **Select Makefile:** dropdown box
 - ✦ Select `Makefile.tau-openmpi-papi-mpi-pdt`

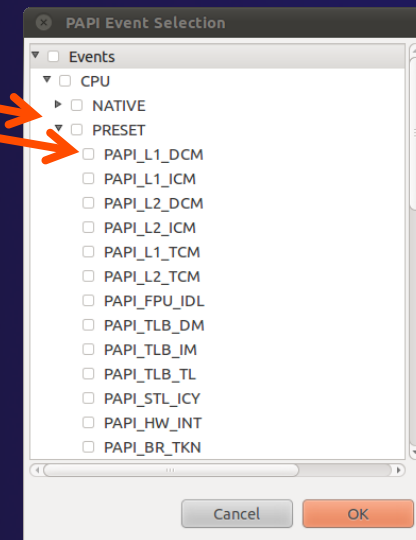


TAU/ETFW Hands-On (5)

Choose PAPI Hardware Counters



- ★ When a PAPI-enabled TAU configuration is selected the PAPI Counter tool becomes available
 - ★ Select the 'Select PAPI Counters' button to open the tool
 - ★ Open the PRESET subtree
 - ★ Select PAPI_L1_DCM (Data cache misses)
 - ★ Scroll down to select PAPI_FP_INS (Floating point instructions)
 - ★ Invalid selections are automatically excluded
 - ★ Select **OK**



TAU/ETFW Hands-On (6)

Compiler and Runtime Options



- ★ Other tab settings are described here but no changes are required...
- ★ TAU Compiler Options
 - ★ Set arguments to TAU compiler scripts
 - ★ Control instrumentation and compilation behavior
- ★ TAU Runtime options
 - ★ Set environment variables used by TAU
 - ★ Control data collection behavior
- ★ All options included context sensitive help

Profile Configurations

Create, manage, and run configurations

Build C, C++ or Fortran applications with performance analysis tools. Launch parallel programs from within Eclipse via the PTP.

Name: shallow

Arguments Environment Performance Analysis TAU Parametric Study

Analysis Options TAU Runtime Tau Compiler Selective Instrumentation Data Collection

-tau_options='-optNoRevert -optComplint'

Verbose

DetectMemoryLeaks

PdtGnuFortranParser

PdtCleanscapeParser

PreProcess

KeepFiles

Shared

PdtDir

PdtF95Opts

PdtF95Reset

PdtCOpts

PdtCReset

PdtCxxOpts

PdtCxxReset

PdtF90Parser

Filter matched 4 of

Profile Configurations

Create, manage, and run configurations

Build C, C++ or Fortran applications with performance analysis tools. Launch parallel programs from within Eclipse via the PTP.

Name: shallow

Arguments Environment Performance Analysis TAU Parametric Study

Analysis Options TAU Runtime Tau Compiler Selective Instrumentation Data Collection

Verbose Output

Enable Profiling

Profile Output Directory

Profile Format

Callpath Profiling

Set Callpath Depth

Overhead Compensation

Track Message

Communication Matrix

Enable Throttling

Throttle Percall

Throttle Numcalls

Enable Tracing

Trace Output Directory

Synchronize Clocks

Track Heap

Track Headroom

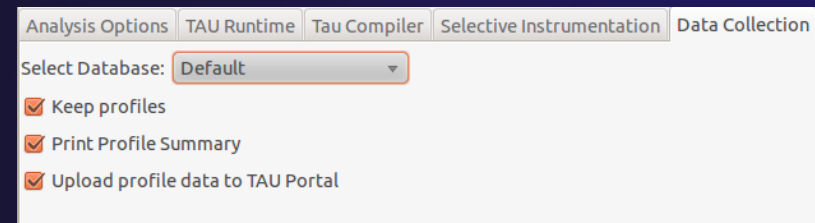
Track memory headroom availability at function entry/exit with a context event

Hover help

TAU/ETFW Hands-On (7)



- ★ If local PerfDMF databases are available you may select one to hold profile output
- ★ A text summary may be printed to the console
- ★ Profiles may be uploaded to the TAU Portal for viewing online
 - ★ tau.nic.uoregon.edu
- ★ Profiles may be copied to your workspace



TAU Profile Output

MULTI_GET_TIME_OF_DAYReading Profile files in profile.*

FUNCTION SUMMARY (total):

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	196	16,797	9	9	1866428 .TAU application
98.8	56	16,601	9	3662	1844640 main
68.7	11,538	11,538	9	0	1282002 MPI_Init()
26.0	204	4,360	8	59684	545009 worker
9.5	803	1,602	80000	160000	20 neighbour_receive
8.3	789	1,402	80000	160000	18 neighbour_send
8.3	234	1,398	8000	64000	175 time_load
7.1	1,188	1,188	81968	0	14 MPI_Recv()
6.5	182	1,085	8000	48000	136 calc_load

TAU Portal URL:

TAU Portal Username:

TAU Portal Password:

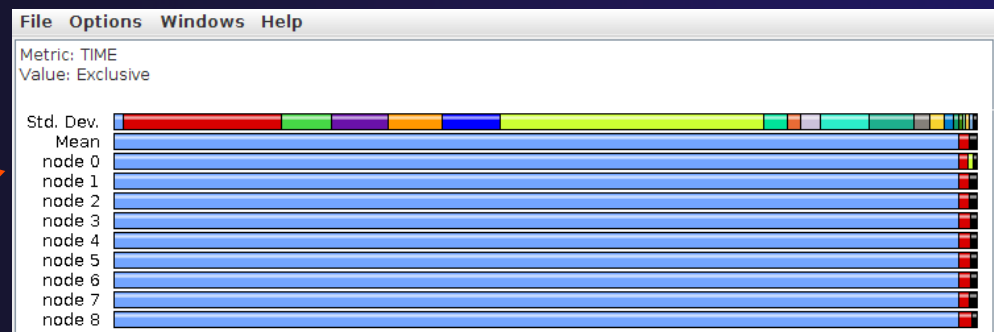
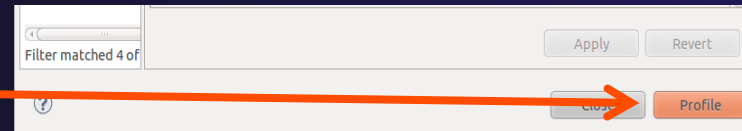
Select a workspace

Select Workspace:



TAU/ETFW Hands-On (8)

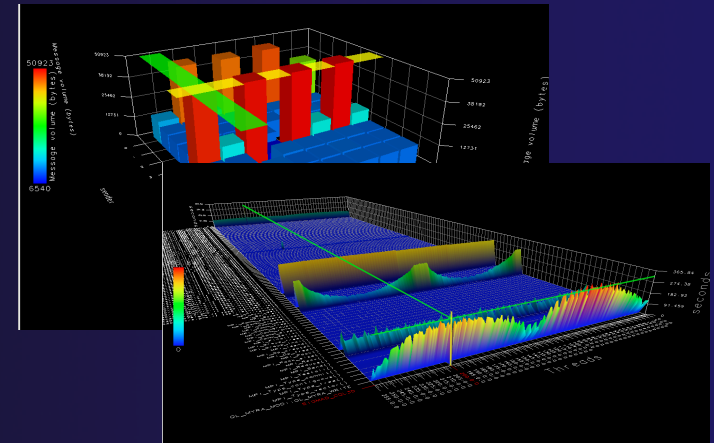
- ★ Once your TAU launch is configured select 'Profile'
- ★ Notice that the project rebuilds with TAU compiler commands
- ★ The project will execute normally but TAU profiles will be generated
- ★ TAU profiles will be processed as specified in the launch configuration.
- ★ If you have a local PerfDMF database the run will show up in the Performance Data Management view
 - ★ Double click the new entry to view in ParaProf
 - ★ Right click on a function bar and select **Show Source Code** for source callback to Eclipse





TAU/ETFW Hands-On (9)

- ★ Use ParaProf for profile visualization to identify performance hotspots
 - ★ Inefficient sequential computation
 - ★ Communication overhead
 - ★ IO/Memory bottlenecks
 - ★ Load imbalance
 - ★ Suboptimal cache performance
- ★ Compare multiple trials in PerfExplorer to identify performance regressions and scaling issues
- ★ To use ParaProf, install TAU from tau.uoregon.edu or use Java webstart from tau.uoregon.edu/paraprof



GEM

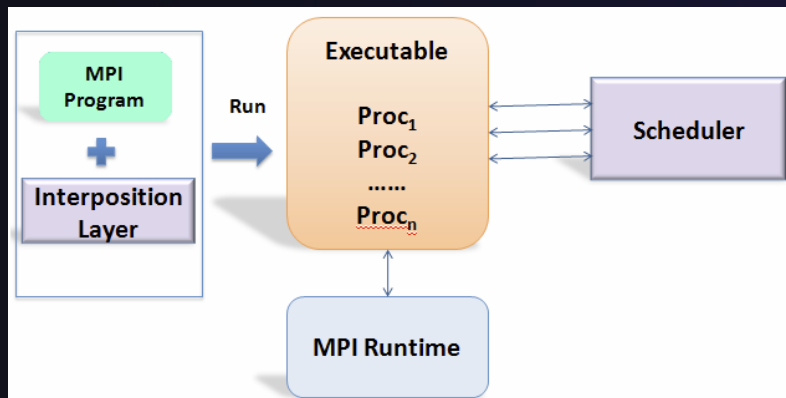
Graphical Explorer of MPI Programs

GEM

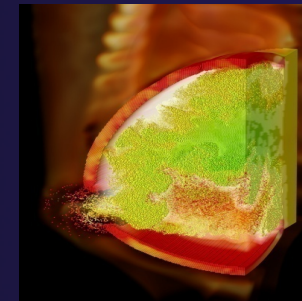
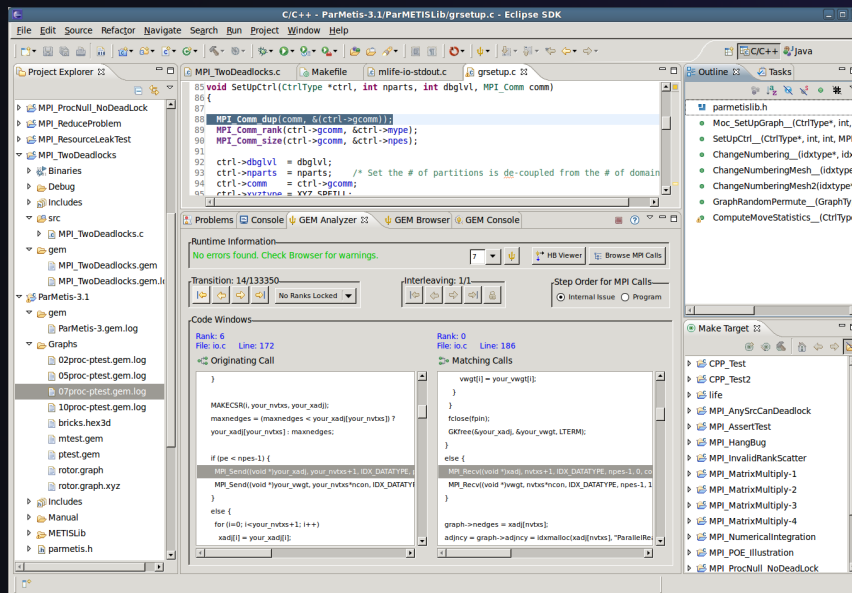
Graphical Explorer of MPI Programs

- ★ Dynamic verification for MPI C/C++ that detects:
 - ★ Deadlocks
 - ★ MPI object leaks (communicators, requests, etc)
 - ★ Functionally irrelevant barriers
 - ★ Local assertion violations
 - ★ MPI Send/Recv Type Mismatches
- ★ Offers rigorous coverage guarantees
 - ★ Complete nondeterministic coverage for MPI (MPI_ANY_SOURCE)
 - ★ Determines relevant interleavings, replaying as necessary
- ★ Examines communication / synchronization behaviors

GEM - Overview



- ★ Front-end for In-situ Partial Order (ISP) developed at University of Utah
- ★ Contributes “push-button” C/C++ MPI verification and analysis to the development cycle
- ★ Automatically instruments and runs user code, displaying post verification results
- ★ Variety of views & tools to facilitate debugging and MPI runtime understanding

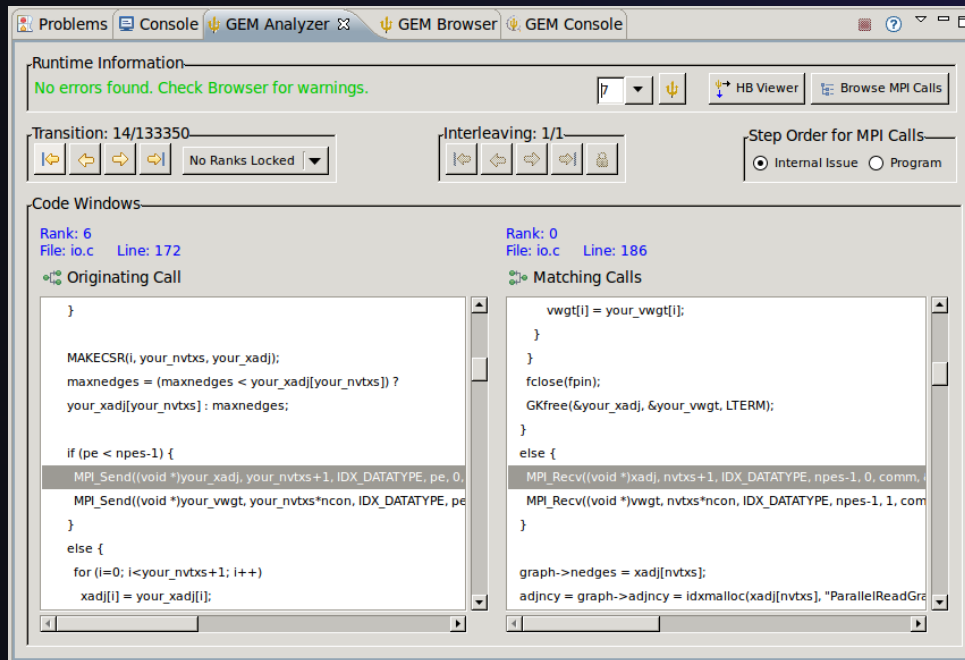


(Image courtesy of Steve Parker, U of Utah)

GEM – Views & Tools

Analyzer View

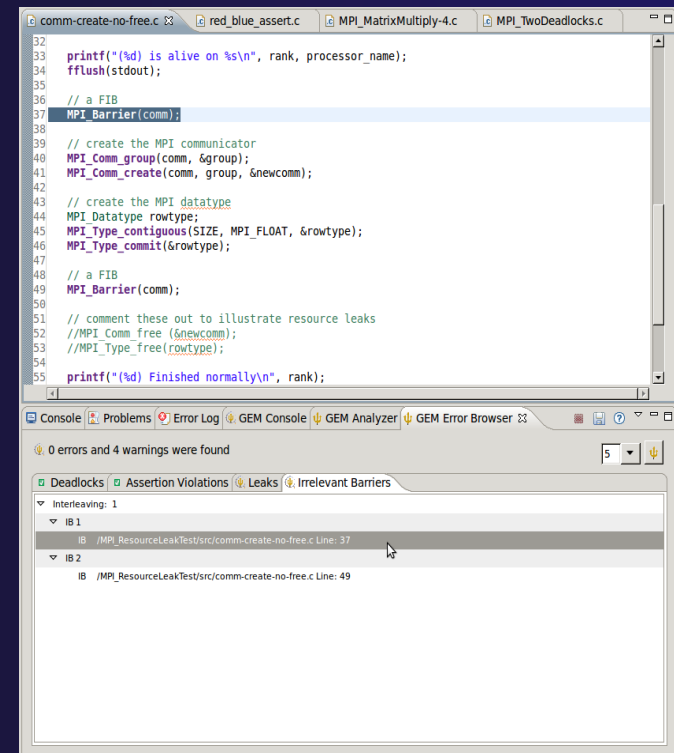
Highlights bugs, and facilitates post-verification review / debugging



Module 5

Browser View

Groups and localizes MPI problems. Maps errors to source code in Eclipse editor

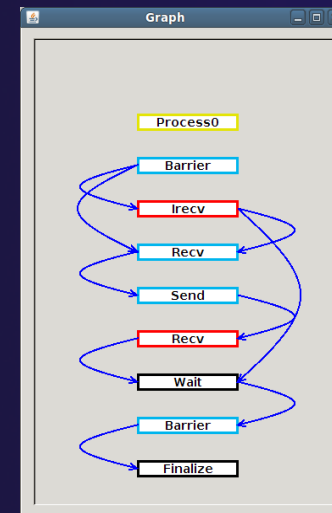
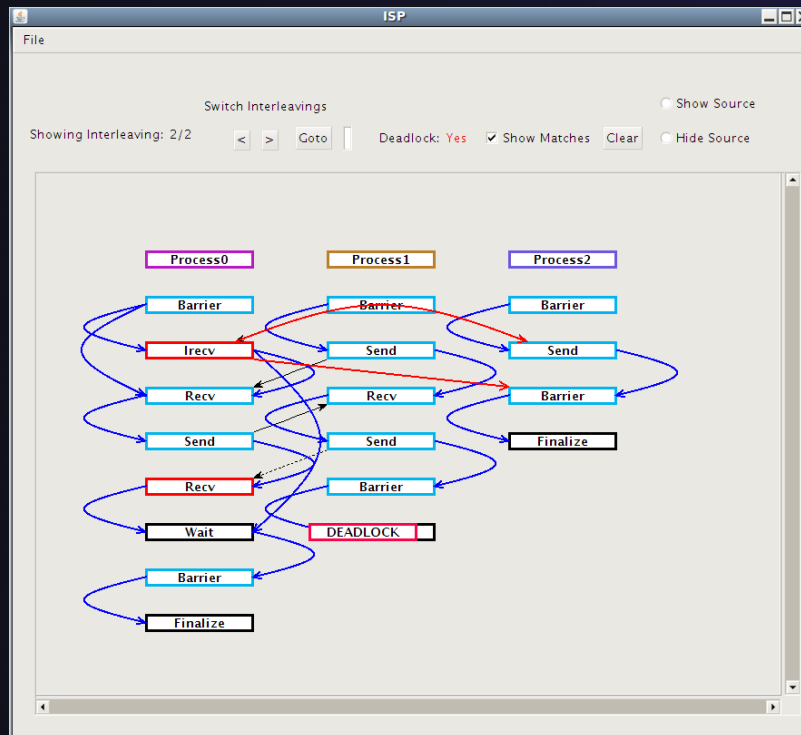


5-20

GEM – Views & Tools (cont.)

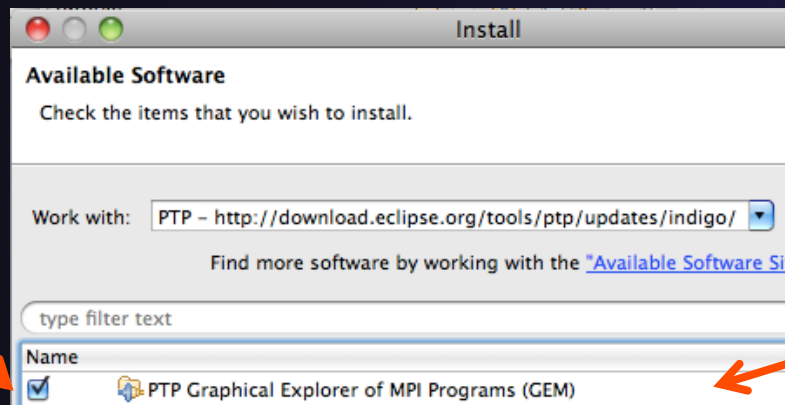
Happens-Before Viewer

Shows required orderings and communication matches
(currently an external tool)



GEM PTP Installation

- ✦ This tutorial assumes that the GEM extensions for PTP are installed – they are not included in the “Eclipse IDE for Parallel Application Developers”
- ✦ The installation section (Module 1) shows how to install GEM and other features from the PTP update site – be sure GEM was selected



To confirm:

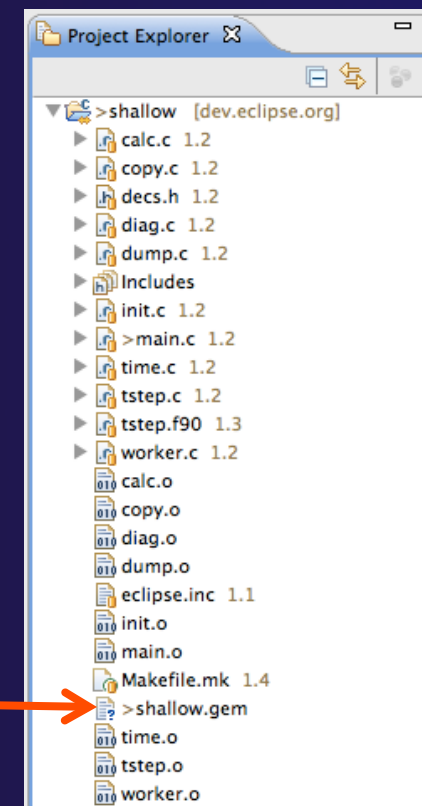
- ✦ Help>Install New Software...
- ✦ Select the link “What is already installed” at the bottom of the dialog
- ✦ You should see the GEM Extension

GEM Hands-on (0) Assumptions




- ★ Obtain and install ISP on the target machine *
 - ★ Download ISP at <http://www.cs.utah.edu/fv/ISP>
 - ★ Web page has installation instructions
- ★ Configure GEM/ISP paths for local/remote machine *
 - ★ On forge.ncsa.illinois.edu, the necessary bin directories have been placed in the default path for your training account
- ★ General GEM configuration necessary for the shallow project *

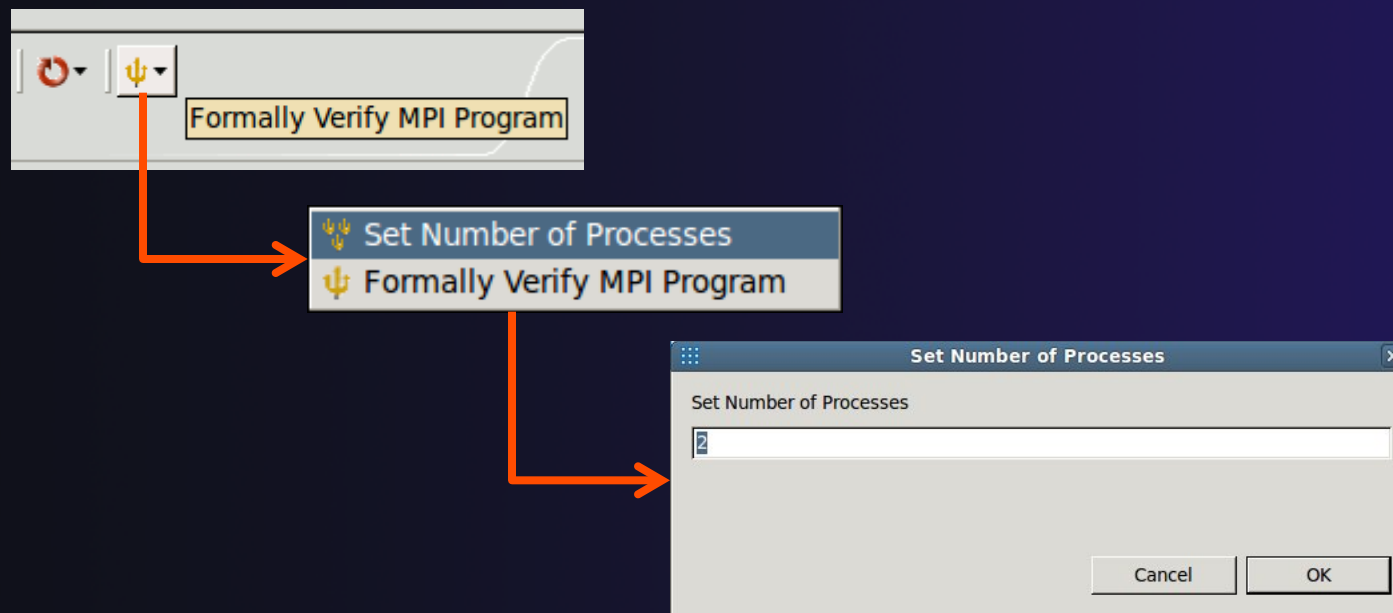
Instrumented
shallow.gem
should be built





GEM Hands-on (1)

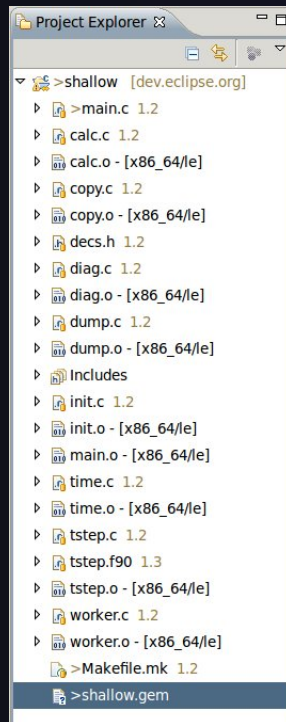
- ✦ No Launch configuration, run GEM directly
- ✦ Locate the trident  Icon on the toolbar
- ✦ From pull-down menu, set number of processes to **2**



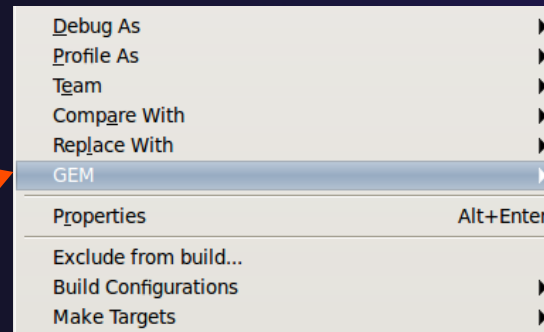


GEM Hands-on (2)

- ★ Make sure the project is built remotely (on Forge)
- ★ Locate file **shallow.gem** in the Project Explorer view



- ★ Right click on **shallow.gem** and select:
 - ★ GEM -> Formally Verify Profiled MPI Executable



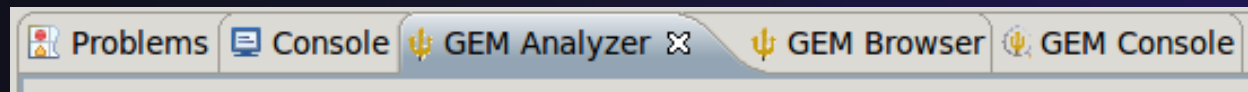
Formally Verify MPI Program

Then be patient, analysis can take several seconds...

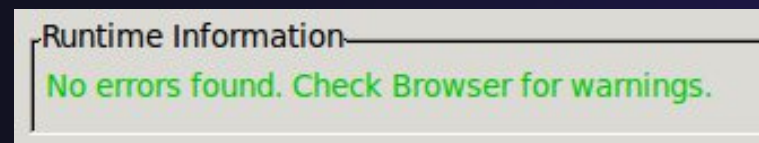


GEM Hands-on (3)

- ★ After the analysis, notice that three new Eclipse views have opened:
 - ★ GEM Analyzer, GEM Browser, GEM Console



- ★ In the **GEM Analyzer View** (in focus by default), notice there were no MPI errors, e.g. deadlock, assertion violation, but warnings were issued

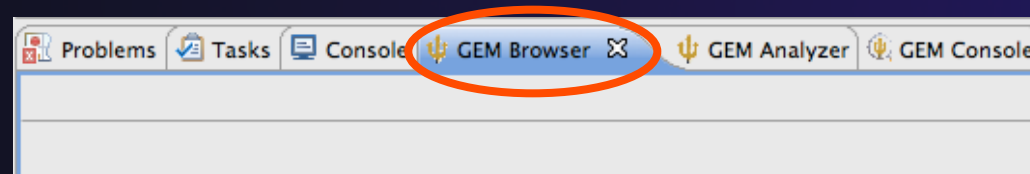


- ★ Examine these warnings in the **GEM Browser View**

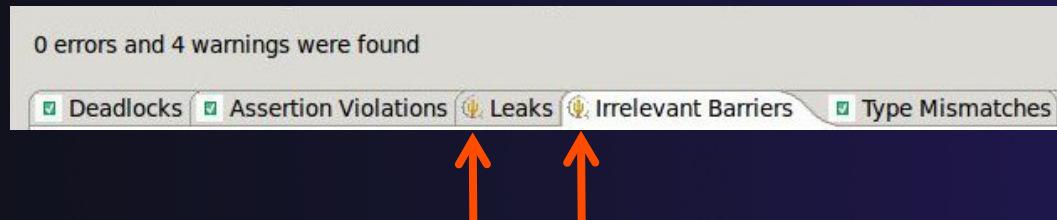


GEM Hands-on (4)

- ★ Click on the tab for the **GEM Browser View**



- ★ GEM has found two types of errors:
 - ★ Functionally Irrelevant Barriers (unnecessary synchronization)
 - ★ Resource Leaks (allocated MPI data types and request objects not freed)



- ★ GEM maps these warnings/errors to the offending line of source code within the Eclipse editor



GEM Hands-on (5)

- ★ Click on **Irrelevant Barriers** tab
- ★ Expand the **Interleaving 1** tree
- ★ Select each **MPI_Barrier** call separately

```

80     if ( tid == 0 )
81         fprintf(stderr, "(number of processes - 1) must be a multiple of %d\n", n);
82
83     MPI_Finalize();
84     return 1;
85 }
86
87 if (tid != 0) {
88     worker();
89     MPI_Barrier(MPI_COMM_WORLD);
90     MPI_Finalize();
91 } else {
92
93     /* master process */
94
95     chunk_size = n / (proc_cnt - 1);
96
97     for (i = 1; i < proc_cnt; i++) {
98         /* calculate each worker's boundary */
99         master_packet[JSTART] = (i - 1) * chunk_size;
100
101         if (i == proc_cnt - 1)
102             master_packet[JEND] = n - 1;
103         else
104             master_packet[JEND] = i * chunk_size - 1;
105
106         if (i == 1)
107             prv = proc_cnt-1;
108         else

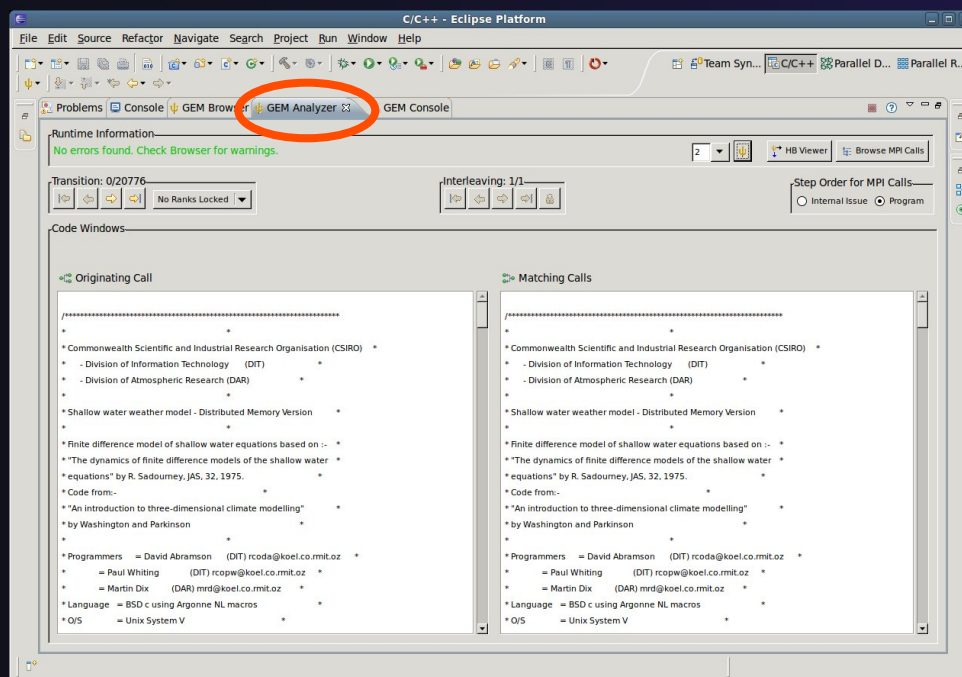
```

- ★ Notice you are taken to the corresponding line in the Eclipse Editor
- ★ This pair (match set) of **MPI_Barrier** calls can be safely removed without changing program behavior
- ★ If you delete only one and run GEM again, you will notice that a deadlock is introduced. The Analyzer View can be used to examine this.



GEM Hands-on (6)

- ✦ Double Click on the **Gem Analyzer View** tab
 - ✦ This makes the view fill the workbench window
- ✦ Examine communication and MPI runtime behavior using the **GEM Analyzer View**



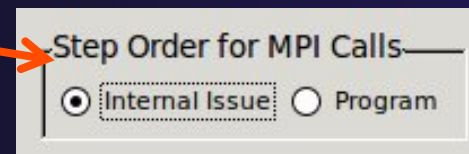
- ✦ The GEM Analyzer View should look similar to this figure

- ✦ **main.c** is open in both code view windows to examine P2P matches

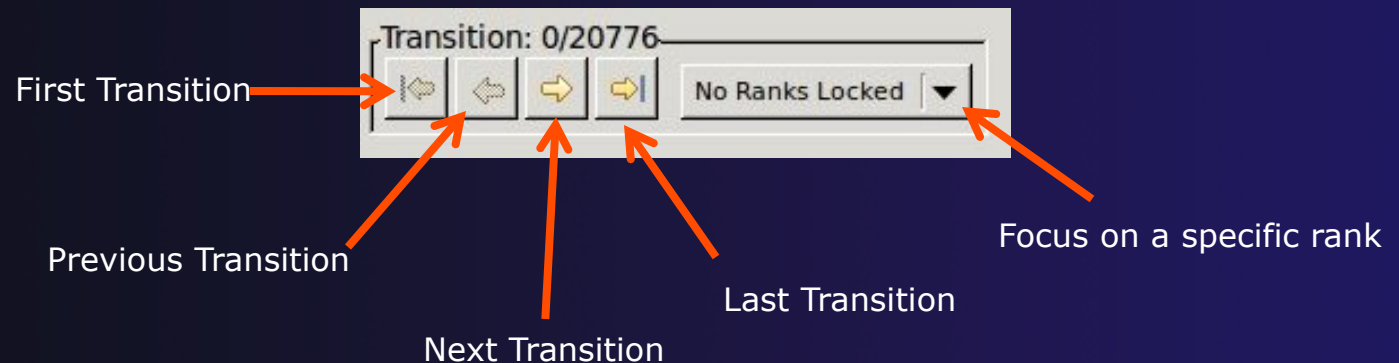


GEM Hands-on (7)

- ★ Under **Step Order for MPI Calls**, select **Internal Issue Order**. This is the order in which GEM issues the intercepted MPI calls to the MPI runtime



- ★ From the **Transition Group**, step through the single interleaving discovered by GEM using the navigation buttons.





GEM Hands-on (8)

★ P2P matches shown in **Code Windows**

Code Windows

Rank: 0
File: main.c Line: 150

Originating Call

```
for (j = 0; j < n; j++) {  
    acopy_two_to_one(p, p_start, j);  
    MPI_Send(&p_start, m, MPI_FLOAT, i, P_ROW,  
            MPI_COMM_WORLD);  
  
    acopy_two_to_one(u, u_start, j);  
    MPI_Send(&u_start, m, MPI_FLOAT, i, U_ROW,  
            MPI_COMM_WORLD);  
  
    acopy_two_to_one(v, v_start, j);  
    MPI_Send(&v_start, m, MPI_FLOAT, i, V_ROW,  
            MPI_COMM_WORLD);  
  
    acopy_two_to_one(psi, psi_start, j);  
    MPI_Send(&psi_start, m, MPI_FLOAT, i, PSI_ROW,  
            MPI_COMM_WORLD);  
  
    acopy_two_to_one(pold, pold_start, j);  
    MPI_Send(&pold_start, m, MPI_FLOAT, i, POLD_ROW,  
            MPI_COMM_WORLD);  
}
```

Rank: 1
File: worker.c Line: 139

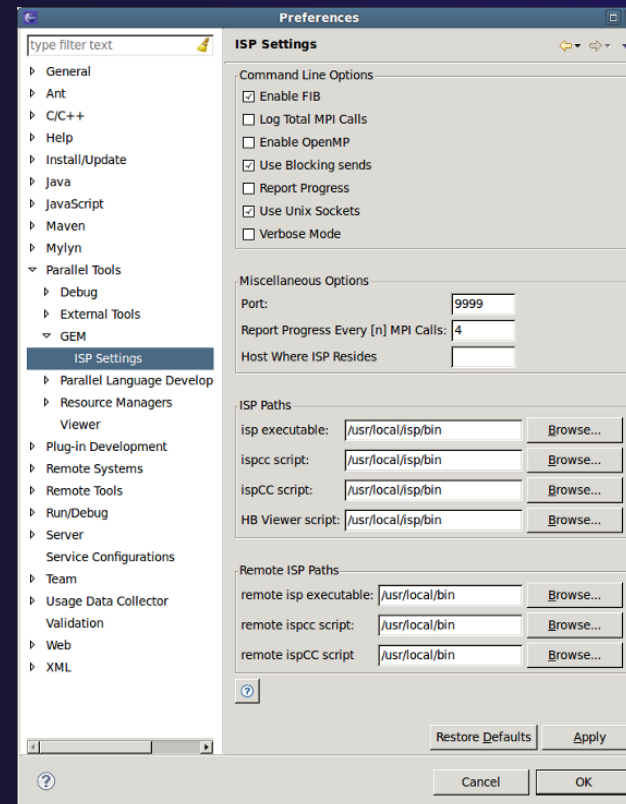
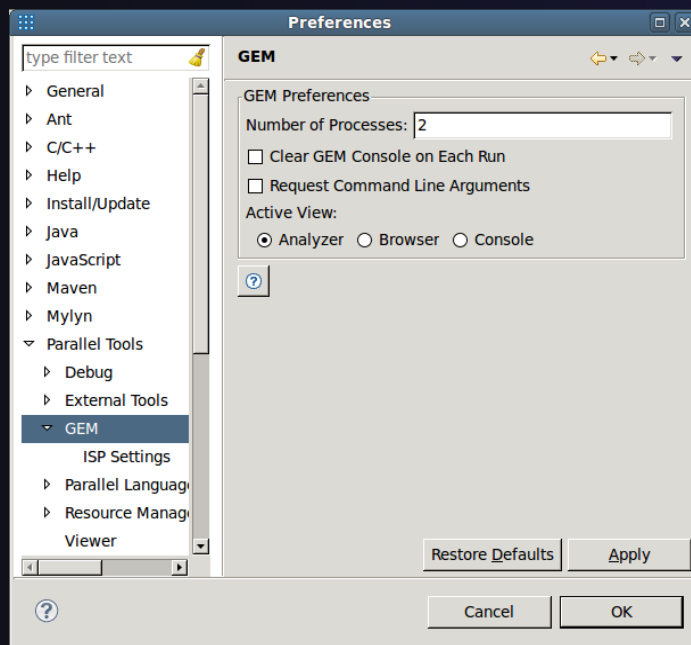
Matching Calls

```
for (i = 0; i < n; i++){  
    MPI_Recv(&p_start, m, MPI_FLOAT, 0, P_ROW,  
            MPI_COMM_WORLD, &status);  
    acopy_one_to_two(p_start, p, i);  
  
    MPI_Recv(&u_start, m, MPI_FLOAT, 0, U_ROW,  
            MPI_COMM_WORLD, &status);  
    acopy_one_to_two(u_start, u, i);  
  
    MPI_Recv(&v_start, m, MPI_FLOAT, 0, V_ROW,  
            MPI_COMM_WORLD, &status);  
    acopy_one_to_two(v_start, v, i);  
  
    MPI_Recv(&psi_start, m, MPI_FLOAT, 0, PSI_ROW,  
            MPI_COMM_WORLD, &status);  
    acopy_one_to_two(psi_start, psi, i);  
  
    MPI_Recv(&pold_start, m, MPI_FLOAT, 0, POLD_ROW,  
            MPI_COMM_WORLD, &status);  
    acopy_one_to_two(pold_start, pold, i);  
}
```

GEM Hands-on (9)

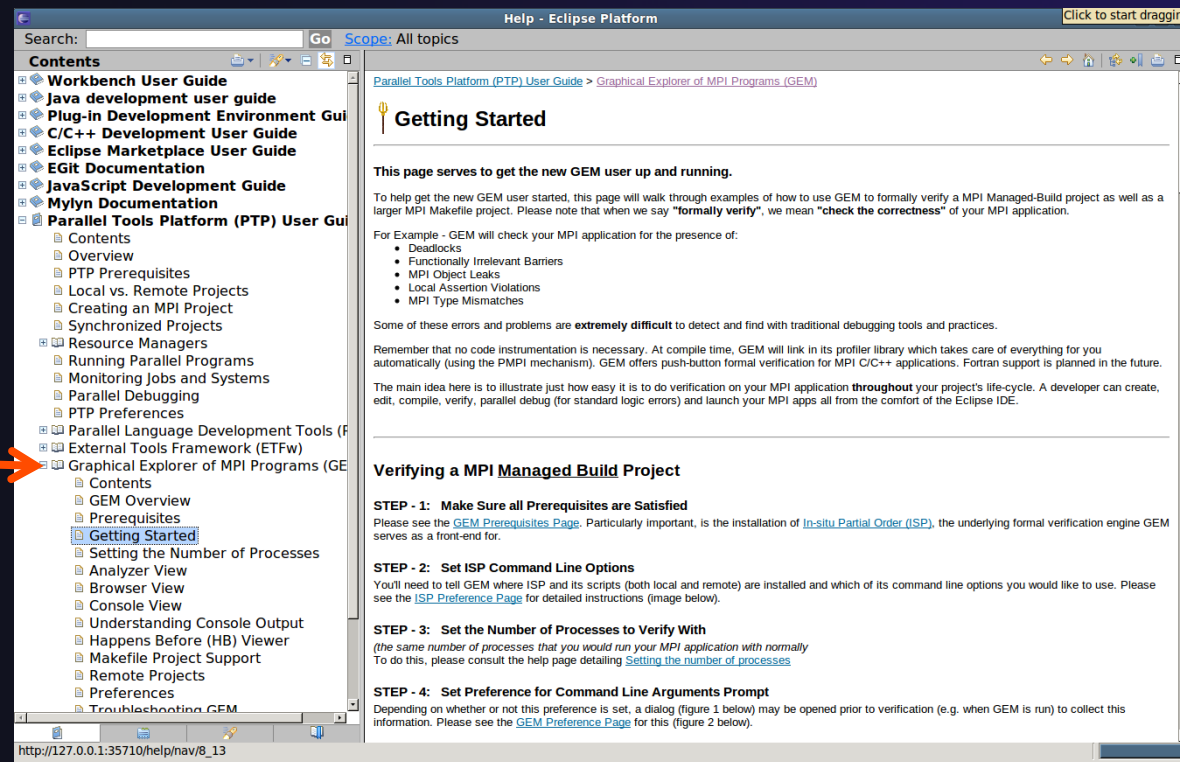


★ Other Preferences for further GEM Configuration



GEM Hands-on (10)

- ★ GEM Help walks through the Managed-Build, PI-C example created by the new project wizard
 - ★ Help > Help Contents



Reference Slides

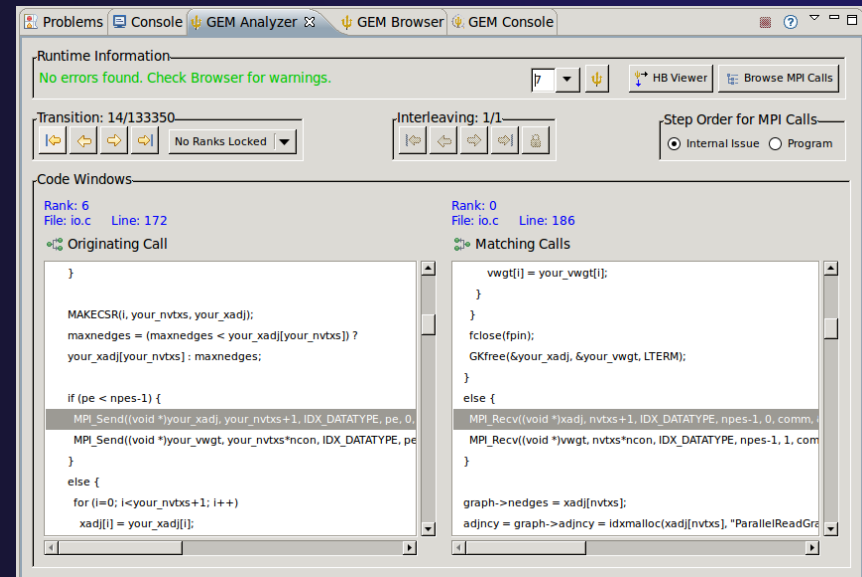
- ★ The following slides are not part of the presentation or hands-on section.
- ★ They are meant to be used as further reference and provide more detailed information on:
 - ★ GEM setup, configuration and preference pages
 - ★ Using GEM views
 - ★ Help contribution
 - ★ Troubleshooting
 - ★ GEM Success stories

ISP Installation

- ★ **ISP itself must be installed prior to using GEM**
- ★ Download ISP at <http://www.cs.utah.edu/fv/ISP>
- ★ Make sure libtool, automake and autoconf are installed.
- ★ Just untar isp-0.3.0.tar.gz into a tmp directory:
 - ★ Configure and install
 - ★ ./configure
 - ★ make
 - ★ make install
 - ★ This installs binaries and necessary scripts

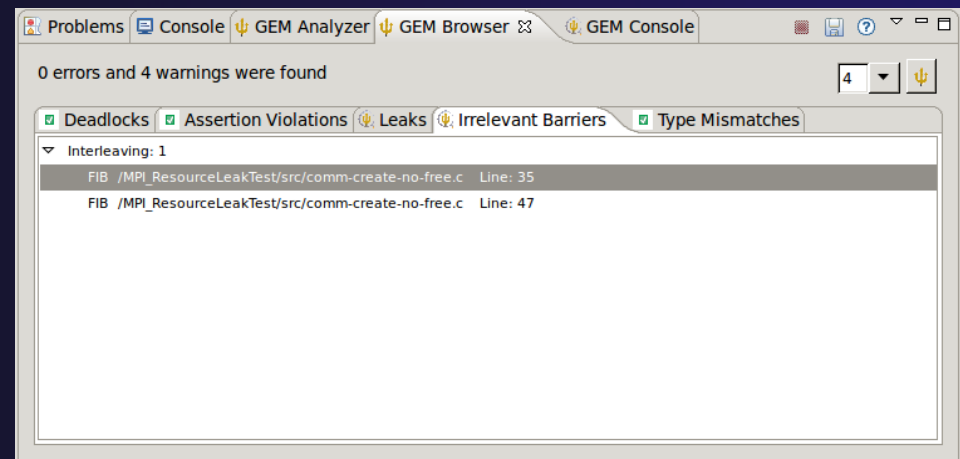
GEM Analyzer View

- ✦ Reports program errors, and runtime statistics
- ✦ Debug-style source code stepping of interleavings
 - ✦ Point-to-point / Collective Operation matches
 - ✦ Internal Issue Order / Program Order views
 - ✦ Rank Lock feature – focus on a particular process
- ✦ Also controls:
 - ✦ Call Browser
 - ✦ Happens Before Viewer launch
 - ✦ Re-launching of GEM



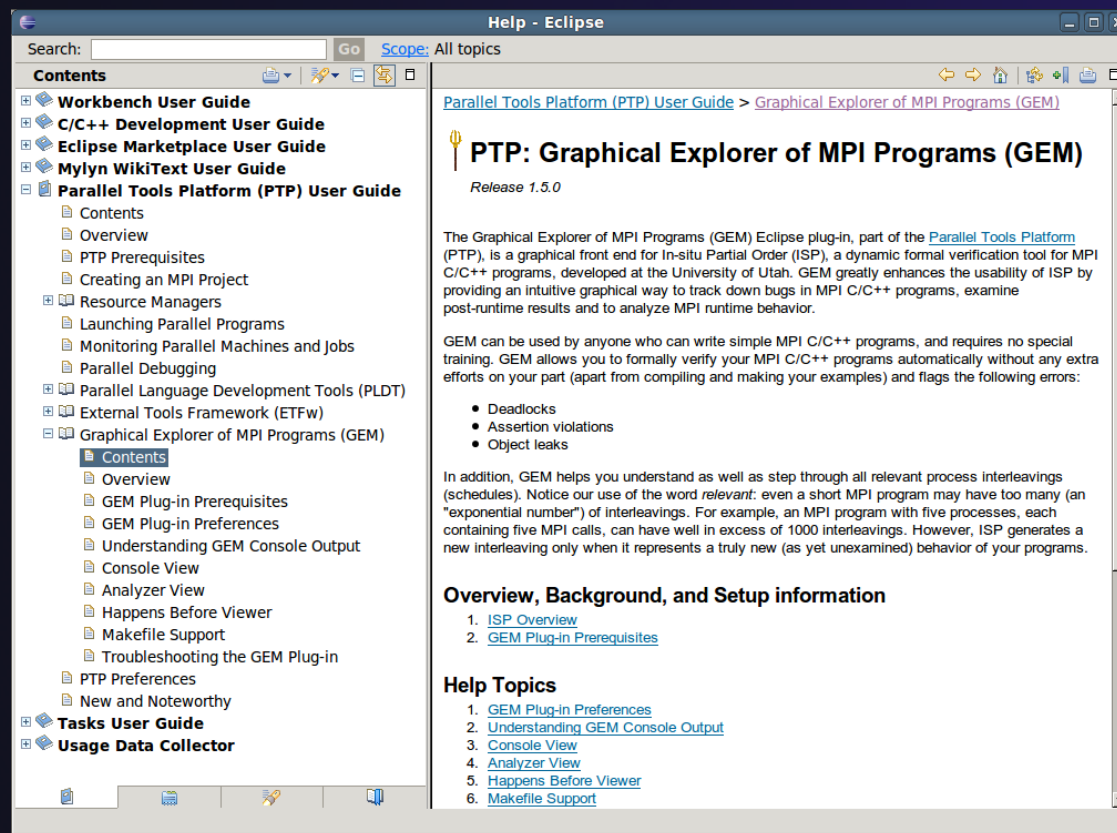
GEM Browser View

- ★ Tabbed browsing for each type of MPI error/warning
- ★ Each error/warning mapped to offending line of source code in Eclipse editor
- ★ One click to visit the Eclipse editor, to examine:
 - ★ Calls involved in deadlock
 - ★ Irrelevant barriers
 - ★ MPI Object Leaks sites
 - ★ MPI type mismatches
 - ★ Local Assertion Violations



GEM – Help Plugin

Extensive how-to sections, graphical aids and trouble shooting section



GEM/ISP Success Stories

★ Umpire Tests

- ★ <http://www.cs.utah.edu/fv/ISP-Tests>
- ★ Documents bugs missed by tests, caught by ISP

★ MADRE (EuroPVM/MPI 2007)

- ★ Previously documented deadlock detected

★ N-Body Simulation Code

- ★ Previously unknown resource leak caught during EuroPVM/MPI 2009 tutorial !

★ Large Case Studies

- ★ ParMETIS, MPI-BLAST, IRS (Sequoia Benchmark), and a few SPEC-MPI benchmarks could be handled

★ Full Tutorial including LiveDVD ISO available

- ★ Visit <http://www.cs.utah.edu/fv/GEM>

Module 6: Other Tools and Wrap-up

✦ Objective

- ✦ How to find more information on PTP
- ✦ Learn about other tools related to PTP
- ✦ See PTP upcoming features

✦ Contents

- ✦ Links to other tools, including performance tools
- ✦ Planned features for new versions of PTP
- ✦ Additional documentation
- ✦ How to get involved



NCSA Blue Waters HPC Workbench

- ★ Tools for NCSA Blue Waters
 - ★ <http://www.ncsa.illinois.edu/BlueWaters/>
 - ★ Sustained Petaflop system
- ★ Based on Eclipse and PTP
- ★ Includes some related tools
 - ★ Performance tools
 - ★ Workflow tools (<https://wiki.ncsa.uiuc.edu/display/MRDPUB/MRD+Public+Space+Home+Page>)
- ★ Part of the enhanced computational environment described at:
<http://www.ncsa.illinois.edu/BlueWaters/ece.html>



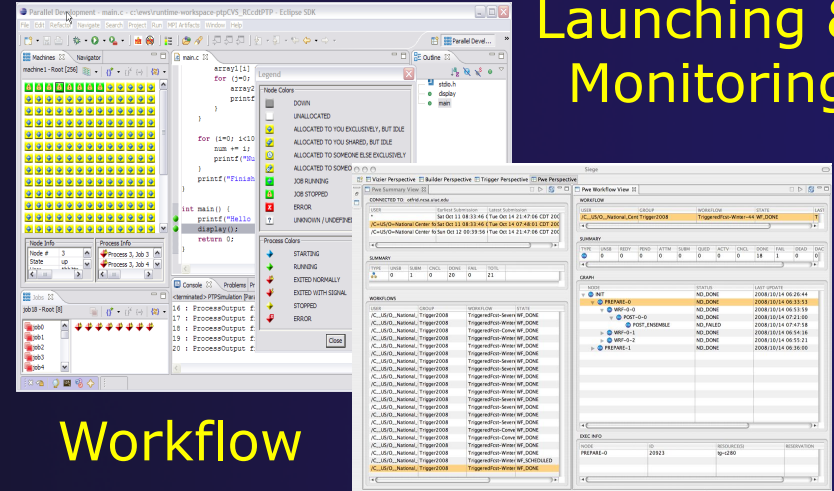
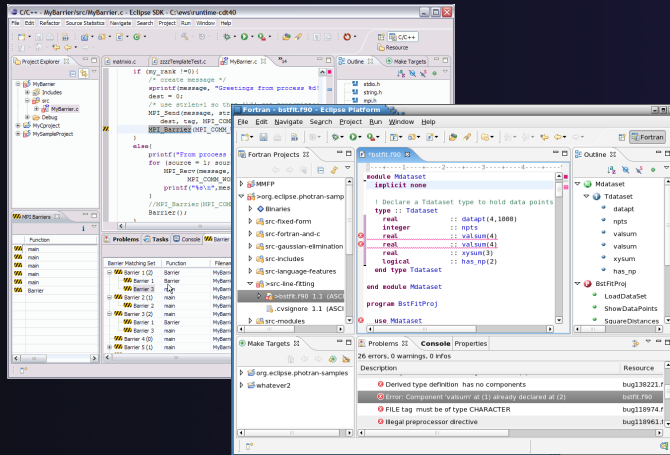
NSF SI2 Workbench for High Performance Computing

- ★ “SI2-SSI Productive and Accessible Development Workbench for HPC Applications”, which is supported by the National Science Foundation under award number OCI 1047956
- ★ Produce a productive and accessible development workbench using Eclipse PTP
- ★ Key Components
 - ★ Determining Requirements, Ensuring Impact
 - ★ Make improvements to Eclipse PTP
 - ★ Engineering Process
 - ★ Metrics
 - ★ Outreach/Training/Education

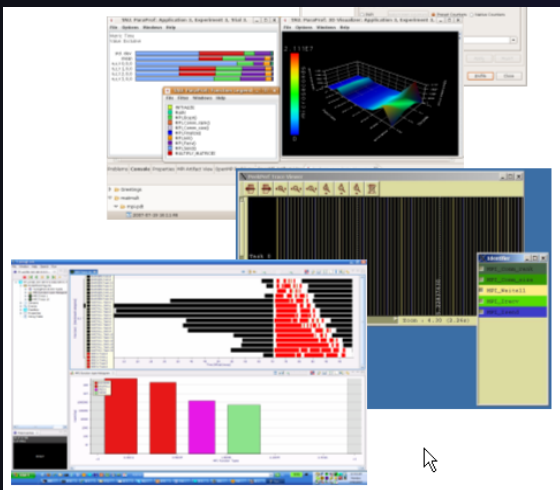
Coding & Analysis (C/C++, Fortran)

NCSA HPC Workbench

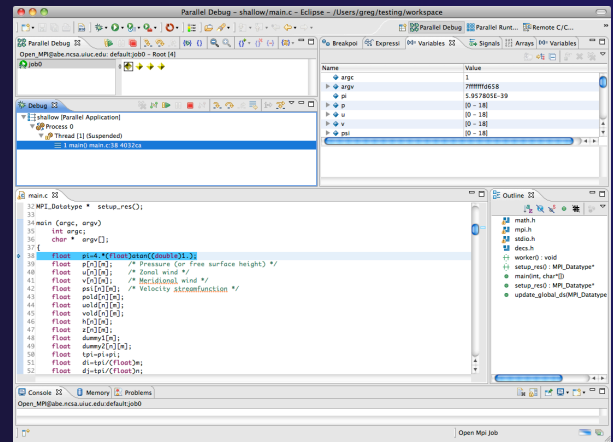
PTP Launching & Monitoring



Workflow



Performance Tuning



Parallel Debugger

Planned PTP Future Work

- ★ Scalability improvements
 - ★ UI to support 1M processes
 - ★ Very large application support
- ★ Usability improvements
 - ★ New wizard to improve setup experience
 - ★ Ability to share configuration information
- ★ Resource Managers
 - ★ More implementations of configurable resource managers
- ★ Synchronized project improvements
 - ★ Conversion wizard
 - ★ Resolving merge conflicts

Useful Eclipse Tools

- ✦ Linux Tools (autotools, valgrind, Oprofile, Gprof)
 - ✦ <http://eclipse.org/linuxtools>
- ✦ Python
 - ✦ <http://pydev.org>
- ✦ Ruby
 - ✦ <http://www.apтана.com/products/radrails>
- ✦ Perl
 - ✦ <http://www.epic-ide.org>
- ✦ Git
 - ✦ <http://www.eclipse.org/egit>
- ✦ VI bindings
 - ✦ Vrapper (open source) - <http://vrapper.sourceforge.net>
 - ✦ viPlugin (commercial) - <http://www.viplugin.com>

Online Information

- ★ Information about PTP
 - ★ Main web site for downloads, documentation, etc.
 - ★ <http://eclipse.org/ptp>
 - ★ Wiki for designs, planning, meetings, etc.
 - ★ <http://wiki.eclipse.org/PTP>
 - ★ Articles and other documents
 - ★ <http://wiki.eclipse.org/PTP/articles>

- ★ Information about Photran
 - ★ Main web site for downloads, documentation, etc.
 - ★ <http://eclipse.org/photran>
 - ★ User's manuals
 - ★ <http://wiki.eclipse.org/PTP/photran/documentation>

Mailing Lists

- ★ PTP Mailing lists
 - ★ Major announcements (new releases, etc.) - low volume
 - ★ <http://dev.eclipse.org/mailman/listinfo/ptp-announce>
 - ★ User discussion and queries - medium volume
 - ★ <http://dev.eclipse.org/mailman/listinfo/ptp-user>
 - ★ Developer discussions - high volume
 - ★ <http://dev.eclipse.org/mailman/listinfo/ptp-dev>
- ★ Photran Mailing lists
 - ★ User discussion and queries
 - ★ <http://dev.eclipse.org/mailman/listinfo/photran>
 - ★ Developer discussions –
 - ★ <http://dev.eclipse.org/mailman/listinfo/photran-dev>

Getting Involved

- ★ See <http://eclipse.org/ptp>
- ★ Read the developer documentation on the wiki
- ★ Join the mailing lists
- ★ Attend the monthly developer meetings
 - ★ Conf Call Monthly: Second Tuesday, 1:00 pm ET
 - ★ Details on the PTP wiki
- ★ Attend the monthly user meetings
 - ★ Teleconference Monthly
 - ★ Each 4th Wednesday, 2:00 pm ET
 - ★ Details on the PTP wiki

PTP will only succeed with your participation!

PTP Tutorial Feedback

- ★ Please complete feedback form
- ★ Your feedback is valuable!

Thanks for attending
We hope you found it useful