Introduction to Application Performance Analysis with CrayPat

Ilias Katsardis
ikatsardis@cray.com

Performance Optimization

We want to get the most science through a supercomputing system as possible

The more efficient codes are the more productive scientists and engineers can be

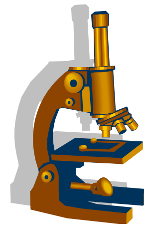

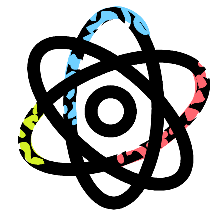

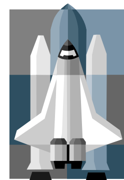

Performance Optimization

- **Adapting the problem to the underlying hardware**
- **Combination of many aspects**
  - Effective algorithms
  - Implementation: Processor utilization & efficient memory use
  - Parallel scalability
- **Important to understand interactions**
  - Algorithm – code – compiler – libraries – hardware
- **Performance is not portable!**

Performance analysis

CRAY

To optimise code we must know *what* is taking the time

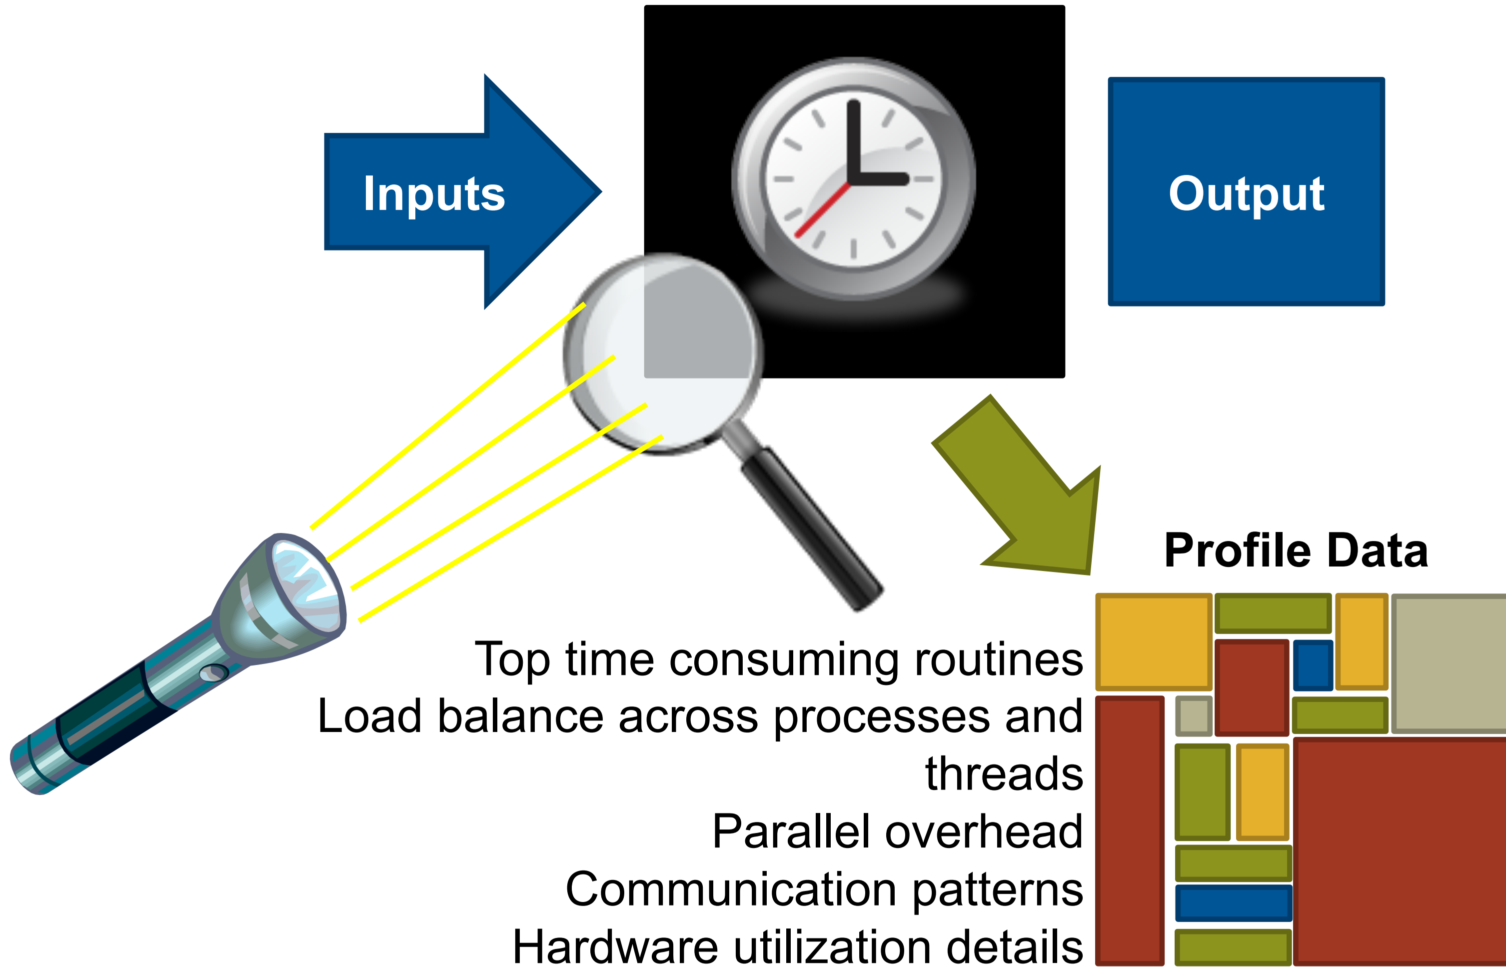

Not going to touch the source code?

- **Find the *compiler* and its *compiler flags* that yield the best performance**
- **Employ *tuned libraries* wherever possible**
- **Find suitable settings for *environment parameters***
- **Mind the *I/O***
  - Do not checkpoint too often
  - Do not ask for the output you do not need

Why does scaling end?

- **Amount of data per process small - computation takes little time compared to communication**
- **Amdahl's law in general**
  - E.g., single-writer or stderr I/O
- **Load imbalance**
- **Communication that scales badly with N_{proc}**
  - E.g., all-to-all collectives
- **Congestion on network – too many messages or lots of data**

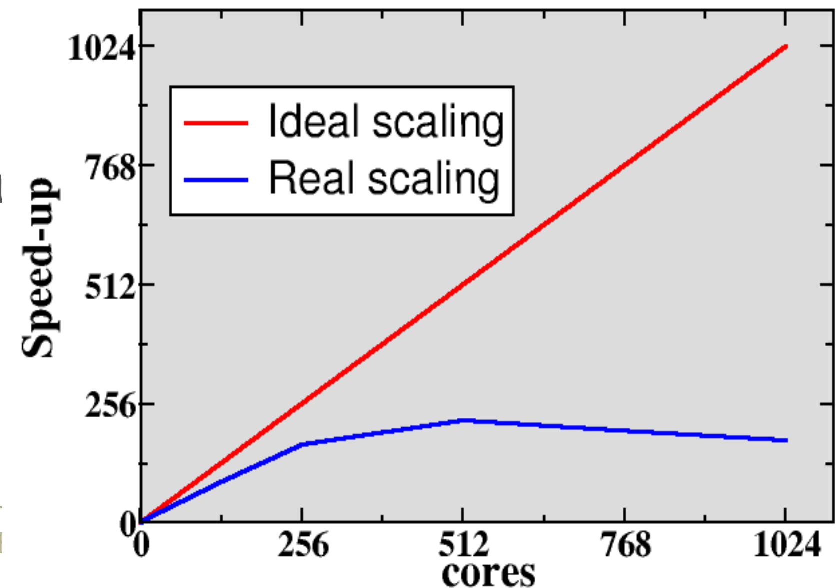

Application timing

- **Most basic information: total wall clock time**
  - Built-in timers in the program (e.g. MPI_Wtime)
  - System commands (e.g. time) or batch system statistics
- **Built-in timers can provide also more fine-grained information**
  - Have to be inserted by hand
  - Typically, no information about hardware related issues e.g. cache utilization
  - Information about load imbalance and communication statistics of parallel program is difficult to obtain

Performance analysis tools

- **Instrumentation of code**
  - Adding special measurement code to binary
    - Special commands, compiler/linker wrappers
    - Automatic or manual
  - Normally all routines do not need to be measured
- **Measurement: running the instrumented binary**
  - Profile: sum of events over time
  - Trace: sequence of events over time
- **Analysis**
  - Text based analysis reports
  - Visualization

CRAY

Sampling

Advantages

- Only need to instrument main routine
- Low Overhead – depends only on sampling frequency
- Smaller volumes of data produced

Disadvantages

- Only statistical averages available
- Limited information from performance counters

Event Tracing

Advantages

- More accurate and more detailed information
- Data collected from every traced function call not statistical averages

Disadvantages

- Increased overheads as number of function calls increases
- Huge volumes of data generated

Guided tracing = trace only program parts that consume a significant portion of the total time

In Cray Performance Analysis Toolkit this is referred to as ”automatic profiling analysis”(APA)

Step 1: Choose a test problem

- **The dataset used in the analysis should**
  - Make scientific sense, i.e. resemble the intended use of the code
  - Be large enough for getting a good view on scalability
  - Be runable in a reasonable time
  - For instance, with simulation codes almost a full-blown model but run only for a few time steps
- **Should be run long enough that initialization/finalization stages are not exaggerated**
  - Alternatively, we can exclude them during the analysis

Step 2: Measure Scalability

- Run the uninstrumented code with different core counts and see where the parallel scaling stops
- Usually we look at strong scaling (fixed problem size)
  - Also weak scaling (fixed amount of work per cpu) is definitely of interest

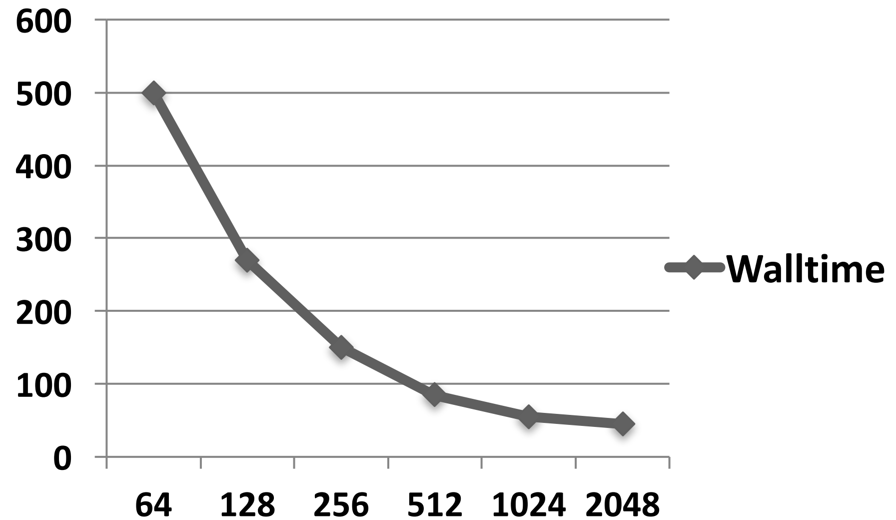

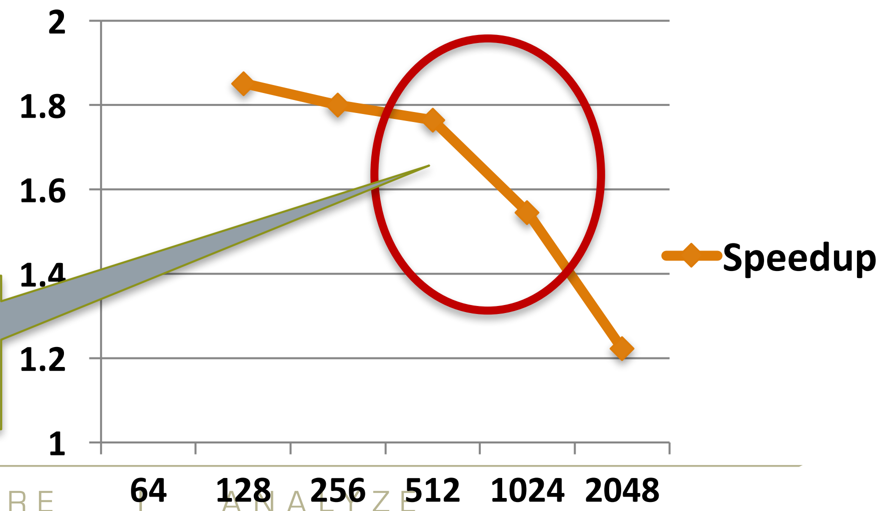

Step 3: Run instrumented version of application

- **Obtain first a sampling profile to find which user functions should be traced**
  - With a large/complex software, one should not trace them all: it causes excessive overhead
- **Make an instrumented exe with tracing time-consuming user functions plus e.g. MPI, I/O and library (BLAS, FFT,...) calls**
- **Execute and record the first analysis with**
  - The core count where the scalability is still ok
  - The core count where the scalability has ended

  and identify the largest differences between these profiles
- **CrayPat has an Automatic Profile Analysis (APA) mode to handle this process:**

Steps to Collect Performance Data

- **Access performance tools software**
  - `module load perftools-base`
  - `module load perftools-lite`
- **Build instrumented version of the application keeping .o files (CCE: `-h keepfiles`)**
  - `make clean`
  - `make`
  - You should get an instrumented version program `a.out`
  - This has been instrumented for sampling (automatic profiling analysis), check with
    - strings a.out | grep 'CrayPat/X'

      CrayPat/X: Version 6.3.0 Revision 14319 09/02/15 13:51:12
- **Run application to get top time consuming routines**
  - `aprun … a.out` *(or* `qsub <pat script>`*)*
  - You should get *.rpt and a *.ap2 files
  - The report in *.rpt is additionally printed to stdout

Example: Sampling report

```
$> make
...
INFO: A maximum of 51 functions from group 'io' will be traced.
INFO: A maximum of 208 functions from group 'mpi' will be traced.
INFO: A maximum of 20 functions from group 'realtime' will be traced.
INFO: A maximum of 56 functions from group 'syscall' will be traced.
INFO: creating the CrayPat-instrumented executable
'/a/certain/dir/cp2k.pdbg' (sample_profile) ...OK

> cat job.out
...
################################################################
#                                                              #
#            CrayPat-lite Performance Statistics               #
#                                                              #
################################################################

CrayPat/X:  Version 6.3.0 Revision 14378 (xf 14041)  09/15/15 10:48:06
Experiment:                  lite  lite/sample_profile
Number of PEs (MPI ranks):     48
Numbers of PEs per Node:       24  PEs on each of  2  Nodes
Numbers of Threads per PE:      1
Number of Cores per Socket:    12
Execution start time:  Wed Oct 14 14:07:17 2015
System name and speed:  mom11 2501 MHz

Avg Process Time:          5.14 secs
High Memory:              2,070 MBytes
MFLOPS:             Not supported (see ob
I/O Read Rate:          4.803892 MBytes/
I/O Write Rate:        88.963763 MBytes/sec
Avg CPU Energy:           1,499 joules 749.50 joules per node
Avg CPU Power:           291.59 watts  145.80 watts per node
```

General job information

```
...
Table 1:  Profile by Function Group and Function (top 8 functions shown)

  Samp% |  Samp |  Imb. |  Imb. |Group
        |       |  Samp | Samp% | Function
        |       |       |       |  PE=HIDE

 100.0% | 263.4 |    -- |    -- |Total
|-----------------------------------------------
|  78.0% | 205.3 |    -- |    -- |MPI
||-----------------------------------------------------------------------
||  62.4% | 164.4 | 115.6 | 42.2% |mpi_bcast
||  10.4% |  27.4 | 186.6 | 89.1% |MPI_ALLREDUCE
||   4.7% |  12.4 |  86.6 | 89.3% |MPI_IPROBE
||=======================================================================
|  13.1% |  34.5 |    -- |    -- |USER
||-----------------------------------------------------------------------
||   3.3% |   8.6 |  61.4 | 89.5% |__message_passing_MOD_mp_probe
||   2.8% |   7.5 |   8.5 | 54.4% |__fist_nonbond_force_MOD_force_nonbond
||   2.0% |   5.2 |   5.8 | 53.6% |__ewalds_MOD_ewald_evaluate
||   1.1% |   2.9 |   3.1 | 52.5% |__splines_methods_MOD_potential_s
||=======================================================================
|   8.2% |  21.5 |    -- |    -- |ETC
||-----------------------------------------------------------------------
||   2.5% |   6.6 |   9.4 | 59.7% |__memmove_ssse3
||   1.7% |   4.4 |   4.6 | 52.7% |__memset_sse2
|=========================================================================..
...
```

Portions of samples

Significant portion of communication

Example: Sampling report (2)

```
...
==================  Observations and suggestions  ==================
Metric-Based Rank Order:

  When the use of a shared resource like memory bandwidth is unbalanced
  across nodes, total execution time may be reduced with a rank order
  that improves the balance.  The metric used here for resource usage
  is: USER Samp

  For each node, the metric values for the ranks on that node are
  summed.  The maximum and average value of those sums are shown below
  for both the current rank order and a custom rank order that seeks
  to reduce the maximum value.

  A file named MPICH_RANK_ORDER.USER_Samp was generated
  along with this report and contains usage instructions and the
  Custom rank order from the following table.

    Rank    Node Reduction    Maximum  Average
    Order  Metric    in Max      Value  Value
             Imb.     Value

   Current  11.17%           9.310e+02  8.270e+02
   Custom   2.59%     8.808%  8.490e+02  8.270e+02
=========================  End Observations  =========================
...
```

Rank reorder suggestions

```
...
Table 2:  File Input Stats by Filename
     Read |      Read |  Read Rate |   Reads |    Bytes/ |File Name[max15]
     Time |    MBytes |  MBytes/sec |         |     Call | PE=HIDE
 0.113291 | 0.544238 |    4.803892 | 2,964.0 |    192.54 |Total
|----------------------------------------------------------------------------
| 0.057170 | 0.214447 |    3.751054 | 1,586.0 |    141.78 |topology_fist_WAT.psf
| 0.026845 | 0.138477 |    5.158328 |   844.0 |    172.04 |H2O_ice.inp
| 0.014117 | 0.000700 |    0.049586 |     3.0 |    244.67 |TMC_NPT.inp
| 0.007784 | 0.098442 |   12.646622 |   176.0 |
| 0.006957 | 0.078669 |   11.307646 |    25.0 |
|==============================================

Table 3:  File Output Stats by Filename
    Write |     Write | Write Rate |  Writes |
     Time |    MBytes | MBytes/sec |         |     Call | PE=HIDE
 0.162883 | 14.490714 |  88.963763 | 5,203.0 | 2,920.36 |Total
|----------------------------------------------------------------------------
| 0.096137 | 13.861026 | 144.179480 | 3,805.0 | 3,819.80 |tmc_traj_T270.xyz
| 0.021800 |  0.064217 |   2.945740 |    18.0 | 3,740.89 |tmc_E_worker_1.out
| 0.016016 |  0.064296 |   4.014441 |    18.0 | 3,745.50 |tmc_E_worker_6.out
| 0.013735 |  0.155310 |  11.307340 |   761.0 |   214.00 |tmc_traj_T270.cell
| 0.004775 |  0.063504 |  13.300140 |    18.0 | 3,699.39 |tmc_E_worker_7.out
| 0.003025 |  0.026007 |   8.596676 |   505.0 |    54.00 |stdout
| 0.001983 |  0.064375 |  32.470347 |    19.0 | 3,552.74 |tmc_E_worker_3.out
| 0.001915 |  0.064375 |  33.624425 |    19.0 | 3,552.74 |tmc_E_worker_2.out
| 0.001905 |  0.063979 |  33.588895 |    18.0 | 3,727.06 |tmc_E_worker_4.out
| 0.001582 |  0.063504 |  40.142573 |    18.0 | 3,699.39 |tmc_E_worker_5.out
| 0.000011 |  0.000122 |  11.053907 |     4.0 |    32.00 |_UnknownFile_
|============================================================================
```

Input/Output analysis

Steps to Collecting Performance Data (2)

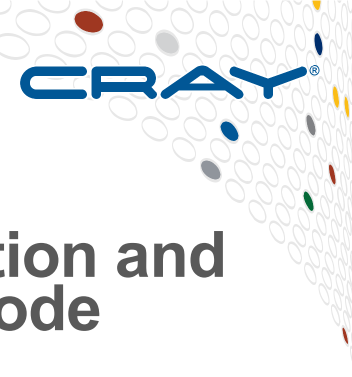

- **At this stage the report gives us useful information and we should get sample hits in time-consuming code sections**

- **We can see more info**
  - pat_report a.out+20199-40s.ap2
  - You should see this printed to stdout
  - Includes
    - job info
    - profile by functions
    - observations and suggestions
    - runtime environment variables
    - hardware performance counter events

- **We can also view graphically with Apprentice^2**

- **We can go further on to tracing and loop profiling**

Example: Tracing report

- Access perftools, then build and run application
  - `module load perftools-base`
  - `module load perftools-lite-event`
  - `make clean; make`
  - `aprun … a.out`
- Comparable to sampling experiment, but now the function are really traced from beginning to end
- Again observations and suggestions are printed
  - E.g. rank reordering
  - And IO observations

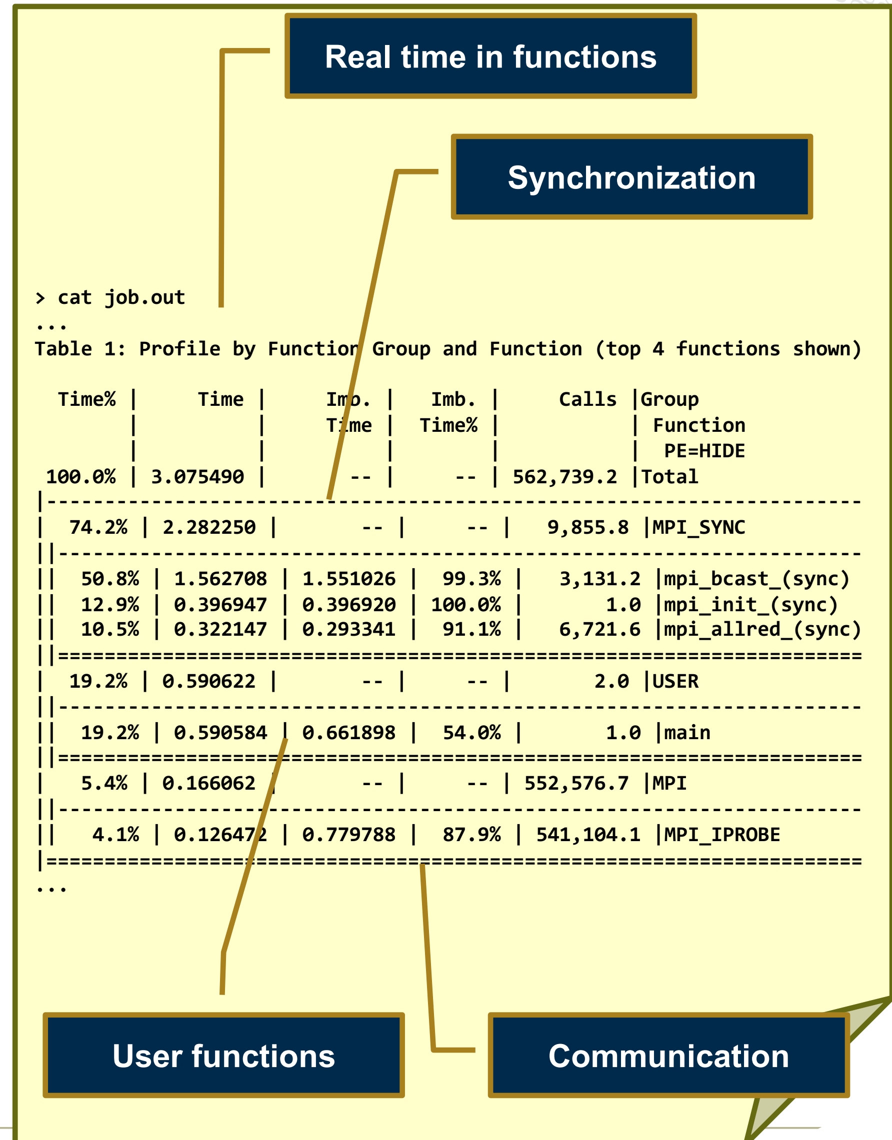


```
> cat job.out
...
Table 1: Profile by Function Group and Function (top 4 functions shown)

  Time% |     Time |     Imb. |  Imb.  |      Calls |Group
        |          |     Time |  Time% |            | Function
        |          |          |        |            |  PE=HIDE
 100.0% | 3.075490 |       -- |     -- | 562,739.2  |Total
|------------------------------------------------------------------------
|  74.2% | 2.282250 |       -- |     -- |    9,855.8 |MPI_SYNC
||-----------------------------------------------------------------------
||  50.8% | 1.562708 | 1.551026 |  99.3% |    3,131.2 |mpi_bcast_(sync)
||  12.9% | 0.396947 | 0.396920 | 100.0% |        1.0 |mpi_init_(sync)
||  10.5% | 0.322147 | 0.293341 |  91.1% |    6,721.6 |mpi_allred_(sync)
||=======================================================================
|  19.2% | 0.590622 |       -- |     -- |        2.0 |USER
||-----------------------------------------------------------------------
||  19.2% | 0.590584 | 0.661898 |  54.0% |        1.0 |main
||=======================================================================
|   5.4% | 0.166062 |       -- |     -- | 552,576.7 |MPI
||-----------------------------------------------------------------------
||   4.1% | 0.126472 | 0.779788 |  87.9% | 541,104.1 |MPI_IPROBE
|========================================================================
...
```

Example: Generate a loop Profile

- **Access performance tools software, provide basic tools and environment settings**
  - `module load perftools-base`

- **Set environment for tracing experiments with loop profiling**
  - `module load perftools-loops`

- **Build instrumented version of the application**
  - `make clean`
  - `make`
  - You should get an instrumented version program `a.out`
  - This has been instrumented for sampling (automatic profiling analysis), check with
    - strings a.out | grep 'CrayPat/X'
      CrayPat/X: Version 6.3.0 Revision 14319 09/02/15 13:51:12

- **Run application to get top time consuming routines**
  - `aprun … a.out` *(or* `qsub <pat script>`*)*
  - You should get *.rpt and a *.ap2 files
  - The report in *.rpt is additionally printed to stdout

Example: Generate a loop Profile

CRAY

Subroutine

Line number

Table 1: Inclusive and Exclusive Time in Loops (from -hprofile_generate)

Loop Incl Time%	Loop Incl Time	Time (Loop Adj.)	Loop Hit	Loop Trips Avg	Loop Trips Min	Loop Trips Max	Function=/.LOOP[.] PE=HIDE
93.0%	19.232051	0.000849	2	26.5	3	50	jacobi.LOOP.1.li.236
77.8%	16.092021	0.001350	53	255.0	255	255	jacobi.LOOP.2.li.240
77.8%	16.090671	0.110827	13515	255.0	255	255	jacobi.LOOP.3.li.241
77.3%	15.979844	15.979844	3446325	511.0	511	511	jacobi.LOOP.4.li.242
14.1%	2.906115	0.001238	53	255.0	255	255	jacobi.LOOP.5.li.263
14.0%	2.904878	0.688611	13515	255.0	255	255	jacobi.LOOP.6.li.264
10.7%	2.216267	2.216267	3446325	511.0	511	511	jacobi.LOOP.7.li.265
4.3%	0.881573	0.000010	1	259.0	259	259	initmt.LOOP.1.li.191
4.3%	0.881563	0.000645	259	259.0	259	259	initmt.LOOP.2.li.192
4.3%	0.880918	0.880918	67081	515.0	515	515	initmt.LOOP.3.li.193
2.7%	0.560499	0.000055	1	257.0	257	257	initmt.LOOP.4.li.210
2.7%	0.560444	0.006603	257	257.0	257	257	initmt.LOOP.5.li.211
2.7[illegible]	[illegible]	[illegible]3842	66049	513.0	513	513	initmt.LOOP.6.li.212

Nested Loops

perftools-lite vs. perftools

- **There are two ways of using CrayPat**

- **perftools-lite**
  - An entry-level approach
  - Aimed at users unfamiliar with the full perftools framework
  - Provides a report automatically at the end of the job
    - Measures the basic set of performance statistics

- **perftools**
  - A more advanced environment
  - Provides full control over the performance statistics collected
  - Requires a few more steps from the user

- **Both generate results as:**
  - a text report
  - a data file (*.ap2) that can be explored using a GUI (Cray Apprentice^2)

Steps to Collect Performance Data with perftools

- **Access performance tools software**
  - `module load perftools-base`
  - `module load perftools`
- **Build application keeping .o files (CCE: `-h keepfiles`)**
  - `make clean`
  - `make`
- **Instrument application for automatic profiling analysis**
  - `pat_build -O apa a.out`
  - You should get an instrumented program `a.out+pat`
  - This has been instrumented for sampling
- **Run application to get top time consuming routines**
  - `aprun … a.out+pat` *(or* `qsub <pat script>`*)*
  - You should get one or more *.xf performance files

Steps to Collecting Performance Data with perftools (2)

- **Run pat_report, on the .xf file or the directory**
  - `pat_report -o <report> <xf file>`
  - `pat_report -o <report> <xf directory>`
  - Generates text report and an .apa instrumentation file
    - We'll discuss pat_report in more detail later
- **At this stage the report gives us useful information and we should get sample hits in time-consuming code sections**
- **We use the .apa file to re-instrument binary for tracing**
  - the most important functions have been identified for tracing
- **We can inspect and edit the .apa file at this point**
  - if we want to tweak the choice of routines to be traced

APA File Example

```
#  You can edit this file, if desired, and use it
#   to reinstrument the program for  tracing like this:
#
#           pat_build -O standard.cray-xt.PE-2.1.56HD.pgi-8.0.amd64.pat-5.0.0.2-Oapa.512.quad.cores.seal.090405.1154.mpi.pat_rt_exp=default.pat_rt_hwpc=none.14999.xf.xf.apa
#
#  These suggested trace options are based on data from:
#
#
/home/users/malice/pat/Runs/Runs.seal.pat5001.2009Apr04/./pat.quad/homme/standard.cray-xt.PE-2.1.56HD.pgi-8.0.amd64.pat-5.0.0.2-Oapa.512.quad.cores.seal.090405.1154.mpi.pat_rt_exp=default.pat_rt_hwpc=none.14999.xf.xf.cdb
# ----------------------------------------------------------------------

#       HWPC group to collect by default.

  -Drtenv=PAT_RT_HWPC=1  # Summary with TLB metrics.

# ----------------------------------------------------------------------

#       Libraries to trace.

  -g mpi

# ----------------------------------------------------------------------

#       User-defined functions to trace, sorted by % of samples.

#       The way these functions are filtered can be controlled with
#       pat_report options (values used for this file are shown):
#
#       -s apa_max_count=200    No more than 200 functions are listed.
#       -s apa_min_size=800     Commented out if text size < 800 bytes.
#       -s apa_min_pct=1        Commented out if it had < 1% of samples.
#       -s apa_max_cum_pct=90   Commented out after cumulative 90%.

#       Local functions are listed for completeness, but cannot be traced.

  -w  # Enable tracing of user-defined functions.
      # Note: -u should NOT be specified as an additional option.
```

```
# 31.29%  38517 bytes
         -T prim_advance_mod_preq_advance_exp_

# 15.07%  14158 bytes
         -T prim_si_mod_prim_diffusion_

#  9.76%  5474 bytes
         -T derivative_mod_gradient_str_nonstag_

. . .

#  2.95%  3067 bytes
         -T forcing_mod_apply_forcing_

#  2.93%  118585 bytes
         -T column_model_mod_applycolumnmodel_

#  Functions below this point account for less than 10% of samples.

#  0.66%  4575 bytes
#         -T bndry_mod_bndry_exchangev_thsave_time_

#  0.10%  46797 bytes
#         -T baroclinic_inst_mod_binst_init_state_

#  0.04%  62214 bytes
#         -T prim_state_mod_prim_printstate_

. . .
#  0.00%  118 bytes
#         -T time_mod_timelevel_update_

# ----------------------------------------------------------------------

  -o preqx.cray-xt.PE-2.1.56HD.pgi-8.0.amd64.pat-5.0.0.2.x+apa
# New instrumented program.

/.AUTO/cray/css.pe_tools/malice/craypat/build/pat/2009Apr03/2.1.56HD/amd64/homme/pgi/pat-5.0.0.2/homme/2005Dec08/build.Linux/preqx.cray-xt.PE-2.1.56HD.pgi-8.0.amd64.pat-5.0.0.2.x  # Original program.
```

Effectively a series of command line arguments to pat_build

Generating Event Traced Profile from APA

- **Re-instrument application for further analysis**
  - `pat_build -O <apa file>`
  - creates new binary: `<exe>+apa`

- **Re-run application**
  - `aprun … a.out+apa` (or `qsub <apa script>`)
  - This generates a new set of .xf data files

- **Generate new text report and visualization file (.ap2)**
  - `pat_report -o <report> <xf file>`
  - `pat_report -o <report> <xf directory>`

- **View report in text and/or with Cray Apprentice2**
  - `app2 <ap2 file>`
  - We'll cover this in more detail later

Steps to Using CrayPat with perftools-lite

Access light version of performance tools software

```
> module load perftools-base
> module load perftools-lite
```

Build program

```
> make
```

a.out (instrumented program)

Run program (no modification to batch script)

```
aprun a.out
```

Condensed report to stdout
a.out*.rpt (same as stdout)
a.out*.ap2
MPICH_RANK_XXX files

Steps to Using CrayPat “classic” with perftools

Access performance tools software

```
> module load perftools-base
> module load perftools
```

Build program, retaining .o files

```
> make
```

a.out

Instrument binary

```
> pat_build –O apa a.out
```

a.out+pat

Modify batch script and run program

```
aprun a.out+pat
```

a.out+pat*.xf

Process raw performance data and create report

```
> pat_report a.out+pat*.xf
```

a.out+pat*.ap2
Text report to stdout
a.out+pat*.apa
MPICH_RANK_XXX

CrayPat (perftools) vs CrayPat (perftools-lite)

- **Both use the same process under the hood**
- **With perftools-lite pat_build runs automatically when the code is linked**
  - but keeps the same executable name
- **The `sample_profile` is equivalent to**
  - `pat_build -O apa a.out`
  - `CRAYPAT_LITE = sample_profile (perftools-lite)`
- **The `event_profile` is equivalent to**
  - `pat_build -u -gmpi a.out`
  - CRAYPAT_LITE = event_profile (perftools-lite-event)
- **It also runs pat_report automatically**
  - at the end of the job

Analysing Data with pat_report

Using pat_report

- **pat_report converts raw profiling data into a profile**
  - Combines .xf data with binary
    - Instrumented binary must still exist when data is converted!
  - Produces a text report and an .ap2 file
  - .ap2 file can be used for further pat_report calls or display in GUI

- **Generates a text report of performance results**
  - Data laid out in tables
  - Many options for sorting, slicing or dicing data in the tables.
    - `pat_report -O <table option> *.ap2`
    - `pat_report -O help` (list of available profiles)
  - Volume and type of information depends upon sampling vs tracing.

Advantages of the .ap2 file

- **.ap2 file is a self contained compressed performance file**
  - Normally it is about 5 times smaller than the .xf file
  - Contains the information needed from the application binary
  - Can be reused
- **Independent of the perftools version used to generate it**
  - The xf files are very version-dependent
- **It is the only input format accepted by Cray Apprentice^2**

- **Once you have the .ap2 file, you can delete:**
  - the .xf files
  - the instrumented binary

Files Generated and the Naming Convention

File Suffix	Description
a.out+pat	Program instrumented for data collection
a.out…s.xf	Raw data from sampling experiment available after application execution
a.out…t.xf	Raw data from trace (summarized or full) experiment available after application execution
a.out….ap2	Processed data, generated by pat_report, contains application symbol information
a.out…s.apa	Automatic profiling analysis template, generated by pat_report (based on pat_build -O apa experiment)
a.out+apa	Program instrumented using .apa file
MPICH_RANK_ORDER.Custom	Rank reorder file generated by pat_report from automatic grid detection an reorder suggestions

Job Execution Information

```
CrayPat/X:  Version 5.2.3.8078 Revision 8078 (xf 8063)  08/25/11 …

Number of PEs (MPI ranks):   16

Numbers of PEs per Node:     16

Numbers of Threads per PE:    1

Number of Cores per Socket:  12

Execution start time:  Thu Aug 25 14:16:51 201

System type and speed:  x86_64 2000 MHz

Current path to data file:
  /lus/scratch/heidi/ted_swim/mpi-openmp/run/swim+pat+27472-34t.ap2

Notes for table 1:
…
```

Sampling Output (Table 1)

```
Notes for table 1:

...

Table 1:  Profile by Function

 Samp % | Samp |  Imb.  |   Imb.  |Group
        |      |  Samp  | Samp %  | Function
        |      |        |         |  PE='HIDE'

 100.0% |  775 |    -- |      -- |Total
|------------------------------------------
|  94.2% |  730 |    -- |      -- |USER
||-----------------------------------------
||  43.4% |  336 |  8.75 |    2.6% |mlwxyz_
||  16.1% |  125 |  6.28 |    4.9% |half_
||   8.0% |   62 |  6.25 |    9.5% |full_
||   6.8% |   53 |  1.88 |    3.5% |artv_
||   4.9% |   38 |  1.34 |    3.6% |bnd_
||   3.6% |   28 |  2.00 |    6.9% |currenf_
||   2.2% |   17 |  1.50 |    8.6% |bndsf_
||   1.7% |   13 |  1.97 |   13.5% |model_
||   1.4% |   11 |  1.53 |   12.2% |cfl_
||   1.3% |   10 |  0.75 |    7.0% |currenh_
||   1.0% |    8 |  5.28 |   41.9% |bndbo_
||   1.0% |    8 |  8.28 |   53.4% |bndto_
||=========================================
|   5.4% |   42 |    -- |      -- |MPI
||-----------------------------------------
||   1.9% |   15 |  4.62 |   23.9% |mpi_sendrecv_
||   1.8% |   14 | 16.53 |   55.0% |mpi_bcast_
||   1.7% |   13 |  5.66 |   30.7% |mpi_barrier_
|==========================================
```

pat_report: Flat Profile

```
Table 1:  Profile by Function Group and Function

 Time % |        Time |Imb. Time |   Imb. | Calls |Group
        |             |          | Time % |       | Function
        |             |          |        |       |  PE='HIDE'

 100.0% | 104.593634 |       -- |     -- | 22649 |Total
|-----------------------------------------------------------
|  71.0% |  74.230520 |       -- |     -- | 10473 |MPI
||----------------------------------------------------------
||  69.7% |  72.905208 | 0.508369 |   0.7% |   125 |mpi_allreduce_
||   1.0% |   1.050931 | 0.030042 |   2.8% |    94 |mpi_alltoall_
||==========================================================
|  25.3% |  26.514029 |       -- |     -- |    73 |USER
||----------------------------------------------------------
||  16.7% |  17.461110 | 0.329532 |   1.9% |    23 |selfgravity_
||   7.7% |   8.078474 | 0.114913 |   1.4% |    48 |ffte4_
||==========================================================
|   2.5% |   2.659429 |       -- |     -- |   435 |MPI_SYNC
||----------------------------------------------------------
||   2.1% |   2.207467 | 0.768347 |  26.2% |   172 |mpi_barrier_(sync)
||==========================================================
|   1.1% |   1.188998 |       -- |     -- | 11608 |HEAP
||----------------------------------------------------------
||   1.1% |   1.166707 | 0.142473 |  11.1% |  5235 |free
```

pat_report: Message Stats by Caller

```
Table 4:  MPI Message Stats by Caller

    MPI Msg |MPI Msg |  MsgSz |  4KB<= |Function
      Bytes |  Count |   <16B |  MsgSz | Caller
            |        |  Count |  <64KB |  PE[mmm]
            |        |        |  Count |

 15138076.0 | 4099.4 |  411.6 | 3687.8 |Total
|-----------------------------------------------
| 15138028.0 | 4093.4 |  405.6 | 3687.8 |MPI_ISEND
||----------------------------------------------
||   8080500.0 | 2062.5 |   93.8 | 1968.8 |calc2_
3|             |        |        |        | MAIN_
|||---------------------------------------------
4||   8216000.0 | 3000.0 | 1000.0 | 2000.0 |pe.0
4||   8208000.0 | 2000.0 |     -- | 2000.0 |pe.9
4||   6160000.0 | 2000.0 |  500.0 | 1500.0 |pe.15
|||=============================================
||   6285250.0 | 1656.2 |  125.0 | 1531.2 |calc1_
3|             |        |        |        | MAIN_
|||---------------------------------------------
4||   8216000.0 | 3000.0 | 1000.0 | 2000.0 |pe.0
4||   6156000.0 | 1500.0 |     -- | 1500.0 |pe.3
4||   6156000.0 | 1500.0 |     -- | 1500.0 |pe.5
|||=============================================
. . .
```

Some important options to pat_report -O

callers	Profile by Function and Callers
callers+hwpc	Profile by Function and Callers
callers+src	Profile by Function and Callers, with Line Numbers
callers+src+hwpc	Profile by Function and Callers, with Line Numbers
calltree	Function Calltree View
heap_hiwater	Heap Stats during Main Program
hwpc	Program HW Performance Counter Data
load_balance_program+hwpc	Load Balance across PEs
load_balance_sm	Load Balance with MPI Sent Message Stats
loop_times	Loop Stats by Function (from -hprofile_generate)
loops	Loop Stats by Inclusive Time (from -hprofile_generate)
mpi_callers	MPI Message Stats by Caller
profile	Profile by Function Group and Function
profile+src+hwpc	Profile by Group, Function, and Line
samp_profile	Profile by Function
samp_profile+hwpc	Profile by Function
samp_profile+src	Profile by Group, Function, and Line

- For a full list see: pat_report -O help

Loop Statistics

- **Just like adding automatic tracing at the function level, we can add tracing to individual loops.**
- **Helps identify candidates for parallelization:**
  - Loop timings approximate how much work exists within a loop
  - Trip counts can be used to understand parallelism potential
    - useful if considering porting to manycore
- **Only available with CCE:**
  - Requires compiler add additional features into the code.
  - Should be done as separate profiling experiment
    - compiler optimizations are restricted with this feature
- **Loop statistics reported by default in `pat_report` table**

Collecting Loop Statistics

- Load PrgEnv-cray module (default on most systems)
- Load perftools module
- Compile **AND** link with CCE flag: `-h profile_generate`
- Instrument binary for tracing
  - All user functions: `pat_build -u my_program`
  - Or even no user functions: `pat_build -w my_program`
    - This is sufficient for loop-level profiling of all loops!
  - Or use an existing apa file.
- Run the application
- Create report with loop statistics
  - `pat_report <xf file> > <report file>`

Default Report Table 2

```
Notes for table 2:
  Table option:
    -O loops
  …
  The Function value for each data item is the avg of the PE values.
    (To specify different aggregations, see:  pat_help report options s1)

  This table shows only lines with Loop Incl Time / Total > 0.0095.
    (To set thresholds to zero, specify:  -T)

Loop instrumentation can interfere with optimizations, so time
  reported here may not reflect time in a fully optimized program.

  Loop stats can safely be used in the compiler directives:
   !PGO$        loop_info est_trips(Avg) min_trips(Min) max_trips(Max)
   #pragma pgo loop_info est_trips(Avg) min_trips(Min) max_trips(Max)

  Explanation of Loop Notes (P=1 is highest priority, P=0 is lowest):
   novec (P=0.5): Loop not vectorized (see compiler messages for reason).
   sunwind (P=1): Loop could be vectorized and unwound.
   vector (P=0.1): Already a vector loop.
```

Profile guided optimization feedback for compiler: see man pgo

Default Report Table 2

```
Table 2:  Loop Stats from -hprofile_generate

   Loop |Loop Incl |Loop Incl |  Loop |  Loop |    Loop |Function=/.LOOP\.
   Incl |      Time |    Time / |   Hit | Trips |   Notes | PE='HIDE'
  Time / |           |       Hit |       |   Avg |         |
  Total |           |           |       |       |         |

|-------------------------------------------------------------------------
|  24.6% | 0.057045 | 0.000570 |   100 |  64.1 |   novec |calc2_.LOOP.0.li.614
|  24.0% | 0.055725 | 0.000009 |  6413 | 512.0 |  vector |calc2_.LOOP.1.li.615
|  18.9% | 0.043875 | 0.000439 |   100 |  64.1 |   novec |calc1_.LOOP.0.li.442
|  18.3% | 0.042549 | 0.000007 |  6413 | 512.0 |  vector |calc1_.LOOP.1.li.443
|  17.1% | 0.039822 | 0.000406 |    98 |  64.1 |   novec |calc3_.LOOP.0.li.787
|  16.7% | 0.038883 | 0.000006 |  6284 | 512.0 |  vector |calc3_.LOOP.1.li.788
|   9.7% | 0.022493 | 0.000230 |    98 | 512.0 |  vector |calc3_.LOOP.2.li.805
|   4.2% | 0.009837 | 0.000098 |   100 | 512.0 |  vector |calc2_.LOOP.2.li.640
|=========================================================================
```

Step 4: Assessing the big picture

- **Profile = Where the most of the time is really being spent?**
  - See also the call-tree view
  - Ignore (from the optimization point-of-view) user routines with less than 5% of the execution time
- **Why does the scaling end: the major differences in these two profiles?**
  - Has the MPI fraction ’blown up’ in the larger run?
  - Have the load imbalances increased dramatically?
  - Has something else emerged to the profile?
  - Has the time spent for user routines decreased as it should (i.e. do they scale independently)?

Example with CrayPat

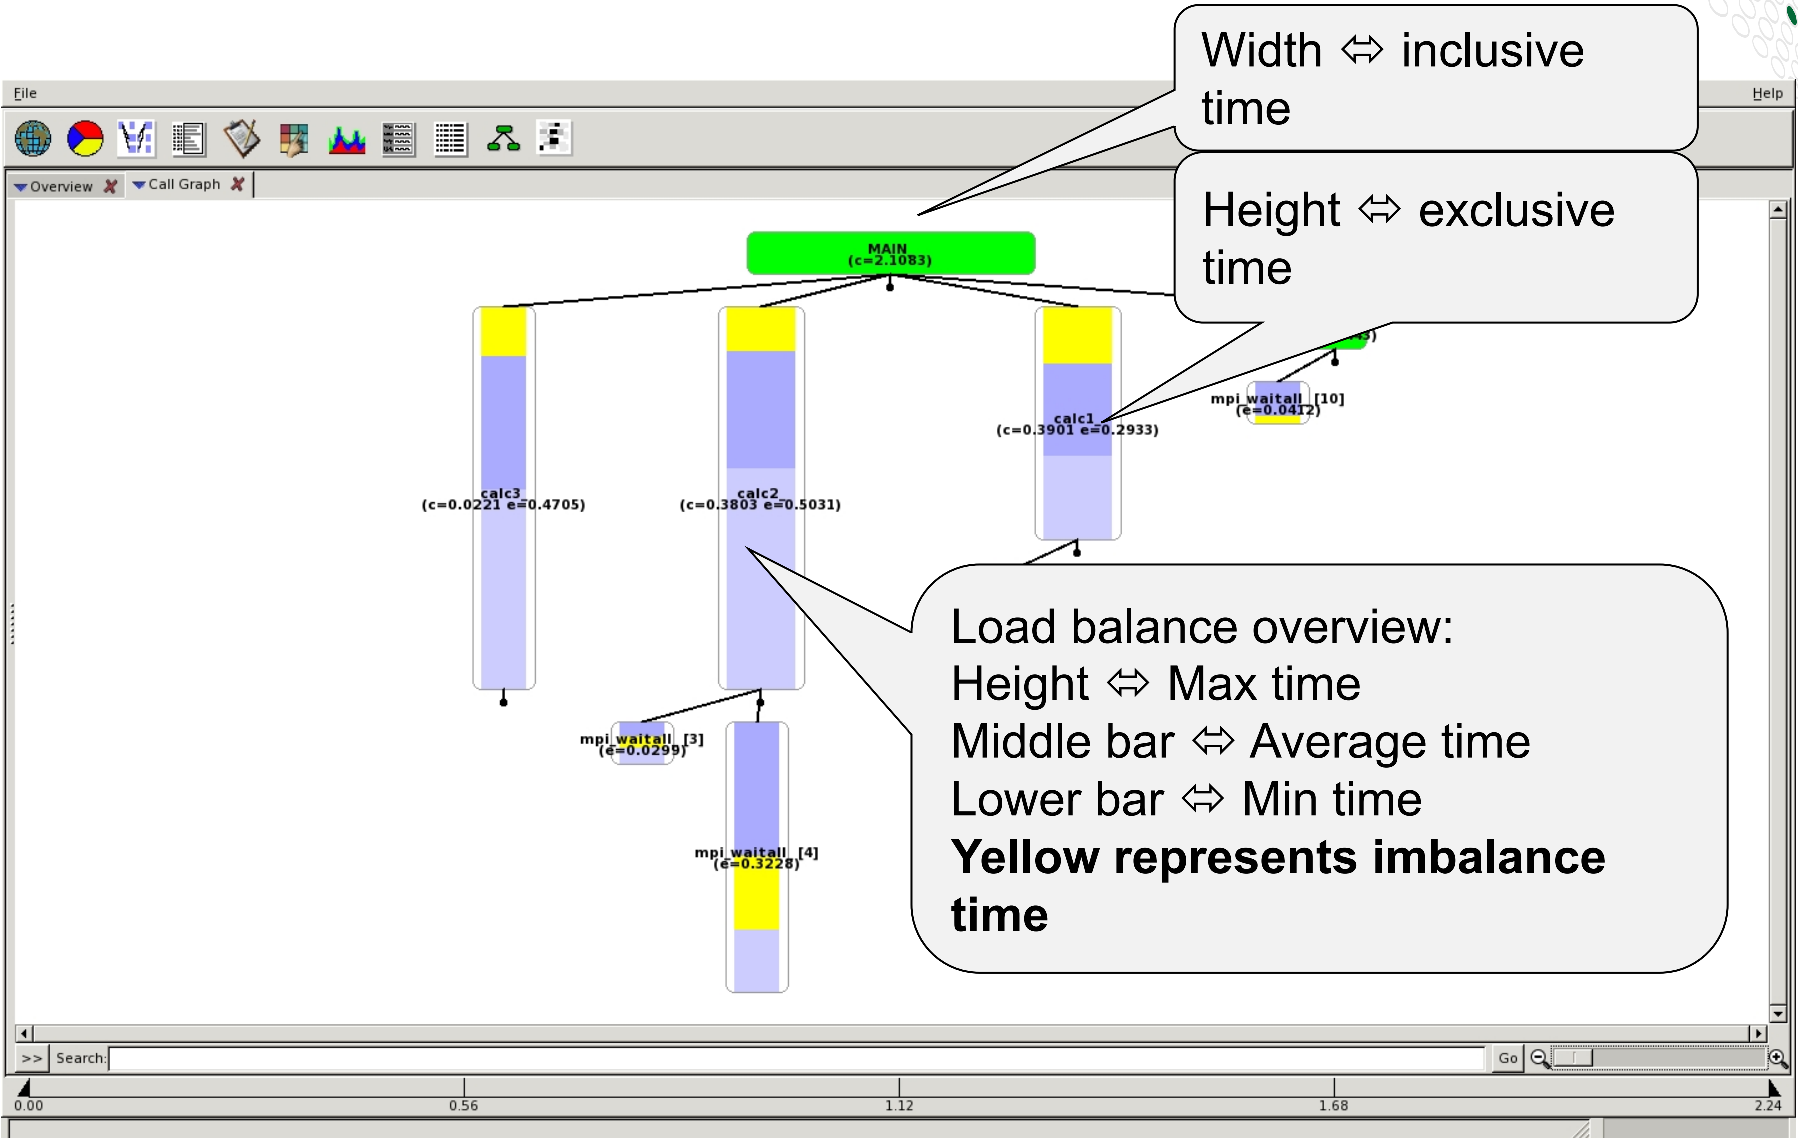

Step 5: Analyze load imbalance

- **What is causing the imbalance?**
- **Computation**
  - Tasks call for computational kernels (user functions, BLAS routines,...) for varying times and/or the execution time varies depending on the input/caller
- **Communication**
  - Large MPI_Sync times
- **I/O**
  - One or more tasks are performing I/O and the others are just waiting for them in order to proceed

Example with CrayPat

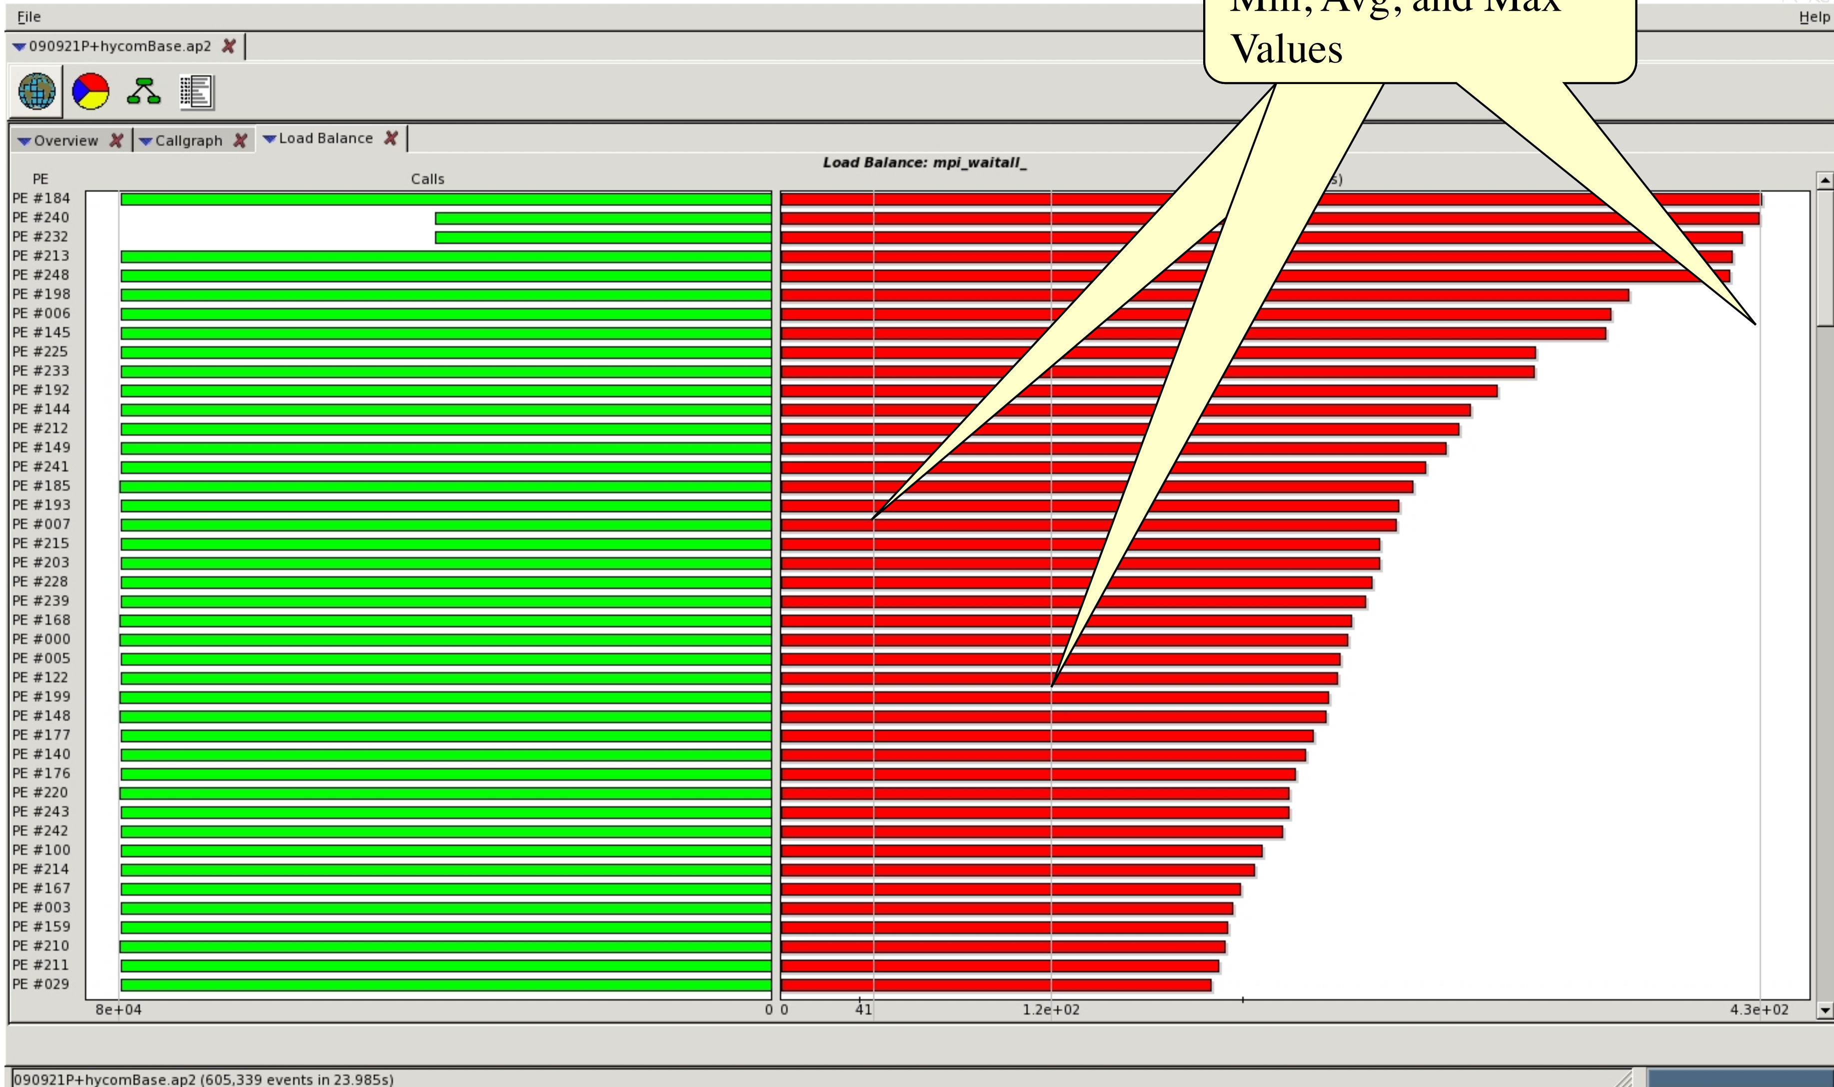

Step 6: Analyze communication

- **What communication pattern is dominating the true time spent for MPI (excluding the sync times)**
  - Refer to the call-tree view on Apprentice2 and the “MPI Message Stats” tables in the text reports produced by pat_report
- **Note that the analysis tools may report load imbalances as ”real” communication**
  - Put an MPI_Barrier before the suspicious routine - load imbalance will aggregate into it in when then analysis is rerun
- **How does the message-size profile look like?**
  - Are there a lot of small messages?

Example with CrayPat report (message stats)

```
Table 4:  MPI Message Stats by Caller

    MPI Msg |MPI Msg |  MsgSz |  4KB<= |Function
      Bytes |   Count |   <16B |  MsgSz | Caller
            |         |  Count |  <64KB |  PE[mmm]
            |         |        |  Count |

 15138076.0 | 4099.4 |  411.6 | 3687.8 |Total
|------------------------------------------------
| 15138028.0 | 4093.4 |  405.6 | 3687.8 |MPI_ISEND
||----------------------------------------------
||  8080500.0 | 2062.5 |   93.8 | 1968.8 |calc2_
3|            |        |        |        | MAIN_
||||--------------------------------------------
4|||  8216000.0 | 3000.0 | 1000.0 | 2000.0 |pe.0
4|||  8208000.0 | 2000.0 |     -- | 2000.0 |pe.9
4|||  6160000.0 | 2000.0 |  500.0 | 1500.0 |pe.15
||||============================================
||  6285250.0 | 1656.2 |  125.0 | 1531.2 |calc1_
3|            |        |        |        | MAIN_
||||--------------------------------------------
4|||  8216000.0 | 3000.0 | 1000.0 | 2000.0 |pe.0
4|||  6156000.0 | 1500.0 |     -- | 1500.0 |pe.3
4|||  6156000.0 | 1500.0 |     -- | 1500.0 |pe.5
||||============================================
. . .
```

Step 7: Analyze I/O

- **Trace POSIX I/O calls (fwrite, fread, write, read,…)**
- **How much I/O?**
  - Do the I/O operations take a significant amount of time?
- **Are some of the load imbalances or communication bottlenecks in fact due to I/O?**
  - Synchronous single writer
  - Insert MPI_Barriers to investigate this

Step 8: Find single-core hotspots

- **Remember: pay attention only to user routines that consume significant portion of the total time**
- **View the key hardware counters, for example**
  - L1 and L2 cache metrics
  - use of vector (SSE/AVX) instructions
  - Computational intensity (= ratio of floating point ops / memory accesses)
- **CrayPat has mechanisms for finding "the" hotspot in a routine (e.g. in case the routine contains several and/or long loops)**
  - CrayPat API
    - Possibility to give labels to "PAT regions"
  - Loop statistics (works only with Cray compiler)
    - Compile & link with CCE using -h profile_generate
    - pat_report will generate loop statistics if the flag is being enabled

Example with CrayPat

```
==========================================================================
USER / conj_grad_.LOOPS
--------------------------------------------------------------------------
  Time%                                               59.5%
  Time                                            73.010370 secs
  Imb. Time                                        3.563452 secs
  Imb. Time%                                           4.7%
  Calls                         1.383 /sec            101.0 calls
  PERF_COUNT_HW_CACHE_L1D:ACCESS                 183909710385
  PERF_COUNT_HW_CACHE_L1D:
    PREFETCH                                       7706793512
  PERF_COUNT_HW_CACHE_L1D:MISS                    21336476999
  ...
  SIMD_FP_256:PACKED_DOUBLE                        1961227352
  User time (approx)           73.042 secs    189983282830 cycles  100.0% Time
  CPU_CLK                       3.454GHz
  HW FP Ops / User time       969.844M/sec     70839736685 ops    9.3%peak(DP)
  Total DP ops                969.844M/sec     70839736685 ops
  Computational intensity         0.37 ops/cycle      0.33 ops/ref
  MFLOPS (aggregate)        124140.04M/sec
  TLB utilization              1058.97 refs/miss     2.068 avg uses
  D1 cache hit,miss ratios       90.0% hits          10.0% misses
  D1 cache utilization (misses)   9.98 refs/miss     1.248 avg hits
  D2 cache hit,miss ratio        17.5% hits          82.5% misses
  D1+D2 cache hit,miss ratio     91.7% hits           8.3% misses
  D1+D2 cache utilization        12.10 refs/miss     1.512 avg hits
  D2 to D1 bandwidth        18350.176MB/sec  1405449334558 bytes
  Average Time per Call                           0.722875 secs
```

Flat profile data

HW counter values

Derived metrics

Example with CrayPat

```
Table 2:  Loop Stats from -hprofile_generate

   Loop |Loop Incl |Loop Incl |  Loop |  Loop |    Loop |Function=/.LOOP\.
   Incl |     Time |   Time / |   Hit | Trips |   Notes | PE='HIDE'
 Time / |          |      Hit |       |   Avg |         |
  Total |          |          |       |       |         |

|-----------------------------------------------------------------------
|  24.6% | 0.057045 | 0.000570 |   100 |  64.1 |   novec |calc2_.LOOP.0.li.614
|  24.0% | 0.055725 | 0.000009 |  6413 | 512.0 |  vector |calc2_.LOOP.1.li.615
|  18.9% | 0.043875 | 0.000439 |   100 |  64.1 |   novec |calc1_.LOOP.0.li.442
|  18.3% | 0.042549 | 0.000007 |  6413 | 512.0 |  vector |calc1_.LOOP.1.li.443
|  17.1% | 0.039822 | 0.000406 |    98 |  64.1 |   novec |calc3_.LOOP.0.li.787
|  16.7% | 0.038883 | 0.000006 |  6284 | 512.0 |  vector |calc3_.LOOP.1.li.788
|   9.7% | 0.022493 | 0.000230 |    98 | 512.0 |  vector |calc3_.LOOP.2.li.805
|   4.2% | 0.009837 | 0.000098 |   100 | 512.0 |  vector |calc2_.LOOP.2.li.640
|=======================================================================
```

Hardware Counter Selection

- **HW counter collection enabled**
  - `export PAT_RT_PERFCTR= <group> | <event list>`

```
$> man hwpc
...
Table 5. Intel Haswell Event Sets
 ----------------------------------------
 Group Description
 ----------------------------------------
 0   D1 with instruction counts
 1   Summary with cache and TLB metrics
 2   D1, D2, and L3 metrics
 6   Micro-op queue stalls
 7   Back-end stalls
 8   Instructions and branches
 9   Instruction cache
 10  Cache hierarchy
 19  Prefetches
 23  Summary with cache and TLB metric
 -----------------------------------
```

```
$> papi_avail
...
=========================================================
PAPI Preset Events
=========================================================
    Name         Code    Avail Deriv Description (Note)
PAPI_L1_DCM  0x80000000  Yes No Level 1 data cache misses
PAPI_L1_ICM  0x80000001  Yes No Level 1 inst cache misses
PAPI_L2_DCM  0x80000002 Yes Yes Level 2 data cache misses
...
PAPI_FP_OPS  0x80000066 Yes Yes Floating point operations
...
```

- **more details:**
  `pat_help counters`

Further Information

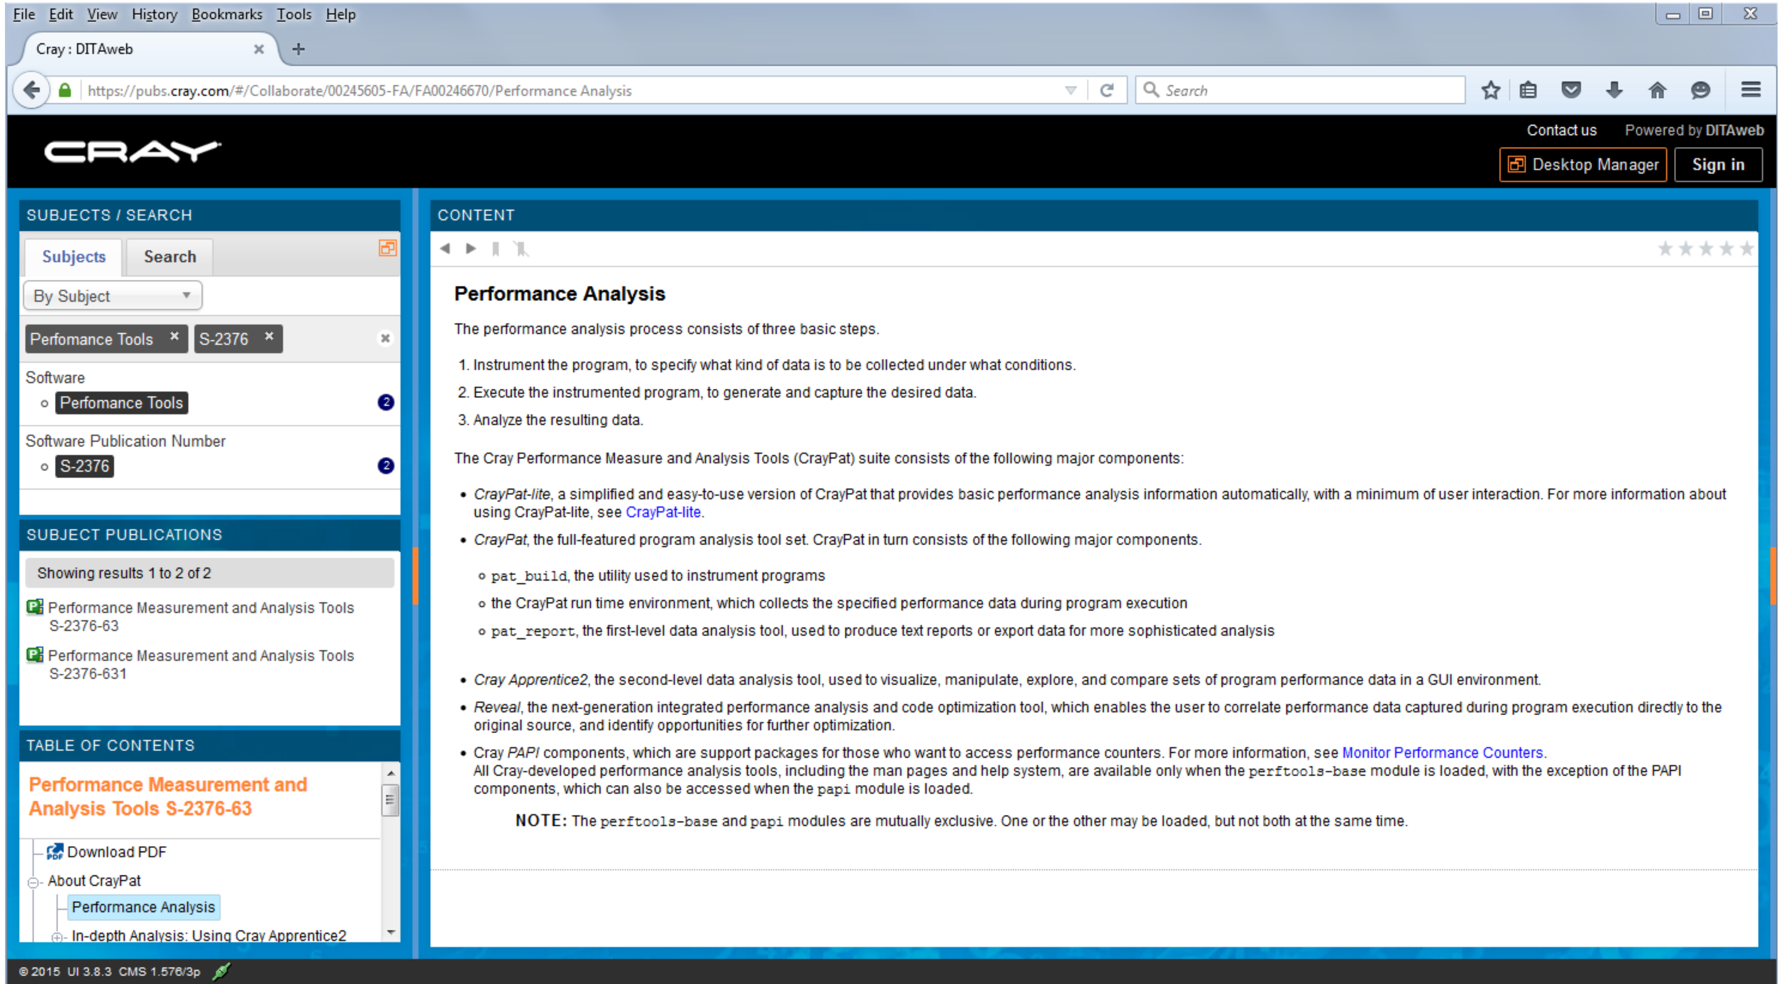

Doesn't the compiler do everything?

- **Not yet...**
  - Standard answer, unchanged for last 50 or so years
- **What does it do**
  - It tries to compile the loops in your application to be as fast as possible
  - Performance depends on reducing memory use and using the best machine instructions (vectorization)
  - This means your code may be significantly transformed
- **What can you do**
  - Work out what you care about (profile)
  - Experiment with alternative source implementations but a lot of expertise is needed here
  - Give the compiler additional information
  - Use compiler output to determine what it is doing and influence it via directives

Loop optimisation techniques

- **Most HPC codes are loop-based**
  - Repeatedly process all the elements of an array
- **There are various optimization techniques for loops**
  - unrolling/unwinding
  - stripmining
  - blocking/tiling
- **We are not going to explain HOW to do this manually but it is useful to be aware of these even if you are not going to optimise source**

- **In many cases, the compiler does these automatically**
  - the material here will help you understand what the compiler did
  - if necessary, you can then step in to assist the compiler

EXAMPLE 1: Loop unrolling/unwinding

- **Unrolling and unwinding are equivalent terms**
- **Replaces a loop by an equivalent set of statements**
  - Removes the overhead of loop control logic
    - incrementing the loop index counter
    - checking if the counter has exceeded the loop bounds

- **Most important for small tripcount/low work loops**
  - Especially when nested inside other loops
  - Full unwinding requires tripcount to be known at compile time

Original code	After unwinding

`do i=1,N`
` a(i)=a(i) + b(i)`
`enddo`
	`a(i)   =a(i)   + b(i)`
	`a(i+1)=a(i+1) + b(i+1)`
	`a(i+2)=a(i+2) + b(i+2)`
	`:`
	`a(N)   =a(N) + b(N)`

Example 2: Loop blocking/tiling

- **Applied to multi-dimensional loopnests**
  - Two or more loops are stripmined
  - Loop interchange moves the strip loops innermost
- **Most often used to preserve memory locality**

Original loopnest	Equivalent explicit code

`do j = 1,Nj`
` do i = 1,Ni`
`  !stencil`
` enddo`
`enddo`
	`do jb = 1,Nj,16`
	` do ib = 1,Ni,16`
	`  do j = jb,jb+16-1`
	`   do i = ib,ib+16-1`
	`    !stencil`
	`   enddo`
	`  enddo`
	` enddo`
	`enddo`

- (strictly, upper strip loop limits should be MIN(Nj,jb+16-1) and similar)

Control: Example blocking with Cray Directives

- CCE blocks well, but it sometimes blocks better with help

Original loopnest	Loopnest with help	Equivalent explicit code

do k = 1,Nk
 	
 do j = 1,Nj
  do i = 1,Ni
   ! *stencil*
  enddo
 enddo
enddo
	!DIR$ BLOCKABLE(j,k)
	!dir$ BLOCKINGSIZE(16)
	do k = 1,Nk
	!dir$ BLOCKINGSIZE(20)
	 do j = 1,Nj
	  do i = 1,Ni
	   ! *stencil*
	  enddo
	 enddo
	enddo
		do kb = 1,Nk,16
		 do jb = 1,Nj,20
		  do k = kb,kb+16-1
		   do j = jb,jb+20-1
		    do i = 6, nx-5
		     ! *stencil*
		    enddo
		   enddo
		  enddo
		 enddo
		enddo

  - (again, upper limits should be MIN(Nk,kb+16-1) and similar)

- **Get the loopmark listing**
  - Identifies which loops were blocked
  - Gives the block size the compiler chose

Example 3: Loop interchange

- **One of the simplest cache optimisations**
  - aim to access consecutive elements of arrays in order
- **If multi-dimensional arrays addressed in wrong order**
  - causes a lot of cache misses = bad performance
- **Order loops in loopnest with fastest innermost**
  - Fortran is column-major (LH array index moves fastest)
  - C/C++ is row-major (RH array index moves fastest)

- **Compiler may re-order loops automatically (see loopmark)**

Original loopnest	interchanged code

`do i = 1,N`
`  do j = 1,N`
`    tot = tot + a(i,j)`
`  enddo`
`enddo`
	`do j = 1,N`
	`  do i = 1,N`
	`    tot = tot + a(i,j)`
	`  enddo`
	`enddo`

Optimization for memory access, huge pages

- **Various loop transformations we have seen**
  - Help with memory access order
  - This makes more efficient use of cache
    - Use as much cache as possible
    - Reuse data when it is in cache

- **There is a level beyond cache size to consider**
- **We have virtual memory pages which map to physical pages**
- **The OS keeps track of this in hardware (TLB) and software**
- **As a result we should try to reuse memory within a page**

Using hugepages

- **Load chosen `craype-hugepages*` module**
  - See `module avail craype-hugepages` for list of available options
- **Compile as before**
- **Execute as before, but**
  - Make sure this module is also loaded in PBS jobscript
    - It sets various environment variables

- **Which pagesize is best?**
  - You should try different settings
  - 2M or 8M are usually most successful on Cray XC systems

- **Quick cheat:**
  - no need to rebuild to try a different pagesize
  - can load different hugepages module at runtime
    - compared to that used at compile-time
  - compile-time module enables hugepages in the application
  - runtime module determines the actual size that is used

- **See `man intro_hugepages` for more details**

Vectorisation

- **The most important optimization is for memory access**
- **Then we can think of optimising computation**
- **This will be in loops**
- **Usually only one loop is vectorisable in loopnest**
  - And most compilers (not CCE) only consider inner loop
- **Optimising compilers will use vector instructions**
  - Relies on code being vectorisable
  - Or in a form that the compiler can convert to be vectorisable
    - Some compilers are better at this than others
- **Check the compiler output listing and/or assembler listing**
  - Look for packed SSE/AVX instructions

Helping vectorisation

- **Is there a good reason for this?**
  - There is an overhead in setting up vectorisation; maybe it's not worth it
    - Could you unroll inner (or outer) loop to provide more work?

- **Does the loop have dependencies?**
  - information carried between iterations
    - e.g. counter: `total = total + a(i)`

- **If there are no loop dependencies:**
  - Tell the compiler that it is safe to vectorise
    - `IVDEP` directive above loop (CCE, but works with most compilers)
    - C99: `restrict` keyword (or compile with `-hrestrict=a` with CCE)
  - Perhaps the dependencies are between iterations i and i+8
    - Then it is safe to vectorise with vectors of length 8 or less
    - Use directive: `IVDEP SAFEVL=8`
  - see `man ivdep` for more details

Inhibitors to vectorisation

- **loop dependencies:**
  - The loop cannot be executed in any order
  - Might be hard to rewrite code to fix this
- **Code is not a loop (do while)**
- **Indirect addressing**
- **Non-vectorisable functions**
- **Unknown loop trip count**
- **Function calls in loop need to be inlined**

- **Check the compiler output to see what it did**
  - CCE: `-hlist=a`
  - Intel: `-vec-report[0..5]`
  - GNU: `-ftree-vectorizer-verbose=5`

Final points on vectorisation

- **Strided loops will not currently vectorise**
  - AVX-512F introduces vector instructions for strided memory
- **The compiler won't vectorise loops if it <u>thinks</u> the memory access might strided**
  - For instance:
    - `SUBROUTINE sub1(b(N))` ! argument appears contiguous
    - `CALL sub1(a(1:2*N:2))` ! but really it was strided
  - Loops in `sub1` will then be (at best) partially vectorised
  - Can tell the compiler that the passed arrays will always be contiguous
    - Use `CONTIGUOUS` attribute (Fortran2008) in declaration of b in `sub1()`, or
    - Compile `sub1.f` using CCE flag: `-h contiguous`

CCE directives

Some useful CCE directives

- **Compiler directives avoid the need for explicit coding**
  - They are compiler-specific but should be ignored as comments by:
    - other compilers
    - the same compiler, if overridden by compiler options

- **CCE has a large set of optimisation directives**
  - Fortran: `!DIR$ <directive>`
  - C/C++: `#pragma _CRI <directive>`
    - `_CRI` optional; include it so compiler warns about unrecognised directives

- **Some useful ones are listed on the next few slides**

- **For more information:**
  - `man directives`
  - `man <directive name>`
  - Fortran, C/C++ Reference Manuals on docs.cray.com

Selected CCE scalar optimisation directives

- **INTERCHANGE (i,j...), NOINTERCHANGE**
  - Specified loops should be interchanged, e.g. (i,j,k) -> (k,j,i)
  - NOINTERCHANGE directive suppresses loop interchange
- **UNROLL [n], NOUNROLL**
  - Specify unrolling of next loop, with optional unroll factor
- **BLOCKABLE (i,j...)**
  - Specified loops can be blocked
  - NOBLOCKING directive prevents blocking
- **BLOCKINGSIZE (n)**
  - Apply blocking factor n to next loop
  - Use separate BLOCKINGSIZE directives for each loop to be blocked
- **FUSION, NOFUSION, NOFISSION**
  - Control loop fusion and fission of specified loop

Selected CCE vectorisation directives (1)

- **IVDEP**
  - Ignore dependencies in the next loop that might inhibit vectorisation
- **NEXTSCALAR**
  - Do not vectorise the next loop
- **PREFERVECTOR**
  - If more than one loop in nest can be vectorised, indicates preference
  - Has the same effect as VECTOR ALWAYS directive
- **NOVECTOR**
  - Disable vectorisation for rest of program unit;
  - reset behaviour with VECTOR directive

Selected CCE vectorisation directives (2)

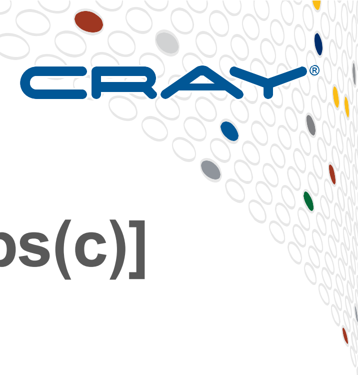

- **LOOP_INFO [min_trips(c)] [est_trips(c)] [max_trips(c)]**
  - Provide information on min/mean/max tripcounts for loop
- **PROBABILITY**
  - Indicate probability of a conditional being true
  - May suggest compiler uses gather/scatter methods to vectorise loop
- **PERMUTATION**
  - The specified integer array does not have repeated values
  - Useful for index array used in indirect addressing
- **CONCURRENT**
  - Stronger than IVDEP
    - IVDEP says loop iterations independent in current order
    - CONCURRENT says independent in any order
  - Both CONCURRENT and IVDEP should allow (possible) vectorisation

Concluding remarks

- **Compilers are good at optimising code, but not perfect**
- **If you do nothing else with your code**
  - Make sure you address arrays in the "right" order
  - Check the compiler feedback to see its not doing anything foolish
- **To go further:**
  - Understand what the compiler does
    - Look at the compiler feedback in more detail
    - Use profiling and hardware counters to see if these optimisations work
  - Help the compiler to understand your code
    - Simpler code is usually a good place to start
    - Use directives to give the compiler more information about your code
  - Only start hand-coding optimisations as a last resort
- **And remember to keep profiling your code**
  - optimise the things that take most time

The Golden Rules of profiling:

- **Profile your code**
  - The compiler/runtime will not do all the optimisation for you.
- **Profile your code yourself**
  - Don't believe what anyone tells you. They're wrong.
- **Profile on the hardware you want to run on**
  - Don't profile on your laptop if you plan to run on a Cray system
- **Profile your code running the full-sized problem**
  - The profile will almost certainly be qualitatively different for a test case.
- **Keep profiling your code as you optimize**
  - Concentrate your efforts on the thing that slows your code down.
  - This will change as you optimise.
  - So keep on profiling.