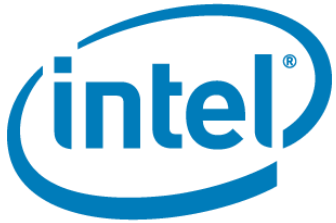

Performance Analysis and SW optimization lab for CERN

Presenter: David Levinthal
Principal Engineer

Business Group, Division: DPD, SSG

Version 1
July 17, 2010

Legal Disclaimer

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

  UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

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

  The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

  Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.
- All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Customers, licensees, and other third parties are not authorized by Intel to use Intel code names in advertising, promotion or marketing of any product or service.
- Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations
-

Risk Factors

The above statements and any others in this document that refer to plans and expectations for the first quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Many factors could affect Intel's actual results, and variances from Intel's current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be the important factors that could cause actual results to differ materially from the corporation's expectations. Current uncertainty in global economic conditions pose a risk to the overall economy as consumers and businesses may defer purchases in response to tighter credit and negative financial news, which could negatively affect product demand and other related matters. Consequently, demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including conditions in the credit market that could affect consumer confidence; customer acceptance of Intel's and competitors' products; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of new Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel's response to such actions; Intel's ability to respond quickly to technological developments and to incorporate new features into its products; and the availability of sufficient supply of components from suppliers to meet demand. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; capacity utilization; excess or obsolete inventory; product mix and pricing; variations in inventory valuation, including variations related to the timing of qualifying products for sale; manufacturing yields; changes in unit costs; impairments of long-lived assets, including manufacturing, assembly/test and intangible assets; and the timing and execution of the manufacturing ramp and associated costs, including start-up costs. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel's products and the level of revenue and profits. The recent financial crisis affecting the banking system and financial markets and the going concern threats to investment banks and other financial institutions have resulted in a tightening in the credit markets, a reduced level of liquidity in many financial markets, and extreme volatility in fixed income, credit and equity markets. There could be a number of follow-on effects from the credit crisis on Intel's business, including insolvency of key suppliers resulting in product delays; inability of customers to obtain credit to finance purchases of our products and/or customer insolvencies; counterparty failures negatively impacting our treasury operations; increased expense or inability to obtain short-term financing of Intel's operations from the issuance of commercial paper; and increased impairments from the inability of investee companies to obtain financing. Intel's results could be impacted by adverse economic, social, political and physical/infrastructure conditions in the countries in which Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel's SEC reports.

Lab Overview

- **Objective: gain familiarity with PTU and events critical to CERN LHC applications**
- **Geant4**
  - **10 GeV photon into simple parallel plate geometry**
- **Synthetic benchmarks for call execution**
  - **Statically linked direct calls**
  - **Dynamically linked direct calls**
  - **Statically linked indirect calls**
  - **Dynamically linked indirect calls**

Installing PTU

- **Installation:**
  THIS IS ALREADY DONE FOR THE LAB
  - Untar ptu package into a commonly accessed directory on a non network mounted drive (for performance reasons)
    - Ex /home/user/ptu
  - cd ptu/sepdk/src/pax
  - ./build-driver
  - ./insmod-pax –g my_user_group
  - cd ..
  - ./build-driver
  - ./insmod-sep3 –g my_user_group
- **The insmod scripts will prompt you for the root PW**
- **You must have a valid vtune license**

Using PTU

- Login through VNCserver or whatever you use
- Account user, pw user
- ptu2/eclipse/eclipse
- Select workspace /home/user/workspace2
- Kill the welcome window
- Start a project called triad, that will run the triad
  - /home/user/lab/triad/triad
  - Arguments 0 3 40
  - Look at project_create.pdf for gui interaction example
  - Start at slide 8 of the pdf
  - In another command shell run lab/triad/triad 0 3 40 to see output

Using PTU

- **Collect basic sampling profile**
  - **Right click on project triad**
  - **Select "profile as"**
  - **Select "basic sampling"**

Watching PTU run is as exciting as watching paint dry

Viewing data

- **Source view**
  - **When the data returns double click on the row called "triad"**
    - **You may need to do this at the extreme left to avoid the tool tip**
  - **Enable the control flow graph (CFG) with the button in the upper left near the assembler and source buttons**
  - **The disassembly is organized as basic blocks**
  - **Double click on a basic block in the CFG**
  - **Use the red arrows in the tool bar to navigate through the hotspots (basic blocks)**

Viewing data

- Click on the tab for the basic sampling to go back to the hotspot view
- Change the granularity to source file
- Change the cpu display to all cores (each)
- Change the limit to 100%
- Change display to show data as
  events = sample_count*SAV
  - In the spreadsheet right click to get a menu
    - Do this on the left to avoid the tool tip
    - Select "view data as"
    - Check events
- Maximize the docking window by double clicking the folder tab

Add a predefined profile

- **Shut down the gui**
- **Cp CERN_NHM.launch ~/ptu2/workspace2/.metadata/.plugins/org.eclipse.core.debug/.launches**
- **Relaunch the gui**
  - **~/ptu2/eclipse/eclipse**

CERN profile on geant4

- **Create a new project, called geant, to run /home/user/lab/geant/G4-local/dirCaloStudies/run.sh**
- **Select multiplexing!!!!**
  - **Expand the bottom of the project creation**
  - **See slide 18 of the installation and usage pdf**
- **Running an application from a script can be the easiest way to ensure that the application is launched with the environment it requires**
- **Collect the CERN_NHM profile**

The Results (while we wait for the paint to dry)

- **transfer the excel spreadsheet geant_lab_nhm.xlsx to a system that supports excel for viewing**

Critical events for OOP

- **Look at the lab.docx document**
  - **It describes a list of events that can be used to study branch execution analysis**
  - **These are the events that identify the dominant issue in the CERN LHC applications:**

Large binary

Too many calls

Too few instructions/call

Create a profile

- Open the profile editor (aka configuration editor)
  - Right click on project and select “profile”
  - See slide 45 of the pdf
- Create a profile with the NHM events in the lab.docx file
  - User the free field filter box (upper right) to help find the events
  - Ex: type demand_ifetch in the filter box

Synthetic benchmarks

- **CERN applications retire 35 -> 65 instructions/call**
- **Objective of the synthetic benchmark**
  - **Make a large binary (10K functions, 100 shared objects, 2MB binary)**
  - **Vary the call invocation style**
  - **Vary the linking options**
- **Measure the cost per call for various function invocation calling style**
  - **Direct, indirect, function pointers, cross shared object, etc**
  - **Maintain size per function**

Processing the PTU results

- Set the granularity to source files
- Right click, show data as events
- Export to csv
- Transfer the csv file to a system that supports excel for viewing
- Transfer the cern_ptu_template_nhm.xls to a system that supports excel for viewing
- Open both spreadsheets in excel
- Copy the data from the csv to the template
  - Ctrl A Ctrl C go to the other Ctrl V
- Click the button, Watch the magic

Code generation

- A program that writes C functions
  - Asm block of __asm{xorq %rdx, %rdx }
  - Followed by some C to invoke the call to the next function
- Functions created as a 3 dimensional matrix of functions, with the calls executing down the third dimension
  - Or executed following a pattern of non repeating random numbers
- generator*.c

Lab exercise

- **Create a project for the binary lab/call_test/sources_att/FOO_static 100000**
  - **Select multiplexing**
- **Collect the CERN_NHM profile without multiplexing**
- **Export the spreadsheet to a csv**
- **Open the csv and import it into the standard CERN analysis excel spreadsheet**

Lab exercise (if there is time)

- **Create a project for the binary lab/call_test/sources_att/run_dyn.sh**
  - **Select multiplexing**
  - **Cat run_dyn.sh**
- **Collect the CERN_NHM profile without multiplexing**
- **What do you notice that changed**
- **Export to csv**
- **Import into Cern excel template**
- **What do you see?**

Summary

- Event based sampling performance analysis is extremely powerful on Intel® Core™ i7, XEON™ 5500 and 5600 Processor Families
- Correct methodology is essential
- Correct usage of events is essential
- Intel® PTU simplifies task