

The ISP RAS activities of improving GCC for Itanium

Arutyun I. Avetisyan
Gelato ICE Singapore
October 4th, 2006

Workflow of our GCC activity

- The ISP RAS is working on GCC under contract with HP Company
- Our primary goal is improving GCC on Itanium Processor Family
- We work in contact with GCC developers and with engineers from HP, Intel, and Red Hat
- Results of our work are contributed back to the GCC community
- We provide full support for accepted patches

Current GCC project

- Implement new aggressive interblock instruction scheduler for GCC
- Choose *selective scheduling* approach
 - focuses on VLIW architectures, but general enough for others
 - aims at exploiting and increasing ILP as much as possible
 - is able to clone instructions, rename and substitute registers, unify instructions during code motion
 - can be used in a software pipeliner implementation
- Implement software pipelining using the scheduler
- 15 months project, started September 2005
- Planned for a team of four
- Consulting is provided by Vladimir Makarov (Red Hat)

Current state of the project

- Major pieces of infrastructure are ready
 - bookkeeping code generation
 - register renaming using existing code from `regrename.c`
 - forward substitution
 - unification
- All pieces are merged in a single source tree
- The scheduler is being tuned on small benchmarks
- Implementation of software pipelining has been started
- Not yet supported:
 - IA-64 speculation (need to port our implementation for the existing scheduler to the new one)
 - code motion of conditional branches

First results of selective scheduling

- Tested on small open-source widely used benchmarks
- Tested on a branch with support for bookkeeping code generation (register renaming is turned off)
- Results are not yet analyzed

Whetstone	-0.03%
LinpackC double precision	+2.9%
Dhrystone 1.1	+0.3%
FFT size=1000000	+2.4%
Scimark	-2.9%
Queens	+25.6%
Blowfish	-1.99%

Future work to be done

- Finish the scheduler project (by December 2006)
 - Test and tune the implementation on IA-64
 - Clean up the code
 - Open the FSF GCC branch
 - Finish the software pipelining implementation
- Complete the contribution cycle (by September 2007)
 - Test the scheduler on other platforms
 - Speed up the implementation
 - Submit for public review
 - Address the issues found
 - Prepare the code for Stage 1 of GCC 4.4 development

GCC contribution cycle

- Create and test your patches
 - using a variety of platforms (x86, x86_64, ia64, powerpc)
- Sort out intellectual property issues
 - sign FSF copyright papers
- Send your patches for a public review
 - large patches should be put on a branch in SVN repository
- Collaborate with other developers
 - Participate in development mailing lists
 - Make a presentation on the GCC Summit
- Support your code
 - Fix the bugs you've introduced or exposed
 - Apply for a SVN write access

Completed work

- Add speculation support for IA-64 to the GCC instruction scheduler
 - done under 6 months contract with HP with a team of three
 - gives ~0.5% improvement on SPEC INT 2000 (up to 6% on selected tests) and ~1.4% improvement on SPEC FP 2000 (up to 11.7% on selected tests)
 - 6 large patches and some minor fixes are approved
 - committed to GCC mainline in March 2006
 - will be included in GCC 4.2 release
 - since then 6 Bugzilla PRs related to the code were fixed
- Propagate alias information from the Tree SSA to the RTL level
 - done under the same contract
 - will be included in the ia64-improvements GCC branch

Gelato GCC Workshop in Moscow

- Moscow, Russia, August 7th-8th, 2006
- Organized by Gelato and ISP RAS
- Sponsored by Intel and HP
- Brought together experts from Intel and HP, members of GCC community, and Gelato members
- Included technical talks, brainstorming sessions, and tutorials

See <http://gcc.gelato.org/MoscowMeeting> for details



Gelato GCC Workshop Sessions

- Memory Disambiguation (Diego Novillo, Red Hat)
- Software Pipelining (Vladimir Makarov, Red Hat)
- Data Prefetching (Zdenek Dvorak, SuSE CR)
- Performance Metrics and Measurement (Shin-Ming Liu, HP)
- Instruction Scheduling (Vladimir Makarov, Red Hat)
- Itanium and ICC Tutorial (Mark Davis, Intel)
- GCC Tutorial (Diego Novillo, Red Hat)
 - Link-time optimizations (Kenneth Zadeck, NaturalBridge)
 - Loop-nest optimizations (Sebastian Pop, Ecole des Mines)
 - Profile-directed optimizations (Jan Hubicka, SuSE CR)
 - IA-64 backend (James Wilson, Specifix)
 - Control flow graph and RTL (Steven Bosscher)

GCC Workshop chosen projects

- Alias analysis (Diego Novillo, ISP RAS)
 - Export to RTL level
 - Keep up to date on all levels
 - Provide test cases to verify disambiguation
- Instruction scheduling (UIUC, ISP RAS, Vladimir Makarov)
 - Continue work on existing scheduling projects
 - Export data dependency information to RTL level
 - Use better data dependency in modulo scheduling
 - Evaluate register pressure in modulo scheduling
- Link time optimizations (Kenneth Zadeck)
- Prefetching (Zdenek Dvorak, NUDT China)
- Profile directed optimizations (Jan Hubicka, UIUC)

ISP RAS future activities

- Prepare the scheduler patches for inclusion in GCC
 - Tune the scheduler for IA-64
 - Open the FSF GCC branch
 - Test on other platforms
- Improve modulo scheduling in GCC
 - Make it work for IA-64
 - Propagate data dependency information to it
 - Will be sponsored by Intel Corp.
- Continue our work on alias analysis
 - Keep propagated information up to date in RTL
 - Cooperate with Diego Novillo on fixing other issues
- Participate in link-time optimizations project

Timeline of our GCC work - 2005

- January : start of IA-64 speculation project
- March : control speculation is supported
- April : data speculation is supported
- April : preliminary report on Gelato meeting
- June : speculation with recovery code is supported
- June : a presentation on GCC Summit 2005
- July : end of speculation project
- August : tuning of speculation patches on SPEC CPU2000
- September : start of the new scheduler work
- October : alias propagation patches are posted for review
- November : speculation patches are being prepared for posting
- December : speculation patches are posted on gcc-patches@ list

Timeline of our GCC work - 2006

- February : aliasing patch is fixed for mainline using the review
- March : speculation patches are approved
- March : basic routines for code motion are ready
- April : Gelato ICE talk on the scheduler project
- May : further work and discussion on aliasing patch
- June : a presentation on GCC Summit 2006
- July : initial support for instruction cloning
- August : Gelato GCC Workshop in Moscow
- September : instruction cloning and register renaming are merged
- October : testing and tuning the scheduler on SPEC CPU2000

Planned timeline of our work

2006:

- November : new branch for the project in the FSF GCC repository
- December : finish the implementation of software pipelining
- December : the new scheduler project completion

2007:

- February : test the scheduler on IA-64 and x86_64 platforms
- March : propagate data dependency information to RTL
- April : report preliminary results on Gelato ICE
- May : tune the scheduler for SPEC CPU2000 and 2006
- June : report the results on GCC Summit 2007
- August : use better data dependency in modulo scheduling
- September : prepare the scheduler for inclusion in GCC 4.4

Acknowledgments

- Gelato
 - Mark Smith
 - Sverre Jarp
- Hewlett-Packard
 - Shin-Ming Liu
 - Al Stone
- Intel
 - Mark Davis
- GCC community
 - Vladimir Makarov
 - Diego Novillo
 - Sebastian Pop
 - James Wilson