

15-745 Class Project Proposal

Brendan Meeder, Jamie Morgenstern, Richard Peng

April 13, 2011

1 Group Info

Brendan Meeder (bmeeder@cs.cmu.edu)
Jamie Morgenstern (jamiemmt@cs.cmu.edu)
Richard Peng (richard.peng@gmail.com)

2 Project Webpage

<http://www.cs.cmu.edu/~bmeeder/15745/proposal.html>

3 Major Changes

We do not have any major changes to our initial proposal.

4 Accomplishments

We have downloaded and built the SPEC 2006 framework. We have investigated and decided which LLVM optimizations to use and which analysis passes to run. Integration of the SPEC 2006 test suite with the LLVM test suite has not been fully completed, and this is the major bottleneck at the moment.

5 Meeting the Milestone

From our proposal document:

1. Select dependent passes
2. Select relevant features regarding these passes

We have selected the dependent passes from those provided with LLVM and have categorized them into optimizations that mostly restructure and move code, and those that reduce operations. We have also looked into

the analysis passes provided by LLVM and have selected several that will be appropriate for the optimizations we have selected. We were able to modify code from previous homeworks to output useful statistics on register count/pressure, as well as obtaining information about the operation counts/block structures from existing analysis passes. However, we are not sure at the moment which of the existing LLVM Analysis passes can be adapted. Richard has been having trouble extracting data from Loop-Info, although he feels that those should be the only set of data left that need to be collected to run the experiments on the benchmarks.

3. Design profiling passes to collect these features

This task is partially completed; we have the commands required to run the passes. However, this isn't yet integrated into the LLVM-SPEC testing framework, so it's not entirely completed.

4. Profile or training and test code

Because we haven't completely figured out custom pass integration, this step has not been completed.

6 Surprises

Getting SPEC and LLVM to work together has been more difficult than we anticipated. This is in no part due to a lack of effort; Jamie has invested a great amount of time wrestling with configuration issues and understanding the interaction between the make system and the benchmark framework.

We are still trying to figure out how to specify custom optimization passes and pre/post-optimization feature extraction. We will be in touch if this issue doesn't resolve itself very soon (by end of day Friday).

7 Revised Schedule

Jamie will finalize integration of LLVM and SPEC; this includes integrating analysis passes and compilation of SPEC with various optimization settings. Concurrently, Richard and Brendan will work on completing the set of analysis passes required. This will be done by the weekend (April 17), and we will start collecting data. At this point, all two of the three members will work on building different classifiers in Matlab while the third member works on writing an evaluation framework. We want to have all training data collected no later than the 22nd of April, which leaves time to build the models and evaluate the system performance.

8 Resources Needed

We have gotten ahold of SPEC and at this time we believe this satisfies our resource requirements. Brendan thinks there are some machines he can gather up in the case that we require more computational facilities to run compilation and benchmark passes.