

Predicting sequences of parameterized transform passes using Machine Learning: Midterm Report

Ashique Khudabukhsh (khudabukhsh@gmail.com)
Jayant Krishnamurthy (jayantk@cs.cmu.edu)

April 18, 2012

Our project is progressing as planned, with no major changes. Our milestone objectives were to implement our data-generation code and experiment with simple machine learning techniques. So far, we have implemented (1) a framework for running optimization passes and generating timing measurements / compiled object files to train the machine learning component, and (2) feature-generation code for the compiled programs. We are currently working on applying machine learning algorithms to the generated features, but do not yet have results for this portion. We have additionally obtained the the SPEC 2006 benchmarks that we plan to use in our final evaluation.

Overall, our current progress is close to our milestone objective, except that we have not yet been able to produce preliminary results for the machine learning. We are slightly delayed on this front because generating features from compiled programs turned out to be somewhat more complicated than initially anticipated. However, this component has been completed and we are close to producing preliminary results.

Revised schedule:

1. April 19th - April 26th:
 - Finish implementing machine learning algorithms (Ashique).
 - Finalize the data set for the evaluation (Jayant).
 - Implement search to determine the best optimization sequence using the learned models (Jayant).

2. April 26th - May 2nd:

- Tune machine learning models to improve prediction performance (Ashique).
- Run final evaluation and write project report (Ashique and Jayant).