



Does Training Input Selection Matter for Feedback-Directed Optimizations?

Paul Berube

`berube@cs.ualberta.ca`

University of Alberta

CDP05, October 17, 2005




Outline

- Background and motivation
- *Aestimo*: an FDO evaluation tool
- Workload Selection
- Results



What Is FDO?

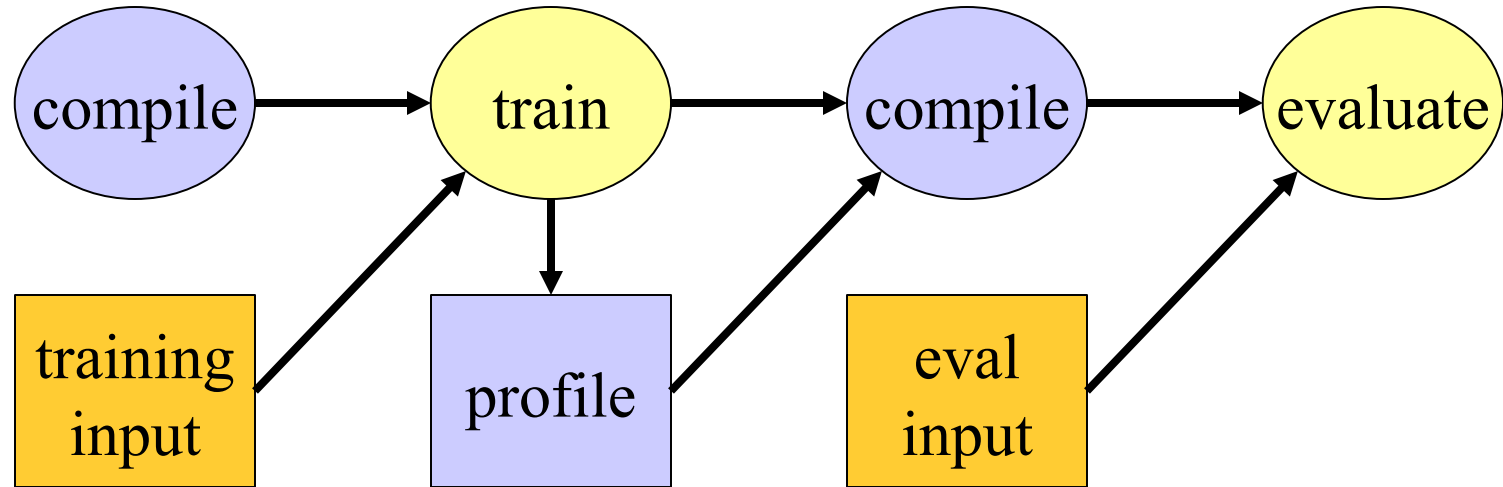
Feedback-Directed Optimization:





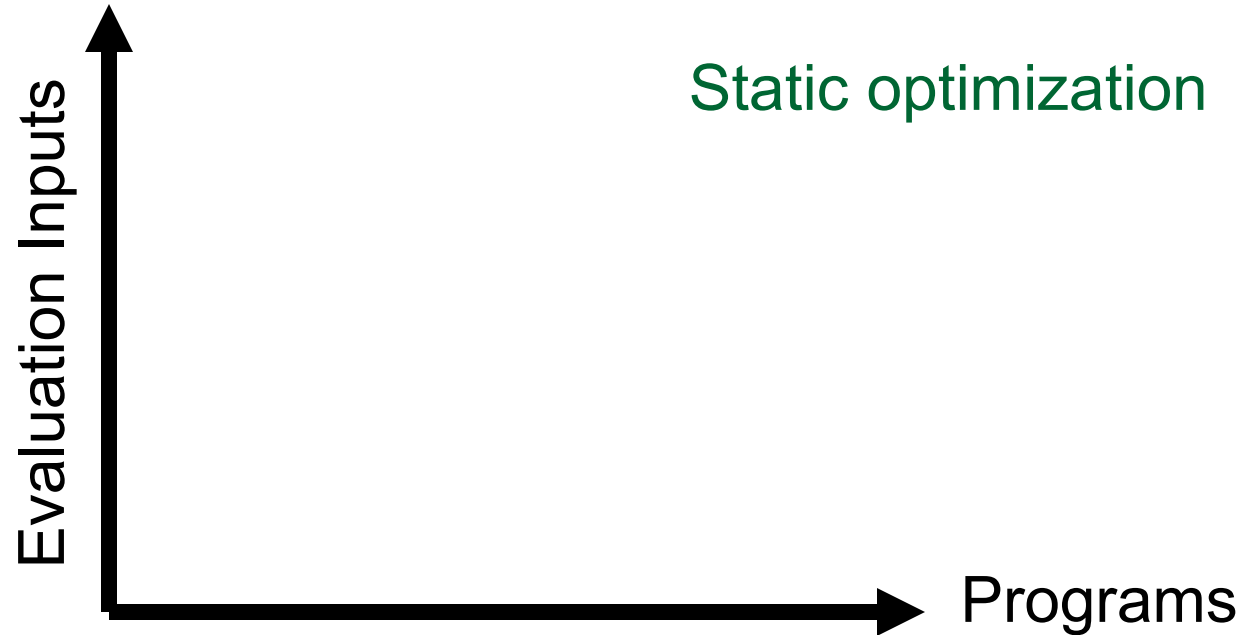
What Is FDO?

Feedback-Directed Optimization:



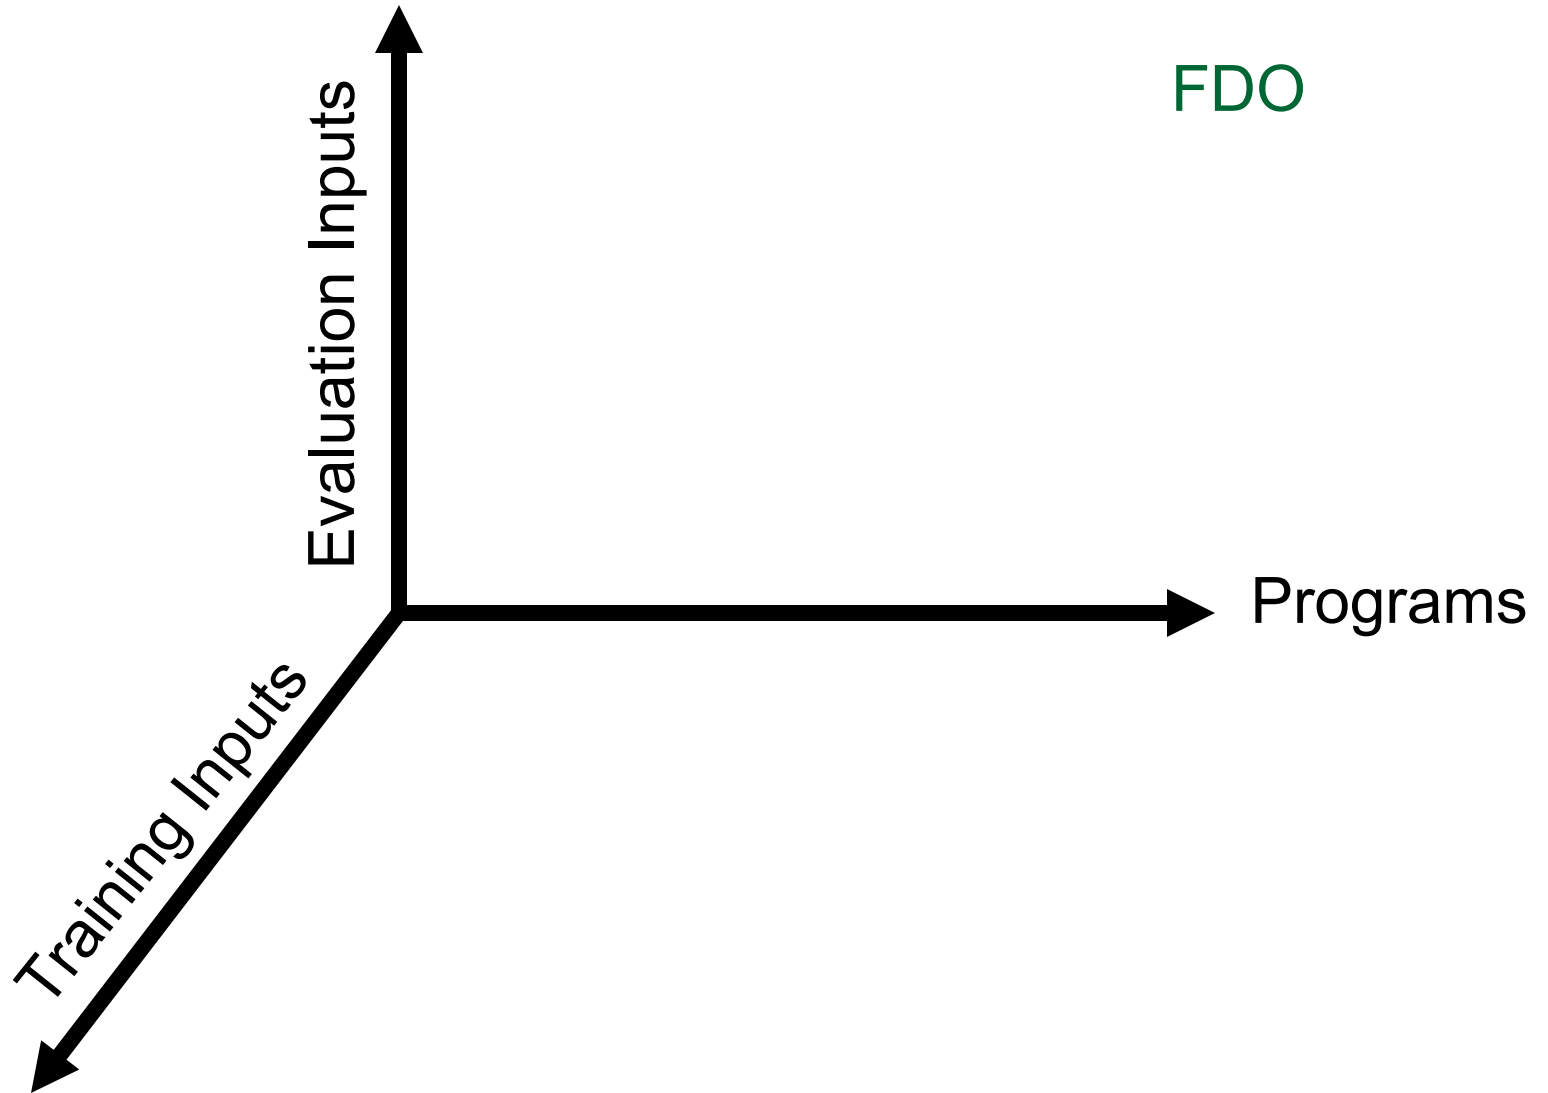


Performance Evaluation Space



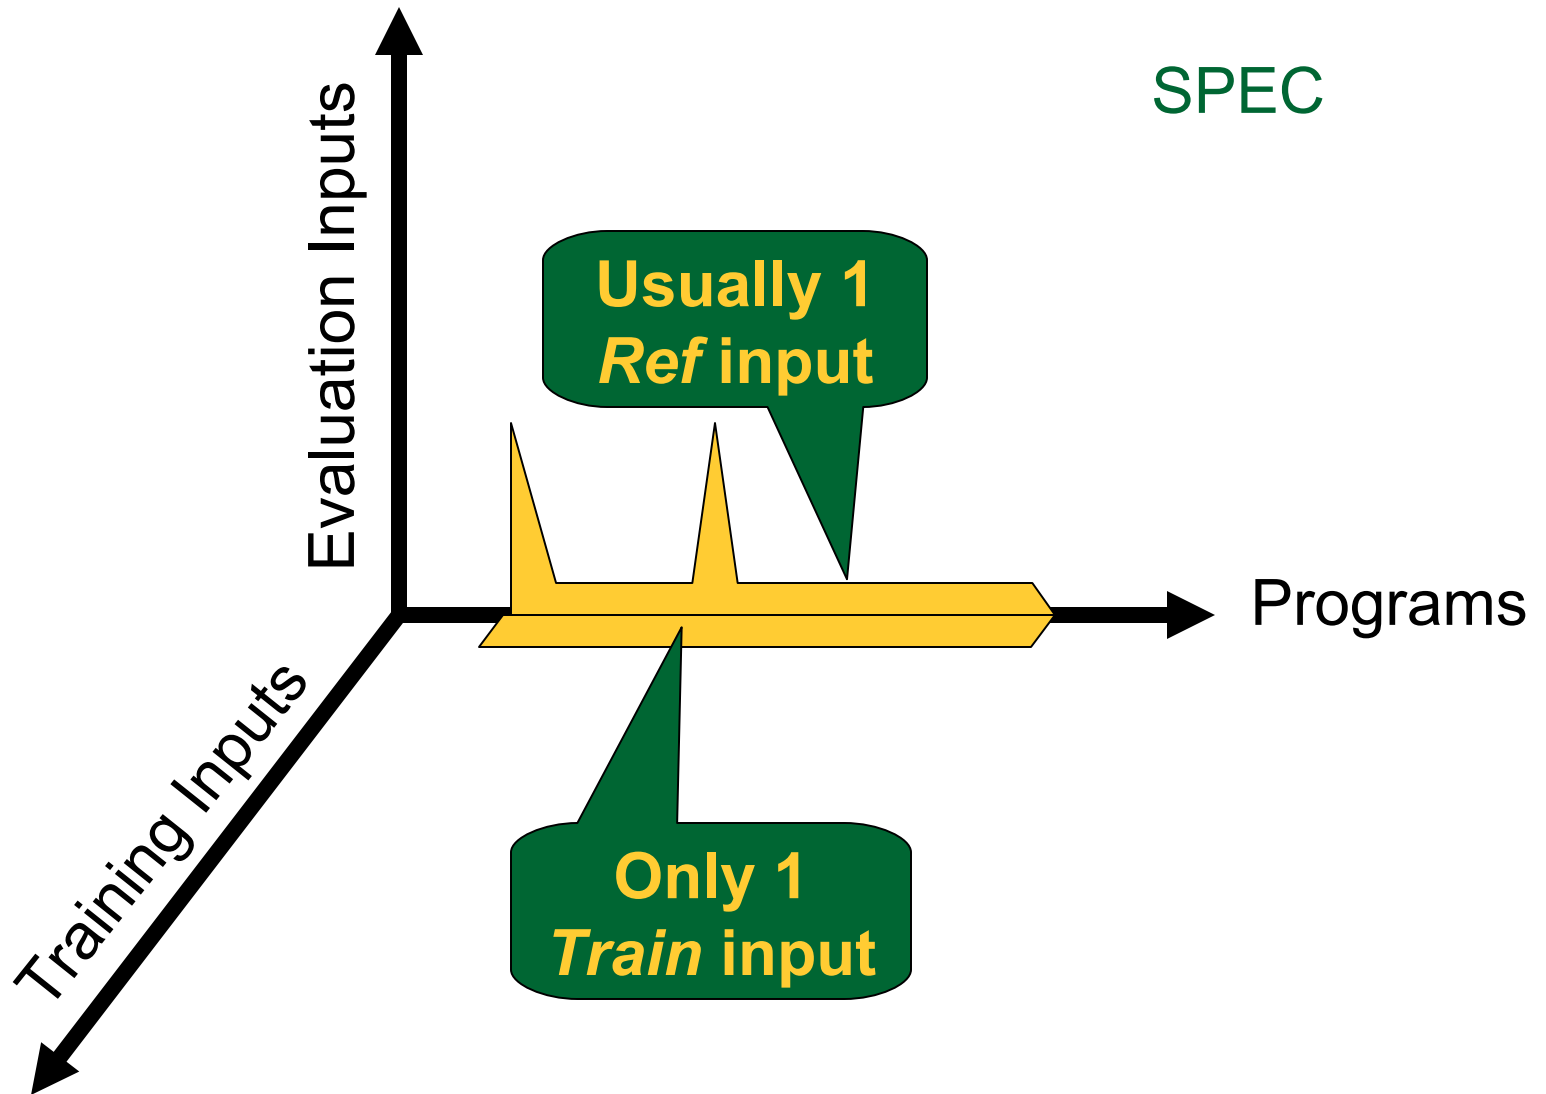


Performance Evaluation Space





Performance Evaluation Space





The Big Question

- Does the selection of training inputs matter for feedback-directed optimization?
 - Different transformation decisions?
 - Different performance?

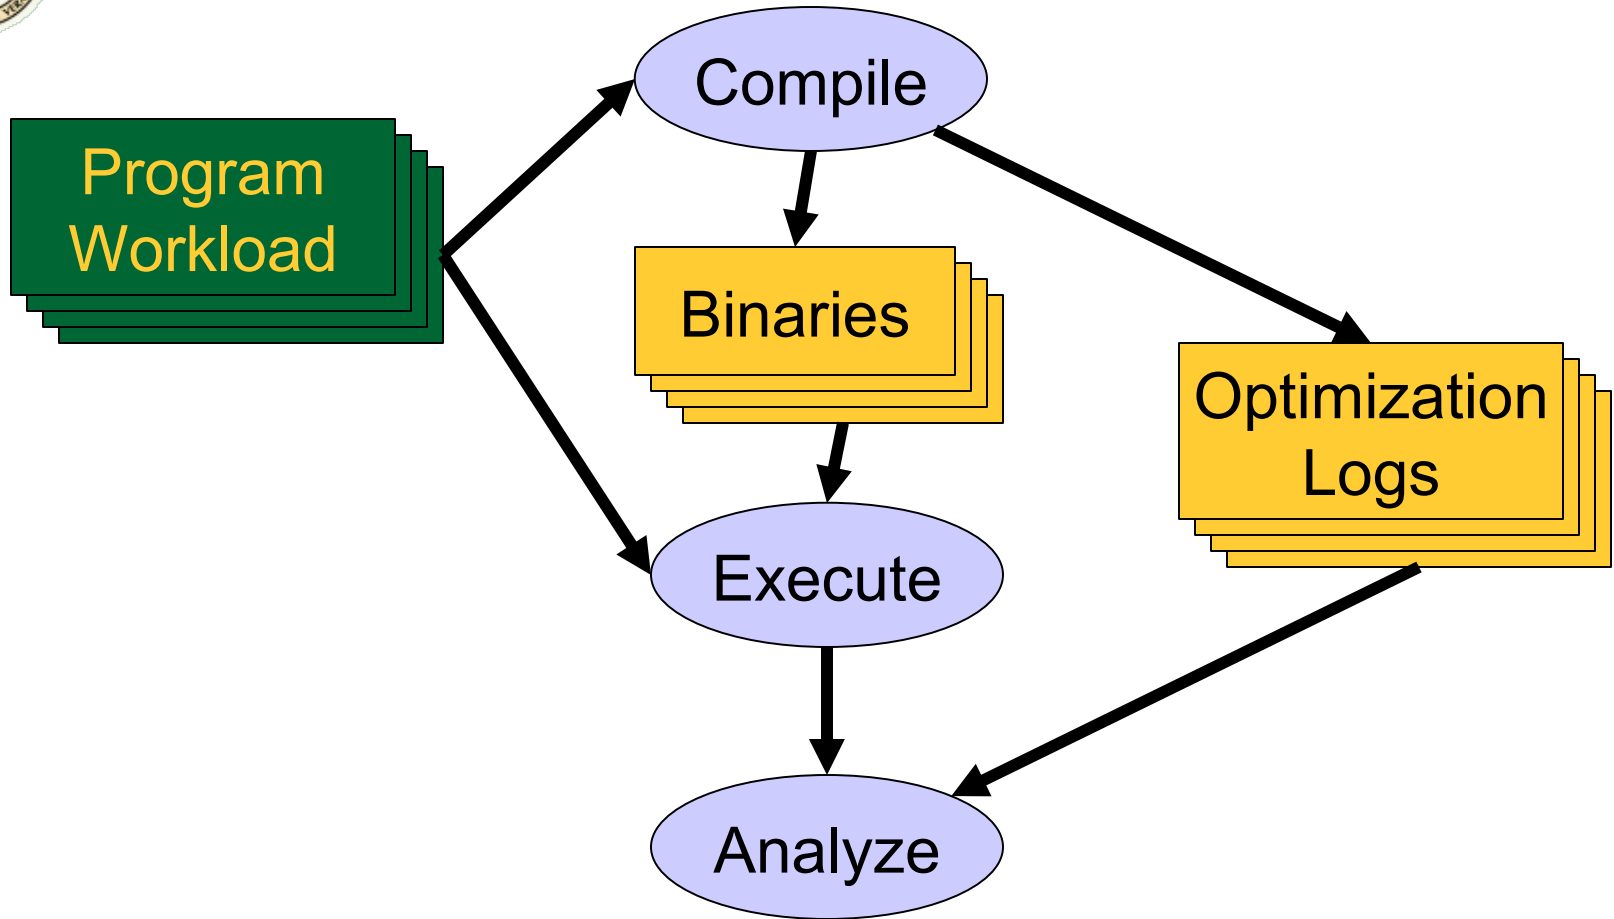


Aestimo

- An FDO evaluation tool
- Automates training and evaluating on a large number of inputs
- Isolates individual transformations
 - Fewer experiment variables
 - Results vary by transformation
- Measures:
 - Differences in transformation decisions
 - Performance differences




An Overview of *Aestimo*



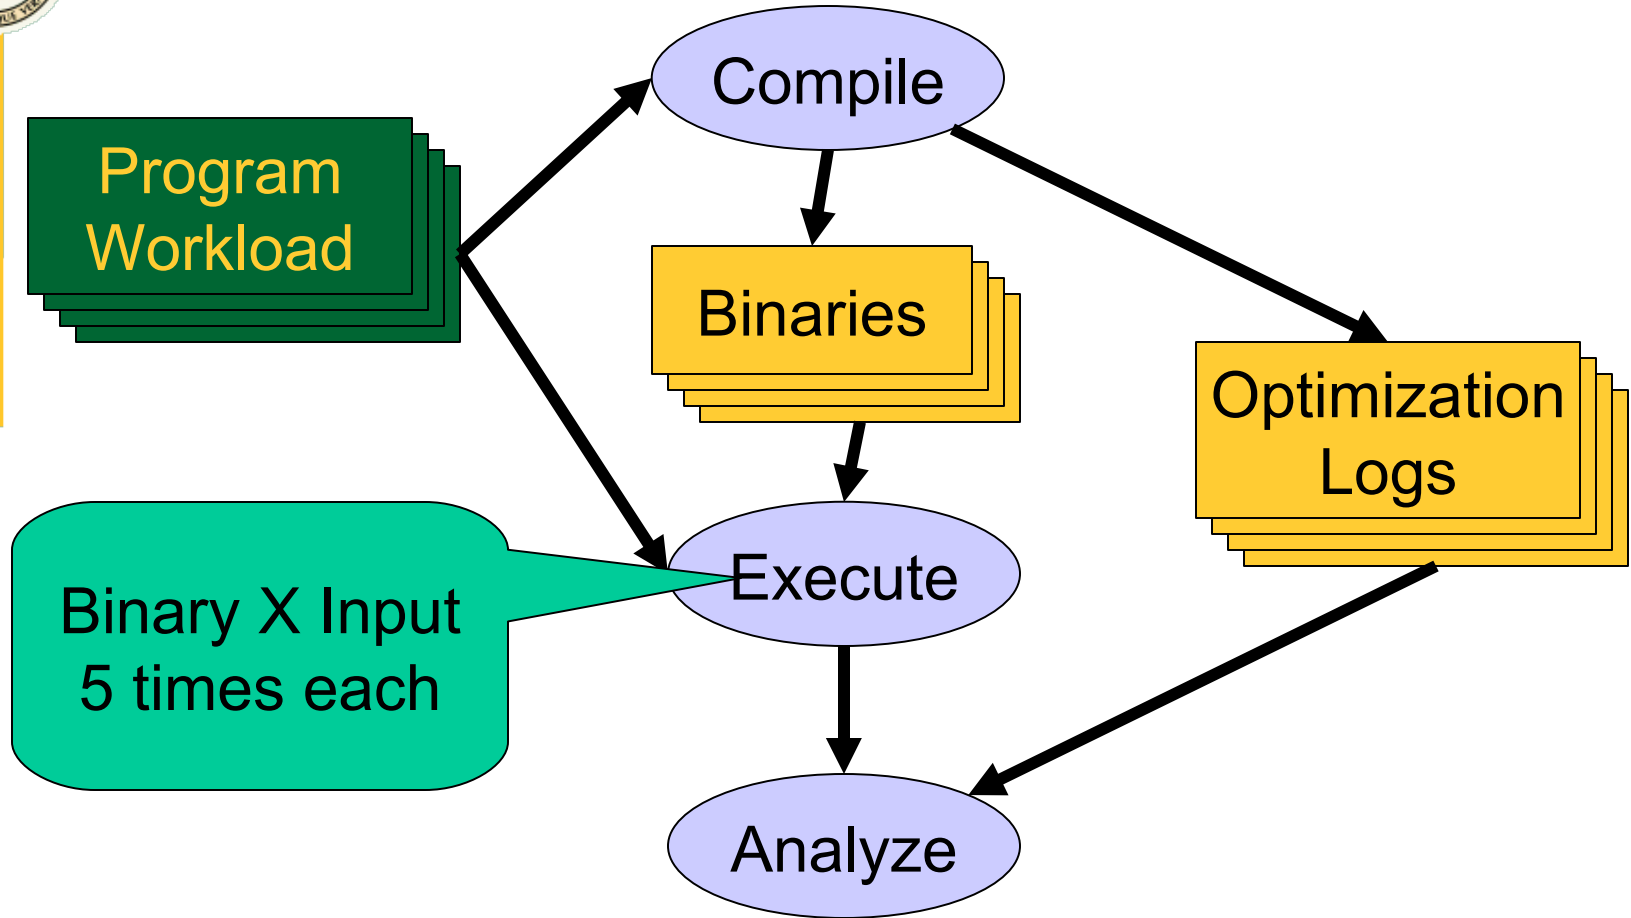


An Overview of *Aestimo*



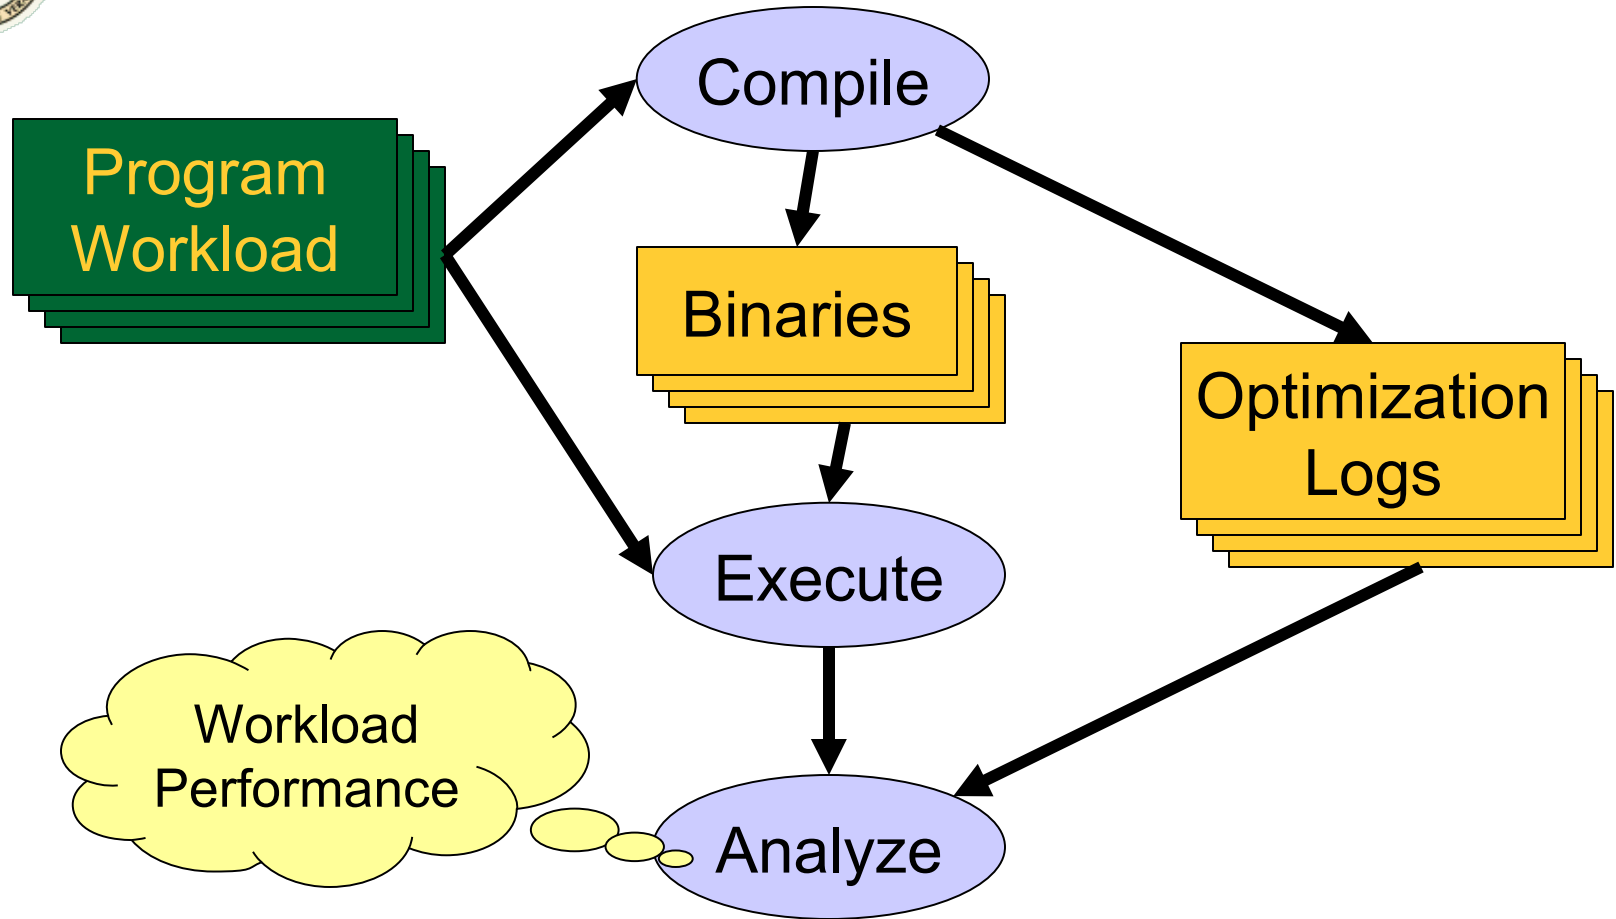


An Overview of *Aestimo*



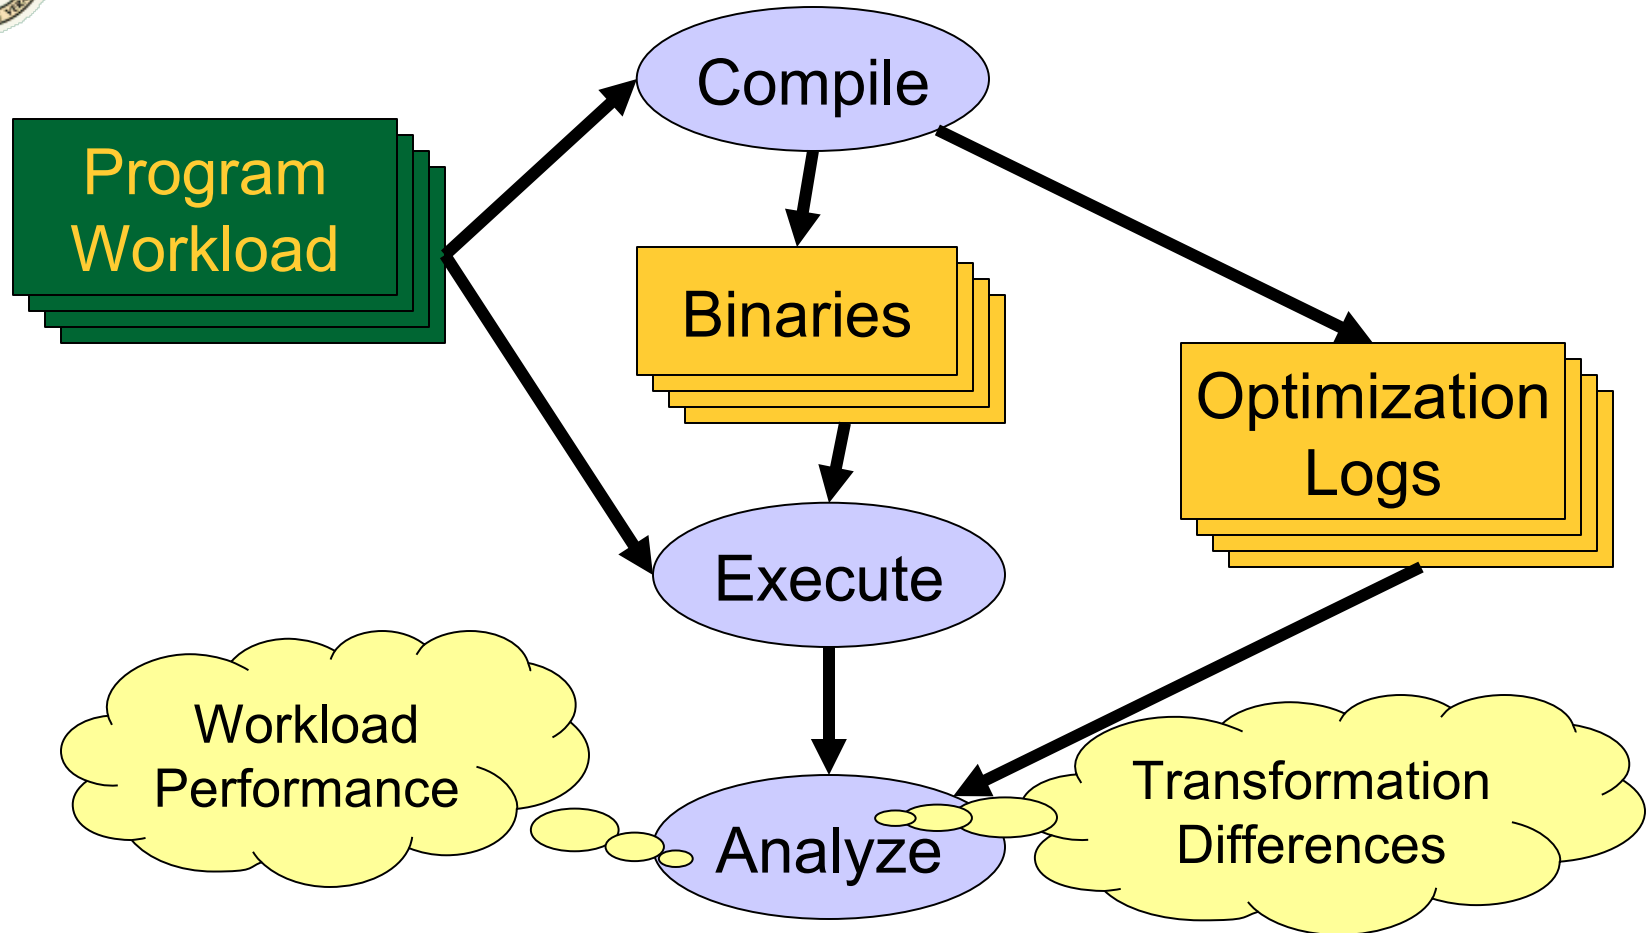


An Overview of *Aestimo*



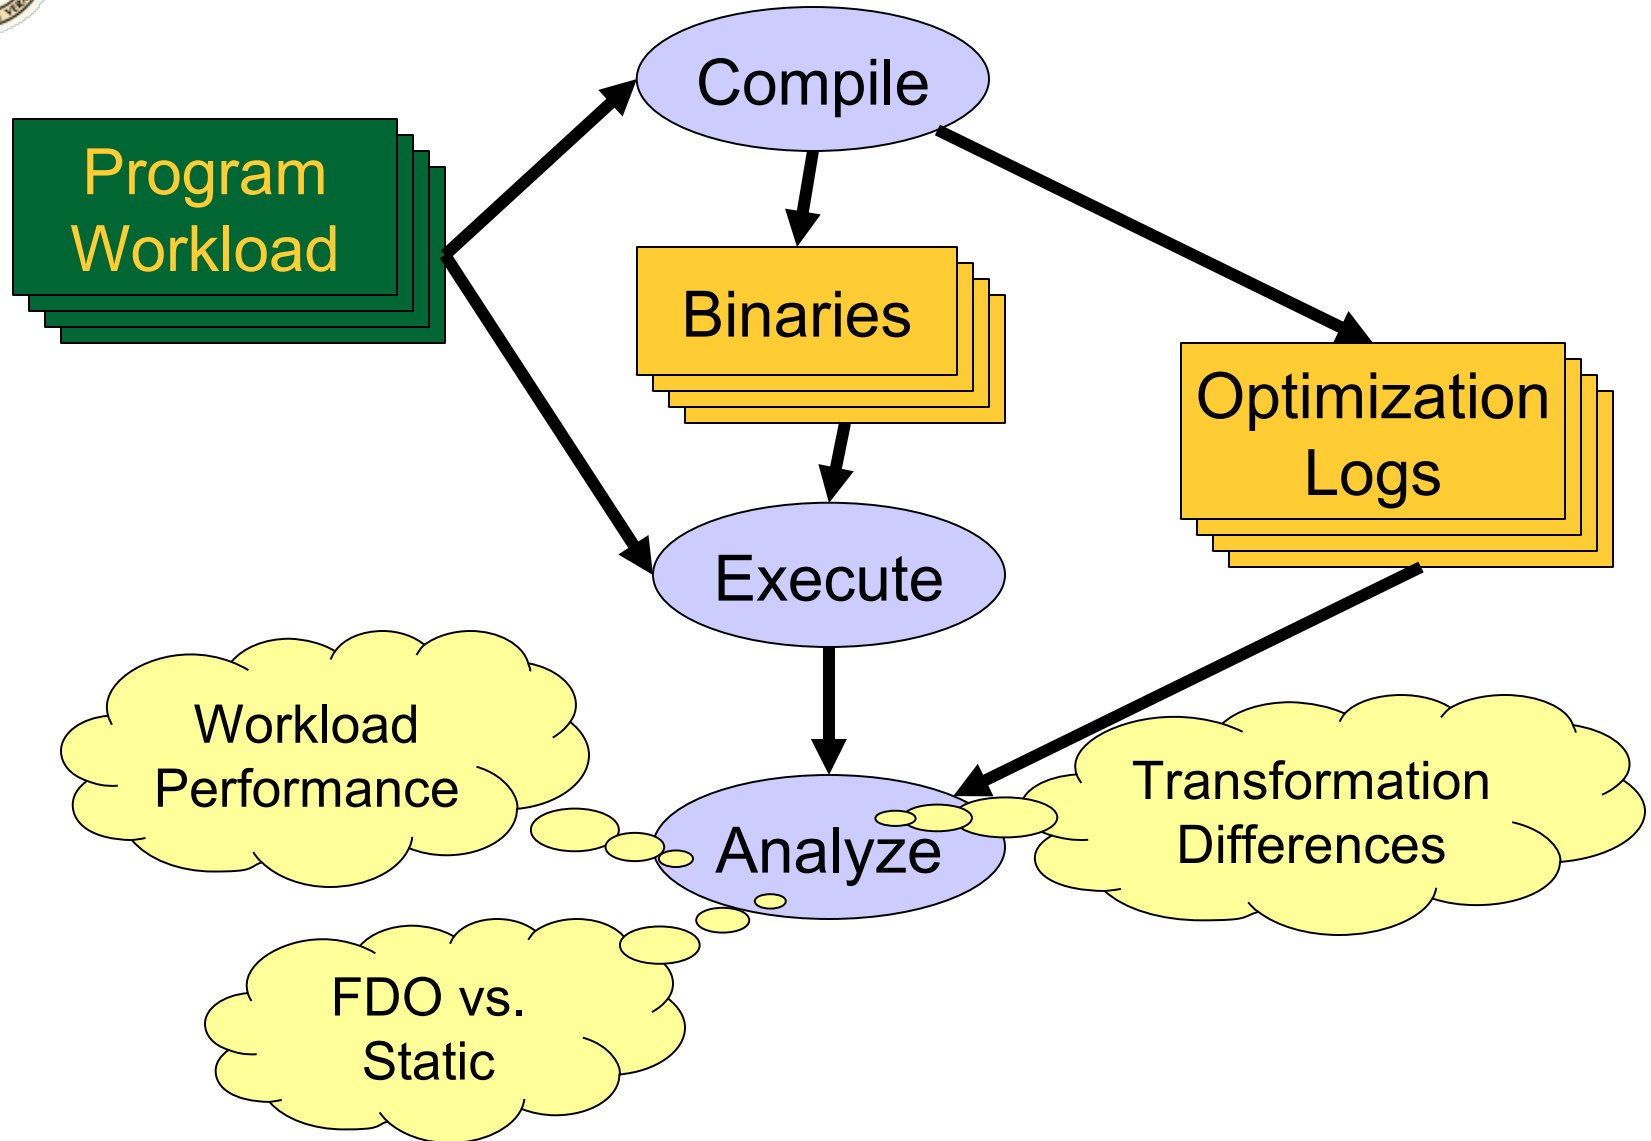


An Overview of *Aestimo*



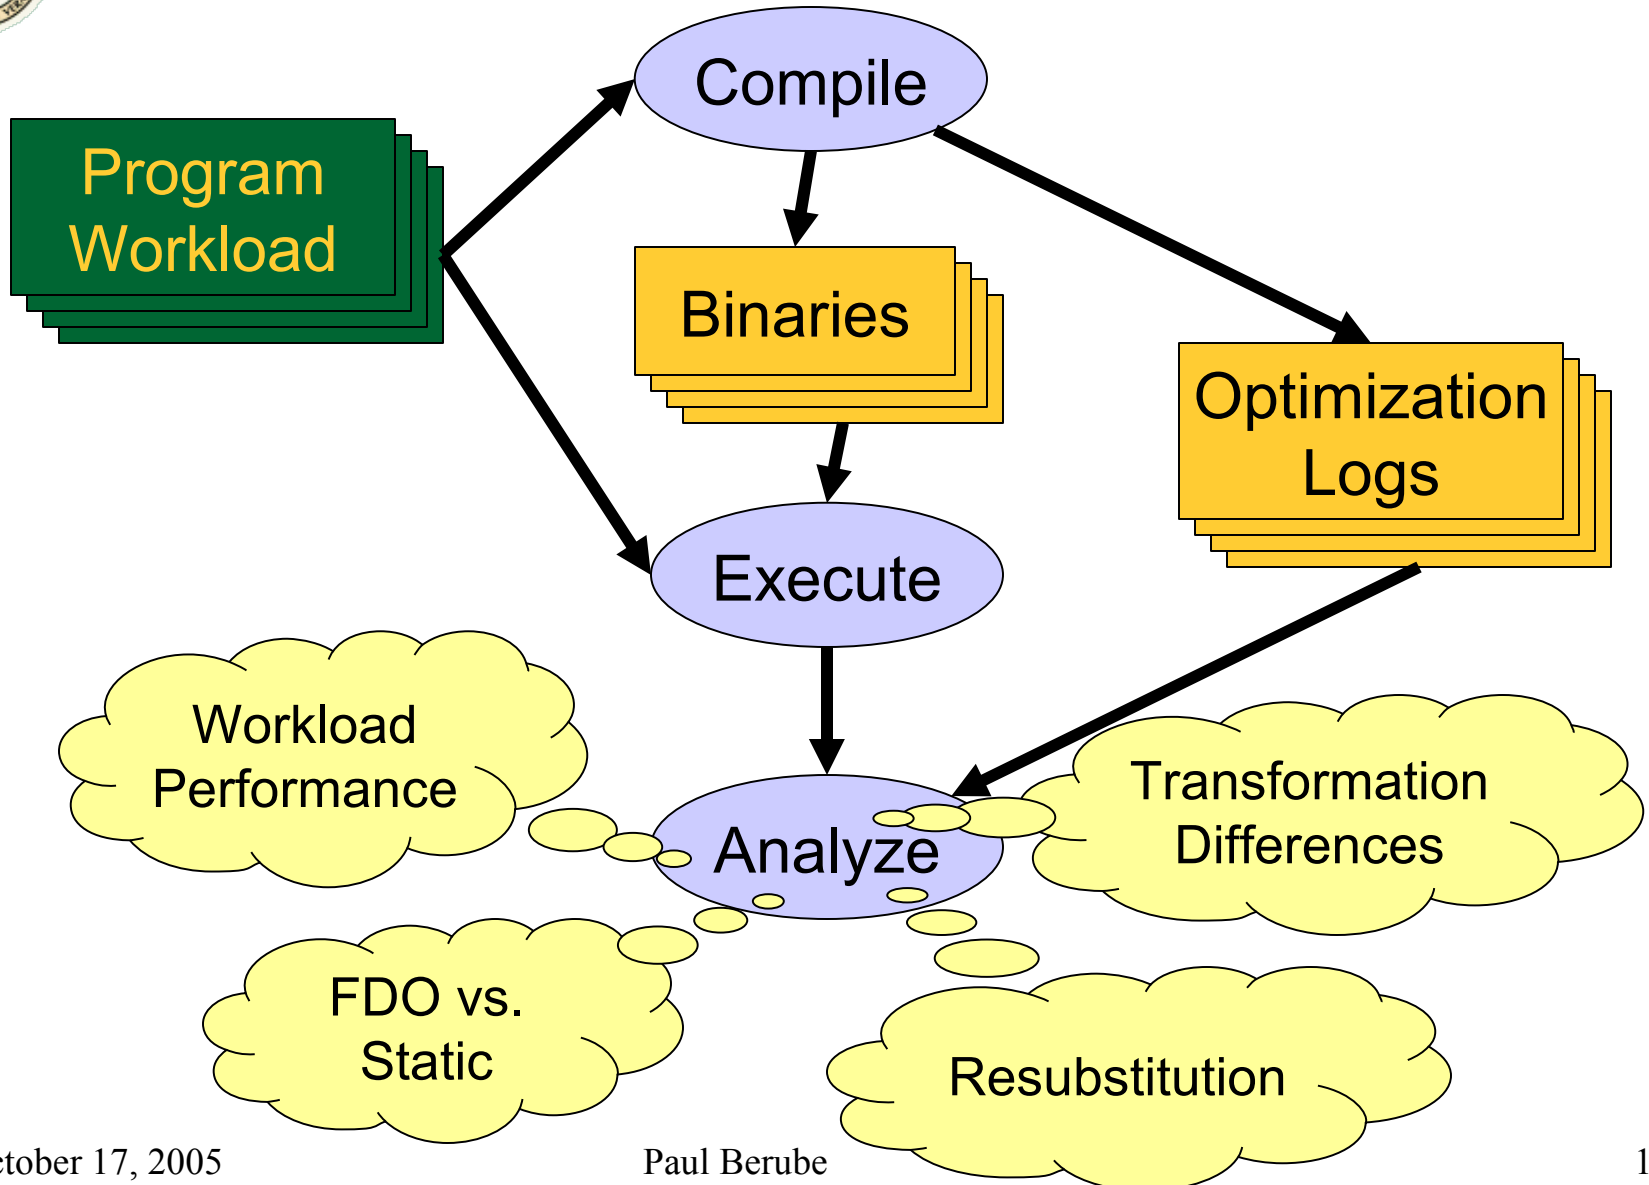


An Overview of *Aestimo*



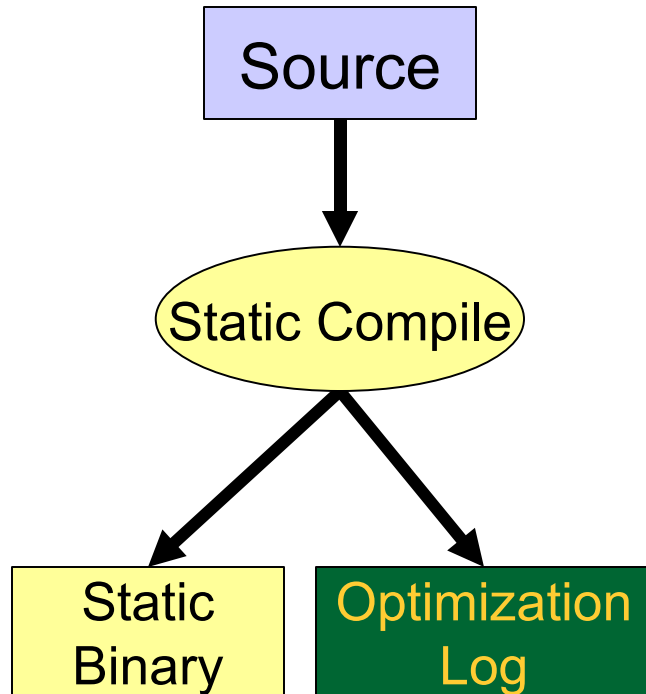


An Overview of *Aestimo*



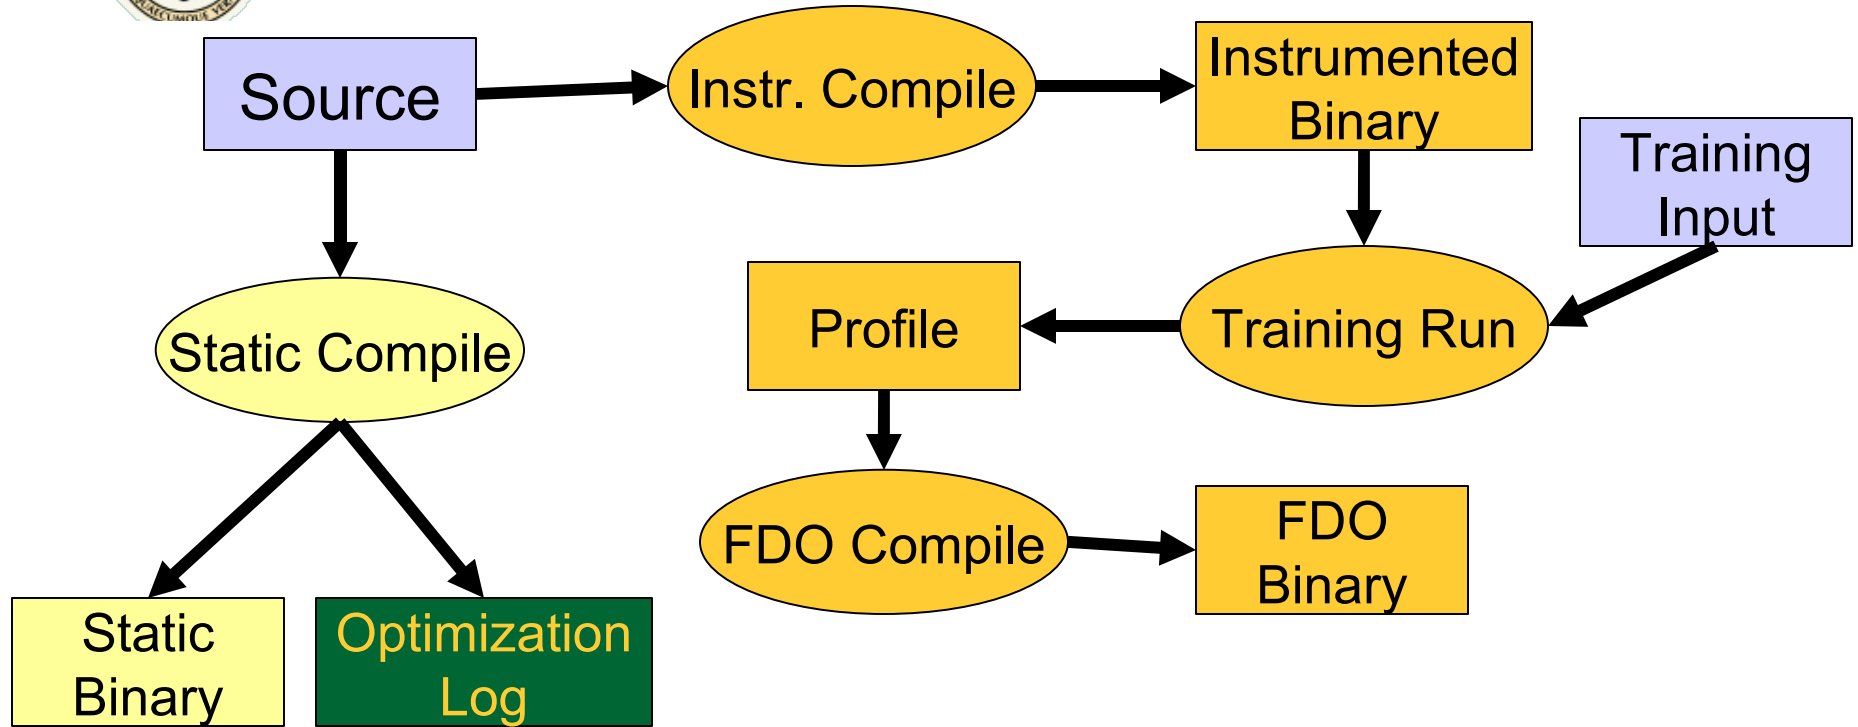


Compilation Process



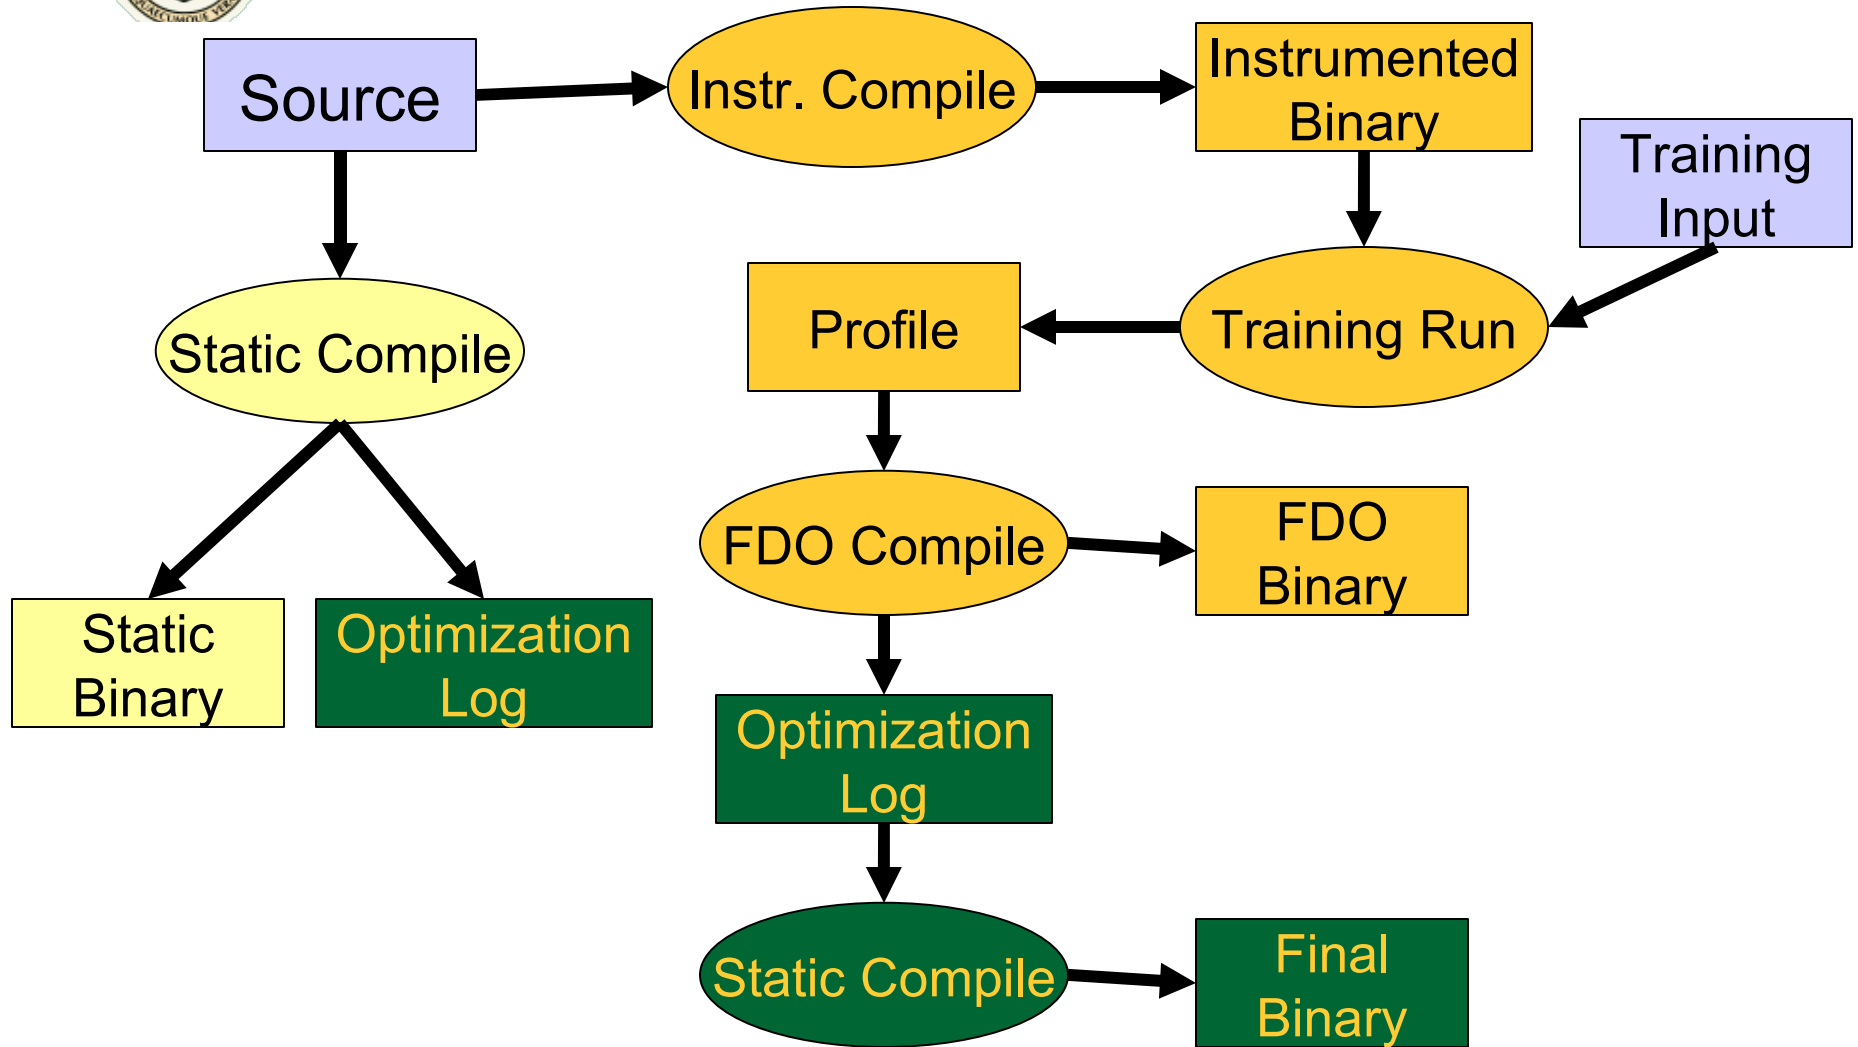


Compilation Process



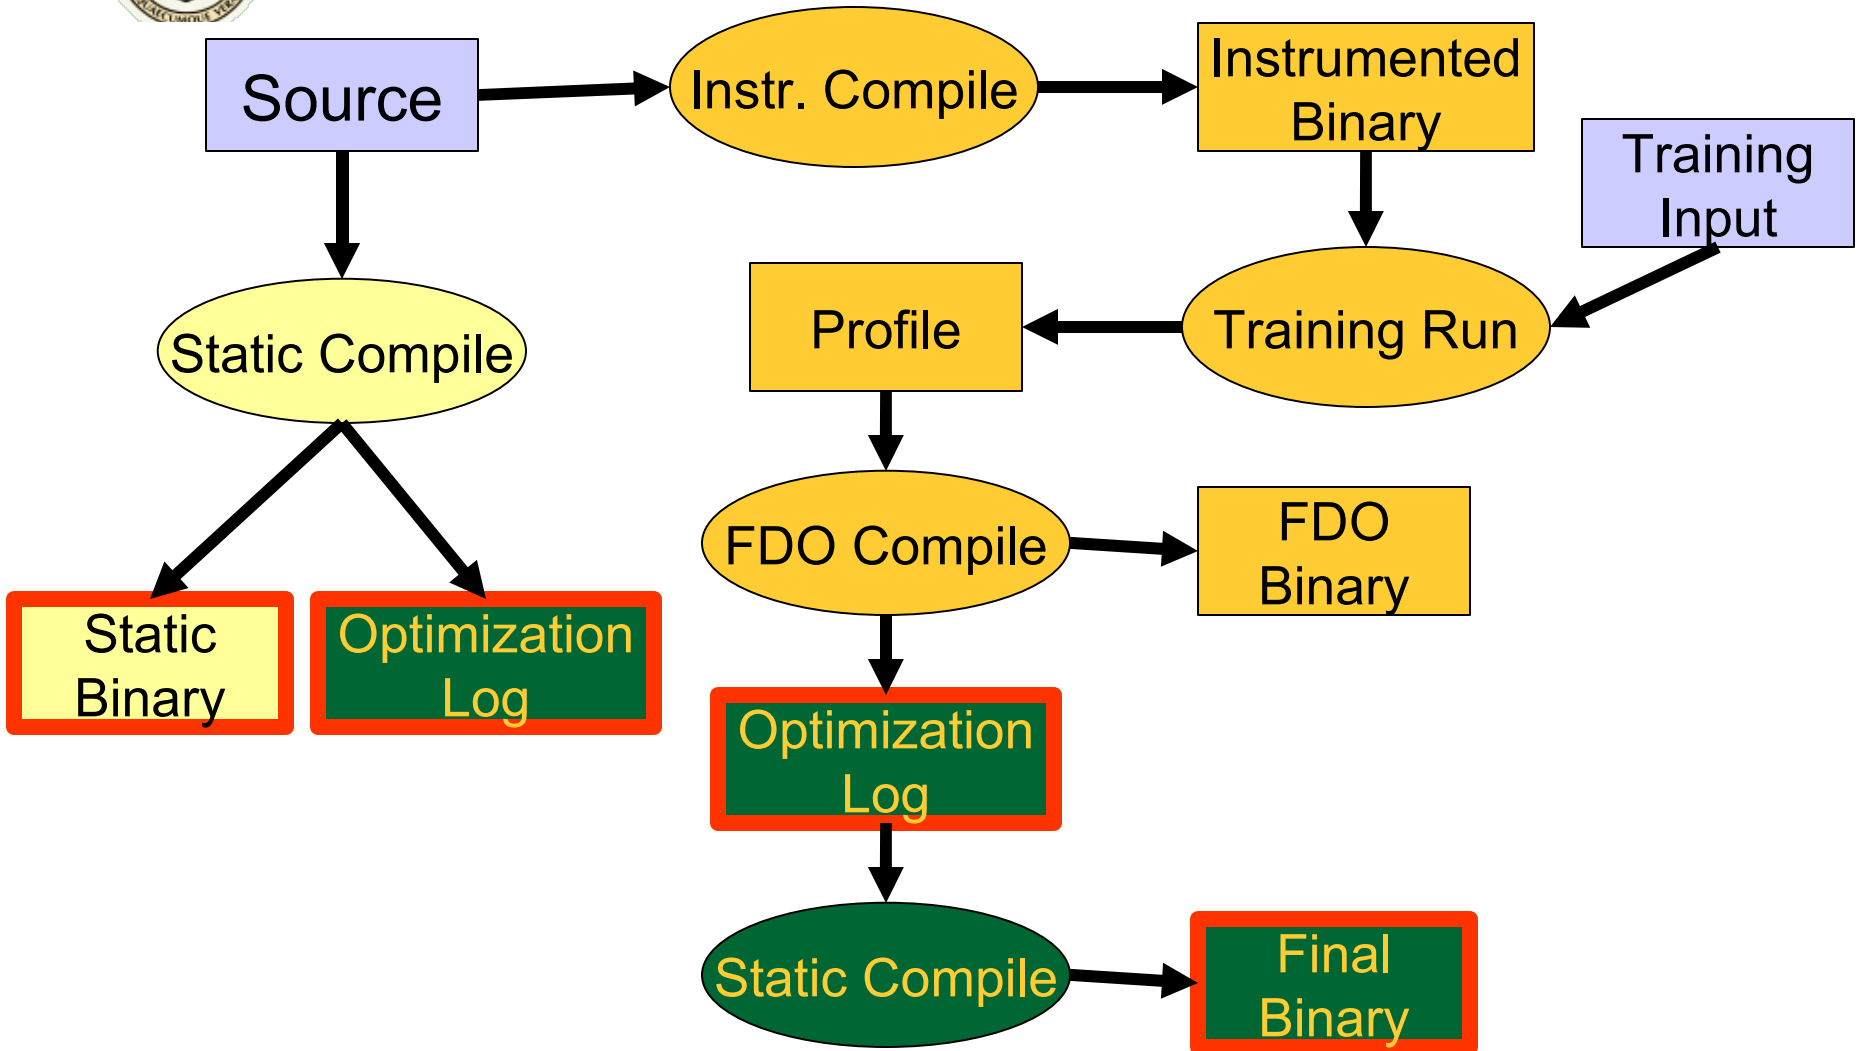


Compilation Process





Compilation Process





Workload Selection

- SPEC CINT2000 Benchmark inputs
 - 8 programs, 32 input
- 84 Additional Inputs
 - Contacted benchmark authors
 - Varied representative inputs
 - Existing collections
 - Synthetic input generator

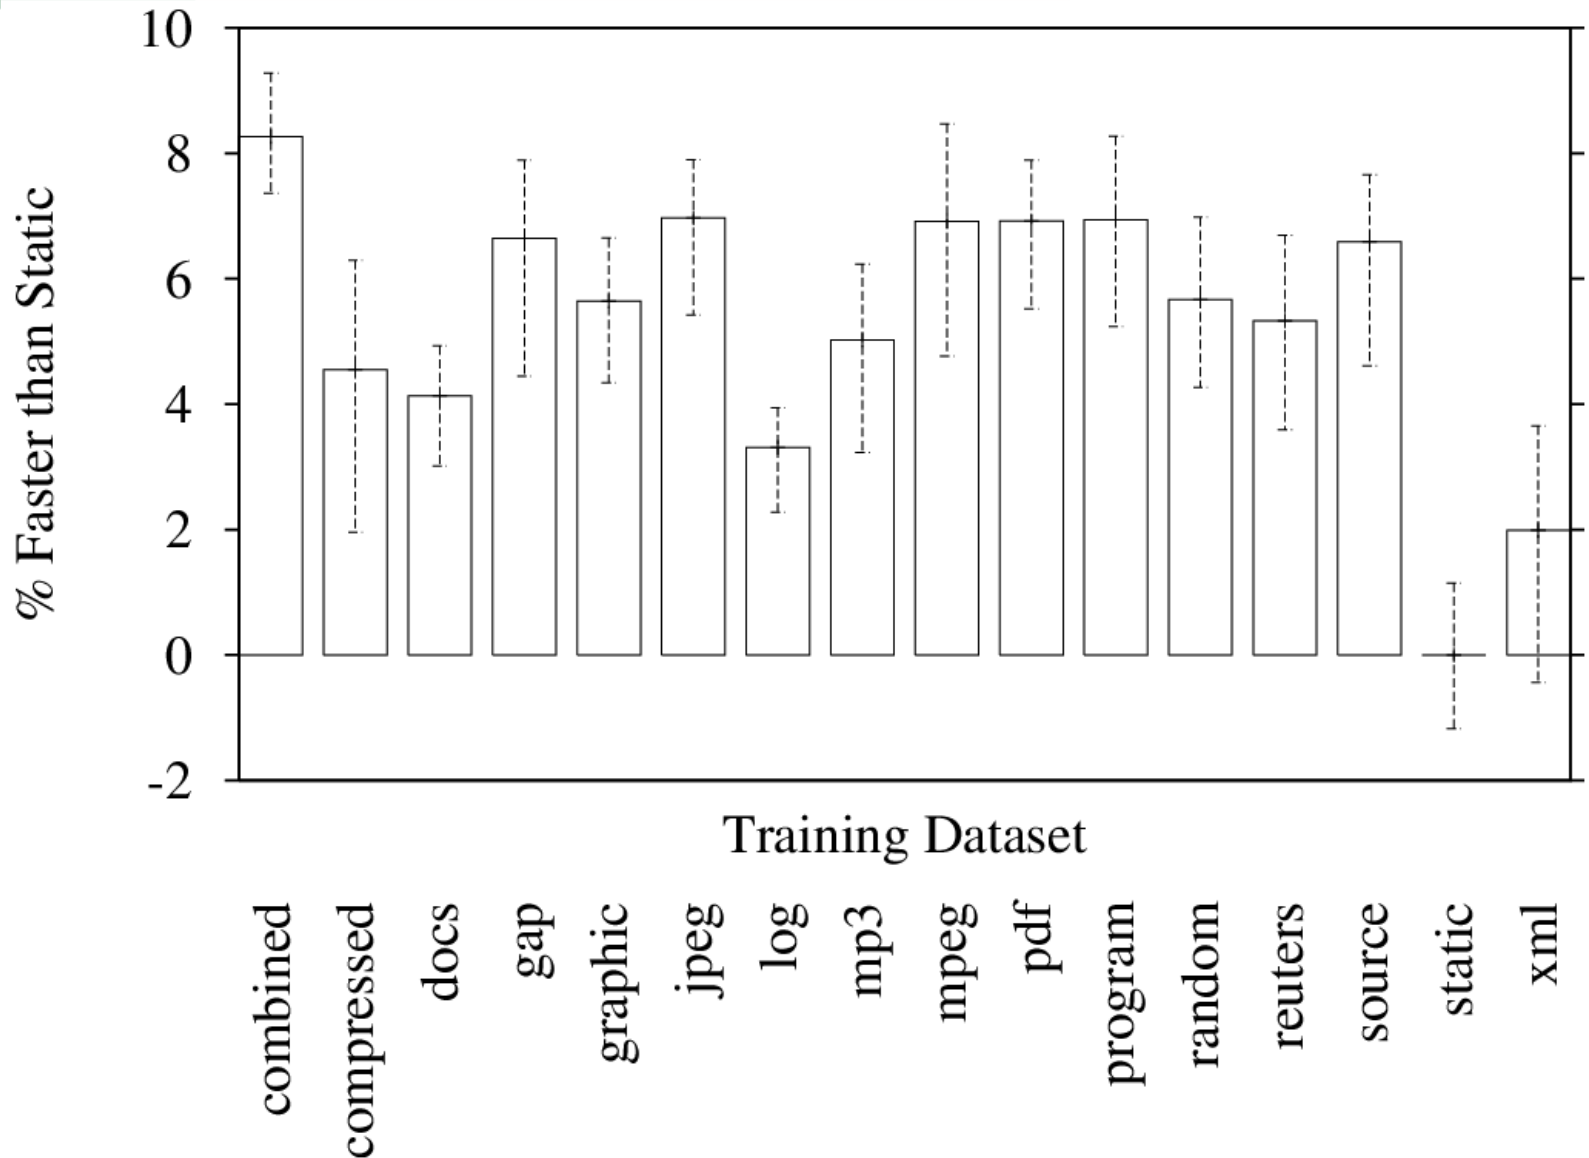


Results

- ORC compiler
- **Inlining** and if conversion
- **Itanium** and Itanium 2 processors



Workload Performance: bzip2

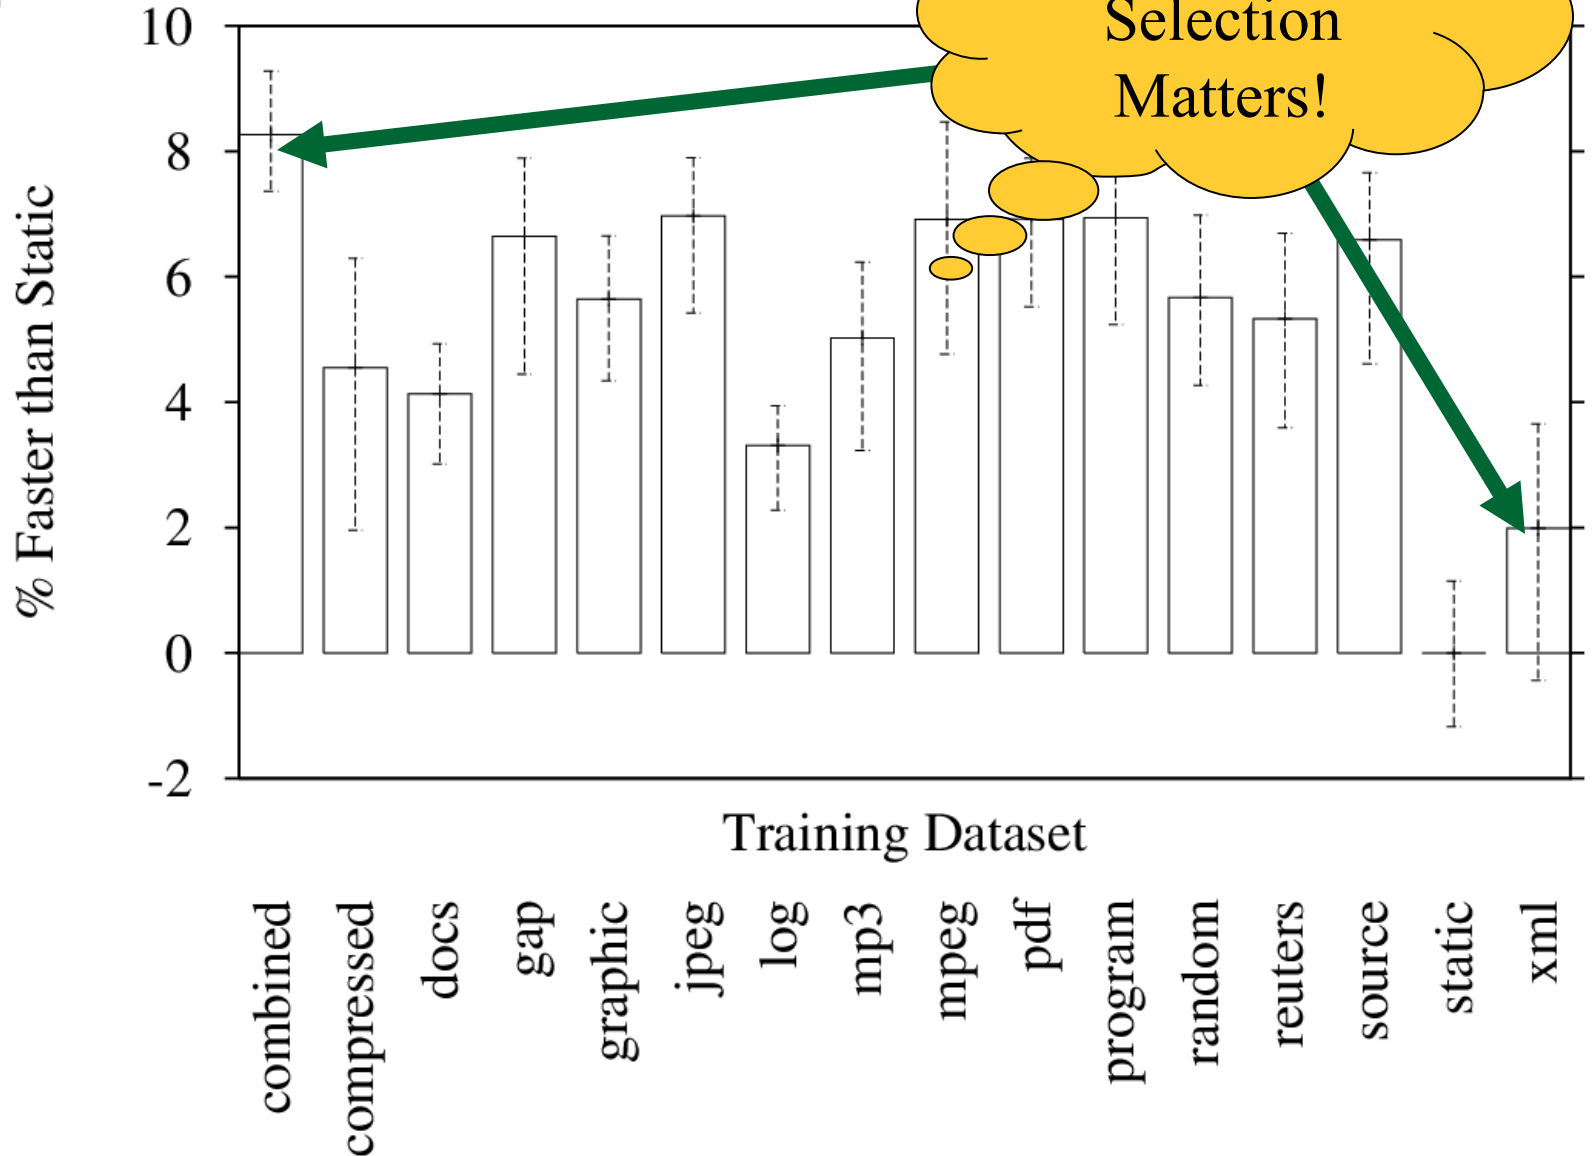


Inlining
Itanium



Workload Performance

Training Input Selection Matters!



Inlining
Itanium



Summary of Contributions

- Training input selection does impact optimization decisions and performance
- *Aestimo*:
 - Automates training and evaluating on a large number of inputs
 - Isolates individual transformations
- A large collection of representative inputs for SPEC CINT2000 programs

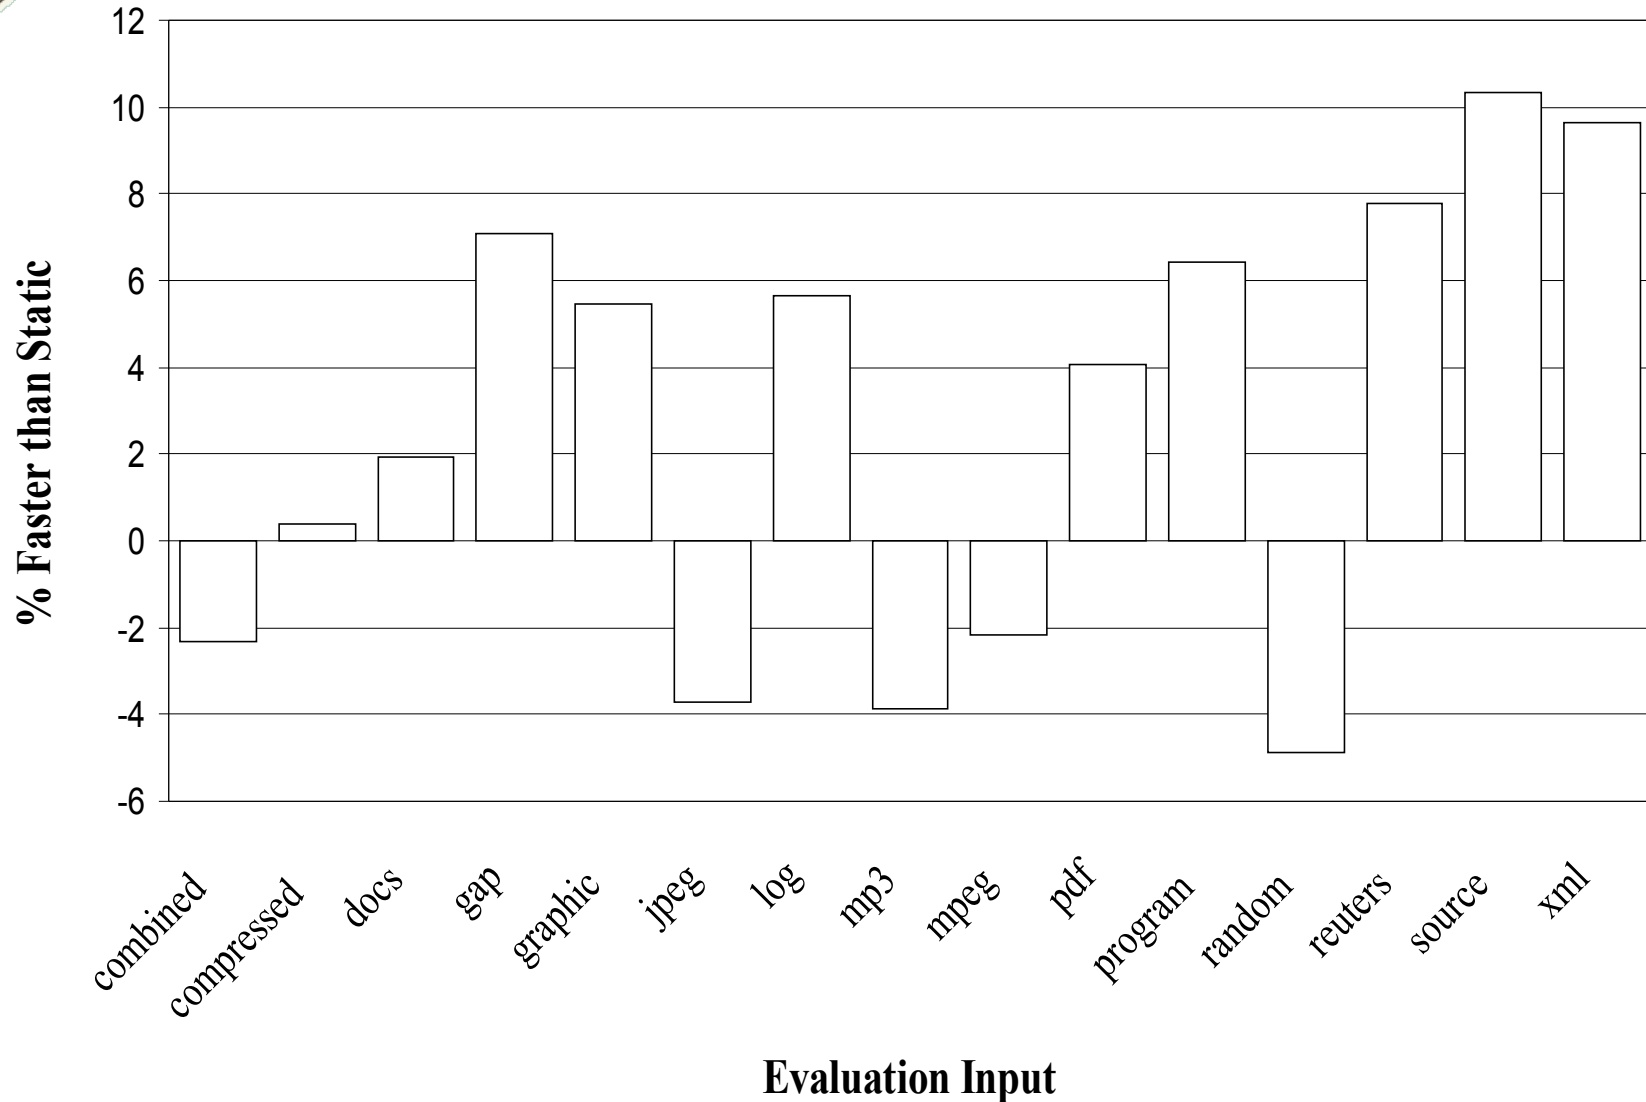


Thank You

Questions?




Performance: bzip2 trained on xml



Inlining
Itanium



Performance: bzip2.combined



Inlining
Itanium