

# Improving Inlining Decisions in the Open Research Compiler

Peng Zhao  
José Nelson Amaral

University of Alberta  
Edmonton, AB, Canada

*<http://www.cs.ualberta.ca/~amaral>  
[amaral@cs.ualberta.ca](mailto:amaral@cs.ualberta.ca)*

Compiler-Driven Performance,  
CASCON, Markham, Oct. 2003



Making  
IT  
happen **Computing Science**

# Yet Another Paper on Inlining?

What is new?

**Adapt Decisions to Benchmark Sizes**

Aggressive for small benchmarks, careful for large ones.

**Use Cycle Density to Control Code Bloat**

A correction to the temperature heuristic in ORC.

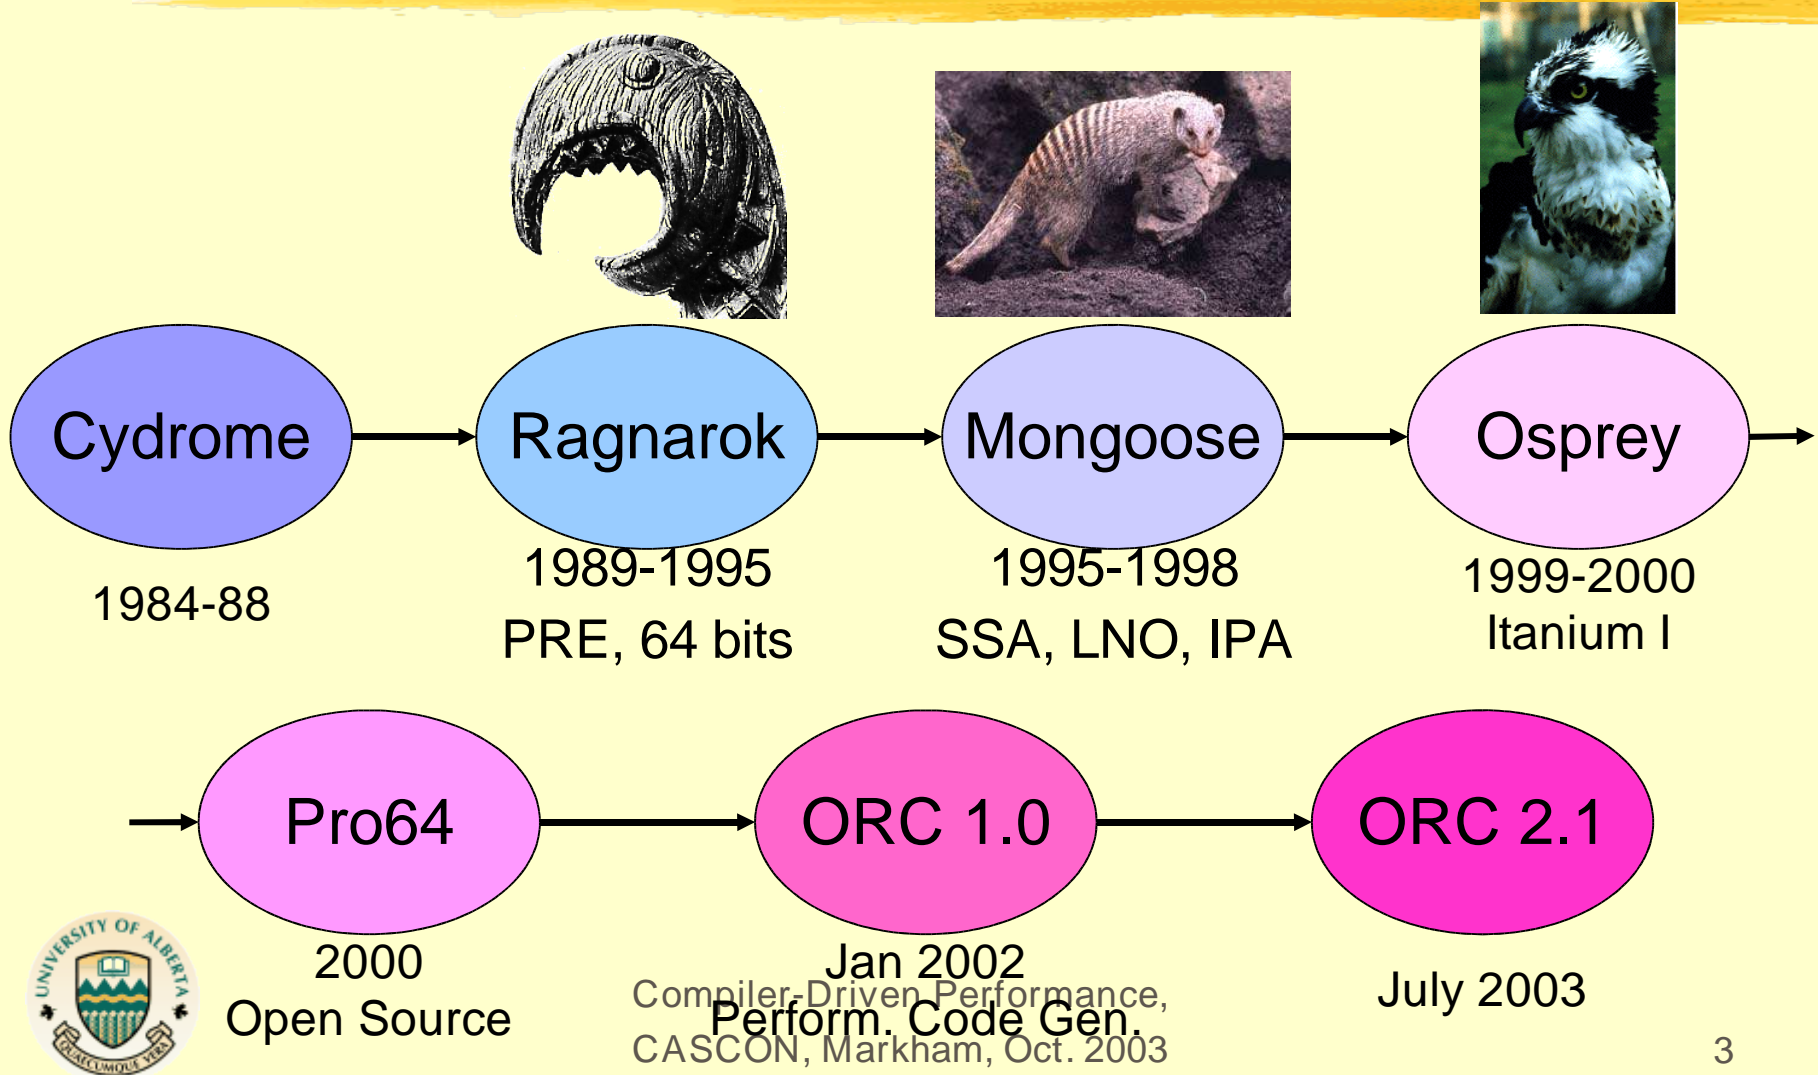
**What is left to do?**

Investigated why remainder procedures not inlined.

Next: **Partial Inlining** and **Recursive Procedure Inlining**.



# Open Research Compiler



# Why inline?

## Eliminate function call overhead

Building stack frame, passing parameters, ...

## Increase scope for code analysis

Better identification for loop optimizations

## Improve code placement

Affine code can be placed nearby



# Why not to inline?

## Code bloat

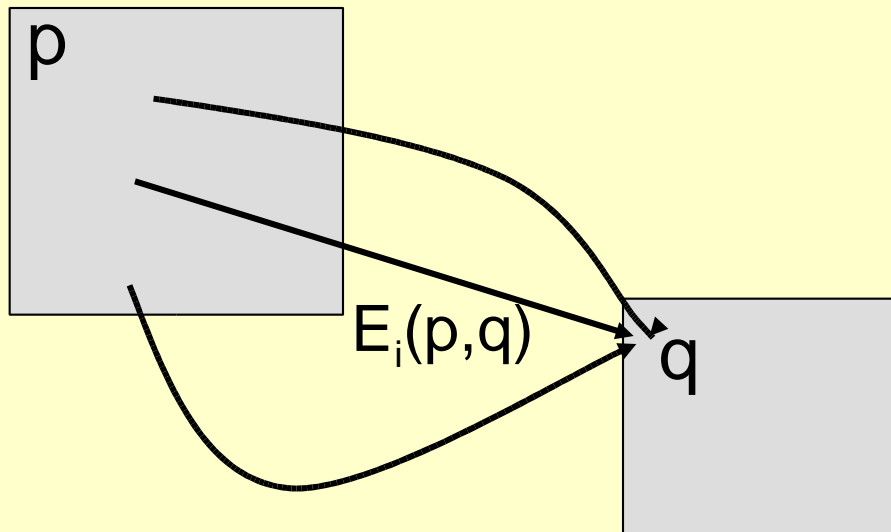
Negative instruction cache effects

## Compiler resources

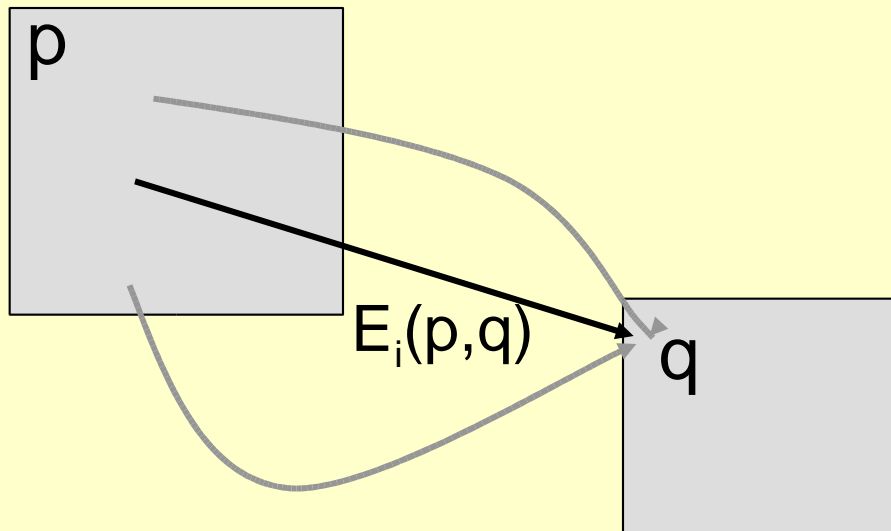
Some analysis may choke on large procedures



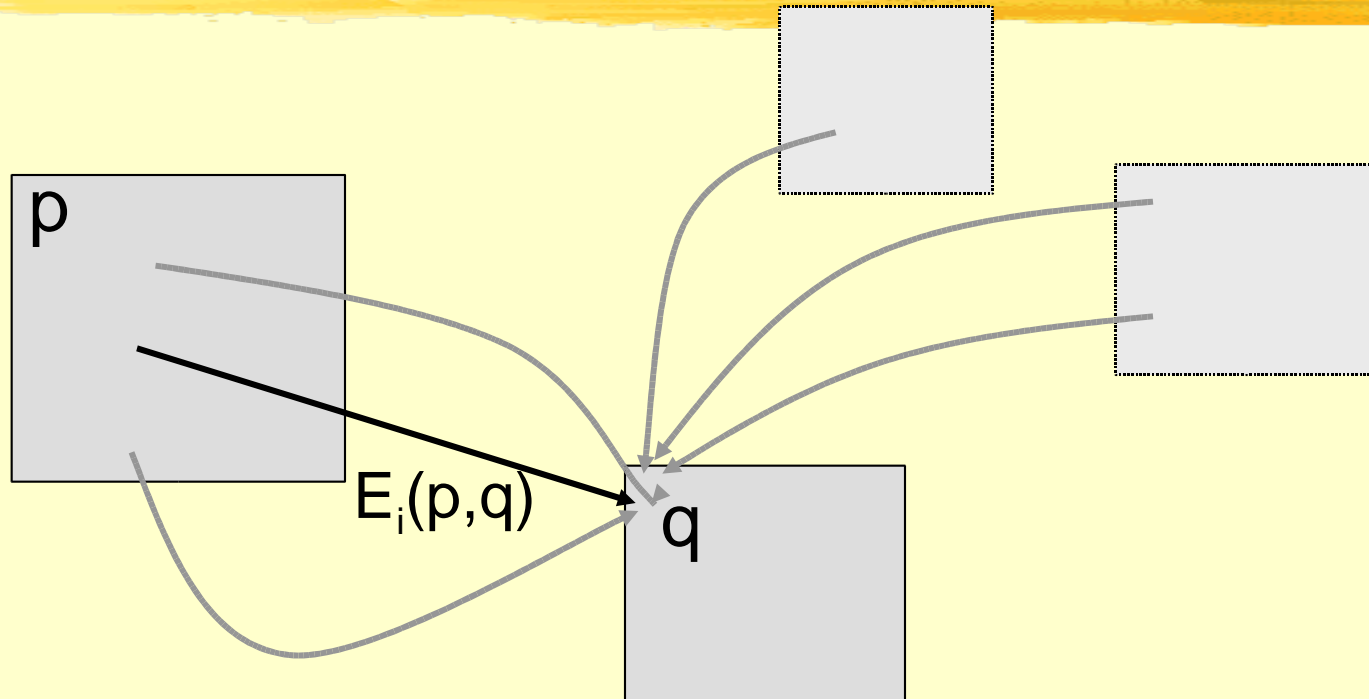
# Original ORC Inlining Heuristic



# Original ORC Inlining Heuristic



# Original ORC Inlining Heuristic

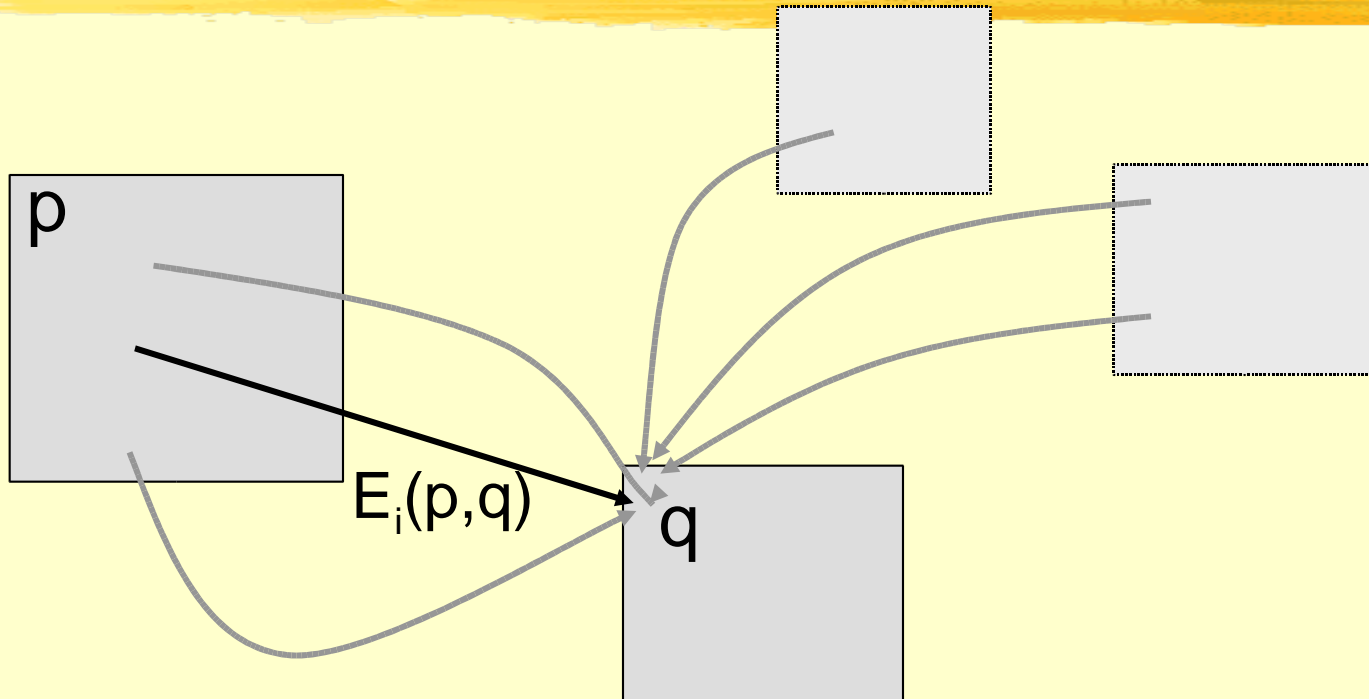


Temperature[ $E_i(p,q)$ ] = ?





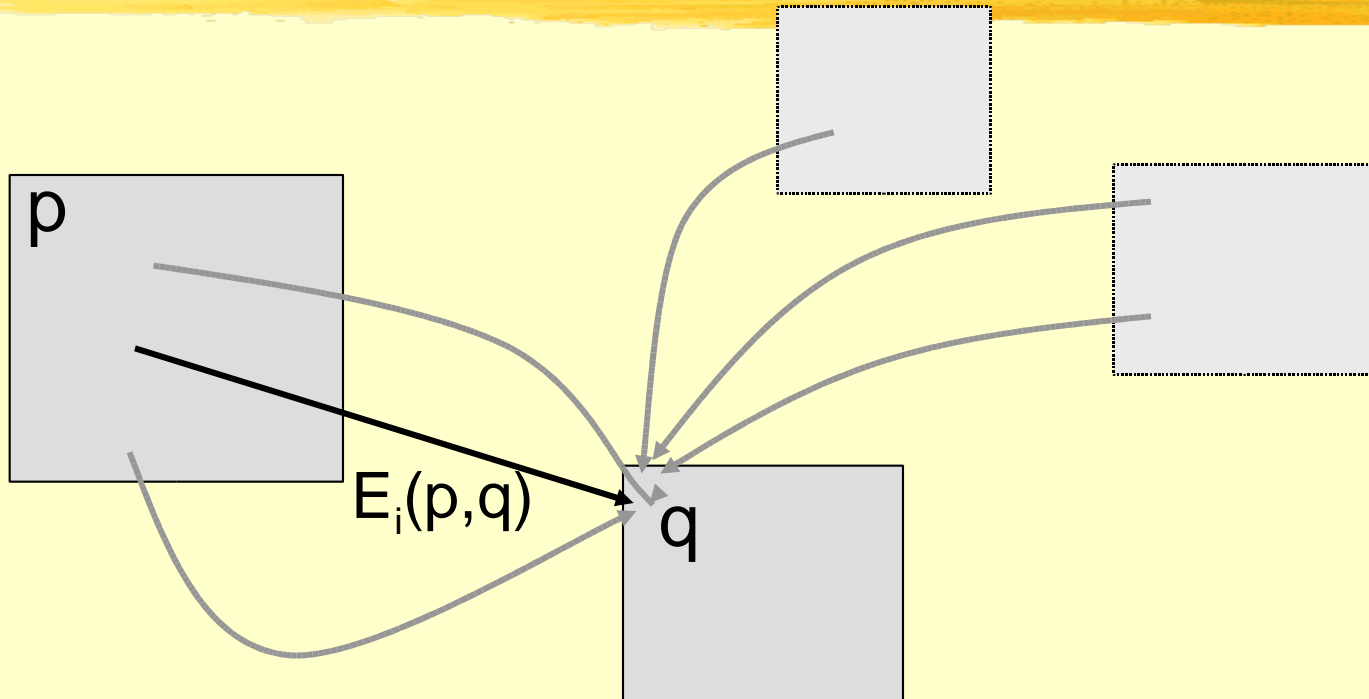
# Original ORC Inlining Heuristic



$$\text{cycle\_ratio}[E_i(p,q)] = \frac{\text{cycle\_count}(q)}{\text{Total\_cycle\_count}} \times \frac{\text{call\_freq}[E_i(p,q)]}{\text{call\_freq}(q)}$$



# Original ORC Inlining Heuristic



$$\text{size\_ratio}(q) = \frac{\text{size}(q)}{\text{Total\_application\_size}}$$



# Original ORC Inlining Heuristic

$$\text{Temperature}[E_i(p,q)] = \frac{\text{cycle\_ratio}[E_i(p,q)]}{\text{size\_ratio}(q)}$$

$$\text{cycle\_ratio}[E_i(p,q)] = \frac{\text{cycle\_count}(q)}{\text{Total\_cycle\_count}} \times \frac{\text{call\_freq}[E_i(p,q)]}{\text{call\_freq}(q)}$$

$$\text{size\_ratio}(q) = \frac{\text{size}(q)}{\text{Total\_application\_size}}$$



# Original ORC Inlining Heuristic

$$\text{Temperature}[E_i(p,q)] = \frac{\text{cycle\_ratio}[E_i(p,q)]}{\text{size\_ratio}(q)}$$

Edges called often are hot. **Good!**

$$\text{cycle\_ratio}[E_i(p,q)] = \frac{\text{cycle\_count}(q)}{\text{Total\_cycle\_count}} \times \frac{\text{call\_freq}[E_i(p,q)]}{\text{call\_freq}(q)}$$

$$\text{size\_ratio}(q) = \frac{\text{size}(q)}{\text{Total\_application\_size}}$$



# Original ORC Inlining Heuristic

$$\text{Temperature}[E_i(p,q)] = \frac{\text{cycle\_ratio}[E_i(p,q)]}{\text{size\_ratio}(q)}$$

Functions that execute longer are hot. **Good!**

$$\text{cycle\_ratio}[E_i(p,q)] = \frac{\text{cycle\_count}(q)}{\text{Total\_cycle\_count}} \times \frac{\text{call\_freq}[E_i(p,q)]}{\text{call\_freq}(q)}$$

$$\text{size\_ratio}(q) = \frac{\text{size}(q)}{\text{Total\_application\_size}}$$



# Original ORC Inlining Heuristic

$$\text{Temperature}[E_i(p,q)] = \frac{\text{cycle\_ratio}[E_i(p,q)]}{\text{size\_ratio}(q)}$$

Even if they are not called often. **Bad!**

$$\text{cycle\_ratio}[E_i(p,q)] = \frac{\text{cycle\_count}(q)}{\text{Total\_cycle\_count}} \times \frac{\text{call\_freq}[E_i(p,q)]}{\text{call\_freq}(q)}$$

$$\text{size\_ratio}(q) = \frac{\text{size}(q)}{\text{Total\_application\_size}}$$



# Original ORC Inlining Heuristic

$$\text{Temperature}[E_i(p,q)] = \frac{\text{cycle\_ratio}[E_i(p,q)]}{\text{size\_ratio}(q)}$$

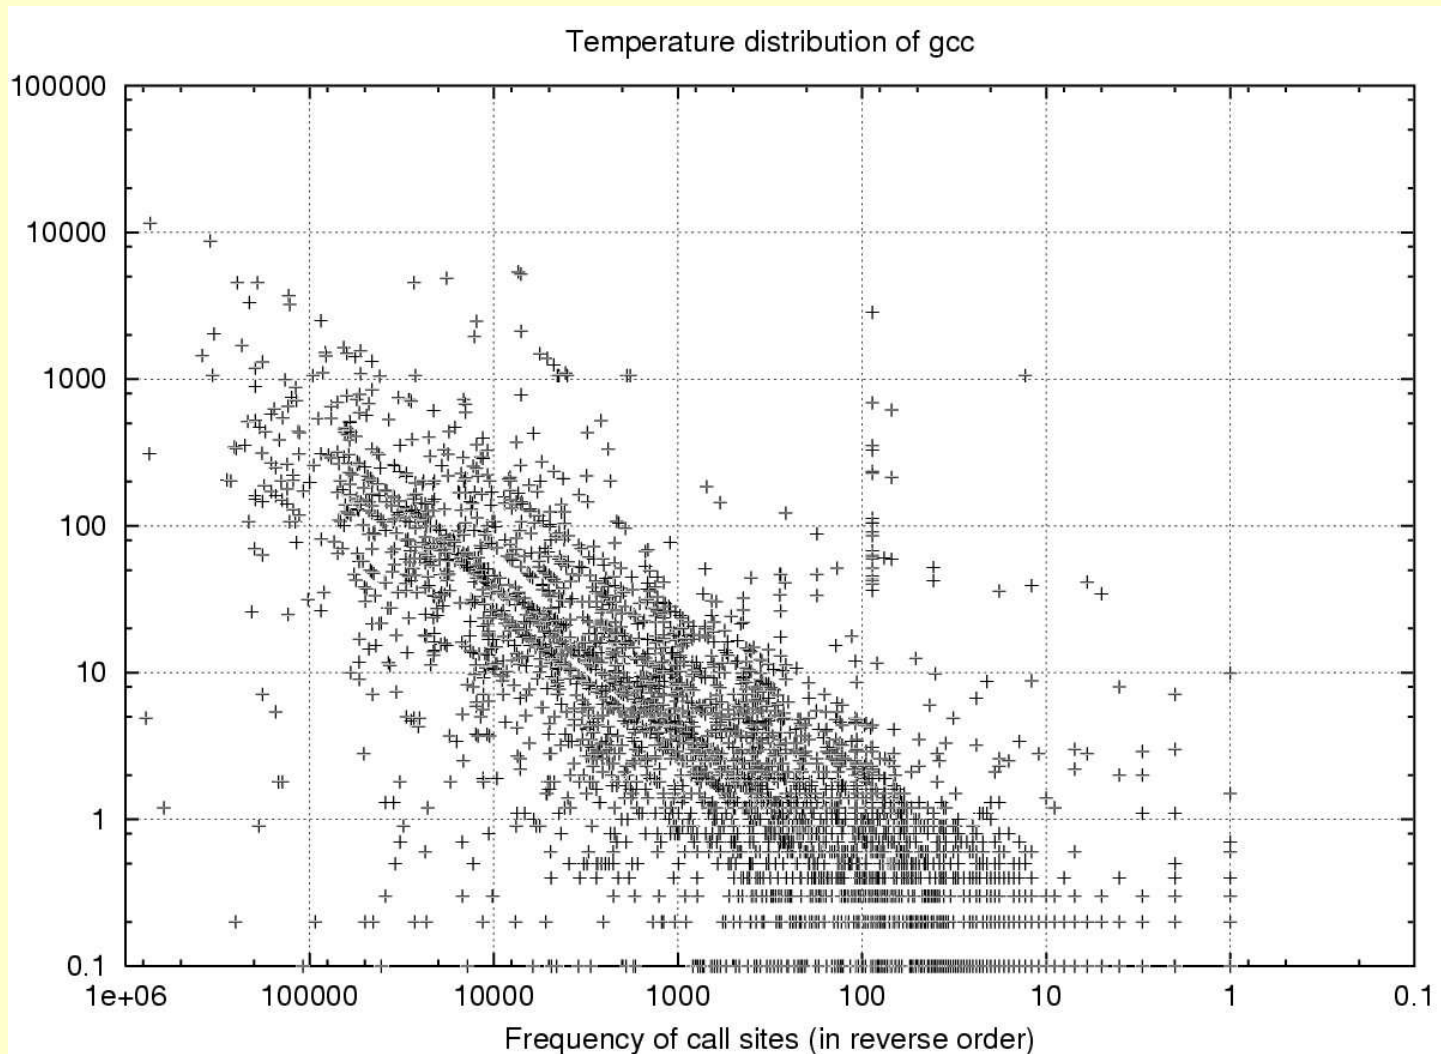
Small Code ==> small functions are cold. **Bad!**

$$\text{cycle\_ratio}[E_i(p,q)] = \frac{\text{cycle\_count}(q)}{\text{Total\_cycle\_count}} \times \frac{\text{call\_freq}[E_i(p,q)]}{\text{call\_freq}(q)}$$

$$\text{size\_ratio}(q) = \frac{\text{size}(q)}{\text{Total\_application\_size}}$$

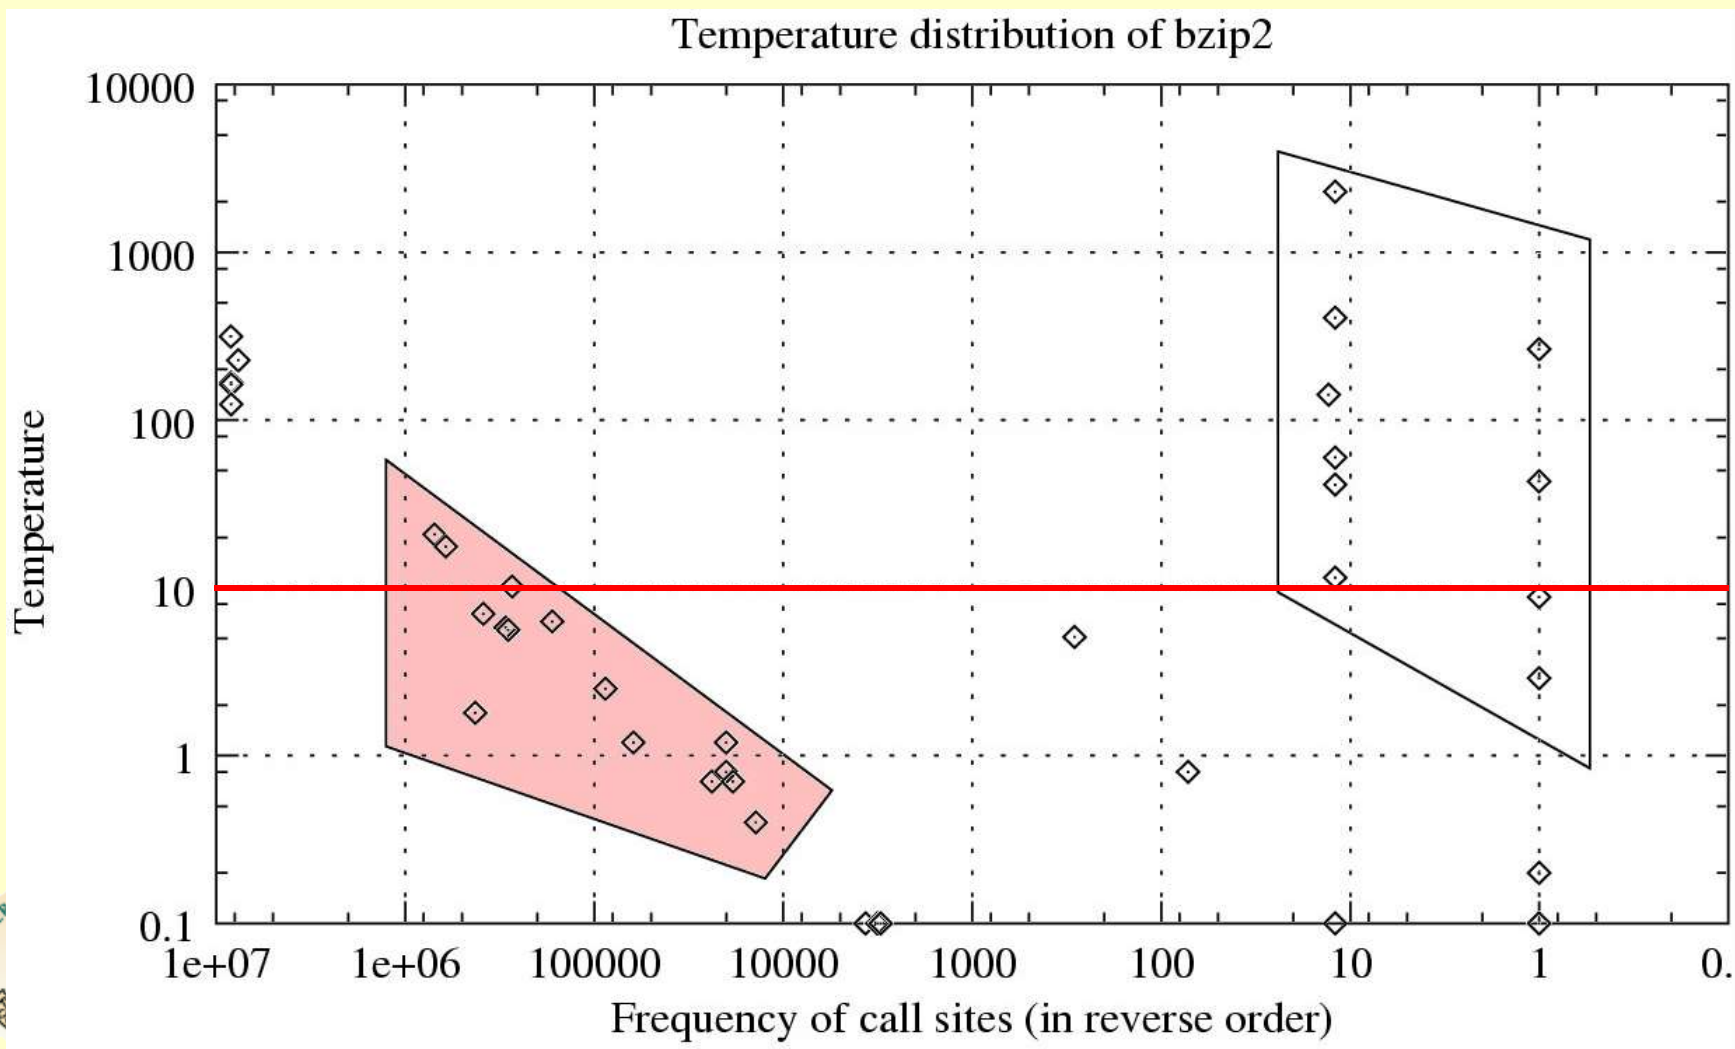


# Temperature Distribution for gcc





# Temperature Distribution for BZIP2



# Adapting the ORC Heuristic to Benchmark Size

Empirical classification of benchmarks  
(based on SPEC):

		Temperature Threshold
<b>Small:</b> < 10,000 AST Nodes	→	1
<b>Medium:</b> Anything in between	→	50
<b>Large:</b> > 250,000 AST Nodes	→	120



# “ Heavy” Procedures

A procedure call that:

- is hot in the original ORC heuristic
- but that is not called often

must have high trip count loops.

We call these **heavy procedures**.

We introduce the **cycle density heuristic** to fix the ORC inlining decisions for heavy procedures.



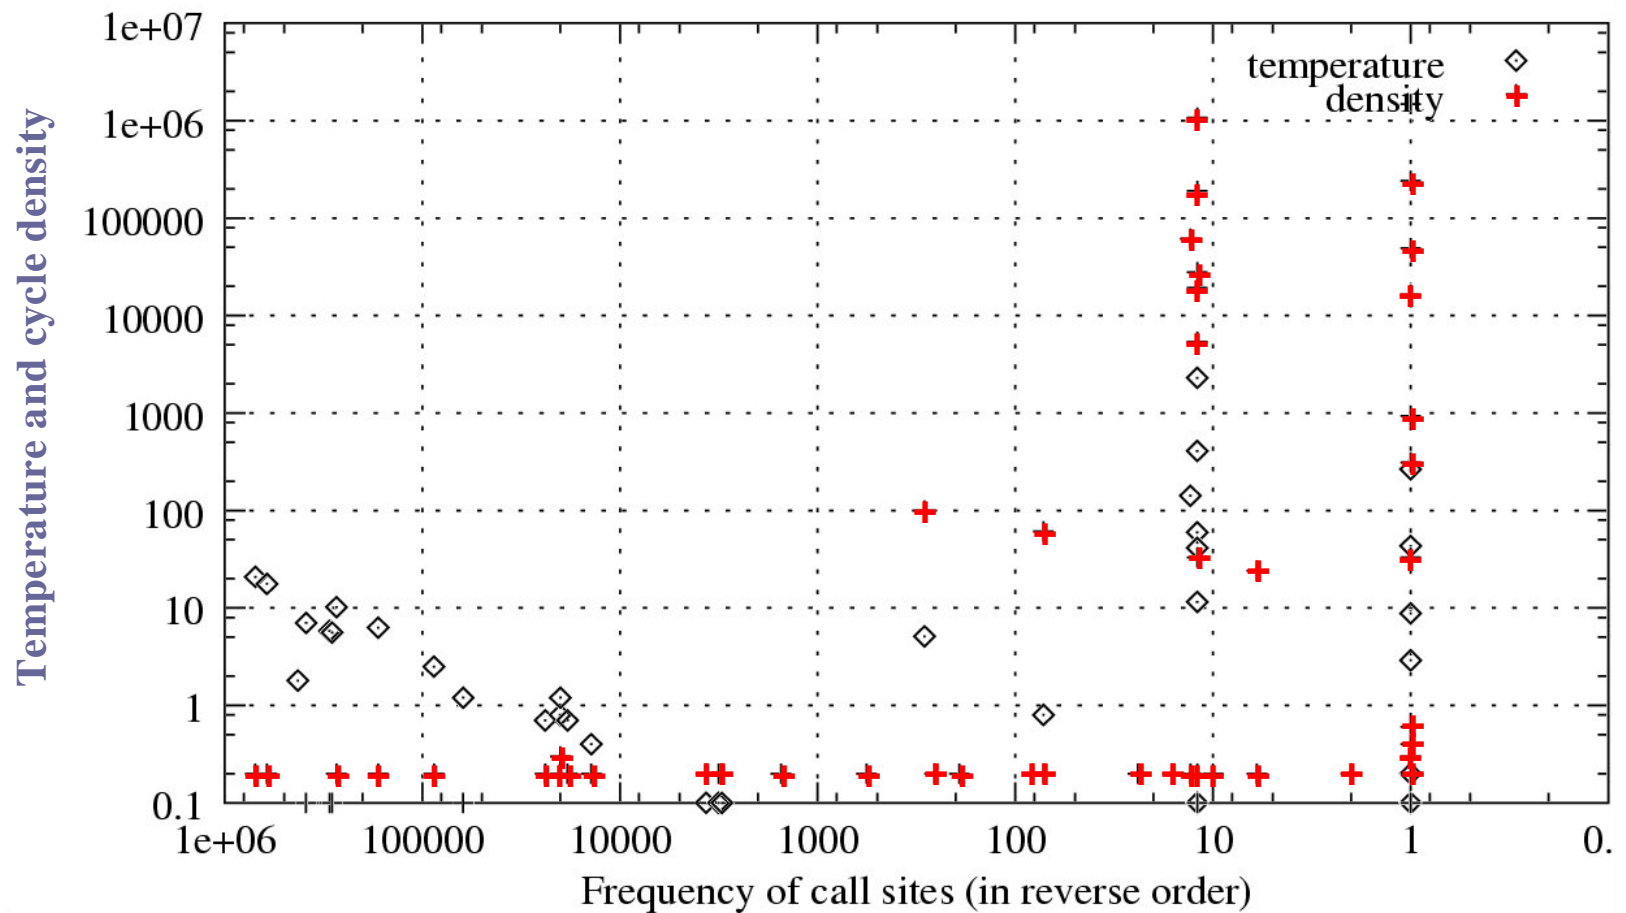
# Cycle Density

$$\text{cycle\_density}(q) = \frac{\text{cycle\_count}(q)}{\text{frequency}(q)}$$

High cycle density indicates a heavy procedure.



# Temp. $\times$ Cycle Density (BZIP2)



# Experimental Study

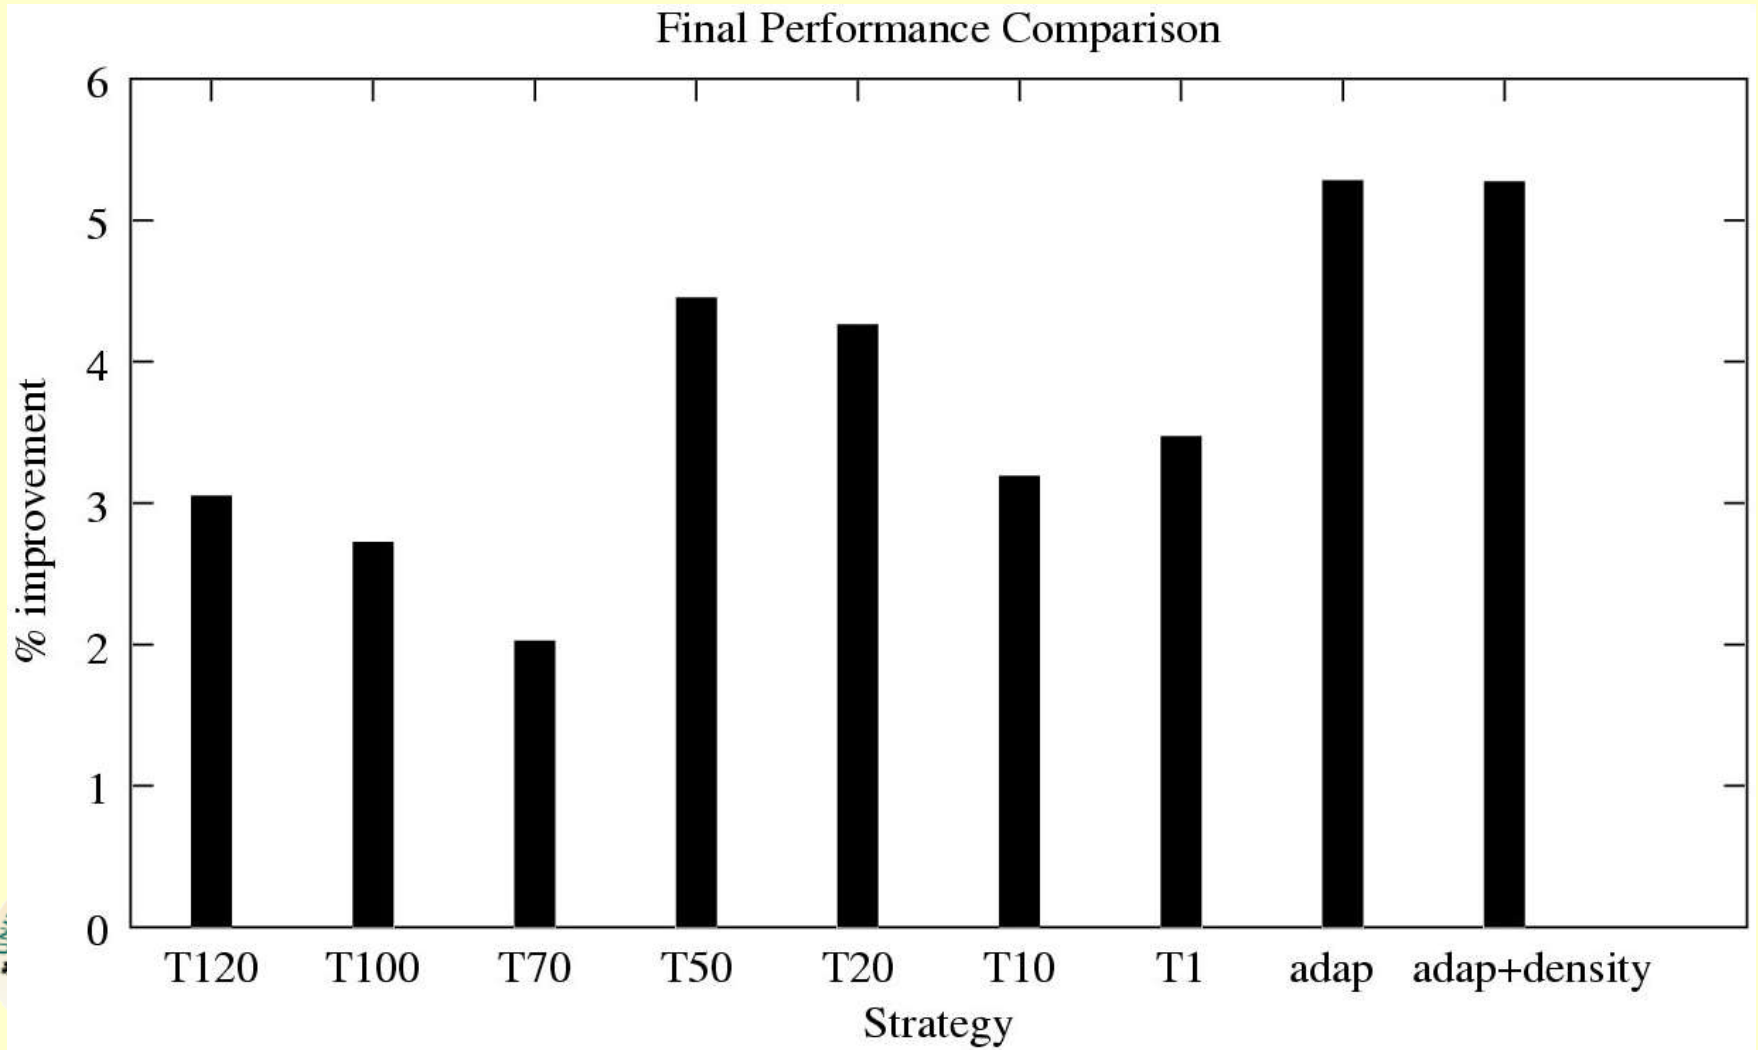
SPEC2000 benchmarks, except EON.

Runtime on an Itanium-I (733 MHz,  
1 GB mem, RH-Linux 7.1)

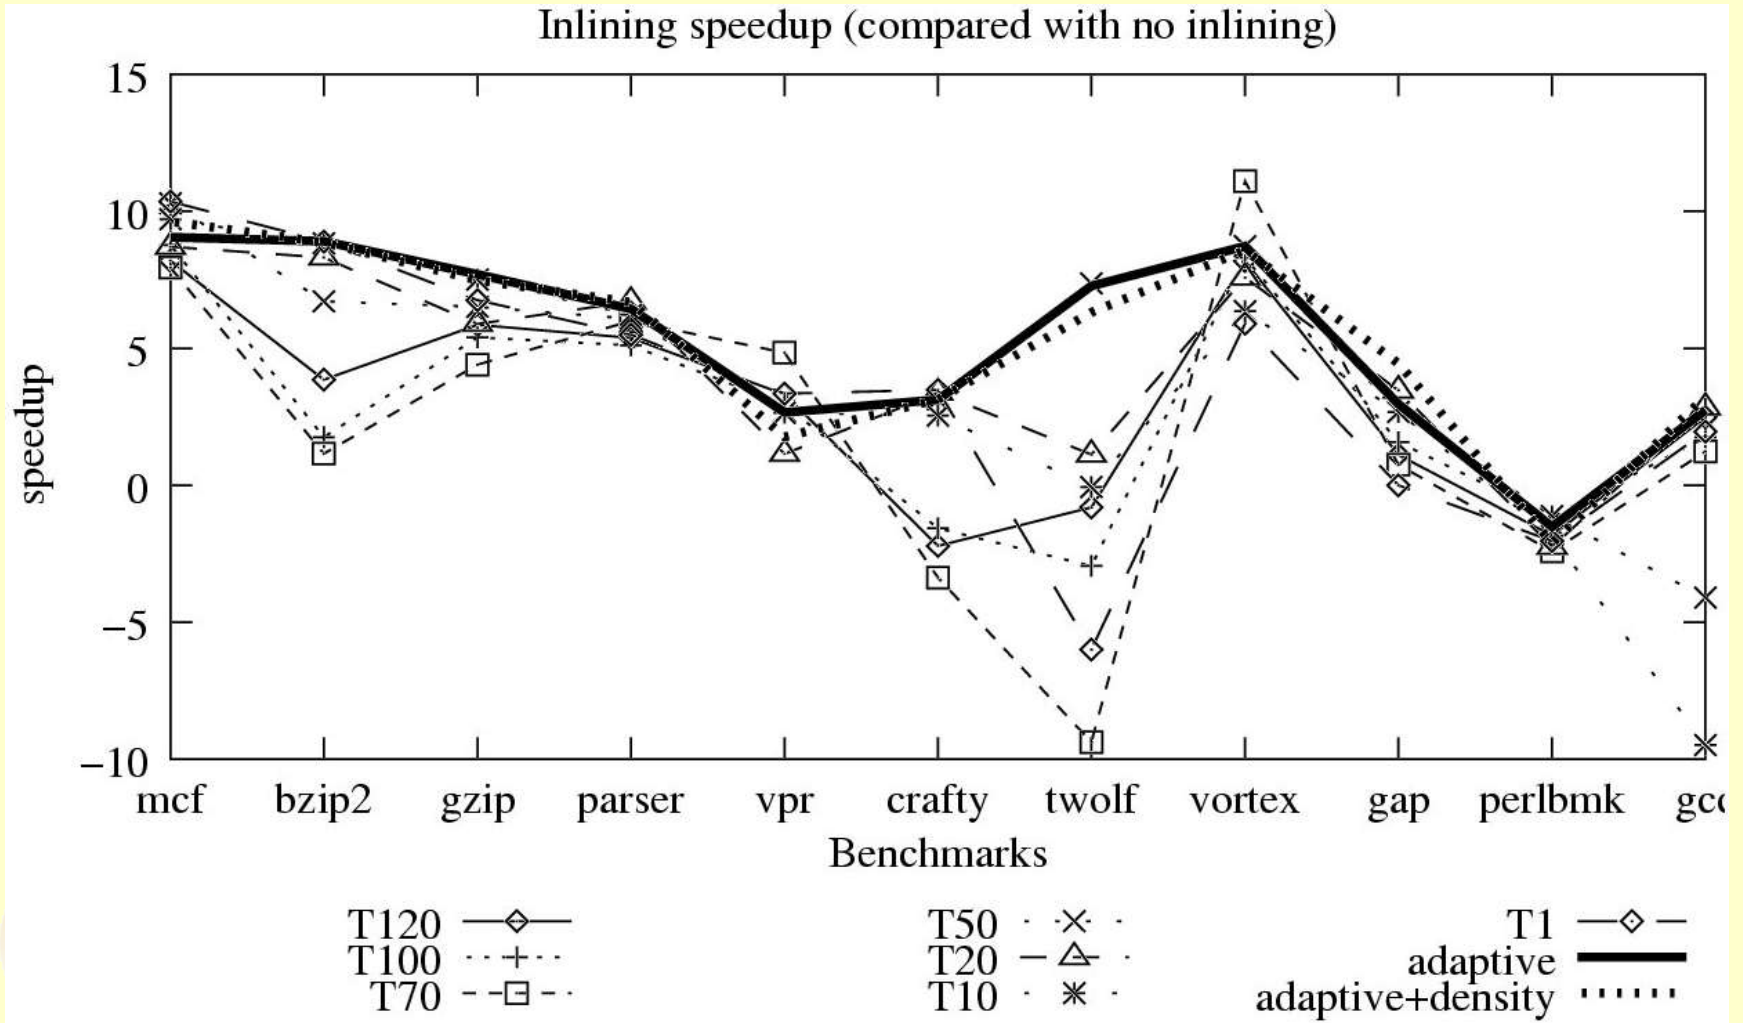
Compilation on a dual Pentium III (600 MHz,  
512 MB mem, RH-Linux 7.2)



# Average Performance Improvement on SPEC2000

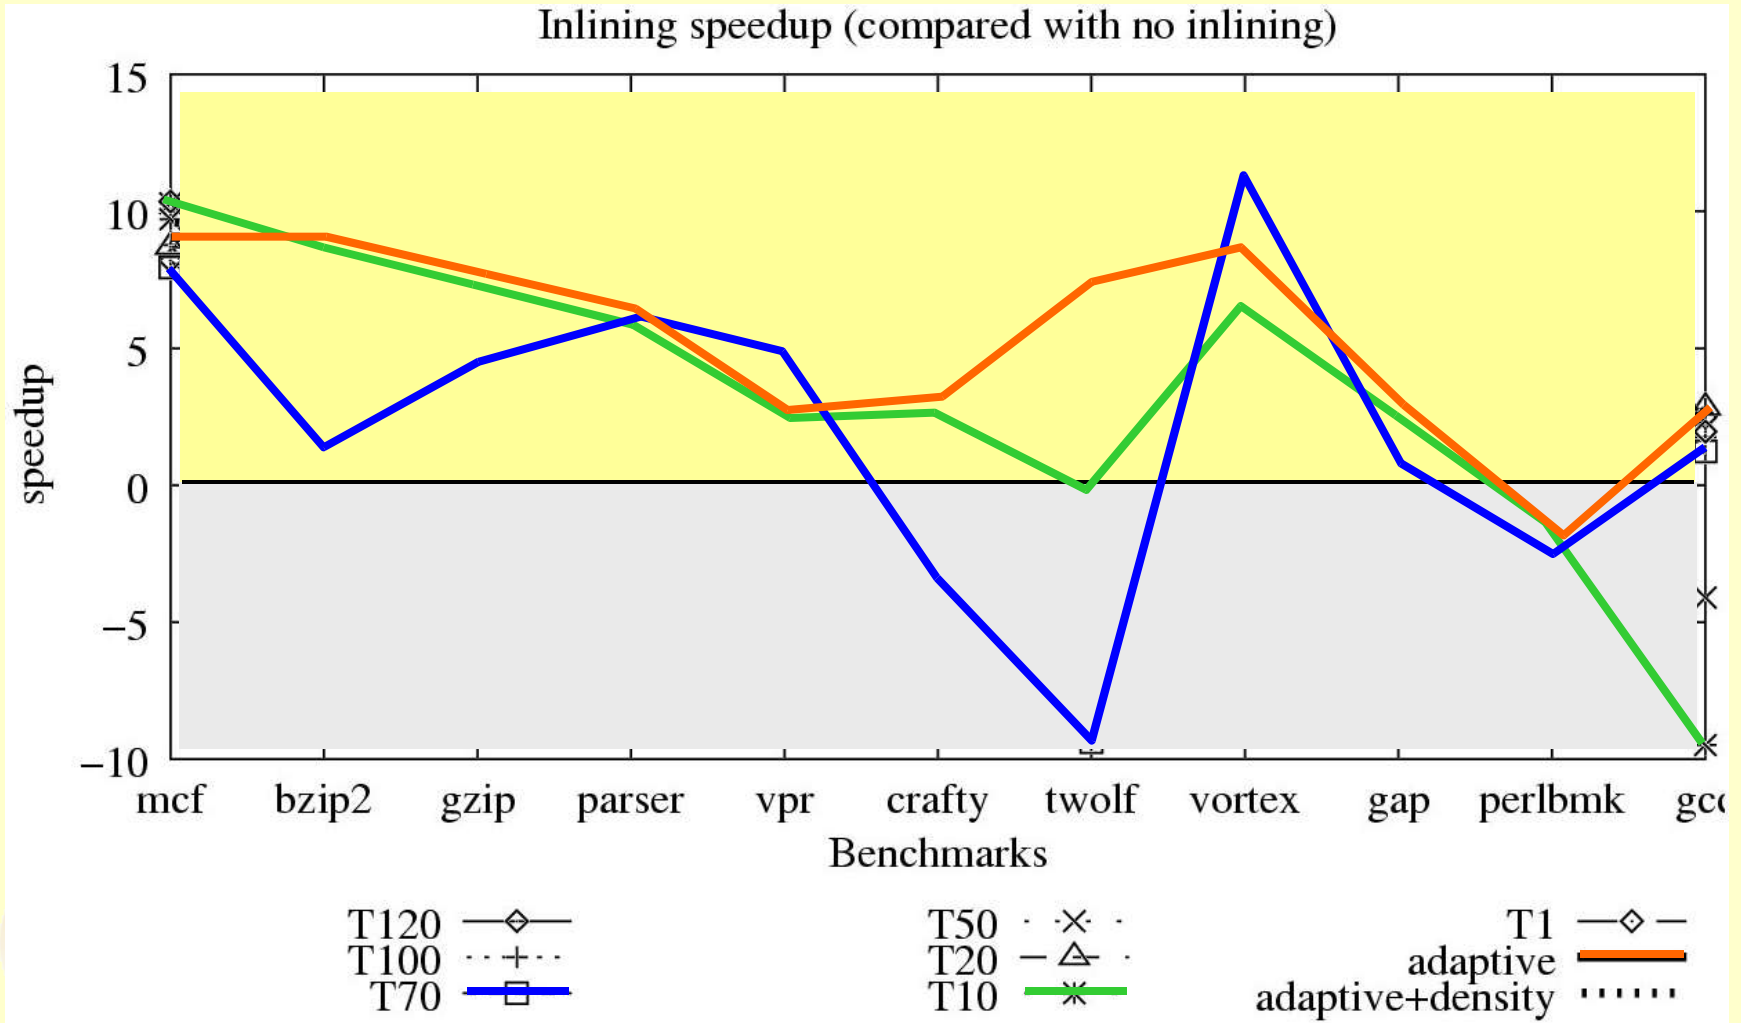


# Inlining Speedup

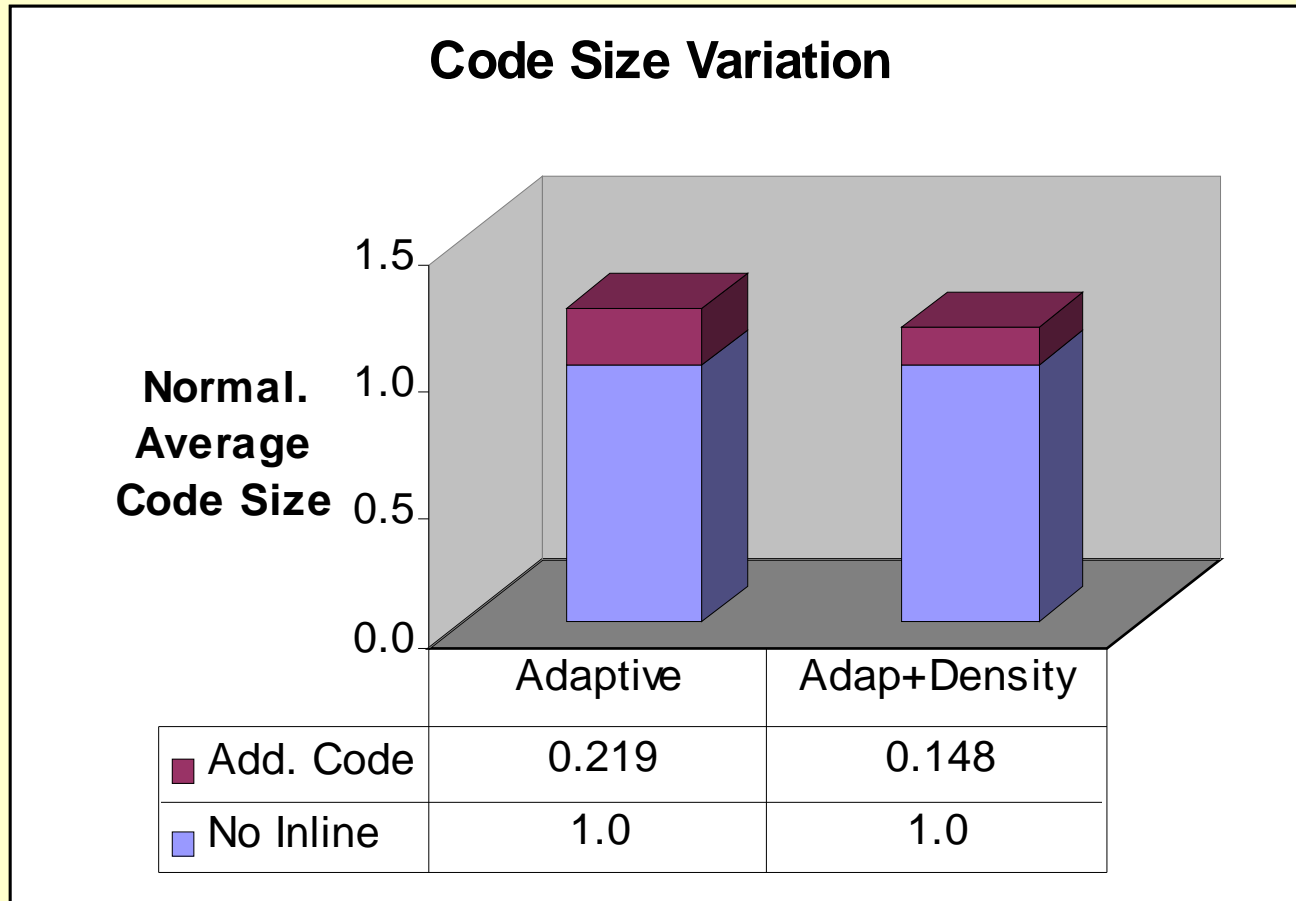




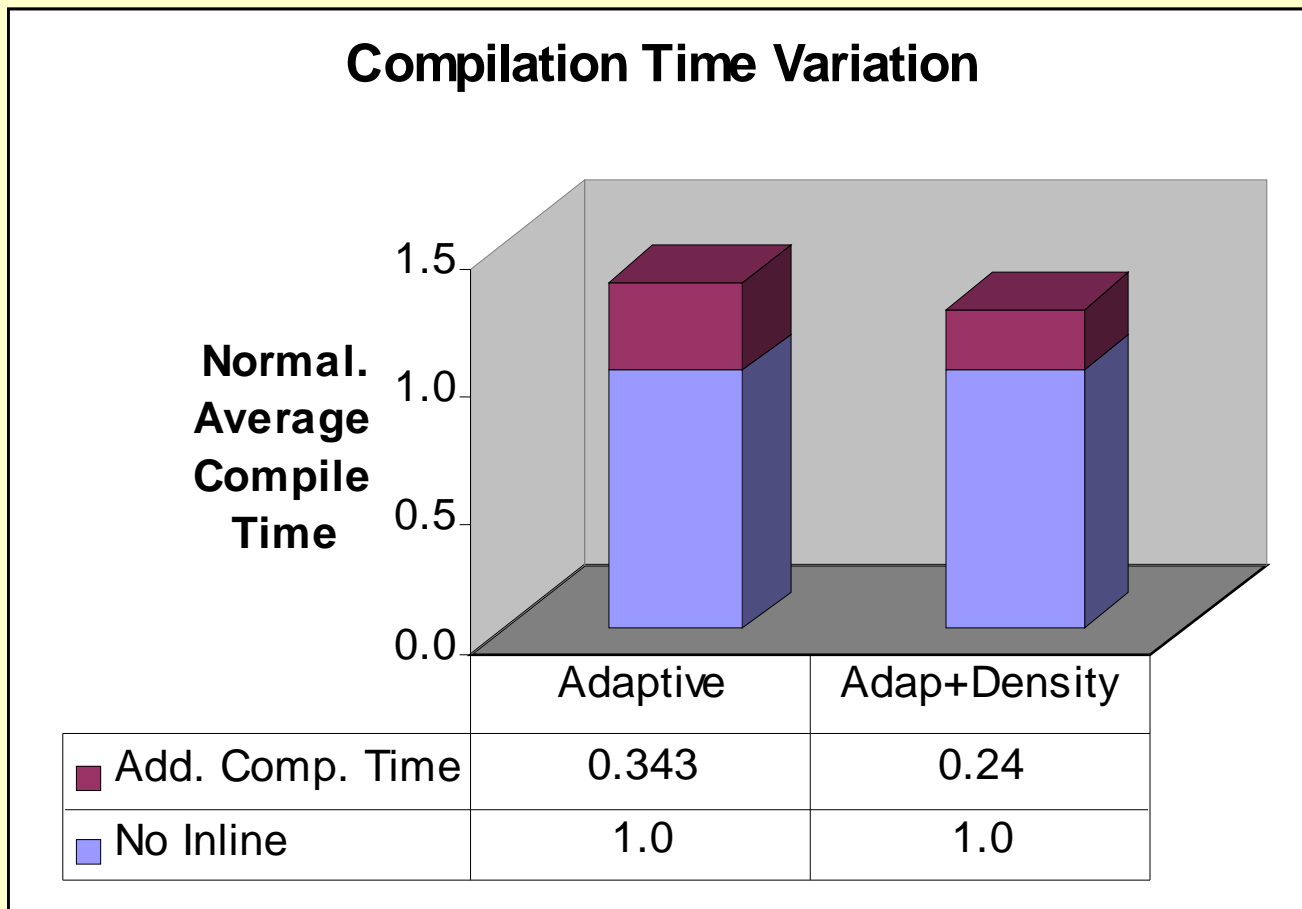
# Inlining Speedup



# Effect of Cycle Density on Code Size



# Effect of Cycle Density on Compilation Time



# Can We Inline Further?

