

Cache-Conscious Structure Definition

T. Chilimbi, University of Wisconsin-Madison

B. Davidson, J. Larus, Microsoft

PLDI '99

Outline

- Motivation
- Contributions
- Class Splitting
 - Algorithm Description
 - Empirical Results
- Field Reordering
 - Algorithm Description
 - Empirical Results
- Conclusions

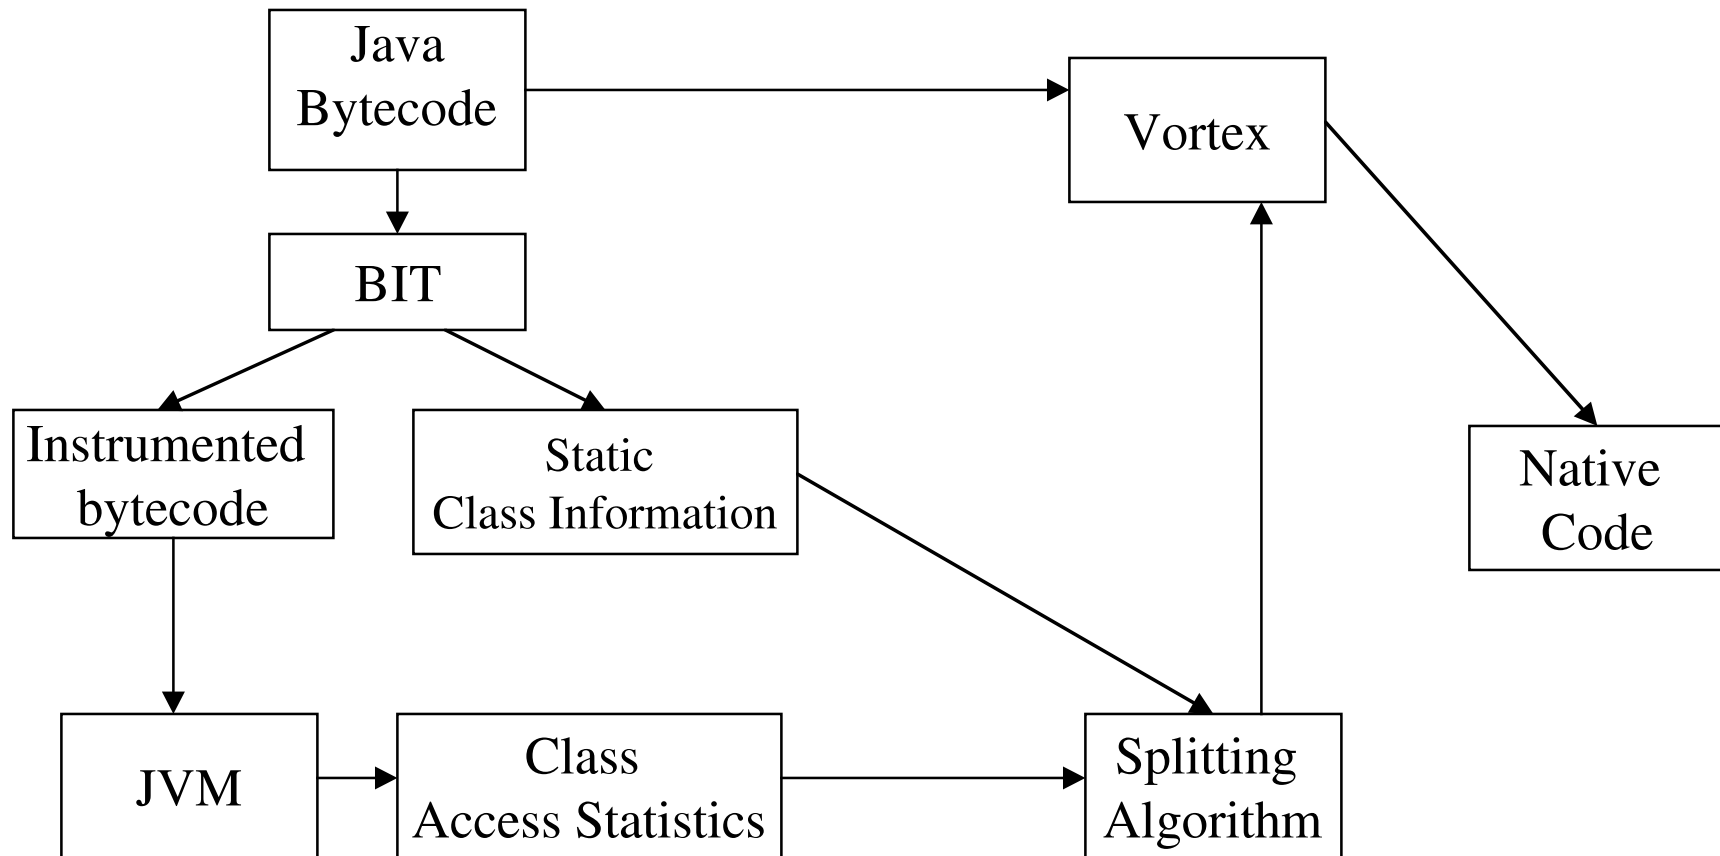
Motivation

- Processor - memory performance gap
- Data reference locality
- Improve performance of layout tools
- Small number of fields frequently accessed

Contribution

- Automatic class splitting
 - Two small objects fit into one cache block
 - Implicit pre-fetch
 - Dynamic co-location is improved
 - Faster execution time
- Field reordering recommendations
 - Time-related fields put in one cache block
 - Better cache utilization, less cache pressure

Class Splitting



October 18, 1999

Algorithm

- Step 1 - filter classes

Frequently accessed, Size > 8 bytes, at least 2 fields

- Step 2 - mark cold fields aggressively

If field access count $< A_i / (2 * F_i)$, mark cold

- Step 3 - do split

If cold portion > 8 bytes

If *Temperature Differential*, then split

Otherwise, remark cold fields

If cold portion > 8 bytes, split.

Temperature Differential

$\max(\text{hot}(\text{class}_i)) - 2 \sum \text{cold}(\text{class}_i) \gg 0$ Why?

- Assume:

$$\max(a_1, a_2 \dots a_n) < \text{cost}(o_1) < \sum (a_1, a_2 \dots a_n)$$

- Benefit from locality:

$$\text{cost}(o_1) + \text{cost}(o_2) <$$

$$(\max(\max(\text{hot}(\text{class}_1), \max(\text{hot}(\text{class}_2))) + e) + 2(\sum \text{cold}(\text{class}_1) + \sum \text{cold}(\text{class}_2)))$$

Temperature Differential

- To benefit from splitting:

$$\begin{aligned} & \max(a_1, a_2 \dots a_n) + \max(b_1, b_2 \dots b_m) > \\ & (\max(\max(\text{hot}(\text{class}_1), \max(\text{hot}(\text{class}_2))) + e) + \\ & 2(\sum \text{cold}(\text{class}_1) + \sum \text{cold}(\text{class}_2))) \end{aligned}$$

- The best they can do?

$$\begin{aligned} & \text{For every } i \max(\text{hot}(\text{class}_i)) - 2 \sum \text{cold}(\text{class}_i) \\ & \gg 0 \end{aligned}$$

Program Transformation

- Add a cold class
 - Contains public cold fields
 - Only has constructor
- Add reference in the hot class
- Transform program
 - Include reference to new class for every cold field access
 - Create cold class instance

Program Transformation

October 18, 1999

Empirical Results

- UltraSPARC, 167MHz, 2Gb, 1Mb L2
- 5 Java programs - 3K - 28K LOC
- Optimizations:
 - Vortex with aggressive optimizations (base)
 - CL object co-location (only)
 - Class Splitting + CL
- Metrics
 - L2 miss rate reduction
 - Execution time

Field Reordering

- bocache recommends better structure field orders in C programs.
- Structures bigger than cache block
- Might be unsafe:

```
struct bar {int x; float *f; int  
    y;} *p;  
int *a;  
a = (int *)p;
```

Bbcache

- Build structure access database (static)
 - Hash table of structures
 - For each structure, list of all instances
 - For each instance, list all accessed fields
 - For each field, list all access sites.
- Process trace (dynamic)
 - Count field accesses
 - Count contemporaneous use < 100 ms

Bbcache (cont.)

- Structure field orders
 - Build per instance affinity graphs
 - Combine into per class affinity graphs
 - Greedy field layout algorithm:
 - Step 1 - Select highest weight edge first
 - Step 2 - Append the field with highest configuration-locality increase
 - Repeat Step 2

Bbcache (cont.)

- Evaluation - qualitative metrics
 - Cache block pressure:

$$\sum(b_1, b_2, \dots, b_n) / n$$

- Cache block utilization:

$$\sum(f_{11}, f_{12}, \dots, f_{nbn}) / \sum(b_1, b_2, \dots, b_n)$$

Empirical Results

- Pentium II Xeon, 4Gb, 1Mb L2
- MS SQL server running TPC-C
- 5 active SQL server structures reordered
- Performance improved by 2-3%

Conclusion

- Class Splitting works because
 - Field access profiles have bimodal division
 - Splitting insensitive to input data
- Benefits on hot objects co-location
- Disadvantages
 - Another level of indirection
 - Increase objects in memory
 - Code bloat - opposite of Jax transformations