

Precomputing Memory Locations for Parametric Allocations

Jörg Herter Sebastian Altmeyer

Department of Computer Science
Saarland University

WCET Workshop, July 2010

What we have ...

- 1 Sound and precise WCET analysis
- 2 Dynamic Memory Allocation
 - ▶ often clearer program structure
 - ▶ easy memory reuse (e.g. in-situ transformations)

... but can we have both together?

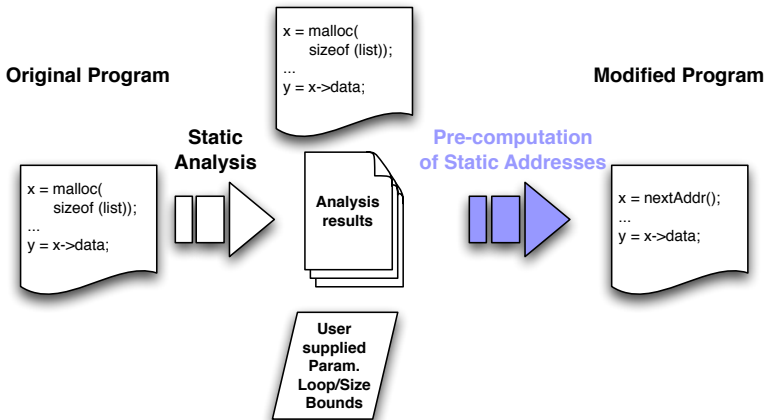
Why is DSA Problematic? An Example ...

Consider a program that ...

- 1 builds-up an internal data structure A
- 2 works on A
- 3 transforms this structure to another data structure B
- 4 works on B

What are the challenges for a WCET analysis?

- 1 builds-up an internal data structure A
→ unknown cache state after allocation
- 2 works on A
→ unknown cache set mappings of list elements
- 3 transforms this structure to another data structure B
→ unknown cache state after (de-)allocation
- 4 works on B
→ unknown cache set mappings of list elements



What are *Good* Memory Addresses?

What do we consider good memory addresses for heap allocated objects?

- Good addresses enable a subsequent WCET analysis to calculate tight WCET bounds,
- exhibit optimal cache performance by
- minimal memory consumption.

What are *Good* Memory Addresses?

Unfortunately, optimizing ...

- for tight WCET bounds requires knowledge about subsequently applied analyses
- for cache performance too costly to compute¹
- for memory consumption is (still) NP hard

¹Petranc and Rawitz: The hardness of cache conscious data placement, 2005

We can use

- simple heuristics² for achieving *good* cache performance
- and only approximate minimum memory consumption.

²Chilimbi et al: putting temporarily accessed objects adjacent in memory

An Example: Intuitive Solution

Assume A and B to be a singly- and a doubly-linked list, respectively.

An intuitive, memory and cache optimal memory placement is:




Assume a target hardware with cache line size 32 bytes and let the algorithm


- consider 4 consecutive elements of A as a single object
- consider 8 consecutive elements of B as a single object

According to our heuristics, all cache benefits are now already exploited!


Computed patterns (chunks) for our list-copy example:



1 repetition



$(p/8 - 1)$ repetitions
 $i \in [2; p/8]$



2 repetitions
 $i \in [p/4 - 1; p/4]$

Putting these chunks consecutively in memory yields ...



This allocation scheme is

- almost memory optimal and
- cache optimal.

How Does Our Algorithm Work?

- start with formal description of a program's allocation behavior
- normalize according to target hardware
- compute conflict free 'patterns'
- suitable allocation scheme is any concatenation of these patterns

- novel algorithm to statically precompute memory addresses
- less limitations than previous approach
- promising results for benchmarks

- Future Work:
 - ▶ full evaluation of the approach
 - ▶ fully automatize whole work chain

An Example: What was gained?

Program	Memory Allocation	WCET Bound (cycles)
In-Situ List-Copy	TLSF	111,698,519
	TLSF*	1,710,656
	Pre-computed	830,202