

# Amortised memory analysis using the depth of data structures

Brian Campbell

`Brian.Campbell@ed.ac.uk`

Laboratory for Foundations of Computer Science  
University of Edinburgh

March 26, 2009

# Overview

Interested in efficient memory bounds inference.

- ▶ Hofmann-Jost [POPL'03] amortised analysis gives **linear heap** space bounds in **total size** of the input.
- ▶ Such bounds are a poor fit for stack space.  
Especially for functional programming, tree structured data.
- ▶ Want *maxima* and bounds in terms of *depth*.

Keep close to efficient Linear Programming inference.

# Principles of Hofmann-Jost analyses

- ▶ The type system *certifies* bounds;
- ▶ annotations describe bounding function in terms of input sizes:

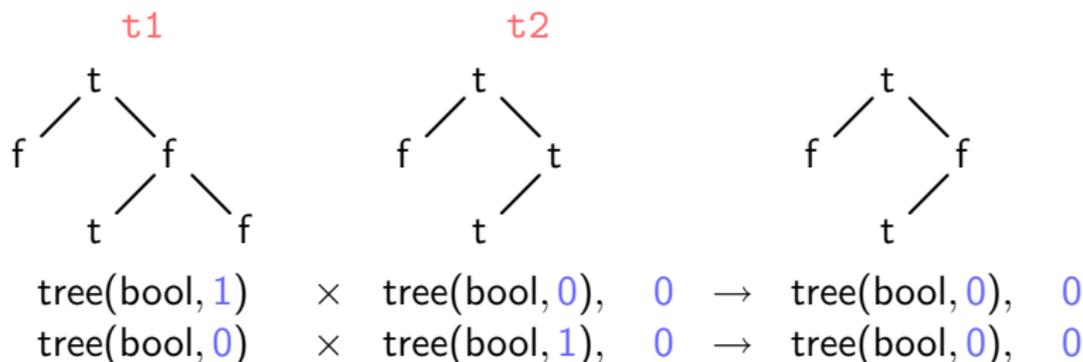
$$\begin{array}{l} x : \text{tree}(\text{bool}, 2), y : \text{tree}(\text{bool}, 3), 5 \vdash e : \text{tree}(\text{bool}, 2), 1 \\ 2 \times |x| \quad + \quad 3 \times |y| \quad + 5 \qquad \qquad \qquad \rightarrow \quad 2 \times |\text{result}| + 1 \end{array}$$

- ▶ Side conditions in rules guarantee the soundness of the bounding functions.

- Inference:
- construct 'plain' typing
  - collect side conditions together
  - solve the resulting LP.

## Heap memory example

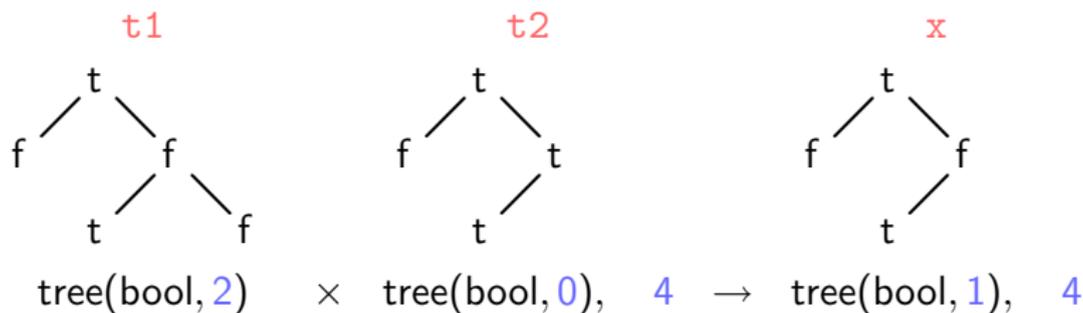
The `andtrees` function computes the pointwise ‘and’ of two boolean trees (up to the largest common subtree):



- ▶ means `andtrees t1 t2` uses no more than  $|t1|$  units of space.
- ▶ The typings (and bounds) are not unique.  $|t2|$  is also sufficient.

## Heap memory example

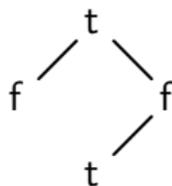
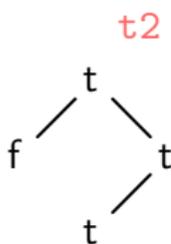
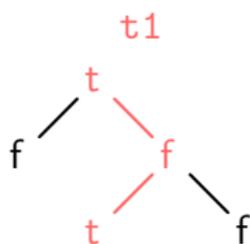
The `andtrees` function computes the pointwise ‘and’ of two boolean trees (up to the largest common subtree):



let  $x = \text{andtrees } t1 \ t2$  in  $e_2$

- ▶ Signatures also ‘translate’ requirements:
- ▶ If  $e_2$  requires  $|x| + 4$  units, then  $2 \times |y| + 4$  is sufficient for both allocation and  $|x| + 4$  later.

## Example with stack space



$\text{tree}(\text{bool}, 1) \times \text{tree}(\text{bool}, 0), 0 \rightarrow \text{tree}(\text{bool}, 1), 0$   
 $\text{tree}(\text{bool}, 0) \times \text{tree}(\text{bool}, 1), 0 \rightarrow \text{tree}(\text{bool}, 1), 0$

- ▶ means and trees  $t_1$   $t_2$  uses at most  $|t_1|$  (or  $|t_2|$ ) units of stack space.
- ▶ Stack space is reusable.
- ▶ But now we want to use the depth to get a better bound (i.e.,  $|t_1|_d$ ).

## Developing an analysis with maximums

Previously we just added all the contributions from the context:

$$l:\text{tree}(\text{bool}, k), r:\text{tree}(\text{bool}, k), v:\text{bool}, n \vdash e : \dots \\ |l| \times k + |r| \times k + 0 + n$$

Now we introduce a second context former to denote 'max' (;):

$$(l:\text{tree}(\text{bool}, k); r:\text{tree}(\text{bool}, k); v:\text{bool}), n \vdash e : \dots \\ \max\{ |l|_d \times k, |r|_d \times k, 0 \} + n$$

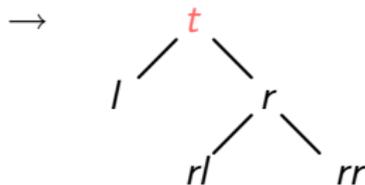
- ▶ Note that contexts are now trees.
- ▶ Treat tree types as 'folded up' version of above context.
- ▶ So  $t:\text{tree}(\text{bool}, k)$  denotes  $|t|_d \times k$ .

Inspired by O'Hearn's Bunched Typing.

## Unfolding trees in the context

$\Gamma(\cdot)$  is a context with a 'hole'.

$$\frac{\Gamma(\cdot) \vdash e_1 : T, k' \quad \Gamma((l:\text{tree}(T, k); r:\text{tree}(T, k); v:T), k) \vdash e_2 : T, k'}{\Gamma(t:\text{tree}(T, k)) \vdash \text{match } t \text{ with } \begin{array}{l} \text{leaf} \rightarrow e_1 \\ \text{node}(l, v, r) \rightarrow e_2 : T, k' \end{array} \quad (\text{MATCH})}$$

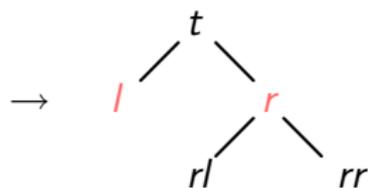


$t:\text{tree}(\text{bool}, k)$

## Unfolding trees in the context

$\Gamma(\cdot)$  is a context with a 'hole'.

$$\frac{\Gamma(\cdot) \vdash e_1 : T, k' \quad \Gamma((l : \text{tree}(T, k); r : \text{tree}(T, k); v : T), k) \vdash e_2 : T, k'}{\Gamma(t : \text{tree}(T, k)) \vdash \text{match } t \text{ with } \begin{array}{l} \text{leaf} \rightarrow e_1 \\ \text{node}(l, v, r) \rightarrow e_2 : T, k' \end{array} \quad (\text{MATCH})}$$

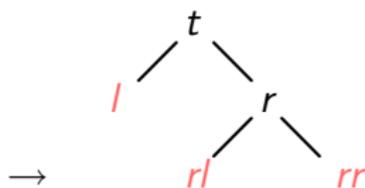


$$t : \text{tree}(\text{bool}, k)$$
$$(l : \text{tree}(\text{bool}, k); r : \text{tree}(\text{bool}, k); v : \text{bool}), k$$

## Unfolding trees in the context

$\Gamma(\cdot)$  is a context with a 'hole'.

$$\frac{\Gamma(\cdot) \vdash e_1 : T, k' \quad \Gamma((l : \text{tree}(T, k); r : \text{tree}(T, k); v : T), k) \vdash e_2 : T, k'}{\Gamma(t : \text{tree}(T, k)) \vdash \text{match } t \text{ with } \begin{array}{l} \text{leaf} \rightarrow e_1 \\ \text{node}(l, v, r) \rightarrow e_2 : T, k' \end{array} \quad (\text{MATCH})}$$



$$t : \text{tree}(\text{bool}, k)$$
$$(l : \text{tree}(\text{bool}, k); r : \text{tree}(\text{bool}, k); v : \text{bool}), k$$
$$\left( l : \text{tree}(\text{bool}, k); \left( (rl : \text{tree}(\text{bool}, k); rr : \text{tree}(\text{bool}, k); rv : \text{bool}), k \right); v : \text{bool} \right), k$$

## let expressions

let  $x = e_1$  in  $e_2$

- ▶ Overall bound is  $\max\{\text{bound for } e_1, \text{bound for } e_2\}$ .
- ▶ But we also need to translate bound for  $e_2$ .

Instead replace subcontext for  $x$  with that needed to produce it,  $\Delta$ :

$$\frac{\Delta \vdash e_1 : T_0, n_0 \quad \Gamma(x : T_0, n_0) \vdash e_2 : T, k'}{\Gamma(\Delta) \vdash \text{let } x = e_1 \text{ in } e_2 : T, k'} \quad (\text{LET})$$

(Only sound for stack discipline.)

## New rules

We need to be able to manipulate contexts to get the right shape.  
Hence new rules such as:

$$\frac{\Gamma(\Delta') \vdash e : T, n' \quad \Delta \cong \Delta'}{\Gamma(\Delta) \vdash e : T, n'} \quad (\equiv)$$

$$\Gamma, (\Delta; \Delta') \cong (\Gamma, \Delta); (\Gamma, \Delta') \quad (\text{distribution})$$

$$\Gamma \cong \Gamma; \Gamma \quad (\text{max-contraction})$$

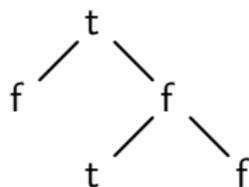
$$\Gamma \cong q\Gamma, (1 - q)\Gamma \quad q \in [0, 1] \quad (\text{plus-contraction})$$

All preserve the bounding functions derived from the context.

Also: weakening and a max-to-plus approximate conversion.

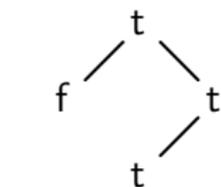
## Example with stack space

Function signatures are also 'structured'.



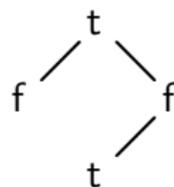
`tree(bool, 1)`

`tree(bool, 0)`



`tree(bool, 0)`

`tree(bool, 1)`



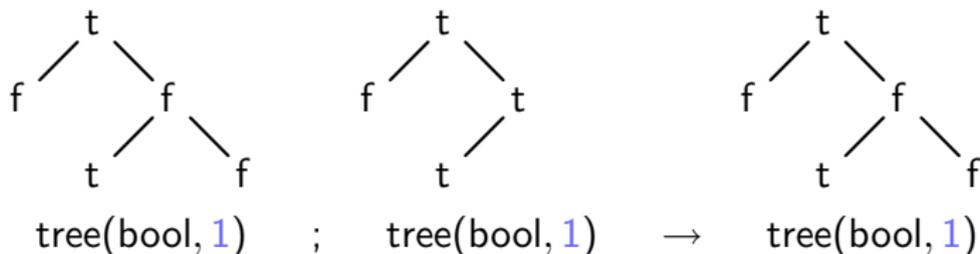
`tree(bool, 1)`

`tree(bool, 1)`

- ▶ means `andtrees t1 t2` uses at most  $|t1|_d$  or  $|t2|_d$  units of stack space.

## Example with stack space

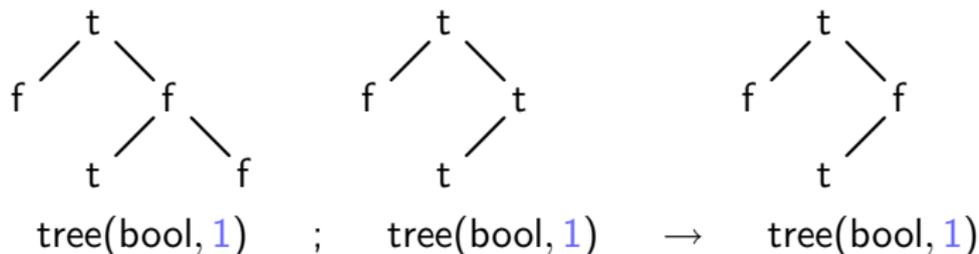
Function signatures are also 'structured'.



- ▶ means `andtreesmax t1 t2` uses at most  $\max\{|t1|_d, |t2|_d\}$  units of stack space.
- ▶ We can now also type a version of `andtrees` which returns false for all the nodes which are only in one of the arguments.

## Example with stack space

Function signatures are also 'structured'.



$$\frac{k \geq \text{stack}(f) \quad k + k'_1 \geq k' \quad \Sigma(f) = \Gamma \rightarrow T, k'_1}{\Gamma[x_1, \dots, x_p / \text{namesof}(\Gamma)], k \vdash f(x_1, \dots, x_p) : T, k'} \text{(FUN)}$$

## Extra benefit from maxima

```
let maybeleft(t,b) =  
  match t with leaf -> leaf | node(l,r,v) ->  
    if b then l else t
```

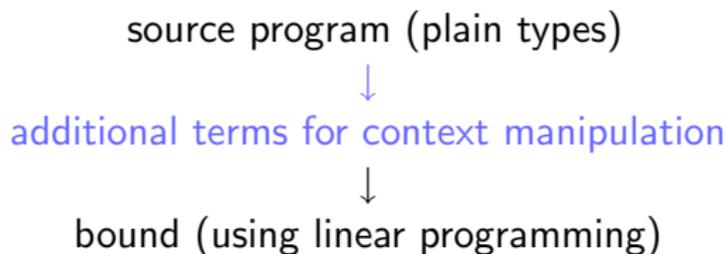
In heap analysis need to sum requirements because of use of contraction at `match`. Doubles the bound unnecessarily.

- ▶ In depth type system we can use max-contraction.
- ▶ So requirement goes  $|t| \Rightarrow \max\{|t|, |t|\} \Rightarrow \max\{|t|, |l|\}$ .
- ▶ Context goes  $t : \text{tree} \Rightarrow t : \text{tree}; t : \text{tree} \Rightarrow t : \text{tree}; l : \text{tree}$

# Stack space inference

- ▶ Would like to take advantage of linear programming again
- ▶ But new context manipulation rules are not syntax-directed

We add an **extra stage** to the inference process:



Assume context structure given for function signatures to make problem more tractable.

## Basic ideas for inference

- ▶ Work from the leaves of the expression outwards.
- ▶ At every stage, keep track of a generated context derived from subexpressions and the typing rule.

$$\frac{\Gamma \vdash e_1 \mapsto \Gamma_1 \quad \Gamma \vdash e_2 \mapsto \Gamma_2}{\Gamma \vdash \text{if } x \text{ then } e_1 \text{ else } e_2 \mapsto \Gamma_1; \Gamma_2; x:\text{bool}}$$

Need to add context manipulation at two points:

1. where binding occurs, to deal with contraction, etc;
2. to make the generated context match the function signature.

Expand context to max-of-sums form, factor out bound variables.

# The full analysis

- ▶ Have algebraic data types, not just trees.
  - ▶ Can specify the form of bounds:
    - in terms of *depth*, *total size*, or a *mixture*.
- Bounds w.r.t. total size useful when depth analysis fails.
- ▶ Resource polymorphism (different function signatures at different points).

Implementation in Standard ML.

## Conclusion and Further work

Type-based amortized analysis can provide good bounds on stack space for programs using tree-structured data.

- ▶ Nested containers don't behave that well due to the definition of depth.  
May be possible to separate structure of containers from their contents.
- ▶ Inferring the structure of function signatures.
- ▶ Deal with construction of log-depth trees (e.g, heap sort).

Potentially the most interesting:

Find sound `LET` rule for heap space, get max-plus bounds for heap.