

Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

Heap Recycling for Lazy Languages

Jurriaan Hage (Joint work with Stefan Holdermans)

Dept. of Information and Computing Sciences, Utrecht University P.O. Box 80.089, 3508 TB Utrecht, The Netherlands E-mail: jur@cs.uu.nl Web pages: http://people.cs.uu.nl/jur/

> May 15, 2008 Fun In The AfterNoon at Hertfordshire Presented earlier at PEPM '08, San Francisco

Announcement

- The Helium Haskell compiler is being made ready for release 1.7.
- New website (already contains the sources of pre-1.7)
- Improved usability and standardization
- Extended logging facilities
- Bugfixes
- But what about those type classes?



Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

< 日 > < 目 > < 目 > < 目 > < 日 > < 日 > < 日 > < 0 < 0

Lazy languages better be pure

Functional languages can be classified along several axes:

pure vs. impure;

strict (eager) vs. nonstrict (lazy).

Provide the sense: reasoning about side-effects in a nonstrict context is hard.



Universiteit Utrecht

Referential transparency

- Pure languages are referential transparent: each term can always be safely replaced by its value.
- Referential transparency enables equational reasoning.
- Referential transparency enables memoization, common subexpression elimination, parallel evaluation strategies, etc.
- Provide the set of the



Universiteit Utrecht

Monads can do the job

Referential transparency requires us to either ban side-effects or deal with them in some special way.

Example: monadic encapsulation of side-effects in Haskell.

main :: IO() $main = \mathbf{do} \ input \leftarrow readFile "in"$ writeFile "out" (reverse input)

1P

Monads come with their own programming style. Reasoning about monadic code can be hard. F



Universiteit Utrecht

Faculty of Science Information and Computing Sciences

< 日 > < 目 > < 目 > < 目 > < 日 > < 日 > < 日 > < 0 < 0

Don't overdo

- Combining the monadic and "ordinary" functional style is okay if side-effects are fundamental to the program.
- If side-effects are only peripheral, a purely functional look and feel is preferred.
- Example: use of an I/O monad makes sense for programs that are indeed about I/O, but not for the occasional debug statement.

 $revSort :: [Int] \rightarrow [Int]$ $revSort = (trace "applying revSort") (reverse \circ sort)$



Universiteit Utrecht

Don't overdo

- Combining the monadic and "ordinary" functional style is okay if side-effects are fundamental to the program.
- If side-effects are only peripheral, a purely functional look and feel is preferred.
- Example: use of an I/O monad makes sense for programs that are indeed about I/O, but not for the occasional debug statement.

 $revSort :: [Int] \rightarrow [Int]$ $revSort = (trace "applying revSort") (reverse \circ sort)$

Similarly, monadic in-place updates make sense for the union-find algorithm, but not for the occasional performance tweak.



Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

<日 > < 同 > < 目 > < 目 > < 目 > < 目 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

Idiomatic list reversal

Idiomatic list reversal needs linear space to run:

 $reverse :: [a] \rightarrow [a]$ $reverse \quad l = rev \ l \ []$ where $rev \ [] \quad acc = acc$ $rev \ (x : xs) \ acc = rev \ xs \ (x : acc)$

rev constructs a new heap cell for every node in the input.

If the input list is used only once, we would like to reuse its cons-nodes and only use constant space.



Universiteit Utrecht

Monadic in-place list reversal

In-place list reversal can be implemented with lazy state threads (Lauchbury and Peyton Jones, PLDI'94):

```
\begin{array}{l} \textbf{type } STList \; s \; a = STRef \; s \; (L \; s \; a) \\ \textbf{data } L \; s \; a = STNil \; | \; STCons \; (STRef \; s \; a) \; (STRef \; s \; (STList \; s \; a)) \\ reverse' \; :: \; STList \; s \; a \; \rightarrow \; ST \; s \; (STList \; s \; a) \\ reverse' \; r \; &= \; \textbf{do} \; acc \; \leftarrow \; newSTRef \; STNil \\ rev \; r \; acc \\ \textbf{where} \\ rev \; r \; acc \; = \; \textbf{do} \; l \; \leftarrow \; readSTRef \; r \\ \textbf{case} \; l \; \textbf{of} \; STNil \; \rightarrow \; return \; acc \\ & \; STCons \; hd \; tl \; \rightarrow \; \textbf{do} \; r' \; \leftarrow \; readSTRef \; tl \\ & \; writeSTRef \; tl \; acc \\ rev \; r' \; r \end{array}
```

A lot of work for a simple performance tweak!

Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

< 日 > < 目 > < 目 > < 目 > < 日 > < 日 > < 日 > < 0 < 0

Idiomatic in-place list reversal

We propose a small language extension:

 $reverse'' :: [a] \rightarrow [a]$ $reverse'' \quad l = rev \ l \ []$ where $rev [] \qquad acc = acc$ $rev \ l@(x:xs) \ acc = rev \ xs \ l@(x:acc)$

- IP We allow the @-construct not only at the left-hand side of a function definition, but also at the right-hand side, where it denotes explicit reuse of a heap node.



Universiteit Utrecht

Faculty of Science Information and Computing Sciences < 日 > < 目 > < 目 > < 目 > < 日 > < 日 > < 日 > < 0 < 0

Challenges

- Q How do we ensure that in-place updates do not compromise referential transparency?
- Q How do we ensure that in-place updates make sense with respect to the underlying memory model?



Universiteit Utrecht

Challenges

- Q How do we ensure that in-place updates do not compromise referential transparency?
- Q How do we ensure that in-place updates make sense with respect to the underlying memory model?
- A We put statically enforced restrictions on the contexts in which updates occur.



Universiteit Utrecht

Referential transparency at stake

In-place filter:

$$\begin{array}{cccc} filter' :: (a \rightarrow Bool) \rightarrow [a] & \rightarrow [a] \\ filter' & p & [] & = [] \\ filter' & p & l@(x:xs) = \mathbf{if} & p x \\ & \mathbf{then} \ l@(x:filter' \ p \ xs) \\ & \mathbf{else} \ filter' \ p \ xs \end{array}$$

Putting odd numbers before even numbers:

let l = [1..10]in filter' odd l + filter' even l

 \mathbb{CP} Yields [1,3,5,7,9]! What happened to [2,4,6,8,10]?



Universiteit Utrecht

Keeping track of single-threadedness

- We only allow in-place updates of values that are passed around single-threadedly.
- Single-threadedness in enforced through type-based uniqueness analysis.
- We annotate typing judgements with uniqueness annotations φ: 1 for single-threaded terms, ω for multi-threaded terms (with 1 ⊑ ω).
- ► For example: l ::¹ [Int^ω] indicates that the list l is passed around single-threadedly, but its elements may be used multi-threadedly.



Universiteit Utrecht

Possible analysis for *filter'* even:

 $filter' \ even \ ::^{\omega} \quad [Int^{\omega}]^1 \ \rightarrow^{\omega} \ [Int^{\omega}]^{\omega}$



Universiteit Utrecht

Possible analysis for *filter'* even:

 $filter' \ even \ ::_{\boxed{1}}^{\omega} \ [Int^{\omega}]^1 \ \rightarrow^{\omega} \ [Int^{\omega}]^{\omega}$

1 The filter may be passed around multi-threadedly.



Universiteit Utrecht



filter' even :: $\frac{\omega}{[1]} [Int_{[2]}^{\omega}]^1 \rightarrow^{\omega} [Int_{[2]}^{\omega}]^{\omega}$

1 The filter may be passed around multi-threadedly.

2 The elements of the argument list may be passed around multi-threadedly.



Possible analysis for *filter'* even:

 $filter' \ even \ ::_{\boxed{1}}^{\omega} \ [Int_{\boxed{2}}^{\omega}]_{\boxed{3}}^{1} \rightarrow^{\omega} \ [Int_{\boxed{2}}^{\omega}]_{\omega}^{\omega}$

- 1 The filter may be passed around multi-threadedly.
- 2 The elements of the argument list may be passed around multi-threadedly.
- 3 The argument list is required to be single-threaded!



Possible analysis for *filter'* even:

filter' even :: $\frac{\omega}{[1]}$ $[Int_{[2]}^{\omega}]_{[3]}^{1} \rightarrow \frac{\omega}{[4]} [Int_{[2]}^{\omega}]^{\omega}$

- 1 The filter may be passed around multi-threadedly.
- 2 The elements of the argument list may be passed around multi-threadedly.
- 3 The argument list is required to be single-threaded!
- The filter is not subjected to any containment restriction (see paper).



Universiteit Utrecht

Possible analysis for *filter'* even:

 $filter' \ even \ ::_{\boxed{1}}^{\omega} \ [Int_{\boxed{2}}^{\omega}]_{\boxed{3}}^{1} \ \rightarrow_{\boxed{4}}^{\omega} \ [Int_{\boxed{5}}^{\omega}]^{\omega}$

- 1 The filter may be passed around multi-threadedly.
- 2 The elements of the argument list may be passed around multi-threadedly.
- 3 The argument list is required to be single-threaded!
- The filter is not subjected to any containment restriction (see paper).
- 5 The elements of the result list may be passed around multi-threadedly.



Universiteit Utrecht

Possible analysis for *filter'* even:

filter' even :: $\overset{\omega}{[1]}$ $[Int^{\omega}_{[2]}]^{1}_{[3]} \rightarrow^{\omega}_{[4]} [Int^{\omega}_{[5]}]^{\omega}_{[6]}$

- 1 The filter may be passed around multi-threadedly.
- 2 The elements of the argument list may be passed around multi-threadedly.
- 3 The argument list is required to be single-threaded!
- The filter is not subjected to any containment restriction (see paper).
- 5 The elements of the result list may be passed around multi-threadedly.
- 6 The result list may be passed around multi-threadedly.



Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

(日)

Judgements for uniqueness analysis

The typing rules for uniqueness analysis are of the form $\Gamma \vdash t :: \varphi \sigma$, where σ can contain annotations.

$$\frac{\Gamma = \Gamma_1 \bowtie \Gamma_2 \qquad \Gamma_1 \vdash t_1 \ ::^{\varphi_1} \ \tau_2^{\varphi_2} \rightarrow^{\varphi_0} \ \tau^{\varphi}}{\Gamma \Vdash \varphi_1 \sqsubseteq \varphi_0 \qquad \Gamma_2 \vdash t_2 \ ::^{\varphi_2} \ \tau_2}$$
$$\frac{\Gamma \vdash t_1 \ t_2 \ ::^{\varphi} \ \tau}{\Gamma \vdash t_1 \ t_2 \ ::^{\varphi} \ \tau}$$

The auxiliary judgement $\Gamma = \Gamma_1 \bowtie \Gamma_2$ ensures that single-threaded variables are not passed down to multiple subterms.

 $\mathfrak{P} \Gamma \Vdash \varphi_1 \sqsubseteq \varphi_0$ enforces a containment restriction.

The analysis allows for both type polymorphism and uniqueness polymorphism (cf. Hage et al., ICFP 2007).



Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

(日)

Fitting the memory model

- Often, a language specification does not prescribe a particular memory model: so, we only allow updates that are likely to be implementable in all implementations of lazy languages.
- For example: replacing a nil-cell by a cons-cell will in most cases be problematic and should therefore be prohibited.
- The scheme we adopt only allows updates with values built by the same constructor.
- To keep track the constructors values are built by, we store them in the typing context Γ in bindings of the form x ::^{φ|ψ} σ, where ψ is either a constructor C or ε.



Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

Rule for updates

Both aspects (referential transparency and the memory model) show up in the typing rule for in-place updates:

$$\Gamma = \Gamma_1 \Join \Gamma_2$$

$$\Gamma_1(x) = {}^{1|C} \sigma_0 \qquad \Gamma_2 \vdash C \ t_1 \dots t_n :: {}^{\varphi} \sigma$$

$$\Gamma \vdash x @ (C \ t_1 \dots t_n) :: {}^{\varphi} \sigma$$

x = x is required to be passed around single-threadedly. x = x is required to be built by C.



Universiteit Utrecht

Properties

Using an instrumented natural semantics, with judgements of the form

 $H;\eta;t \Downarrow_n H';\eta';w$

(with H a heap, η a mapping from variables to heap locations, w a weak-head normal form, and n the number of heap cells allocated),

we can demonstrate a subject-reduction result.

Furthermore, we can show that adding well-behaved updates to a program preserves the meaning of the original program and the new program requires at most the same amount of space.



Universiteit Utrecht

[Faculty of Science Information and Computing Sciences]

(日)

Assessment

- Should update instructions be inferred?
- Do we need two versions of *reverse*? Do we need two versions of *filter*? What about *zip*?
- Do we expose annotated types to the programmer?
- How does our system relate to Clean?

