

# Type-Based Resource Analysis on Haskell

Franz Sigmüller

Ludwig Maximilian University  
Munich, Germany  
frsiglmu@web.de

We propose an amortized analysis that approximates the resource usage of a Haskell expression. Using the plugin API of GHC, we convert the Haskell code into a simplified representation called GHC Core. We then apply a type-based system which derives linear upper bounds on the resource usage. This setup allows us to analyze actual Haskell code, whereas previous implementations of similar analyses do not support any commonly used lazy functional programming languages.

## 1 Introduction

It can be difficult to estimate the runtime costs of a program written in a lazy programming language. For example, consider the following two versions of a “repeat” function, which returns an infinite list with a given element:

```
repeat x = let xs = x : xs in xs
it = repeat 1 :: [Int]
repeat' x = x : repeat' x
it' = repeat' 1 :: [Int]
```

To a novice, it may not be obvious that the memory usage of this function is constant for the first version and linear for the second version. To assist in detecting these differences, we want to automatically generate annotated types such as the following:

$$\begin{array}{l} \frac{4}{0} \text{ it} : \mu X. \{ [] : (0, []) \mid (:): (0, [\mathbb{T}^0(\text{Int}), \mathbb{T}^0(X)]) \} \\ \frac{0}{0} \text{ it}' : \mu X. \{ [] : (0, []) \mid (:): (0, [\mathbb{T}^0(\text{Int}), \mathbb{T}^2(X)]) \} \end{array}$$

In these types, the second field of the  $(:)$  constructor is wrapped in a thunk type that is annotated with costs 0 and 2, respectively. These values represent additional costs – in this case, memory allocations – that arise from accessing subsequent nodes of the linked list. As this cost is 0 in the type of the first version, this means that no additional allocations will take place and the list therefore requires constant space. For the second version, the value is greater than 0, indicating linear costs.

In this paper, we describe our implementation of a tool that can derive these annotated types automatically. Given a Haskell module, our tool will use the GHC compiler to parse the code and translate it into an intermediate representation which is more easy to analyze. We then apply our own custom type-based system on each expression to derive an annotated typing which depicts an upper bound on the number of resources required for evaluation. Depending on the cost model used, this can be used to approximate memory usage or runtime duration, for example.

This paper is a summary of the master’s thesis “Implementation of an Automated Amortized Analysis on GHC Core as Compiler Plugin”. [12] Note that sections of this paper may be similar or identical to that thesis.

## 2 Related Work

Our work is rooted in Steffen Jost’s PhD thesis “Automated Amortized Analysis” [8]. In his work, Jost introduced a custom language Arthur which can be analyzed statically using a custom type system. For any given Arthur expression, this system can derive a type annotated with numeric variables, as well as a linear program over these variables. After feeding this linear program to an LP solver, an annotated type can be generated which represents linear upper bounds on the memory usage.

A variation of this system was later integrated into the OCaml compiler by Jan Hoffmann et al. [5, 6] This allows them to analyze real OCaml code by using the existing compilation pipeline to parse the original code, and then applying their own analysis on the intermediate code generated by the compiler.

However, these previous systems are designed around strict programming languages and cannot be applied to lazy languages such as Haskell in a meaningful way. While delayed evaluation is supported to some extent – for example, using function abstraction and application –, this needs to be encoded explicitly within the analyzed expression. In lazy languages such as Haskell, however, it is implicit that we need to treat every sub-expression as lazy. As an example of this distinction, consider the `repeat`’ function provided in the previous section: This definition is only meaningful in a lazy language, where only the required prefix of the infinite list will be evaluated; While a strict interpretation of the same implementation will attempt to evaluate the full infinite list at once, triggering an infinite loop.

A new type-based system designed around lazy evaluation was proposed in “Type-based Cost Analysis for Lazy Functional Languages” by Steffen Jost, Pedro Vasconcelos, Mário Florido and Kevin Hammond. [9] This is achieved by introducing “think types” which are used for sub-expressions that may be evaluated lazily. Think types are annotated with the costs required for evaluating the corresponding sub-expression; This cost has to be paid only when the expression is actually evaluated, and can be discarded otherwise. This system has later been examined and extended as part of several bachelors’ theses. [7, 10, 11] Most notably, these extended the language and the type system to introduce support for mutual recursion.

The major drawback of this system is the programming language which is analyzed. Just like Arthur before, Jost et al. introduced their own (unnamed) custom programming language which we have dubbed “JV FH”. This language is unwieldy to use, as the primary goal in designing this language was being convenient to analyze rather than to write programs with. For example, in order to analyze the `repeat`’ function given before, we first have to translate it into JV FH as follows:

```
let repeat' = \x -> let xs = repeat' x
                    in letcons xt = Cons (x,xs)
                       in xt
in let one = 1
   in let it' = repeat' one
      in it'
```

While it is possible to use the JV FH analysis for Haskell in this way, performing this translation step manually is time-consuming and error-prone. In our work, we will circumvent this issue by adapting the existing JV FH type system to one of the intermediate languages used within the GHC compiler. This means that the translation can be performed automatically by using the existing GHC compilation pipeline.

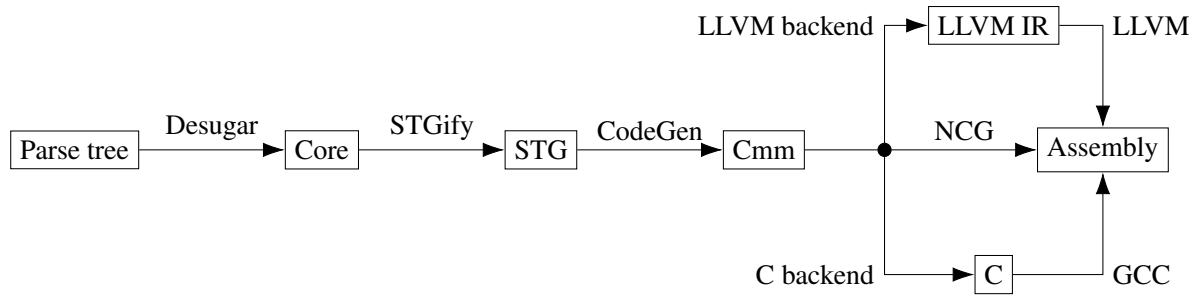


Figure 1: Basic overview of the GHC compilation pipeline. Based on a graphic in [4].

### 3 GHC Architecture

In this section, we provide a basic summary of the compilation pipeline used within GHC, depicted in Figure 1. The original Haskell code is first parsed into an abstract syntax tree, which is then type-checked. It is then desugared and translated to GHC Core, a lazy functional language based on System F. At this point, the compiler will apply most of its code optimizations. Afterwards, the code is translated to STG, a subset of GHC Core minus certain type information which is not needed anymore after this point. It is then translated into a low-level imperative language called Cmm, before one of several different backends can be used to translate it into Assembler and, finally, the finished binary.

For our analysis, we will be looking at GHC Core, as it is the most similar to the JVFH language. Additionally, the other intermediate languages used by GHC Core have a number of drawbacks that make them unsuitable for our analysis:

The syntax of the parse tree is a considerably more redundant than that of GHC Core; In fact, the type used in GHC to represent a Core expression consists of only 10 different constructors [1], while the Haskell syntax is represented by around 50 different constructors [2]. While it is possible to implement a type system on the parse tree, it is tedious and does not provide any advantage for our purpose. Additionally, most of the compiler optimizations are applied only after the translation to GHC Core. It is preferable to analyze this optimized code, as it is more representative of the final compiled binary.

The STG code is unsuitable for our purpose, because it lacks certain type information such as type abstraction and application. As GHC Core is an extension to System F and STG is a subset of Core missing type information, we assume that the undecidability of type inference in System F [15] also applies to STG; This is not an issue for the compiler, as type checking has already completed at this point and the compiler does not need the discarded type information anymore. However, as our analysis is based on a type system, this information is essential to us in fully supporting the language, particularly polymorphism. While our current implementation does not support polymorphism yet and therefore may also work on STG, choosing STG over Core would severely hamper any future work on supporting the full language.

Finally, Cmm and any of the languages used thereafter are strict imperative languages. Thus, they are unsuitable for an adaption of the JVFH type system, which is designed around a lazy functional language. Analyzing these languages may be feasible, but this would require a completely different approach.

The GHC compiler has an API for custom plugins that can be used to interact with the compilation pipeline. Plugins can be implemented by writing a Haskell module that exports variables with specific names and specific types, and they can be used by calling GHC with specific command line arguments. [3, 14]

$$\begin{aligned}
e ::= & x \mid c \mid l \mid \tau \mid \lambda x. e \mid \Lambda \tau. e \mid e_1 e_2 \\
& \mid \text{let } x = e_1 \text{ in } e_2 \mid \text{letrec } x_1 = e_1; \dots; x_n = e_n \text{ in } e \\
& \mid \text{case } x = e_0 \text{ of } \{\text{default} \rightarrow e_d \mid c_1(\mathbf{x}_1) \rightarrow e_1 \mid \dots \mid c_n(\mathbf{x}_n) \rightarrow e_n\} \\
& \mid \text{case } x = e_0 \text{ of } \{\text{default} \rightarrow e_d \mid l_1 \rightarrow e_1 \mid \dots \mid l_n \rightarrow e_n\}
\end{aligned}$$

Figure 2: Simplified representation of the GHC Core syntax

$$A ::= X \mid A \xrightarrow{p} B \mid \top^p(A) \mid \mu X. \{c_1 : (q_1, \mathbf{A}_1) \mid \dots \mid c_n : (q_n, \mathbf{A}_n)\}$$

Figure 3: Type syntax, as originally defined in [9]

When we started our implementation, three different types of plugins were supported: Core plugins can modify the optimization pipeline applied to the Core intermediate code; Frontend plugins can take full control over the entire compilation pipeline; And typechecking plugins can enhance the capabilities of the constraint solver used during type checking. The latter has proven unsuitable for our purpose, as we need to establish our own type system, not enhance the existing one. We have implemented both a Core and a frontend plugin which serve to obtain a copy of the Core intermediate code and pass it to our analysis. This future-proofs our tool in case one of the APIs is removed in the future.

## 4 Language & Type Syntax

In the next section, we will discuss how we had to adapt the existing JVFH type rules for GHC Core. For this purpose, we will be using a streamlined representation of the Core syntax as depicted in Figure 2. This representation is slightly different from the one actually used within the compiler [1]. For example, we distinguish more clearly between variables and constructors, both of which are represented as `Var` within the GHC compiler. On the other hand, we also omit several expressions which are currently unsupported in our analysis, namely `Cast`, `Coercion` and `Tick`.

For our types, we are reusing the type syntax from the original JVFH system without modifications. This syntax, as depicted in Figure 3, covers type variables, functions, thunks and algebraic types. Most of these contain annotations that signify the resource usage of the respective expression. Depending on the cost model used, this can refer to memory usage or runtime duration, for example.

The type  $A \xrightarrow{q} B$  represents a function, that takes an argument of type  $A$ , and returns a result of type  $B$ . The type annotation  $q$  signifies the cost of executing this function, i.e., evaluating the function result to weak head normal form.

The algebraic type  $\mu X. \{c_1 : (q_1, \mathbf{A}_1) \mid \dots \mid c_n : (q_n, \mathbf{A}_n)\}$  is a possibly recursive type consisting of several constructors  $c_i$ . For each constructor, we specify a potential  $q_i$ , which are resources “reserved” during structure allocation for future use. During a pattern match, this potential can be redeemed to pay for any upcoming calculations. Additionally, each constructor is also associated with a list of constructor fields  $\mathbf{A}_i$ . The type variable  $X$  can be used within these constructor fields as a recursive reference to the algebraic type. As an example, consider the following representation of a list type:

$$\mu X. \{\text{Nil} : (1, []) \mid \text{Cons} : (2, [\top^3(\text{Int}), \top^4(X)])\}$$

This example type consists of two constructors: The first, `Nil`, has potential 1 and no fields; and the second, `Cons`, has potential 2 and two fields, namely a list element of type `Int`, and a reference to the next list node, represented recursively with a type variable  $X$ . Note that type variables in our system may only be used as recursive references, and therefore must always be contained within algebraic types that bind them.

Finally, the thunk type  $T^q(A)$  is a wrapper that represents lazy evaluation; It signifies that an expression of this type possibly has not been evaluated to weak head normal form yet, and doing so will use up  $q$  resources. Thunk types can only (and will always) appear in specific places: Namely, as the argument of any function type, as the constructor fields of any algebraic type, and as the type of every variable in the context of any typing judgment.

We apply some workarounds for Haskell/Core types that do not have any direct equivalent in this syntax: Primitive types such as “`Int#`” or “`Float#`” are represented as an empty algebraic type “ $\mu X.\{\}$ ” in our system; and type abstractions “`forall x. T`” are replaced with an “artificial” function “ $\mu X.\{\} \xrightarrow{0} T$ ”. However, note that polymorphism is currently unsupported in our implementation; Therefore, this artificial function does not actually have any meaningful effect within the type system, and only serves as an informatory annotation for the user. Also note that this representation breaks the requirement that all function arguments are wrapped in thunk types; This can be used to easily distinguish “real” functions from an artificial function representing type abstraction.

Applying these workarounds instead of extending the type syntax may seem restricting, but it offers one major advantage: The JVFH type system defines several relations between different types, which we can reuse without any modifications. This would not be possible if we modified the type syntax or its semantics.

## 5 Type Rules

In this section, we will give an overview over the adapted type rules used in our implementation. However, note that we abstain from reciting a large portion of the rules and definitions which we reuse from the original JVFH system without modification. Instead, please refer to the original publication by Jost et al. [9]

The following types of judgments are used in both the JVFH and our type system:

- The typing judgment  $\Gamma \vdash_{\frac{p}{q}} e : T$  states that, given context  $\Gamma$ , expression  $e$  has type  $T$ . Additionally,  $p$  is an upper bound on the number of resources required to evaluate  $e$  to weak head normal form; and  $q$  is a lower bound on the number of resources that are available after this evaluation.

Among the type rules that can be applied to these judgments, we distinguish between syntax driven rules and structural rules. The latter serve various difference purposes; For example, the `PREPAY` rule can be used to pay the thunk cost of any variable only once, even if the variable is used multiple times in the expression; And the `RELAX` rule allows us to increase the over-approximation of the cost upper bounds in a typing judgment on purpose, thereby relaxing linear constraints which otherwise might render the linear program unsolvable. While the purpose of the various structural type rules vary widely, they have in common that they can be applied to any typing judgment, regardless of the expression. This means that we can reuse these rules in our system without any modifications.

The syntax driven rules can only be applied to typing judgments with specific expressions. Many of these rules contain cost constants, which can be set by the user; The basic goal of the type system

is to count how often certain kinds of expressions are executed, and to add up the cost constants associated with these expressions. For example, if we set the cost constants associated with any expressions involving variable bindings to 1, and the remaining constants to 0, our analysis will derive an upper bound on the number of memory allocations required to evaluate an expression. However, as these type rules directly depend on the available expressions of the language and their syntax, the original JVFH rules cannot be applied to GHC Core without modifications. We will elaborate on how we adapted these rules for Core in the second half of this section.

- The type equality judgment  $T_1 = T_2$  is used to unify two types  $T_1$  and  $T_2$ .
- The Sharing judgment  $T \checkmark \{T_1, T_2\}$  enforces that all types have the same structure, but may have different type annotations within certain limitations. Most importantly, we enforce that the potential in the left hand side is greater or equal the sum of the potentials in the right hand side; In other words, the potential of  $T$  is “shared” between  $T_1$  and  $T_2$ .

This property is used to ensure that potential can be redeemed only once. For example, consider a typing judgment  $\Gamma, x:X \vdash_T^p e : T$ , where the variable  $x$  appears in expression  $e$  multiple times. In a non-affine system, this would mean that any potential in  $X$  could be redeemed once for each occurrence. But in our system, we instead use Sharing to split up the potential of  $X$  between several subtypes  $X_i$  which are then assigned to the various instances of  $x$  in  $e$ .

Note that the Sharing relation is defined for any number of types on the right hand side; In fact, the subtype relation  $T_1 <: T_2$  is defined as a shorthand for  $T_1 \checkmark \{T_2\}$ , a Sharing relation with only one type on the right hand side.

As the Sharing judgment is a relation between different types and does not depend on the language syntax at all, we reuse the inference rules provided for JVFH [9] without any modifications.

- The judgment  $T_1 \triangleleft T_2$  is another relation between different types and therefore can also be reused in our implementation without modifications. Just like Sharing, this judgment forces both  $T_1$  and  $T_2$  to have the same structure while giving some leeway to the type annotations. In particular, if  $T_1$  and  $T_2$  are some algebraic type  $\mu X. \{\dots\}$ , then all recursive references  $\top^q(X)$  within  $T_1$  may have reduced thunk costs – even zero – compared to their counterparts in  $T_2$ . This is used as an optimization when analyzing recursive variable definitions, where we can assume that these thunks have already been evaluated. However, as we will discuss shortly, there were concerns about the correctness of this assumption and the operator is currently not used in our implementation.
- The final kind of judgment in our system is linear constraints such as  $v_1 \cdot c_1 + \dots + v_n \cdot c_n \leq c_0$ , where  $v_i$  are numerical variables and  $c_i$  are constants. We collect every linear constraint that comes up during the analysis into a single linear program, which is then passed to the LP solver.

As illustrated in this list, only the syntax driven type rules actually require modifications for GHC Core. In the remainder of this section, we will list our adapted type rules and explain why and how they were modified. For the original type rules, please refer to the original paper on JVFH [9] and the bachelor’s thesis that introduced the type rule for mutually recursive definitions [10].

Even among the syntax driven type rules, several rules could be reused with minor or no modifications, due to overlaps between the JVFH and GHC Core languages. As depicted in Figure 4, these encompass variables, abstraction, applying functions to variables, and recursive variable bindings. Functionally, these four rules are identical to their respective counterparts in the JVFH systems. However, the rules were renamed from VAR, ABS, APP and LETAND $\dagger$  to VAR<sub>GC</sub>, ABS<sub>GC</sub>, APPVAR<sub>GC</sub> and LETREC<sub>GC</sub>; And the latter has been adjusted for the minor syntactical differences between the JVFH letand and the Core letrec expressions.

$$\begin{array}{c}
\frac{}{x:\mathbb{T}^q(A) \vdash_{\frac{q+\text{Kvar}}{0}} x:A} \quad (\text{VARGC}) \\
\\
\frac{\Gamma, x:A \vdash_{\frac{q}{0}} e:C \quad x \notin \Gamma \quad \Gamma \checkmark \{\Gamma, \Gamma\}}{\Gamma \vdash_{\frac{0}{0}} \lambda x. e : A \xrightarrow{-q} C} \quad (\text{ABSGC}) \\
\\
\frac{\Gamma \vdash_{\frac{p}{p}} e:A \xrightarrow{-q} C}{\Gamma, x:A \vdash_{\frac{p+q+\text{Kapp}}{p}} ex:C} \quad (\text{APPVARGC}) \\
\\
\frac{\begin{array}{l} \Gamma_1, x_1:\mathbb{T}^0(A''_1), x_2:\mathbb{T}^{q'_2}(A''_2), \dots, x_n:\mathbb{T}^{q'_n}(A''_n) \vdash_{\frac{q_1}{0}} e_1:A_1 \\ \vdots \\ \Gamma_n, x_1:\mathbb{T}^{q'_1}(A''_1), \dots, x_{n-1}:\mathbb{T}^{q'_{n-1}}(A''_{n-1}), x_n:\mathbb{T}^0(A''_n) \vdash_{\frac{q_n}{0}} e_n:A_n \\ A_1 \checkmark \{A_1, A'_1\} \quad A'_1 \triangleleft A'_1 \quad \dots \quad A_n \checkmark \{A_n, A'_n\} \quad A''_n \triangleleft A'_n \\ q_1 \leq q'_1 + q''_1 \quad \dots \quad q_n \leq q'_n + q''_n \quad p \geq p_0 + q'_1 + \dots + q''_n \\ \Delta, x_1:\mathbb{T}^{q'_1}(A_1), \dots, x_n:\mathbb{T}^{q'_n}(A_n) \vdash_{\frac{p_0}{p}} e:B \quad x_1, \dots, x_n \notin \Gamma_1, \dots, \Gamma_n, \Delta \end{array}}{\Gamma_1, \dots, \Gamma_n, \Delta \vdash_{\frac{p+n \cdot \text{Kletrec}}{p}} \text{letrec } x_1 = e_1 ; \dots ; x_n = e_n \text{ in } e : B} \quad (\text{LETREC}_{\text{GC}})
\end{array}$$

Figure 4: Unmodified syntax driven type rules

However,  $\text{LETREC}_{\text{GC}}$  is also noteworthy as it is the only inference rule in our adapted type system that makes use of the  $\triangleleft$  operator. There exists a soundness proof of this rule [10], indicating that the usage of this operator should be correct; But during our tests, we encountered a counter-example where this rule does allow for under-approximation:

```

letrec zipWith = ... in let one = ... in let zero = ... in
letrec fibs = Cons zero fibs1 ;
      fibs1 = Cons one fibs2 ;
      fibs2 = zipWith (+)Int fibs fibs1
in fibs

```

For this (abridged) expression, which evaluates to an infinite list of Fibonacci numbers, our analysis was able to derive a list type with constant costs. This is due to the  $\triangleleft$  operator in the  $\text{LETREC}_{\text{GC}}$  type rule, which allows us to treat the reference to  $\text{fibs}_2$  in the definition of  $\text{fibs}_1$  as if it had constant cost. However, this is clearly incorrect –  $\text{fibs}_2$  is an infinite list which can never be fully evaluated, so the cost of the list has to be linear at least. To work around this issue, we have disabled the  $\triangleleft$  operator for now and simply use type unification in its place. With this modification, our implementation will fail to analyze the code above, as the generated linear program is now unsolvable. However, conceding failure is preferable to returning incorrect results.

In Figure 5, we list several rules which required larger modifications from the JVFH type system:

The application rule  $\text{APP}_{\text{GC}}$  is a generalization of the original APP rule. In JVFH, the arguments of function applications always have to be variables, but in GHC Core, this is not necessarily the case. Therefore, we adapted this rule to allow for any expression in the argument. However, if the argument actually is a variable, using this generalized rule can artificially increase the thunk cost of the variable. In those cases, we instead use the previously mentioned  $\text{APPVARGC}$  rule, which is an unmodified copy of the original APP rule, to prevent this undesired side-effect.

$\text{CONSGC}$  can be considered a replacement of the original LETCONS and CONS rules. JVFH handles

$$\begin{array}{c}
\frac{\Gamma \vdash_{p'}^p e_1 : \mathbb{T}^r(A) \xrightarrow{q} C \quad \Delta \vdash_0^r e_2 : A}{\Gamma, \Delta \vdash_{p'}^{\frac{p+q+\text{Kapp}}{p}} e_1 e_2 : C} \quad (\text{APP}_{\text{GC}}) \\
\\
\frac{C = \mu X. \{ \dots \mid c : (q, \mathbf{A}) \mid \dots \} \quad \mathbf{A}[X \mapsto C] = \mathbf{B}}{\frac{\vdash_{\frac{q+\text{Kcons}}{0}}^{q+\text{Kcons}} c : \mu T_1. \{ \} \xrightarrow{0} \dots \xrightarrow{0} \mu T_{|\mathbf{T}|}. \{ \} \xrightarrow{0} B_1 \xrightarrow{0} \dots \xrightarrow{0} B_{|\mathbf{B}|} \xrightarrow{0} C}{\Gamma \vdash_0^q e_1 : A \quad \Delta, x : \mathbb{T}^q(A) \vdash_{p'}^p e_2 : C \quad x \notin \Gamma, \Delta}}{\Gamma, \Delta \vdash_{p'}^{\frac{p+\text{Klet}}{p}} \text{let } x = e_1 \text{ in } e_2 : C} \quad (\text{LET}_{\text{GC}})} \\
\\
\frac{}{\vdash_0^0 l : \mu X. \{ \}} \quad (\text{LIT}_{\text{GC}})
\end{array}$$

Figure 5: Modified type rules

$$\begin{array}{c}
M \subset \{c_i(\mathbf{x}_i) \rightarrow e_i\}_{i=1}^n \cup \{\text{default} \rightarrow e_d\} \\
|\mathbf{A}_i| = |\mathbf{x}_i| \quad B = \mu X. \{c_i : (q_i, \mathbf{A}_i)\}_{i=1}^n \quad \Gamma \vdash_{p'}^p e_0 : B \\
\Delta, \mathbf{x}_i : \mathbf{A}_i[X \mapsto B], y : \mathbb{T}^0(B) \vdash_{p''}^{\frac{p'+q_i}{p''}} e_i : C \quad \text{for all } i, \text{ if } (c_i(\mathbf{x}_i) \rightarrow e_i) \in M \\
\Delta, y : \mathbb{T}^0(B) \vdash_{p''}^{\frac{p'+q_i}{p''}} e_d : C \quad \text{for all } i, \text{ if } (c_i(\mathbf{x}_i) \rightarrow e_i) \notin M \text{ and } \{\text{default} \rightarrow e_d\} \in M \\
\frac{}{\Gamma, \Delta \vdash_{p''}^{\frac{p+\text{Kmatch}}{p''}} \text{case } y = e_0 \text{ of } M : C} \quad (\text{CASEALG}_{\text{GC}}) \\
\\
M \subset \{l_i \rightarrow e_i\}_i \cup \{\text{default} \rightarrow e_d\} \\
\Gamma \vdash_{p'}^p e_0 : B \quad \Delta, y : \mathbb{T}^0(B) \vdash_{p''}^p e_i : C \quad \text{for all } (l_i \rightarrow e_i) \in M \\
\frac{}{\Gamma, \Delta \vdash_{p''}^{\frac{p+\text{Kmatch}}{p''}} \text{case } y = e_0 \text{ of } M : C} \quad (\text{CASELIT}_{\text{GC}})
\end{array}$$

Figure 6: Adapted type rules for case expressions

structure allocation via a letcons expression, which fulfills multiple purposes at once: First, it fully applies the constructor to all of its arguments; This is simple, because all arguments have to be variables, so this can be done by looking up their types in the context. In GHC Core, the arguments can be more complex, so we defer handling the arguments to the  $\text{APP}_{\text{GC}}$  and  $\text{APPVAR}_{\text{GC}}$  rules instead. Second, the resulting structure is immediately bound to a variable in JVFH. However, in GHC Core, a structure can sometimes be used immediately without binding it to a variable first. This is generally the case when the constructor does not have any fields, such as the “empty list” constructor of a list type. Because of these differences, we chose to process constructors individually, instead of aggregating them with surrounding application and binding expressions.

The  $\text{LET}_{\text{GC}}$  rule is a simplified version of the original  $\text{LET}_{\dagger}$  rule. Unlike JVFH, the let expression in GHC Core does not support recursive definitions – the letrec expression has to be used instead. Our rule was adapted to reflect this difference.

$\text{LIT}_{\text{GC}}$  is an entirely new rule introduced to type literals. As we do not distinguish between different primitive types such as  $\text{Int}\#$ ,  $\text{Float}\#$  or  $\text{Char}\#$ , we simply type these as the empty algebraic type  $\mu X. \{ \}$ .

The rules in Figure 6 cover two different kinds of pattern matches, namely on algebraic types and on primitive types:

$\text{CASEALG}_{\text{GC}}$  was extended from the original  $\text{MATCH}$  rule to account for some additional features



$$\frac{\Gamma \vdash_0^q e : C}{\Gamma \vdash_0^q \Lambda \tau. e : \mu T. \{ \} \xrightarrow{0} C} \quad (\text{TYABS}_{\text{GC}})$$

$$\frac{\Gamma \vdash_{p'}^p e : \mu T. \{ \} \xrightarrow{0} C}{\Gamma \vdash_{p'}^p e \tau : C} \quad (\text{TYAPP}_{\text{GC}})$$

$$\frac{\Gamma \vdash_{p'}^p e : C}{\Gamma \vdash_{p'}^p \text{let } \tau_1 = \tau_2 \text{ in } e : C} \quad (\text{TYLET}_{\text{GC}})$$

Figure 7: New type rules for handling type expressions

of Core’s case expression which are not available to JVfH’s match; Namely, binding the value of the scrutinee to a variable, and an additional default case which is used when none of the other cases match.  $\text{CASELIT}_{\text{GC}}$ , then, was heavily simplified from our own  $\text{CASEALG}_{\text{GC}}$  rule, as primitive types cannot have fields and as we do not distinguish between different primitive values.

Finally, we introduced the rules listed in Figure 7, which offer very limited support for type abstraction and type application. Unfortunately, as mentioned earlier, our type system currently does not support polymorphism yet, as this feature is also missing in JVfH, and extending the type system with major new features was not within the scope of our work. Instead, we opted to simply “ignore” type abstraction and application for now. To see why this can be convenient, consider the `repeat` example discussed earlier:

```
repeat x = let xs = x : xs in xs
it = repeat 1 :: [Int]
```

We do not explicitly specify a type for the function `repeat`; Therefore, GHC will automatically infer the polymorphic type “forall a. a -> [a]” and introduce the appropriate type abstractions and applications into the generated Core code. Without our type rules from Figure 7, our analysis would fail when it encounters any of these expressions. However, if we instead simply ignore them, we will then unify the type variable `a` with `Int`, which essentially turns the function into a monomorphic function of type “Int -> [Int]”. An additional wrapper  $\mu X. \{ \} \xrightarrow{0} C$  around the actual type  $C$  is introduced as an annotation to inform the user about any ignored type abstractions.

However, this should be considered a workaround, and will only work if every polymorphic function is used monomorphically. If it is used polymorphically, this will lead to the unification of two incompatible types; And if it is never used at all, the resulting type will contain free type variables, which is not valid in our system.

## 6 Implementation

The architecture of our tool is shown in Figure 8. We initiate the analysis by calling the GHC compiler on a given Haskell module. After the code was compiled to GHC Core, it is then fetched by our GHC plugin and passed to the actual analysis. Note that this Core “program” is actually a list of variable bindings; However, we can convert this list into a single expression by simply nesting `let` and `letrec` expressions. This allows us to generate an initial typing judgment, or “constraint”, which needs to be proven.

A constraint solver is used to apply the appropriate type rules to each judgment, and to keep track of the judgments that already have been or still need to be proven. Using a central “control unit” for manag-

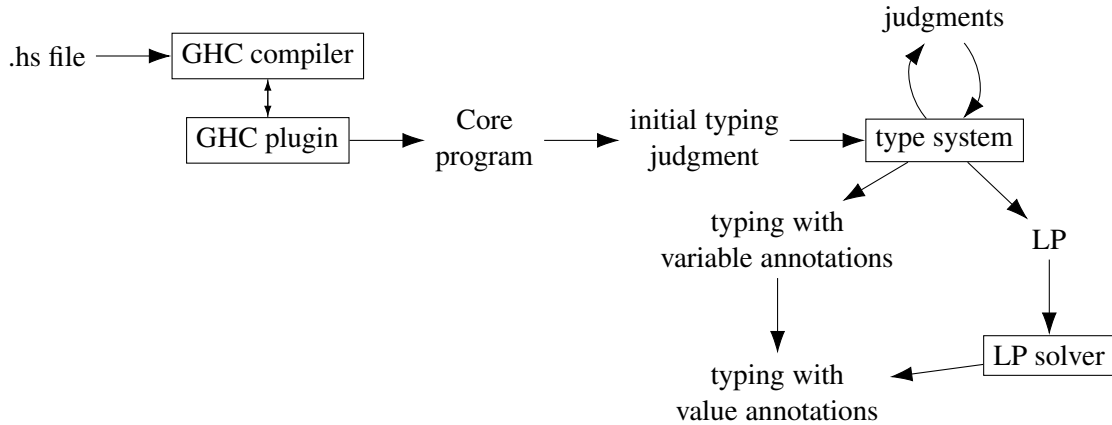


Figure 8: Outline of the architecture of our analysis

ing derived judgments like this has several advantages; For example, this allows us to print a trace when we encounter a typing error, or to output the full derivation tree after the analysis has completed. This can be used as an evidence to support the correctness of the analysis result, or as a tool for debugging.

In most cases, there is no ambiguousness regarding which inference rule to apply on any given judgment; However, structural type rules are an exception, as any of them may be applied to any typing judgment. For our implementation, we therefore need to find a way to make these rule applications deterministic. A naive solution for this problem is to apply every structural once to every generated typing judgment before applying the next syntax-driven type rule. However, this needlessly increases the size of the derivation tree, as most of these rule applications would not add anything to the strength of our type system.

Instead we reused the same heuristics as described in section “4.5 Experimental Results” of the original paper on JVfH. [9] For example, consider the `PREPAY` rule, which allows us to prepay the cost of a variable in the context. This is used to simulate the fact that thunks are evaluated at most once, even when they are referenced multiple times in an expression. According to Jost et al., it suffices to apply this rule once whenever a new variable is inserted into the context of a typing judgment. In our system, this applies to the `ABSGC`, `LETRECGC`, `LETGC` and `CASEALGGC` rules. Note that the `CASELITGC` also introduces a new variable, but its thunk cost is always zero, so prepaying would be redundant.

For the most part, we found the given heuristics to be sufficient, with one exception: In the `CASEALGGC` and `CASELITGC` rules, the number of resources remaining after evaluating the scrutinee is also used as the number of resources available before evaluating the respective cases. However, for most scrutinees, our type system will set this number to a low value which generally does not suffice for evaluating the cases. Therefore, we found it necessary to apply the `RELAX` rule to the scrutinee judgment. This allows us to artificially increase the number of resources available both before and after the evaluation by any value as needed. The intuition is that these additional resources are not used during the evaluation of the scrutinee; Therefore, they are available before its evaluation and also remain available afterwards.

In addition to the syntax-driven type rules mentioned in the previous section, we also had to hard-code the types of some specific variables. Most notably, these include operations on primitive types such as `Int#` or `Float#`. Usually, variables simply resolve to actual Haskell code, which we could analyze to derive an annotated type and an LP over its annotations. However, primitive operators do not correspond

to any Haskell code; Instead, they are detected by the GHC compiler and translated into an appropriate opcode for the target machine of the compilation. Since these operators are hard-coded into the compiler, we also need to hard-code them into our analysis. We introduce an additional cost constant  $K_{\text{prim}}$  which is used for every execution of a primitive operation.

We also hard-coded several other variables from the Prelude, including dictionaries and operators on boxed primitive types such as `Int` and `Float`. This is a temporary solution introduced only because our implementation currently does not support multi-module programs yet. In the future, references to other modules – including the Prelude – should automatically be detected and handled correctly, rendering this workaround obsolete.

Finally, for solving the linear program we make use of `glpk-hs`, which previously was also used in the implementation of the JVfH type system. This package provides Haskell bindings for the GNU Linear Programming Kit. However, we have introduced an additional abstraction layer, so the use of this library is transparent to the majority of our code. Thus, it should be possible to replace the LP solver without much effort if required.

## 7 Evaluation

Our implementation can analyze a substantial subset of Haskell. For example, consider the following simple Haskell module, which contains examples of several concepts that are common to Haskell, but nonexistent in JVfH, such as list comprehension and function applications with non-variable arguments:

```
module Example where
import Prelude hiding (map, repeat)
repeat x = xs where xs = x : xs
map f xs = [f x | x <- xs]
it = map (+1) $ repeat 1 :: [Int]
```

When applied to this module, our analysis will successfully generate the following typing for the variable `it`:

$$\frac{9}{0} \text{ it} : \mu X. \{ [] : (0, []) \mid (:) : (0, [T^1(\text{Int}), T^3(X)]) \}$$

This means that evaluating this expression to weak head normal form will induce a onetime cost of 9 allocations at most, and up to 3 more allocations for each list node accessed. Furthermore, accessing any of the list elements will also evoke one additional allocation.

For comparison, consider the following translation of the previous code to JVfH:

```
let repeat = \x -> letcons xs = Cons(x,xs) in xs
in let map = \f xt -> match xt with
    Nil() -> letcons r = Nil() in r
    | Cons(x,xs) -> let y = f x
                    in let ys = map f xs
                    in letcons r = Cons(y,ys)
                    in r

in let one = 1
    in let ones = repeat one
        in let inc = \x -> let r = one + x in r
            in let it = map inc ones
                in it
```

From this expression, the original JVfH analysis will infer the following typing:

$$\frac{\vdash_0}{\vdash} it : \mu X. \{ [] : (0, []) \mid (:) : (0, [\top^1(\text{Int}), \top^3(X)]) \}$$

Note that the typing is similar, but not identical. The specific values used as type annotations may differ, due to differences in how the code is represented in JVFH and GHC Core, respectively; and the LP solver also has some freedom in how values are assigned to variables.

This result is typical of most comparisons between JVFH and Haskell analyses. As we have not formally proven the correctness of our adapted system, we used JVFH analysis results as a basis for testing the viability of our own analysis. For this purpose, we translated several JVFH example programs published as part of a JVFH online demo<sup>1</sup> to Haskell. We then ran our analysis on this code and compared our results to those from the online demo. While we obtained similar results in most cases, there were a few noteworthy cases where our results were entirely different.

One of these is the following code which is a simplified version of the “list fusion” example from the online demo. For brevity, we will only provide our Haskell translation of the code:

```
map1 f [] = []
map1 f (x:xs) = f x : map1 f xs
map2 f [] = []
map2 f (x:xs) = f x : map2 f xs
lhs f g xs = map1 f $ map2 g xs
```

Note that we provide duplicate definitions of the map function for each usage within the lhs function. This is a workaround commonly used in the JVFH examples, which allows us to assign different types with different annotations to each usage of the function. Using this trick, the original JVFH version of the code above can be analyzed without any issues. However, when we attempt to analyze our Haskell translation, the GHC compiler detects this duplication and removes it, replacing all references to map2 with map1. As a result, our analysis of the lhs function will generate an unsolvable linear program and then fail. This is beyond our control, as we can only edit the Haskell code, and the translation to GHC Core is left entirely to the compiler.

However, this property also facilitated the discovery and elimination of an issue with a type rule that had previously been considered sound. We found this issue while analyzing the following definition of the Fibonacci sequence:

```
zipWith f (x:xs) (y:ys) = f x y : zipWith f xs ys
fibs = 0 : 1 : zipWith (+) fibs fibs' where (_,fibs') = fibs
```

Initially, our analysis erroneously was able to derive a constant cost for this sequence, meaning that it would be possible to fully evaluate the entire infinite list. We have previously detailed the cause and our fix of this issue in our discussion of the LETREC<sub>GC</sub> type rule in section 5. However, this issue was exposed only because GHC translated our Haskell code to Core using letrec for mutual recursion; The original JVFH version of this code did not contain any letand expressions – which is the JVFH equivalent of letrec – and therefore the erroneous type rule was not triggered during analysis. However, we were able to reproduce this error in the JVFH analysis by manually translating the Core code to JVFH, proving that this is an issue with the original rule, and not with our own implementation. While our fix causes the analysis of this Fibonacci code to fail due to an unsolvable linear program, this is preferable over returning an incorrect analysis result.

---

<sup>1</sup>The JVFH online demo is available at: <http://kashmir.dcc.fc.up.pt/cgi/lazy.cgi>

## 8 Outlook

The tool described in this paper was implemented as part of a master’s thesis. [12] The goal of this thesis was to provide an initial implementation, which can then be used as a basis for future improvements. While our analysis already works on a substantial subset of Haskell, there still are several flaws that severely inhibit its usefulness. In this section, we will discuss some of the shortcomings that may be used as a starting point for any future work.

The biggest issue of our current work is the lack of support for polymorphism. As we previously discussed in section 3, we specifically chose GHC Core over STG for our analysis with future support for polymorphism in mind; But as our type system is based on JVFH, which does not feature polymorphism at all, this feature is also missing in our current implementation. While the type syntax does include type variables, these may only be used as recursive references in an algebraic type; And several inference rules in the JVFH type system (which we currently reuse without alterations) were written with this assumption in mind. Therefore, providing full support for polymorphism may require large modifications to our type system.

We currently also do not support the “Cast” and “Coercion” expressions of GHC Core. These are used by the compiler to support types defined via `newtype`, as well as several Haskell language extensions such as GADT, associated types, and functional dependencies. [13] Therefore, none of these features are currently supported in our analysis. As of yet, we have not investigated which modifications to our type system would be necessary to support these two expressions.

Multi-module programs also are not supported yet; Therefore, we cannot analyze any expressions that contain variables imported from different modules, including the `Prelude`. In some cases, however, the compiler may automatically inline function calls, circumventing this limitation. This is the reason why we can use the operators `+` and `$` in the example codes in the previous section.

We also have not formally proven the soundness of our adapted type system. For this purpose, it would be necessary to define the cost-annotated operational semantics for GHC Core; and to show that type annotations calculated by our analysis are always greater or equal to the costs determined by the operational semantics. As we have previously discussed in sections 5 and 7, it may also be necessary to revisit the soundness proof for the original `LETAND†` rule [10].

## 9 Conclusion

In this paper, we have presented an automated amortized analysis on the resource usage of Haskell programs. Our work is based on a previous paper by Jost et al. [9], which proposed a type-based analysis on an artificial language called JVFH. We use the plugin API of the GHC compiler to translate any given Haskell code into a heavily simplified representation called GHC Core. Due to the large similarities to JVFH, we can reuse most of the definitions and inference rules from the JVFH system. Only a subset of the type rules were modified to take minor syntactic and semantic differences between the two languages into account. For the most part, our implementation will return similar results to those produced by the original JVFH analysis; However, there still are several shortcomings, which make our current implementation impractical for everyday use. Most notably, certain features such as polymorphism and `newtype` are not supported yet and were deferred to future work.

## References

- [1] *The GHC API: CoreSyn*. <https://downloads.haskell.org/~ghc/8.2.2/docs/html/libraries/ghc-8.2.2/CoreSyn.html>. Accessed: 2019-04-29.
- [2] *The GHC API: HsExpr*. <https://downloads.haskell.org/~ghc/8.2.2/docs/html/libraries/ghc-8.2.2/HsExpr.html>. Accessed: 2019-04-29.
- [3] *The GHC API: Plugins*. <https://downloads.haskell.org/~ghc/8.2.2/docs/html/libraries/ghc-8.2.2/Plugins.html>. Accessed: 2019-04-29.
- [4] Max Bolingbroke et al.: *The GHC Commentary: I know kung fu: learning STG by example*. [https://gitlab.haskell.org/ghc/ghc/wikis/commentary/compiler/generated-code?version\\_id=33352ae3](https://gitlab.haskell.org/ghc/ghc/wikis/commentary/compiler/generated-code?version_id=33352ae3). Accessed: 2019-04-29.
- [5] Jan Hoffmann, Klaus Aehlig & Martin Hofmann (2012): *Resource Aware ML*. In: *24rd International Conference on Computer Aided Verification (CAV'12), Lecture Notes in Computer Science 7358*, Springer, pp. 781–786, doi:10.1007/978-3-642-31424-7\_64.
- [6] Jan Hoffmann, Ankush Das & Shu-Chun Weng (2017): *Towards Automatic Resource Bound Analysis for OCaml*. In: *ACM SIGPLAN Notices*, 52, ACM, pp. 359–373, doi:10.1145/3009837.3009842.
- [7] Nataliya Irkha (2016): *Evaluation und Erweiterung einer Typ-basierten Kostenanalyse für funktionale Sprachen mit verzögerter Auswertung*. Bachelor's thesis, Ludwig-Maximilian-Universität München.
- [8] Steffen Jost (2010): *Automated Amortised Analysis*. Ph.D. thesis, Ludwig-Maximilian-Universität München.
- [9] Steffen Jost, Pedro Vasconcelos, Mário Florido & Kevin Hammond (2017): *Type-Based Cost Analysis for Lazy Functional Languages*. *Journal of Automated Reasoning* 59(1), pp. 87–120, doi:10.1007/s10817-016-9398-9. Available at <http://www.dcc.fc.up.pt/~pbv/research/JAR2016-draft.pdf>.
- [10] Sophie Kleber (2017): *Mutual Recursive Definitions for a Type-Based Cost Analysis for Lazy Functional Languages*. Bachelor's thesis, Ludwig-Maximilian-Universität München.
- [11] Maksym Redka (2017): *Erweiterung des Prototyps einer Typ-basierten Kostenanalyse für funktionale Sprachen mit verzögerter Auswertung*. Bachelor's thesis, Ludwig-Maximilian-Universität München.
- [12] Franz Sigmüller (2018): *Implementation of an Automated Amortized Analysis on GHC Core as Compiler Plugin*. Master's thesis, Ludwig-Maximilian-Universität München.
- [13] Martin Sulzmann, Manuel Chakravarty, Simon Peyton Jones & Kevin Donnelly (2007): *System F with Type Equality Coercions*. In: *ACM SIGPLAN International Workshop on Types in Language Design and Implementation (TLDI'07)*, pp. 53–66, doi:10.1145/1190315.1190324.
- [14] GHC Team (2017): *GHC User's Guide Documentation: Release 8.2.2*. [https://downloads.haskell.org/~ghc/8.2.2/docs/users\\_guide.pdf](https://downloads.haskell.org/~ghc/8.2.2/docs/users_guide.pdf).
- [15] Joe B Wells (1999): *Typability and type checking in System F are equivalent and undecidable*. *Annals of Pure and Applied Logic* 98(1-3), pp. 111–156, doi:10.1016/S0168-0072(98)00047-5.