Automatic Static Cost Analysis for Parallel Programs

Jan Hoffmann Zhong Shao
Yale University

Abstract

Static analysis of the evaluation cost of programs is an extensively studied problem that has many important applications. However, most automatic methods for static cost analysis are limited to sequential evaluation while programs are increasingly evaluated on modern multicore and multiprocessor hardware.

This article introduces the first automatic analysis for deriving bounds on the worst-case evaluation cost of parallel first-order functional programs. The analysis is performed by a novel type system for amortized resource analysis. The main innovation is a technique that separates the reasoning about sizes of data structures and evaluation cost within the same framework. The cost semantics of parallel programs is based on call-by-value evaluation and the standard cost measures *work* and *depth*. A soundness proof of the type system establishes the correctness of the derived cost bounds with respect to the cost semantics. The derived bounds are multivariate resource polynomials which depend on the sizes of the arguments of a function. Type inference can be reduced to linear programming and is fully automatic. A prototype implementation of the analysis system has been developed to experimentally evaluate the effectiveness of the approach. The experiments show that the analysis infers bounds for realistic example programs such as quick sort for lists of lists, matrix multiplication, and an implementation of sets with lists. The derived bounds are often asymptotically tight and the constant factors are close to the optimal ones.

1. Introduction

Static analysis of the resource cost of programs is a classical subject of computer science. Recently, there has been an increased interest in formally proving cost bounds since they are essential in the verification of safety-critical real-time and embedded systems. As you might recall from books such as Knuth's *The Art of Computer Programming*, it is tedious and error-prone to manually reason about the exact resource cost of programs, as formally defined by a cost semantics. As a result, machine support for statically deriving resource bounds for programs has been extensively studied.

For sequential functional programs there exist many automatic and semi-automatic analysis systems that can statically infer cost bounds. Most of them are based on sized types [34], recurrence relations [14], and amortized resource analysis [24, 31]. The goal of these systems is to automatically compute easily-understood arithmetic expressions in the sizes of the inputs of a program that bound resource cost such as time or space usage. Even though an automatic computation of cost bounds is undecidable in general, novel analysis techniques are able to efficiently compute tight time and space bounds for many non-trivial programs [16, 2, 22, 9, 32].

For functional programs that are evaluated in parallel, on the other hand, no such analysis system exists to support programmers with computer-aided derivation of cost bounds. In particular, there are no type systems that derive cost bounds for parallel programs. This is unsatisfying because parallel evaluation is becoming increasingly important on modern hardware and referential transparency makes functional programs ideal for parallel evaluation.

This article introduces an automatic type-based resource analysis for the derivation of cost bounds for parallel first-order functional programs. Automatic cost analysis for sequential programs is already challenging and it might seem to be a long shot to develop an analysis for parallel evaluation that takes into account low-level features of the underlying hardware such as the number of processors. Fortunately, it has been shown [7, 8] that the cost of parallel functional programs can be analyzed in two steps. First, we derive cost bounds at a high abstraction level where we assume to have an unlimited number of processors at our disposal. Second, we prove once and for all how the cost on the high abstraction level relates to the actual cost on a specific system with limited resources.

In this work, we derive bounds on an abstract cost model that consists of the *work* and the *depth* of an evaluation of a program [7]. Work measures the evaluation time of sequential evaluation and depth measures the evaluation time of parallel evaluation assuming an unlimited number of processors. It is well-known [17] that a program that evaluates to a value using work w and depth d can be evaluated on a shared-memory multiprocessor (SMP) system with p processors in time $O(\max(w/p, d))$ (see Section 2.3). The mechanism that is used to prove this result is comparable to a scheduler in an operating system.

Technically, the analysis computes two separate typing derivations, one for the work and one for the depth. To derive a bound on the work, we use multivariate amortized resource analysis for sequential programs [21]. To derive a bound on the depth, we develop a novel multivariate amortized resource analysis for programs that are evaluated in parallel. The main challenge in the design of this novel parallel analysis is to ensure the same high compositionality as in the sequential analysis. The design and implementation of this novel analysis for bounds on the depth of evaluations is the main contribution of our work. The technical innovation that enables compositionality is an analysis method that separates the static tracking of size changes of data structures from the cost analysis while using the same framework. We envision that this technique will find further applications in the analysis of other non-additive cost such as stack-space usage and recursion depth.

A novelty in the cost semantics in this paper is the definition of work and depth for terminating and non-terminating evaluations. Intuitively, the non-deterministic big-step evaluation judgement that is defined in Section 2 expresses that *there is a (possibly partial) evaluation with work n and depth m*. This statement is used to prove that a typing derivation for bounds on the depth or for bounds on the work ensures termination. Bounds on parallel evaluation also prove termination of the sequential evaluation.

We describe the new type analysis for parallel evaluation for a simple first-order language with lists, pairs, pattern matching, and sequential and parallel composition. This is already sufficient to study the cost analysis of parallel programs. However, we implemented the analysis system in Resource Aware ML (RAML), which

also includes other inductive data types and conditionals [23]. To demonstrate the universality of the approach, we also implemented NESL's [6] parallel list comprehensions as a primitive in RAML (see Section 6). Similarly, we can define other parallel sequence operations of NESL as primitives and correctly specify their work and depth. RAML is currently extended to include higher-order functions, arrays, and user-defined inductive types. This work is orthogonal to the treatment of parallel evaluation.

To evaluate the practicability of the proposed technique, we performed an experimental evaluation of the analysis using the prototype implementation in RAML. Note that the analysis computes worst-case bounds instead of average-case bounds and that the asymptotic behavior of many of the classic examples of Blelloch et al. [7] does not differ in parallel and sequential evaluations. For instance, the depth and work of quick sort are both quadratic in the worst-case. Therefore, we focus on examples that actually have asymptotically different bounds for the work and depth. This includes quick sort for lists of lists in which the comparisons of the inner lists can be performed in parallel, matrix multiplication where matrices are lists of lists, a function that computes the maximal weight of a (continuous) sublist of an integer list, and the standard operations for sets that are implemented as lists. The experimental evaluation can be easily reproduced and extended: RAML and the example programs are publicly available for download and through an user-friendly online interface [1].

In summary we make the following contributions.

1. We introduce the first automatic static analysis for deriving bounds on the depth of parallel functional programs. Being based on multivariate resource polynomials and type-based amortized analysis, the analysis is highly compositional. The computed type derivations are easily checkable certificates of the bound.
2. We prove the soundness of the type-based amortized analysis with respect to an operational big-step semantics that models the work and depth of terminating and non-terminating programs. This allows us to prove that work and depth bounds ensure termination. Our inductively defined big-step semantics is an interesting alternative to coinductive big-step semantics.
3. We implemented the proposed analysis in RAML, an OCaml-like functional language. In addition to the language constructs like lists and pairs that are formally described in this article, the implementation includes binary trees, natural numbers, tuples, Booleans, and NESL's parallel list comprehensions. RAML is publically available for download and through an easy-to-use web interface [1].
4. We evaluated the practicability of the implemented analysis by performing reproducible experiments with typical example programs. Our results show that the analysis is efficient and works for a wide range of examples. The derived bounds are usually asymptotically tight if the tight bound is expressible as a resource polynomial.

The remainder of the article is organized as follows. In Section 2 we define work and depth and discuss its relation to the evaluation of programs on realistic systems. In Section 3 we informally introduce amortized resource analysis and explain the main idea of the novel analysis for the depth. Section 4 introduces multivariate resource polynomials which form the foundation of our analysis. Section 5 contains the type system for the cost analysis of parallel evaluation and the proof of the soundness of the type system with respect to the cost semantics. The implementation and the experimental evaluation are described in Section 7. Finally, we discuss related work (Section 8) and conclude (Section 9).

2. Cost Semantics for Parallel Programs

In this section, we introduce a first-order functional language with parallel and sequential composition. We then define a big-step operational semantics that formalizes the cost measures *work* and *depth* for terminating and non-terminating evaluations. Finally, we prove properties of the cost semantics and discuss the relation of work and depth to the run time on hardware with limited resources.

2.1 Expressions and Programs

Expressions are given in let-normal form. This means that term formers are applied to variables only when this does not restrict the expressivity of the language. Expressions are formed by integers, variables, function applications, lists, pairs, pattern matching, and sequential and parallel composition.

$$\begin{aligned} e, e_1, e_2 ::= {} & n \mid x \mid f(x) \mid \mathsf{nil} \mid \mathsf{cons}(x_1, x_2) \mid (x_1, x_2) \\ & \mid \mathsf{match}\, x \,\mathsf{with}\, \langle \mathsf{nil} \Rightarrow e_1 \mid \mathsf{cons}(x_1, x_2) \Rightarrow e_2 \rangle \\ & \mid \mathsf{match}\, x \,\mathsf{with}\, (x_1, x_2) \Rightarrow e \mid \mathsf{let}\, x = e_1 \,\mathsf{in}\, e_2 \\ & \mid \mathsf{par}\, x_1 = e_1 \,\mathsf{and}\, x_2 = e_2 \,\mathsf{in}\, e \end{aligned}$$

The parallel composition $\mathsf{par}\, x_1 = e_1 \,\mathsf{and}\, x_2 = e_2 \,\mathsf{in}\, e$ is used to evaluate e_1 and e_2 in parallel and bind the resulting values to the names x_1 and x_2 for use in the evaluation of e.

In the prototype, we have implemented other inductive types such as trees, natural numbers, and tuples. Additionally, there are operations for primitive types such as Booleans and integers, and NESL's parallel list comprehensions [6]. Expressions are also transformed automatically into let normal form before the analysis. In the examples in this paper, we use the syntax of our prototype implementation to improve readability.

In the following, we define a standard type system for expressions and programs. Data types A, B and function types F are defined as follows.

$$\begin{aligned} A, B \quad & ::= \quad \mathsf{int} \mid L(A) \mid A * B \\ F \quad & ::= \quad A \to B \end{aligned}$$

Let $\mathcal{A}$ be the set of data types and let $\mathcal{F}$ be the set of function types. A signature $\Sigma : \mathrm{FID} \rightharpoonup \mathcal{F}$ is a partial finite mapping from function identifiers to function types. A context is a partial finite mapping $\Gamma : \mathit{Var} \rightharpoonup \mathcal{A}$ from variable identifiers to data types. A simple type judgement $\Sigma; \Gamma \vdash e : A$ states that the expression e has type A in the context Γ under the signature Σ. The definition of typing rules for this judgement is standard and we omit the rules. Basically, the rules are obtained by erasing the potential annotations of the syntax-directed rules in Section 5.

A *(well-typed) program* consists of a signature Σ and a family $(e_f, y_f)_{f \in \mathrm{dom}(\Sigma)}$ of expressions e_f with a distinguished variable identifier y_f such that $\Sigma; y_f{:}A \vdash e_f{:}B$ if $\Sigma(f) = A \to B$.

2.2 Big-Step Operational Semantics

We now formalize the resource cost of evaluating programs with a big-step operational semantics. The focus of this paper is on time complexity and we only define the cost measures *work* and *depth*. Intuitively, the work measures the time that is needed in a sequential evaluation. The depth measures the time that is needed in a parallel evaluation. In the semantics, time is parameterized by a metric that assigns a non-negative cost to each evaluation step.

Motivation A distinctive feature of our big-step semantics is that it models terminating, failing, and diverging evaluations by inductively describing finite subtrees of (possibly infinite) evaluation trees. By using an inductive judgement for diverging and terminating computations while avoiding intermediate states, it combines the advantages of big-step and small-step semantics. This has two benefits compared to standard big-step semantics.

$\boxed{V, H \overset{M}{\vdash} e \Downarrow \rho \mid (w, d)}$ For stack V and heap H, e evaluates to ρ with work w and depth d in a fixed program $(e_f, y_f)_{f \in \mathrm{dom}(\Sigma)}$.

$$\frac{n \in \mathbb{Z} \qquad H' = H, \ell \mapsto n}{V, H \overset{M}{\vdash} n \Downarrow (\ell, H') \mid (M^{\mathsf{const}}, M^{\mathsf{const}})} \text{(E:CONST)}$$

$$\frac{[y_f \mapsto V(x)], H \overset{M}{\vdash} e_f \Downarrow \rho \mid (w, d)}{V, H \overset{M}{\vdash} f(x) \Downarrow \rho \mid (M^{\mathsf{app}}{+}w, M^{\mathsf{app}}{+}d)} \text{(E:APP)}$$

$$\frac{H(V(x)) = (\ell_1, \ell_2) \qquad V[x_1 \mapsto \ell_1, x_2 \mapsto \ell_2], H \overset{M}{\vdash} e \Downarrow \rho \mid (w, d)}{V, H \overset{M}{\vdash} \mathsf{match}\ x\ \mathsf{with}\ (x_1, x_2) \Rightarrow e \Downarrow \rho \mid (M^{\mathsf{matP}}{+}w, M^{\mathsf{matP}}{+}d)} \text{(E:MATP)}$$

$$\frac{H' = H, \ell \mapsto (V(x_1), V(x_2))}{V, H \overset{M}{\vdash} (x_1, x_2) \Downarrow (\ell, H') \mid (M^{\mathsf{pair}}, M^{\mathsf{pair}})} \text{(E:PAIR)}$$

$$\frac{H(V(x)) = \mathsf{nil} \qquad V, H \overset{M}{\vdash} e_1 \Downarrow \rho \mid (w, d)}{V, H \overset{M}{\vdash} \mathsf{match}\ x\ \mathsf{with}\ \langle \mathsf{nil} \Rightarrow e_1 \mid \mathsf{cons}(x_1, x_2) \Rightarrow e_2 \rangle \Downarrow \rho \mid (M^{\mathsf{matN}}{+}w, M^{\mathsf{matN}}{+}d)} \text{(E:MATN)}$$

$$\frac{}{V, H \overset{M}{\vdash} e \Downarrow \circ \mid (0, 0)} \text{(E:ABORT)}$$

$$\frac{H(V(x)) = (\mathsf{cons}, \ell_1, \ell_2) \qquad V[x_1 \mapsto \ell_1, x_2 \mapsto \ell_2], H \overset{M}{\vdash} e_2 \Downarrow \rho \mid (w, d)}{V, H \overset{M}{\vdash} \mathsf{match}\ x\ \mathsf{with}\ \langle \mathsf{nil} \Rightarrow e_1 \mid \mathsf{cons}(x_1, x_2) \Rightarrow e_2 \rangle \Downarrow \rho \mid (M^{\mathsf{matL}}{+}w, M^{\mathsf{matL}}{+}d)} \text{(E:MATC)}$$

$$\frac{H' = H, \ell \mapsto \mathsf{nil}}{V, H \overset{M}{\vdash} \mathsf{nil} \Downarrow (\ell, H') \mid (M^{\mathsf{nil}}, M^{\mathsf{nil}})} \text{(E:NIL)}$$

$$\frac{V, H \overset{M}{\vdash} e_1 \Downarrow \circ \mid (w, d)}{V, H \overset{M}{\vdash} \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \Downarrow \circ \mid (M^{\mathsf{let}}{+}w, M^{\mathsf{let}}{+}d)} \text{(E:LET1)}$$

$$\frac{H' = H, \ell \mapsto (\mathsf{cons}, V(x_1), V(x_2))}{V, H \overset{M}{\vdash} \mathsf{cons}(x_1, x_2) \Downarrow (\ell, H') \mid (M^{\mathsf{cons}}, M^{\mathsf{cons}})} \text{(E:CONS)}$$

$$\frac{V, H \overset{M}{\vdash} e_1 \Downarrow (\ell, H') \mid (w_1, d_1) \qquad V[x \mapsto \ell], H' \overset{M}{\vdash} e_2 \Downarrow \rho \mid (w_2, d_2)}{V, H \overset{M}{\vdash} \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \Downarrow \rho \mid (M^{\mathsf{let}}{+}w_1{+}w_2, M^{\mathsf{let}}{+}d_1{+}d_2)} \text{(E:LET2)}$$

$$\frac{V(x) = \ell}{V, H \overset{M}{\vdash} x \Downarrow (\ell, H) \mid (M^{\mathsf{var}}, M^{\mathsf{var}})} \text{(E:VAR)}$$

$$\frac{V, H \overset{M}{\vdash} e_1 \Downarrow \rho_1 \mid (w_1, d_1) \qquad V, H \overset{M}{\vdash} e_2 \Downarrow \rho_2 \mid (w_2, d_2) \qquad \rho_1 = \circ \vee \rho_2 = \circ}{V, H \overset{M}{\vdash} \mathsf{par}\ x_1 = e_1\ \mathsf{and}\ x_2 = e_2\ \mathsf{in}\ e \Downarrow \circ \mid (M^{\mathsf{Par}}{+}w_1{+}w_2, M^{\mathsf{Par}} + \max(d_1, d_2))} \text{(E:PAR1)}$$

$$\frac{V, H \overset{M}{\vdash} e_1 \Downarrow (\ell_1, H_1) \mid (w_1, d_1) \qquad V, H \overset{M}{\vdash} e_2 \Downarrow (\ell_2, H_2) \mid (w_2, d_2) \qquad V[x_1 {\mapsto} \ell_1, x_2 {\mapsto} \ell_2], H_1 \uplus H_2 \overset{M}{\vdash} e \Downarrow (\ell, H') \mid (w, d)}{V, H' \overset{M}{\vdash} \mathsf{par}\ x_1 = e_1\ \mathsf{and}\ x_2 = e_2\ \mathsf{in}\ e \Downarrow (\ell, H') \mid (M^{\mathsf{Par}}{+}w_1{+}w_2{+}w, M^{\mathsf{Par}} + \max(d_1, d_2){+}d)} \text{(E:PAR2)}$$

Figure 1: Rules of the operational big-step semantics.

1. We can model the resource consumption of diverging programs and prove that bounds hold for terminating and diverging programs. (In some cost metrics, diverging computations can have finite cost.)
2. For a cost metric in which all diverging computations have infinite cost we are able to show that bounds imply termination.

Note that we cannot achieve this by step-indexing a standard big-step semantics. The available alternatives to our approach are small-step semantics and coinductive big-step semantics. However, it is unclear how to prove the soundness of our type system with respect to these semantics. Small-step semantics is difficult to use because our type-system models an intentional property that goes beyond the classic type preservation: After performing a step, we have to obtain a refined typing that corresponds to a (possibly) smaller bound. Coinductive derivations are hard to relate to type derivations because type derivations are defined inductively.

Our inductive big-step semantics can not only be used to formalize resource cost of diverging computations but also for other effects such as event traces. It is therefore an interesting alternative to recently proposed coinductive operational big-step semantics [11].

Semantic Judgements We formulate the big-step semantics with respect to a stack and a heap. Let *Loc* be an infinite set of *locations* modeling memory addresses on a heap. The set of *values* *Val* is defined as follows.

$$v ::= n \mid (\ell_1, \ell_2) \mid (\mathsf{cons}, \ell_1, \ell_2) \mid \mathsf{nil}$$

A value $v \in \mathit{Val}$ is either an integer $n \in \mathbb{Z}$, a pair of locations (ℓ_1, ℓ_2), a node $(\mathsf{cons}, \ell_1, \ell_2)$ of a list, or nil.

A *heap* is a finite partial mapping $H : \mathit{Loc} \rightharpoonup \mathit{Val}$ that maps locations to values. A *stack* is a finite partial mapping $V : \mathit{Var} \rightharpoonup \mathit{Loc}$ from variable identifiers to locations. Thus we have boxed values. It is not important for the analysis whether values are boxed.

The big-step evaluation rules in Figure 1 are formulated with respect to a resource metric M. They define the evaluation judgment

$$V, H \overset{M}{\vdash} e \Downarrow \rho \mid (w, d) \qquad \text{where} \qquad \rho ::= (\ell, H) \mid \circ\ .$$

It expresses the following. In a fixed program $(e_f, y_f)_{f \in \mathrm{dom}(\Sigma)}$, if the stack V and the initial heap H are given then the expression e evaluates to ρ. Under the metric M, the work of the evaluation of e is w and the depth of the evaluation is d. Unlike standard big-step operational semantics, ρ can be either a pair of a location and a new heap, or $\circ$ (pronounced *busy*) indicating that the evaluation is not finished yet.

It might be sometimes handy to also include the possibility of ρ being a failure value $\perp$ to have an explicit judgement for evaluations that go wrong. However, this is not necessary here since we are mainly concerned with resource consumption rather than safety.

A resource metric $M : K \to \mathbb{Q}_0^+$ defines the resource consumption in each evaluation step of the big-step semantics with a non-negative rational number. We define

$$K = \{\mathsf{const}, \mathsf{var}, \mathsf{app}, \mathsf{cons}, \mathsf{nil}, \mathsf{matN}, \mathsf{matC}, \mathsf{pair}, \mathsf{matP}, \mathsf{let}, \mathsf{par}\}.$$

We write M^k for $M(k)$.

An intuition for the judgement $V, H \overset{M}{\vdash} e \Downarrow \circ \mid (w, d)$ is that there is a partial evaluation of e that runs without failure, has work w and depth d, and has not yet reached a value. This is similar to a small-step judgement.

Rules The rules of the big-step operational semantics are given in Figure 1. For a heap H, we write $H, \ell \mapsto v$ to express that $\ell \notin \mathrm{dom}(H)$ and to denote the heap H' such that $H'(x) = H(x)$ if $x \in \mathrm{dom}(H)$ and $H'(\ell) = v$. In the rule E:PAR2, we write $H_1 \uplus H_2$ to indicate that H_1 and H_2 agree on the values of locations in $\mathrm{dom}(H_1) \cap \mathrm{dom}(H_2)$ and to define the heap H with

$$\dfrac{\dfrac{}{x \mapsto \ell, \ell \mapsto 0 \;\vdash^{M}\; \omega(x) \Downarrow \circ \mid (0,0)}\text{(E:ABORT)}}{x \mapsto \ell, \ell \mapsto 0 \;\vdash^{M}\; \omega(x) \Downarrow \circ \mid (M^{\mathsf{app}}, M^{\mathsf{app}})}\text{(E:APP)}$$

$$\dfrac{\dfrac{\dfrac{}{x \mapsto \ell, \ell \mapsto 0 \;\vdash^{M}\; \omega(x) \Downarrow \circ \mid (0,0)}\text{(E:ABORT)}}{x \mapsto \ell, \ell \mapsto 0 \;\vdash^{M}\; \omega(x) \Downarrow \circ \mid (M^{\mathsf{app}}, M^{\mathsf{app}})}\text{(E:APP)}}{x \mapsto \ell, \ell \mapsto 0 \;\vdash^{M}\; \omega(x) \Downarrow \circ \mid (2M^{\mathsf{app}}, 2M^{\mathsf{app}})}\text{(E:APP)}$$

Figure 2: Two example derivations for the diverging function call $\omega(x)$ where ω is defined through $\omega(x) = \omega(x)$.

$\text{dom}(H) = \text{dom}(H_1) \cup \text{dom}(H_2)$ such that

$$H(x) = \begin{cases} H_1(x) & \text{if } x \in \text{dom}(H_1) \cap \text{dom}(H_2) \wedge H_1(x) = H_2(x) \\ H_1(x) & \text{if } x \in \text{dom}(H_1) \backslash \text{dom}(H_2) \\ H_2(x) & \text{if } x \in \text{dom}(H_2) \backslash \text{dom}(H_1) \end{cases}$$

We assume that the locations that are allocated in parallel evaluations are disjoint. That is easily achievable in an implementation.

The most interesting rules of the semantics are E:ABORT, and the rules for sequential and parallel composition. They allow us to approximate infinite evaluation trees for non-terminating evaluations with finite subtrees. The rule E:ABORT states that we can partially evaluate every expression by doing zero steps. The work w and depth d are then both zero (i.e., $w = d = 0$).

To obtain an evaluation judgement for a sequential composition $\mathsf{let}\, x = e_1 \,\mathsf{in}\, e_2$ we have two options. We can use the rule E:LET1 to partially evaluate e_1 using work w and depth d. Alternatively, we can use the rule E:LET2 to evaluate e_1 until we obtain a location and a heap (ℓ, H') using work w_1 and depth d_1. Then we evaluate e_2 using work w_2 and depth d_2. The total work and depth is then given by $M^{\mathsf{let}}{+}w_1{+}w_2$ and $M^{\mathsf{let}}{+}d_1{+}d_2$, respectively.

Similarly, we can derive evaluation judgements for a parallel composition $\mathsf{par}\, x_1 = e_1 \,\mathsf{and}\, x_2 = e_2 \,\mathsf{in}\, e$ using the rules E:PAR1 and E:PAR2. In the rule E:PAR1, we partially evaluate e_1 or e_2 with evaluation cost (w_1, d_1) and (w_2, d_2). The total work is then $M^{\mathsf{Par}}{+}w_1{+}w_2$ (the cost for the evaluation of the parallel binding plus the cost for the sequential evaluation of e_1 and e_2). The total depth is $M^{\mathsf{Par}}{+}\max(d_1, d_2)$ (the cost for the evaluation of the binding plus the maximum of the cost of the depths of e_1 and e_2). The rule E:PAR2 handles the case in which e_1 and e_2 are fully evaluated. It is similar to E:LET2 and the cost of the evaluation of the expression e is added to both the cost and the depth since e is evaluated after e_1 and e_2.

Figure 2 shows two example derivations for the diverging expressions $\omega(x)$ in the program $(\omega(x), x)_\omega$, that is, the function ω is defined through $\omega(x) = \omega(x)$. We can derive $x \mapsto \ell, \ell \mapsto 0 \vdash^{M} \omega(x) \Downarrow \circ \mid (n{\cdot}M^{\mathsf{app}}, n{\cdot}M^{\mathsf{app}})$ for every $n \in \mathbb{N}$.

An interesting observation is that depth and work form semirings with the operations (+,+) and (+,max), respectively. In general, resource consumption of parallel programs can be described by a semiring in the same way resource consumption of sequential programs can be described by a monoid.

2.3 Properties of the Cost-Semantics

A convenient property of the semantics is that heap-cells are only allocated but never deallocated during evaluation. This is not required for the soundness proof nor for parallel evaluation. However, it simplifies the soundness proof. The following proposition follows directly from an inspection of the evaluation rules.

Proposition 1. If $V, H \vdash^{M} e \Downarrow (\ell', H') \mid (w, d)$ then $H'(\ell) = H(\ell)$ for all $\ell \in \text{dom}(H)$.

$\boxed{H \vDash \ell \mapsto a : A}$ In heap H, ℓ points to the sem. value $a \in [\![A]\!]$.

$$\frac{H(\ell){=}n \qquad n \in \mathbb{Z}}{H \vDash \ell \mapsto n : \text{int}}\text{(V:INT)} \qquad \frac{H(\ell) = \mathsf{nil}}{H \vDash \ell \mapsto [\,] : L(A)}\text{(V:NIL)}$$

$$\frac{H(\ell) = (\ell_1, \ell_2) \qquad H \vDash \ell_1 \mapsto a : A \qquad H \vDash \ell_2 \mapsto b : B}{H \vDash \ell \mapsto (a, b) : A * B}\text{(V:PAIR)}$$

$$\frac{\begin{array}{c} H(\ell) = (\mathsf{cons}, \ell_1, \ell_2) \\ H \vDash \ell_1 \mapsto a_1 : A \qquad H \vDash \ell_2 \mapsto [a_2, \ldots, a_n] : L(A) \end{array}}{H \vDash \ell \mapsto [a_1, \ldots, a_n] : L(A)}\text{(V:CONS)}$$

Figure 3: Relating heap cells to semantic values.

The main theorem of this section states that the resource cost of a partial evaluation is less than or equal to the cost of an evaluation of the same expression that terminates.

Theorem 1. If $V, H \vdash^{M} e \Downarrow (\ell, H') \mid (w, d)$ and $V, H \vdash^{M} e \Downarrow \circ \mid (w', d')$ then $w' \leqslant w$ and $d' \leqslant d$.

Theorem 1 can be proved by a straightforward induction on the derivation of the judgement $V, H \vdash^{M} e \Downarrow (\ell, H') \mid (w, d)$.

Provably Efficient Implementations While work is a realistic cost-model for the sequential execution of programs, depth is not a realistic cost-model for parallel execution. The main reason is that it assumes that an infinite number of processors can be used for parallel evaluation. However, it has been shown [7] that work and depth are closely related to the evaluation time on more realistic abstract machines.

For example, *Brent's Theorem* [17] provides an asymptotic bound on the number of execution steps on the shared-memory multiprocessor (SMP) machine. It states that if $V, H \vdash^{M} e \Downarrow (\ell, H') \mid (w, d)$ then e can be evaluated on a p-processor SMP machine in time $O(\max(w/p, d))$. An SMP machine has a fixed number p of processes and provides constant-time access to a shared memory. The proof of Brent's Theorem can be seen as the description of a so-called *provably efficient implementation*, that is, an implementation for which we can establish an asymptotic bound that depends on the number of processors.

Classically, we are especially interested in non-asymptotic bounds in resource analysis. It would thus be interesting to develop a non-asymptotic version of Brent's Theorem for a specific architecture using more refined models of concurrency [8]. However, such a development is not in the scope of this article.

Well-Formed Environments and Type Soundness For each data type A we inductively define a set $[\![A]\!]$ of values of type A. Lists are interpreted as lists and pairs are interpreted as pairs.

$$\begin{array}{rcl} [\![\text{int}]\!] & = & \mathbb{Z} \\ [\![A * B]\!] & = & [\![A]\!] \times [\![B]\!] \\ [\![L(A)]\!] & = & \{[a_1, \ldots, a_n] \mid n \in \mathbb{N}, a_i \in [\![A]\!]\} \end{array}$$

If H is a heap, ℓ is a location, A is a data type, and $a \in [\![A]\!]$ then we write $H \vDash \ell \mapsto a : A$ to mean that ℓ defines the semantic value $a \in [\![A]\!]$ when pointers are followed in H in the obvious way. The judgment is formally defined in Figure 3.

If we fix a simple type A and a heap H then there exists at most one semantic value a such that $H \vDash \ell \mapsto a : A$.

Proposition 2. Let H be a heap, $\ell \in Loc$, and let A be a simple type. If $H \vDash \ell \mapsto a : A$ and $H \vDash \ell \mapsto a' : A$ then $a = a'$.

Note that if $H \vDash \ell \mapsto a : A$ then ℓ can point to a data structure with aliasing, but circularity is not allowed since this would require values a of infinite size. There is no way of generating such circular values in the functional language we study here.

We write $H \vDash \ell : A$ to indicate that there exists a, necessarily unique, semantic value $a \in [\![A]\!]$ so that $H \vDash \ell \mapsto a : A$. A stack V and a heap H are *well-formed* with respect to a context Γ if $H \vDash V(x) : \Gamma(x)$ holds for every $x \in \mathrm{dom}(\Gamma)$. We then write $H \vDash V : \Gamma$.

Theorem 2 shows that the terminating evaluation of a well-typed expression in a well-formed environment results in a well-formed environment.

Theorem 2 (Preservation)**.** If $\Sigma; \Gamma \vdash e : B$ and $H \vDash V : \Gamma$ and $V, H \overset{M}{\vdash} e \Downarrow (\ell, H') \mid (w, d)$ then $H' \vDash V : \Gamma$ and $H' \vDash \ell : B$.

The Cost-Free Metric A key aspect of the resource-aware type system is based on the *cost-free* resource metric cf, that is, the metric that assigns zero cost to all steps. We have

$$\mathsf{cf}(x) = 0 \text{ for every } x \in K \ .$$

In the cost-free metric, all evaluation judgements have zero work and zero depth. In the type system we use cost-free judgements to relate the sizes of the results of evaluations to the sizes of the values in the context.

Proposition 3. If $V, H \overset{\mathsf{cf}}{\vdash} e \Downarrow (\ell, H') \mid (w, d)$ then $w = d = 0$.

Simple Metrics and Progress In the reminder of this section, we prove some properties of the evaluation judgement under a simple metric. A *simple metric* M assigns the value 1 to every resource constant, that is,

$$M(x) = 1 \text{ for every } x \in K \ .$$

With a simple metric, work counts the number of evaluation steps. Lemma 1 shows that we can always make one partial evaluation step for a well-typed expression in a well-formed environment. It is used in the induction basis of the proof of Theorem 3.

Lemma 1. Let M be a simple metric. If $\Sigma; \Gamma \vdash e : B$ and $H \vDash V : \Gamma$ then $V, H \overset{M}{\vdash} e \Downarrow \rho \mid (1, 1)$ for some ρ.

Lemma 1 is proved by a simple case destination on syntactic forms. For the composed forms such as let expressions we use the rule E:ABORT for inner expressions.

Intuitively, Theorem 3 states that, in a well-formed environment, well-typed expressions either evaluate to a value or the evaluation uses unbounded work and depth.

Theorem 3 (Progress)**.** Let M be a simple metric, $\Sigma; \Gamma \vdash e : B$, and $H \vDash V : \Gamma$. Then $V, H \overset{M}{\vdash} e \Downarrow (\ell, H') \mid (w, d)$ for some $w, d \in \mathbb{N}$ or for every $n \in \mathbb{N}$ there exist $x, y \in \mathbb{N}$ such that $V, H \overset{M}{\vdash} e \Downarrow \circ \mid (x, n)$ and $V, H \overset{M}{\vdash} e \Downarrow \circ \mid (n, y)$.

To prove Theorem 3, we show by induction on n that if $\Sigma; \Gamma \vdash e : B$, $H \vDash V : \Gamma$, $V, H \overset{M}{\vdash} e \Downarrow \rho \mid (x, n)$ for an x then $\rho = (\ell, H')$ or $V, H \overset{M}{\vdash} e \Downarrow \rho' \mid (x', n + 1)$ for some ρ', x'. We then show the corresponding statement for the work.

A direct consequence of Theorem 3 is that bounds on the depth of programs under a simple metric ensure termination.

3. Amortized Analysis and Parallel Programs

In this section, we give a short introduction into amortized resource analysis for sequential programs (for bounding the work) and then informally describe the main contribution of the article: a multivariate amortized resource analysis for parallel programs (for bounding the depth).

Amortized Resource Analysis Amortized resource analysis is a type-based technique for deriving upper bounds on the resource cost of programs [24]. The advantages of amortized resource analysis are compositionality and efficient type inference that is based on linear programming. The idea is that types are decorated with resource annotations that describe a potential function. Such a potential function maps the sizes of typed data structures to a non-negative rational number. The typing rules ensure that the potential defined by a typing context is sufficient to pay for the evaluation cost of the expression that is typed under this context and for the potential of the result of the evaluation.

The basic idea of amortized analysis is best explained by example. Consider the function $\mathsf{mult} : \mathsf{int} * \mathsf{L}(\mathsf{int}) \to \mathsf{L}(\mathsf{int})$ that takes an integer and an integer list and multiplies each element of the list with the integer. The function mult is implemented as follows.

```
mult(x,ys) = match ys with | nil → nil
                    | (y::ys') → x*y::mult(x,ys')
```

For simplicity, we assume a metric M^* that only counts the number of multiplications performed in an evaluation in this section. We then have $V, H \overset{M^*}{\vdash} \mathsf{mult}(\mathsf{x}, \mathsf{ys}) \Downarrow (\ell, H') \mid (n, n)$ for a well-formed stack V and heap H in which ys points to a list of length n. In short, the work and depth of the evaluation of $\mathsf{mult}(\mathsf{x}, \mathsf{ys})$ is $|\mathsf{ys}|$.

To obtain a bound on the work in type-based amortized resource analysis, we derive a type of the following form.

$$\mathsf{x}{:}\mathsf{int}, \mathsf{ys}{:}L(\mathsf{int}); Q \overset{M^*}{\vdash} \mathsf{mult}(\mathsf{x}, \mathsf{ys}) : (L(\mathsf{int}), Q')$$

Here Q and Q' are *coefficients* of multivariate resource polynomials $p_Q : [\![\mathsf{int} * L(\mathsf{int})]\!] \to \mathbb{Q}_0^+$ and $p_{Q'} : [\![L(\mathsf{int})]\!] \to \mathbb{Q}_0^+$ that map semantic values to non-negative rational numbers. The rules of the type system ensure that for every evaluation context (V, H) that maps x to a number m and ys to a list a, the potential $p_Q(m, a)$ is sufficient to cover the evaluation cost of $\mathsf{mult}(\mathsf{x}, \mathsf{ys})$ and the potential $p_{Q'}(a')$ of the returned list a'. More formally, we have $p_Q(m, a) \geqslant w + p_{Q'}(a')$ if $V, H \overset{M^*}{\vdash} \mathsf{mult}(\mathsf{x}, \mathsf{ys}) \Downarrow (\ell, H') \mid (w, d)$ and ℓ points to the list a' in H'.

In our type system we can for instance derive coefficients Q and Q' that represent the potential functions

$$p_Q(n, a) = |a| \qquad \text{and} \qquad p_{Q'}(a) = 0 \ .$$

The intuitive meaning is that we must have the potential $|\mathsf{ys}|$ available when evaluating $\mathsf{mult}(\mathsf{x}, \mathsf{ys})$. During the evaluation, the potential is used to pay for the evaluation cost and we have no potential left after the evaluation. The soundness theorem of the type system (Theorem 4) states that $|\mathsf{ys}|$ is an upper bound on the work of the evaluation of $\mathsf{mult}(\mathsf{x}, \mathsf{ys})$.

To enable compositionality, we also have to be able to pass potential to the result of an evaluation. Another possible instantiation of Q and Q' would for example result in the following potential functions.

$$p_Q(n, a) = 2{\cdot}|a| \qquad \text{and} \qquad p_{Q'}(a) = |a|$$

The resulting typing can be read as follows. To evaluate $\mathsf{mult}(\mathsf{x}, \mathsf{ys})$ we need the potential $2|\mathsf{ys}|$ to pay for the cost of the evaluation. After the evaluation there is the potential $|\mathsf{mult}(\mathsf{x}, \mathsf{ys})|$ left to pay for future cost in a surrounding program. Such an instantiation would be needed to type the inner function application in the expression $\mathsf{mult}(\mathsf{x}, \mathsf{mult}(\mathsf{z}, \mathsf{ys}))$.

Technically, the coefficients Q and Q' are families that are indexed by sets of base polynomials. The set of base polynomials is determined by the type of the corresponding data. For the type $\mathsf{int} * L(\mathsf{int})$, we have for example $Q = \{q_{(*,[])}, q_{(*,[*])}, q_{(*,[*,*])}, \ldots\}$ and $p_Q(n, a) = q_{(*,[])} + q_{(*,[*])}{\cdot}|a| + q_{(*,[*,*])}{\cdot}\binom{|a|}{2} + \ldots$. This allows us to express multivariate functions such as $m \cdot n$ for two lists of length m and n.

The rules of our type system show how to describe the valid instantiations of the coefficients Q and Q' with a set of linear inequalities. As a result, we can use linear programming to infer resource bounds efficiently.

A more in-depth discussion can be found in the literature on automatic amortized resource analysis [24, 18, 22].

Sequential Composition In a sequential composition $\mathsf{let}\, x = e_1 \,\mathsf{in}\, e_2$, the initial potential, defined by a context and a corresponding annotation (Γ, Q), has to be used to pay for the work of the evaluation of e_1 and the work of the evaluation of e_2. Let us consider a concrete example again.

```
mult2(ys) = let xs = mult(496,ys) in
            let zs = mult(8128,ys) in (xs,zs)
```

The work (and depth) of the evaluation of the expression $\mathsf{mult2(ys)}$ is $2|\mathsf{ys}|$ in the metric M^*. In the type judgement, we express this bound as follows. First, we type the two function applications of mult as before using

$$\mathsf{x{:}int}, \mathsf{ys}{:}L(\mathrm{int}); Q \mathrel{\big|\!\!\frac{M^*}{\quad}} \mathsf{mult(x,ys)} : (L(\mathrm{int}), Q')$$

where $p_Q(n,a) = |a|$ and $p_{Q'}(a) = 0$. In the type judgement

$$\mathsf{ys}{:}L(\mathrm{int}); R \mathrel{\big|\!\!\frac{M^*}{\quad}} \mathsf{mult2(ys)} : (L(\mathrm{int}) * L(\mathrm{int}), R')$$

we require that $p_R(a) \geqslant p_Q(a) + p_Q(a)$, that is, the initial potential (defined by the coefficients R) has to be shared in the two sequential branches. Such a sharing can still be expressed with linear constraints such as $r_{[*]} \geqslant q_{(*,[*])} + q_{(*,[*])}$. A valid instantiation of R would thus correspond to the potential function $p_R(a) = 2|a|$. With this instantiation, the previous typing reflects the bound $2|\mathsf{ys}|$ for the evaluation of $\mathsf{mult2(ys)}$.

A slightly more involved example is the function $\mathsf{dyad} : L(\mathrm{int}) * L(\mathrm{int}) \to L(L(\mathrm{int}))$ which computes the dyadic product of two integer lists.

```
dyad (u,v) = match u with | nil  →  nil
     | (x::xs)  →  let x’ = mult(x,v) in
                   let xs’ = dyad(xs,v) in x’::xs’;
```

Using the metric M^* that counts multiplications, multivariate resource analysis for sequential programs derives the bound $|\mathsf{u}|{\cdot}|\mathsf{v}|$. In the cons branch of the pattern match, we have the potential $|\mathsf{xs}|{\cdot}|\mathsf{v}| + |\mathsf{v}|$ which is shared to pay for the cost $|\mathsf{v}|$ of $\mathsf{mult(x,v)}$ and the cost $|\mathsf{xs}|{\cdot}|\mathsf{v}|$ of $\mathsf{dyad(xs,v)}$.

Moving multivariate potential through a program is not trivial; especially in the presence of nested data structures like trees of lists. To give an idea of the challenges, consider the expression e that is defined as follows.

```
let xs = mult(496,ys) in
let zs = append(ys,ys) in dyad(xs,zs)
```

The depth of evaluating e in the metric M^* is bounded by $|\mathsf{ys}| + 2|\mathsf{ys}|^2$. Like in the previous example, we express this in amortized resource analysis with the initial potential $|\mathsf{ys}| + 2|\mathsf{ys}|^2$. This potential has to be shared to pay for the cost of the evaluations of $\mathsf{mult(496,ys)}$ (namely $|\mathsf{ys}|$) and $\mathsf{dyad(xs,zs)}$ (namely $2|\mathsf{ys}|^2$). However, the type of dyad requires the quadratic potential $|\mathsf{xs}|{\cdot}|\mathsf{zs}|$. In this simple example, it is easy to see that $|\mathsf{xs}|{\cdot}|\mathsf{zs}| = 2|\mathsf{ys}|^2$. But in general, it is not straightforward to compute such a conversion of potential in an automatic analysis system, especially for nested data structures and super-linear size changes. The type inference for multivariate amortized resource analysis for sequential programs can analyze such programs efficiently [22].

Parallel Composition The insight of this paper is that the potential method works also well to derive bounds on parallel evaluations. The main challenge in the development of an amortized resource analysis for parallel evaluations is to ensure the same high compositionality as in sequential amortized resource analysis.

The basic idea of our new analysis system is to allow each branch in a parallel evaluation to use all the available potential without sharing. Consider for example the previously defined function $\mathsf{mult2}$ in which we evaluate the two applications of mult in parallel.

```
mult2par(ys) = par xs = mult(496,ys)
               and zs = mult(8128,ys) in (xs,zs)
```

Since the depth of $\mathsf{mult(n,ys)}$ is $|\mathsf{ys}|$ for every n and the two applications of mult are evaluated in parallel, the depth of the evaluation of $\mathsf{mult2par(ys)}$ is $|\mathsf{ys}|$ in the metric M^*.

In the type judgement, we type the two function applications of mult as in the sequential case in which

$$\mathsf{x{:}int}, \mathsf{ys}{:}L(\mathrm{int}); Q \mathrel{\big|\!\!\frac{M^*}{\quad}} \mathsf{mult(x,ys)} : (L(\mathrm{int}), Q')$$

such that $p_Q(n,a) = |a|$ and $p_{Q'}(a) = 0$. In the type judgement

$$\mathsf{ys}{:}L(\mathrm{int}); R \mathrel{\big|\!\!\frac{M^*}{\quad}} \mathsf{mult2par(ys)} : (L(\mathrm{int}) * L(\mathrm{int}), R')$$

for $\mathsf{mult2par}$ we require however only that $p_R(a) \geqslant p_Q(a)$. In this way, we express that the initial potential defined by the coefficients R has to be sufficient to cover the cost of each parallel branch. Consequently, a possible instantiation of R corresponds to the potential function $p_R(a) = |a|$.

In the function dyad, we can replace the sequential computation of the inner lists of the result by a parallel computation in which we perform all calls to the function mult in parallel. The resulting function is $\mathsf{dyad_par}$.

```
dyad_par (u,v) = match u with | nil  →  nil
               | (x::xs)  →  par x’ = mult(x,v)
                             and xs’ = dyad_par(xs,v)
                             in x’::xs’;
```

The depth of $\mathsf{dyad_par}$ is $|\mathsf{v}|$. In the type-based amortized analysis, we hence start with the initial potential $|\mathsf{v}|$. In the cons branch of the pattern match, we can use the initial potential to pay for both, the cost $|\mathsf{v}|$ of $\mathsf{mult(x,v)}$ and the cost $|\mathsf{v}|$ of the recursive call $\mathsf{dyad(xs,v)}$ without sharing the initial potential.

Unfortunately, the compositionality of the sequential system is not preserved by this simple idea. The problem is that the naive reuse of potential that is passed through parallel branches would break the soundness of the system. To see why, consider the following function.

```
mult4(ys) = par xs = mult(496,ys)
  and zs = mult(8128,ys) in (mult(5,xs), mult(10,zs))
```

Recall, that a valid typing for $\mathsf{xs} = \mathsf{mult(496,ys)}$ could take the initial potential $2|\mathsf{ys}|$ and assign the potential $|\mathsf{xs}|$ to the result. If we would simply reuse the potential $2|\mathsf{ys}|$ to type the second application of mult in the same way then we would have the potential $|\mathsf{xs}| + |\mathsf{zs}|$ after the parallel branches. This potential could then be used to pay for the cost of the remaining two applications of mult. We have now verified the unsound bound $2|\mathsf{ys}|$ on the depth of the evaluation of the expression $\mathsf{mult4(ys)}$ but the depth of the evaluation is $3|\mathsf{ys}|$.

The problem in the previous reasoning is that we doubled the part of the initial potential that we passed on for later use in the two parallel branches of the parallel composition. To fix this problem, we need a separate analysis of the sizes of data structures and the cost of parallel evaluations.

In this paper, we propose to use cost-free type judgements to reason about the size changes in parallel branches. Instead of simply using the initial potential in both parallel branches, we share the potential between the two branches but analyze the two branches twice. In the first analysis, we only pay for the resource consumption of the first branch. In the second, analysis we only pay for resource consumption of the second branch.

A cost-free type judgement is like any other type judgement in amortized resource analysis but uses the cost-free metric cf that assigns zero cost to every evaluation step. For example, a cost-free typing of the function $\mathsf{mult(ys)}$ would express that the initial

potential can be passed to the result of the function. In the cost-free typing judgement

$$\mathsf{x}{:}\text{int}, \mathsf{ys}{:}L(\text{int}); Q \overset{\mathsf{cf}}{\vdash} \mathsf{mult}(\mathsf{x}, \mathsf{ys}) : (L(\text{int}), Q')$$

a valid instantiation of Q and Q' would correspond to the potential functions

$$p_Q(n, a) = |a| \qquad \text{and} \qquad p_{Q'}(a) = |a| \; .$$

The intuitive meaning is that in a call $\mathsf{zs} = \mathsf{mult}(\mathsf{x}, \mathsf{ys})$, the initial potential $|ys|$ can be transformed to the potential $|zs|$ of the result.

Using this cost-free typing we can now correctly reason about the depth of the evaluation of mult4. We start with the initial potential $3|\mathsf{ys}|$ and have to consider two cases in the parallel binding. In the first case, we have to pay only for resource cost of $\mathsf{mult}(496, \mathsf{ys})$. So we share the initial potential and use $2|\mathsf{ys}|$: $|\mathsf{ys}|$ to pay the cost of $\mathsf{mult}(496, \mathsf{ys})$ and $|\mathsf{ys}|$ to assign the potential $|\mathsf{xs}|$ to the result of the application. The reminder $|\mathsf{ys}|$ of the initial potential is used in a cost-free typing of $\mathsf{mult}(8128, \mathsf{ys})$ where we simply assign the potential $|\mathsf{zs}|$ to the result of the function without paying any evaluation cost. In the second case, we derive a similar typing in which the roles of the two function applications are switched. In both cases, we started with the potential $3|\mathsf{ys}|$ and ended with the potential $|\mathsf{xs}| + |\mathsf{zs}|$. We can use it to pay for the two remaining calls of mult and have verified the correct bound.

In the univariate case, using the notation from [24, 18], we could formulate the type rule for parallel composition as follows. Here, the coefficients Q are not globally attached to a type or context but appear locally at list types such as $L^q(\text{int})$. The sharing operator $\Gamma \curlyvee (\Gamma_1, \Gamma_2, \Gamma_3)$ requires the sharing of the potential in the context Γ in the contexts Γ_1,Γ_2 and Γ_3. For instance, we have $x{:}L^6(\text{int}) \curlyvee (x{:}L^2(\text{int}), x{:}L^3(\text{int}), x{:}L^1(\text{int}))$.

$$\frac{\begin{array}{c}\Gamma \curlyvee (\Delta_1, \Gamma_2, \Gamma') \qquad \Gamma \curlyvee (\Gamma_1, \Delta_2, \Gamma') \\ \Gamma_1 \overset{M}{\vdash} e_1 : A_1 \qquad \Delta_2 \overset{\mathsf{cf}}{\vdash} e_2 : A_2 \qquad \Delta_1 \overset{\mathsf{cf}}{\vdash} e_1 : A_1 \\ \Gamma_2 \overset{M}{\vdash} e_2 : A_2 \qquad \Gamma', x_1{:}A_1, x_2{:}A_2 \overset{M}{\vdash} e : B\end{array}}{\Gamma \overset{M}{\vdash} \mathsf{par}\, x_1 = e_1 \,\mathsf{and}\, x_2 = e_2 \,\mathsf{in}\, e : B}$$

In the rule, the initial potential Γ is shared twice using the sharing operator $\curlyvee$. First, to pay the cost of evaluating e_2 and e, and to pass potential to x_1 using the cost-free type judgement $\Delta_1 \overset{\mathsf{cf}}{\vdash} e_1 : A_1$. Second, to pay the cost of evaluation e_1 and e, and to pass potential to x_2 via the judgement $\Delta_2 \overset{\mathsf{cf}}{\vdash} e_2 : A_2$.

In this work, we generalize the idea to multivariate resource polynomials for which we also have to deal with mixed potential such as $|x_1|{\cdot}|x_2|$. The approach features the same high compositionality as the sequential version of the analysis. As the experiments in Section 7 show, the analysis works well for many typical example programs.

The use of cost-free typings to separate the reasoning about size changes of data structures and resource cost in amortized analysis has applications that go beyond parallel evaluations. Similar problems arise in sequential (and parallel) programs when deriving bounds for non-additive cost such as stack-space usage or recursion depth. We envision that the developed technique can be used to derive bounds for these cost measures too.

Other Forms of Parallelism The binary parallel binding is a simple yet powerful form of parallelism. However, it is (for example) not possible to directly implement NESL's model of sequences that allows to perform an operation for every element in the sequence in constant depth. The reason is that the parallel binding would introduce a linear overhead.

Nevertheless it is possible to introduce another binary parallel binding that is semantically equivalent except that it has zero depth cost. We can then analyze more powerful parallelism primitives by translating them into code that uses this cost-free parallel binding. To demonstrate such a translation, we implemented NESL's [6] parallel sequence comprehensions in RAML. For instance, the following RAML expression computes the list that contains the squares of all negative numbers in the list ℓ.

$$\{\, x * x \,:\, x \,\mathsf{in}\, \ell \mid x < 0 \,\}$$

The work of the evaluation is linear in ℓ and the depth of the evaluation is constant. The analysis automatically computes precise bounds. More explanations can be found in Section 6.

4. Resource Polynomials and Annotated Types

In this section, we introduce multivariate resource polynomials and annotated types. Our goal is to systematically describe the potential functions that map data structures to non-negative rational numbers. Multivariate resource polynomials are a generalization of non-negative linear combinations of binomial coefficients. They have properties that make them ideal for the generation of succinct linear constraint systems in an automatic amortized analysis. The presentation might appear quite low level but this level of detail is necessary to describe the linear constraints in the type rules.

Two main advantages of resource polynomials are that they can express more precise bounds than non-negative linear-combinations of standard polynomials and that they can succinctly describe common size changes of data that appear in construction and destruction of data. More explanations can be found in the previous literature on multivariate amortized resource analysis [21, 22].

4.1 Resource Polynomials

A resource polynomial maps a value of some data type to a nonnegative rational number. Potential functions and thus resource bounds are always resource polynomials.

Base Polynomials For each data type A we first define a set $P(A)$ of functions $p : [\![A]\!] \to \mathbb{N}$ that map values of type A to natural numbers. These *base polynomials* form a basis (in the sense of linear algebra) of the resource polynomials for type A. The resource polynomials for type A are then given as nonnegative rational linear combinations of the base polynomials. We define $P(A)$ as follows.

$$\begin{aligned} P(\text{int}) &= \{a \mapsto 1\} \\ P(A_1 * A_2) &= \{(a_1, a_2) \mapsto p_1(a_1) \cdot p_2(a_2) \mid p_i \in P(A_i)\} \\ P(L(A)) &= \{\Sigma\Pi[p_1, \ldots, p_k] \mid k \in \mathbb{N}, p_i \in P(A)\} \end{aligned}$$

In the last clause we have

$$\Sigma\Pi[p_1, \ldots, p_k]([a_1, \ldots, a_n]) = \sum_{1 \leqslant j_1 < \cdots < j_k \leqslant n} \; \prod_{1 \leqslant i \leqslant k} p_i(a_{j_i}) \; .$$

Every set $P(A)$ contains the constant function $v \mapsto 1$. For lists $L(A)$ this arises for $k = 0$ (one element sum, empty product).

For example, the function $\ell \mapsto \binom{|\ell|}{k}$ is in $P(L(A))$ for every $k \in \mathbb{N}$; simply take $p_1 = \ldots = p_k = 1$ in the definition of $P(L(A))$. The function $(\ell_1, \ell_2) \mapsto \binom{|\ell_1|}{k_1} \cdot \binom{|\ell_2|}{k_2}$ is in $P(L(A) * L(B))$ for every $k_1, k_2 \in \mathbb{N}$ and $[\ell_1, \ldots, \ell_n] \mapsto \sum_{1 \leqslant i < j \leqslant n} \binom{|\ell_i|}{k_1} \cdot \binom{|\ell_j|}{k_2} \in P(L(L(A)))$ for every $k_1, k_2 \in \mathbb{N}$.

Resource Polynomials A *resource polynomial* $p : [\![A]\!] \to \mathbb{Q}_0^+$ for a data type A is a non-negative linear combination of base polynomials, i.e.,

$$p = \sum_{i=1,\ldots,m} q_i \cdot p_i$$

for $q_i \in \mathbb{Q}_0^+$ and $p_i \in P(A)$. We write $R(A)$ for the set of resource polynomials for A.

An instructive, but not exhaustive, example is given by $R_n = R(L(\text{int}) * \cdots * L(\text{int}))$. The set R_n is the set of linear combinations

of products of binomial coefficients over variables $x_1, \ldots, x_n$, that is, $R_n = \{\sum_{i=1}^{m} q_i \prod_{j=1}^{n} \binom{x_j}{k_{ij}} \mid q_i \in \mathbb{Q}_0^+, m \in \mathbb{N}, k_{ij} \in \mathbb{N}\}$. Concrete examples that illustrate the definitions follow in the next subsection.

4.2 Annotated Types

To relate type annotations in the type system to resource polynomials, we introduce names (or indices) for base polynomials. These names are also helpful to intuitively explain the base polynomials of a given type.

Names For Base Polynomials To assign a unique name to each base polynomial we define the *index set* $\mathcal{I}(A)$ to denote resource polynomials for a given data type A. Essentially, $\mathcal{I}(A)$ is the meaning of A with every atomic type replaced by the *unit index* $\circ$.

$$\mathcal{I}(\text{int}) = \{\circ\}$$
$$\mathcal{I}(A_1 * A_2) = \{(i_1, i_2) \mid i_1 \in \mathcal{I}(A_1) \text{ and } i_2 \in \mathcal{I}(A_2)\}$$
$$\mathcal{I}(L(A)) = \{[i_1, \ldots, i_k] \mid k \geqslant 0, i_j \in \mathcal{I}(A)\}$$

The *degree* $\deg(i)$ of an index $i \in \mathcal{I}(A)$ is defined as follows.

$$\deg(\circ) = 0$$
$$\deg(i_1, i_2) = \deg(i_1) + \deg(i_2)$$
$$\deg([i_1, \ldots, i_k]) = k + \deg(i_1) + \cdots + \deg(i_k)$$

Let $\mathcal{I}_k(A) = \{i \in \mathcal{I}(A) \mid \deg(i) \leqslant k\}$. The indices $i \in \mathcal{I}_k(A)$ are an enumeration of the base polyonomials $p_i \in P(A)$ of degree at most k. For each $i \in \mathcal{I}(A)$, we define a base polynomial $p_i \in P(A)$ as follows: If $A = \text{int}$ then

$$p_\circ(v) = 1 \ .$$

If $A = (A_1 * A_2)$ is a pair type and $v = (v_1, v_2)$ then

$$p_{(i_1,i_2)}(v) = p_{i_1}(v_1) \cdot p_{i_2}(v_2) \ .$$

If $A = L(B)$ is a list type and $v \in [\![L(B)]\!]$ then

$$p_{[i_1,\ldots,i_m]}(v) = \Sigma\Pi[p_{i_1}, \ldots, p_{i_m}](v) \ .$$

We use the notation 0_A (or just 0) for the index in $\mathcal{I}(A)$ such that $p_{0_A}(a) = 1$ for all a. We have $0_{\text{int}} = \circ$ and $0_{(A_1 * A_2)} = (0_{A_1}, 0_{A_2})$ and $0_{L(B)} = []$. If $A = L(B)$ for a data type B then the index $[0, \ldots, 0] \in \mathcal{I}(A)$ of length n is denoted by just n. We identify the index (i_1, i_2, i_3, i_4) with the index $(i_1, (i_2, (i_3, i_4)))$.

For a list $i = [i_1, \ldots, i_k]$ we write $i_0 :: i$ to denote the list $[i_0, i_1, \ldots, i_k]$. Furthermore, we write ii' for the concatenation of two lists i and i'.

Lemma 2. If $p, p' \in R(A)$ then $p + p', p \cdot p' \in R(A)$, and $\deg(p + p') = \max\{\deg(p), \deg(p')\}$ and $\deg(p \cdot p') = \deg(p) + \deg(p')$.

By linearity it suffices to show this lemma for base polynomials. This is done by induction on A.

Lemma 3. Let $a \in [\![A]\!]$ and $\ell \in [\![L(A)]\!]$. Let $i_0, \ldots, i_k \in \mathcal{I}(A)$ and $k \geqslant 0$. Then $p_{[i_0,i_1,\ldots,i_k]}([]) = 0$ and $p_{[i_0,i_1,\ldots,i_k]}(a :: \ell) = p_{i_0}(a) \cdot p_{[i_1,\ldots,i_k]}(\ell) + p_0(a) \cdot p_{[i_0,i_1,\ldots,i_k]}(\ell)$.

To prove this, one decomposes the sum in the definition of $p_{[i_0,i_1,\ldots,i_k]}(a :: \ell)$ into two summands, one corresponding to the case where the first position j_1 equals one, thus hits a and where it is greater than one, thus a is not considered. Note that $p_0(a) = 1$; this factor is there to achieve the format of the resource polynomials for types like $A * L(A)$.

Lemma 4 characterizes concatenations of lists (written as juxtaposition) as they will occur in the construction of lists. Note that, e.g., $\text{elems}(\text{node}(a, t_1, t_2)) = a :: \text{elems}(t_1)\,\text{elems}(t_2)$.

Lemma 4. Let $\ell_1, \ell_2 \in [\![L(A)]\!]$. Then it holds that $\ell_1 \ell_2 \in [\![L(A)]\!]$ and $p_{[i_1,\ldots,i_k]}(\ell_1 \ell_2) = \sum_{t=0}^{k} p_{[i_1,\ldots,i_t]}(\ell_1) \cdot p_{[i_{t+1},\ldots,i_k]}(\ell_2)$.

This can be proved by induction on the length of ℓ_1 using Lemma 3 or else by a decomposition of the defining sum according to which indices hit the first list and which ones hit the second.

Examples First consider the type int. The index set $\mathcal{I}(\text{int}) = \{\circ\}$ only contains the unit element because the only base polynomial for the type int is the constant polynomial $p_\circ : \mathbb{Z} \to \mathbb{N}$ that maps every integer to 1, that is, $p_\circ(n) = 1$ for all $n \in \mathbb{Z}$. In terms of resource-cost analysis this implies that the resource polynomials can not represent cost that depends on the value of an integer.

Now consider the type $L(\text{int})$. The index set for lists of integers is $\mathcal{I}(L(\text{int})) = \{[], [\circ], [\circ, \circ], \ldots\}$, the set of lists of unit indices $\circ$. The base polynomial $p_{[]} : [\![L(\text{int})]\!] \to \mathbb{N}$ is defined as $p_{[]}([a_1, \ldots, a_n]) = 1$ (one element sum and empty product). More interestingly, we have $p_{[\circ]}([a_1, \ldots, a_n]) = \sum_{1 \leqslant j \leqslant n} 1 = n$ and $p_{[\circ,\circ]}([a_1, \ldots, a_n]) = \sum_{1 \leqslant j_1 < j_2 \leqslant n} 1 = \binom{n}{2}$. In general, if $i_k = [\circ, \ldots, \circ]$ is as list with k unit indices then $p_{i_k}([a_1, \ldots, a_n]) = \sum_{1 \leqslant j_1 < \cdots < j_k \leqslant n} 1 = \binom{n}{k}$. The intuition is that the base polynomial $p_{i_k}([a_1, \ldots, a_n])$ describes a constant resource cost that arises for every ordered k-tuple $(a_{j_1}, \ldots, a_{j_n})$.

Finally, consider the type $L(L(\text{int}))$ of lists of lists of integers. The corresponding index set is $\mathcal{I}(L(L(\text{int}))) = \{[]\} \cup \{[i] \mid i \in \mathcal{I}(L(\text{int}))\} \cup \{[i_1, i_2] \mid i_1, i_2 \in \mathcal{I}(L(\text{int}))\} \cup \cdots$. Again we have $p_{[]} : [\![L(L(\text{int}))]\!] \to \mathbb{N}$ and $p_{[]}([a_1, \ldots, a_n]) = 1$. Moreover we also get the binomial coefficients again: If the index $i_k = [[], \ldots, []]$ is as list of k empty lists then $p_{i_k}([a_1, \ldots, a_n]) = \sum_{1 \leqslant j_1 < \cdots < j_k \leqslant n} 1 = \binom{n}{k}$. This describes a cost that would arise in a program that computes something of constant cost for tuples of inner lists (e.g., sorting with respect to the smallest head elements). However, the base polynomials can also refer to the lengths of the inner lists. For instance, we have $p[[\circ, \circ]]([a_1, \ldots, a_n]) = \sum_{1 \leqslant i \leqslant n} \binom{|a_i|}{2}$, which represents a quadratic cost for every inner list (e.g, sorting the inner lists). This is not to be confused with the base polynomial $p_{[\circ,\circ]}([a_1, \ldots, a_n]) = \sum_{1 \leqslant i < j \leqslant n} |a_i||a_j|$, which can be used to account for the cost of the comparisons in a lexicographic sorting of the outer list.

Annotated Types and Potential Functions We use the indices and base polynomials to define type annotations and resource polynomials. We then give examples to illustrate the definitions.

A *type annotation* for a data type A is defined to be a family

$$Q_A = (q_i)_{i \in \mathcal{I}(A)} \text{ with } q_i \in \mathbb{Q}_0^+$$

We say Q_A is of *degree (at most)* k if $q_i = 0$ for every $i \in \mathcal{I}(A)$ with $\deg(i) > k$. An *annotated data type* is a pair (A, Q_A) of a data type A and a type annotation Q_A of some degree k.

Let H be a heap and let ℓ be a location with $H \vDash \ell \mapsto a : A$ for a data type A. Then the type annotation Q_A defines the *potential*

$$\Phi_H(\ell{:}(A, Q_A)) = \sum_{i \in \mathcal{I}(A)} q_i \cdot p_i(a)$$

If $a \in [\![A]\!]$ and Q is a type annotation for A then we also write $\Phi(a : (A, Q))$ for $\sum_i q_i p_i(a)$.

Let for example, $Q = (q_i)_{i \in L(\text{int})}$ be an annotation for the type $L(\text{int})$ and let $q_{[]} = 2$, $q_{[\circ]} = 2.5$, $q_{[\circ,\circ,\circ]} = 8$, and $q_i = 0$ for all other $i \in \mathcal{I}(L(\text{int}))$. The we have $\Phi([a_1, \ldots, a_n] : (L(\text{int}), Q)) = 2 + 2.5n + 8\binom{n}{3}$.

Examples The simplest annotated types are those for atomic data types like integers. The indices for int are $\mathcal{I}(\text{int}) = \{\circ\}$ and thus each type annotation has the form (int, q_0) for a $q_0 \in \mathbb{Q}_0^+$. It defines the constant potential function $\Phi_H(v{:}(\text{int}, q_0)) = q_0$. Similarly, tuples of atomic types feature a single index of the form $(\circ, \ldots, \circ)$ and a constant potential function defined by some $q_{(\circ,\ldots,\circ)} \in \mathbb{Q}_0^+$.

More interesting examples are lists of atomic types like, for example, $L(\text{int})$. The set of indices of degree k is then $\mathcal{I}_k(L(\text{int})) =$

$\{[], [\circ], [\circ,\circ], \ldots, [\circ, ..., \circ]\}$ where the last list contains k unit elements. Since we identify a list of i unit elements with the integer i we have $\mathcal{I}_k(L(\text{int})) = \{0, 1, \ldots, k\}$. Consequently, annotated types have the form $(L(\text{int}), (q_0, \ldots, q_k))$ for $q_i \in \mathbb{Q}_0^+$. The defined potential function is $\Phi([a_1, \ldots, a_n]{:}(L(\text{int}), (q_0, \ldots, q_n)) = \sum_{0 \leqslant i \leqslant k} q_i \binom{n}{i}$.

The next example is the type $(L(\text{int}) * L(\text{int}))$ of pairs of lists of integers. The set of indices of degree k is $\mathcal{I}_k(L(\text{int}) * L(\text{int})) = \{(i,j) \mid i + j \leqslant k\}$ if we identify lists of units with their lengths as usual. Annotated types are then of the from $((L(\text{int}) * L(\text{int})), Q)$ for a triangular $k \times k$ matrix Q with non-negative rational entries. If $\ell_1 = [a_1, \ldots, a_n]$, $\ell_2 = [b_1, \ldots, b_m]$ are two lists then the potential function is $\Phi((\ell_1, \ell_2), ((L(\text{int}) * L(\text{int})), (q_{(i,j)}))) = \sum_{0 \leqslant i+j \leqslant k} q_{(i,j)} \binom{n}{i}\binom{m}{j}$.

The Potential of a Context For use in the type system we need to extend the definition of resource polynomials to typing contexts. We treat a context like a tuple type.

Let $\Gamma = x_1{:}A_1, \ldots, x_n{:}A_n$ be a typing context and let $k \in \mathbb{N}$. The index set $\mathcal{I}(\Gamma)$ is defined through

$$\mathcal{I}(\Gamma) = \{(i_1, \ldots, i_n) \mid i_j \in \mathcal{I}(A_j)\}.$$

The degree of $i = (i_1, \ldots, i_n) \in \mathcal{I}(\Gamma)$ is defined through $\deg(i) = \deg(i_1) + \cdots + \deg(i_n)$. As for data types, we define $I_k(\Gamma) = \{i \in I(\Gamma) \mid \deg(i) \leqslant k\}$. A *type annotation* Q for Γ is a family

$$Q = (q_i)_{i \in \mathcal{I}_k(\Gamma)} \text{ with } q_i \in \mathbb{Q}_0^+.$$

We denote a *resource-annotated context* with $\Gamma; Q$. Let H be a heap and V be a stack with $H \vDash V : \Gamma$ where $H \vDash V(x_j) {\mapsto} a_{x_j} : \Gamma(x_j)$. The potential of $\Gamma; Q$ with respect to H and V is

$$\Phi_{V,H}(\Gamma; Q) = \sum_{(i_1,\ldots,i_n) \in \mathcal{I}_k(\Gamma)} q_{\vec{\imath}} \prod_{j=1}^{n} p_{i_j}(a_{x_j})$$

In particular, if $\Gamma = \varnothing$ then $\mathcal{I}_k(\Gamma) = \{()\}$ and $\Phi_{V,H}(\Gamma; q_{()}) = q_{()}$. We sometimes also write q_0 for $q_{()}$.

5. Type System for Bounds on the Depth

In this section, we formally describe the novel resource-aware type system. We focus on the type judgement and explain the rules that are most important for handling parallel evaluation. .

The main theorem of this section proves the soundness of the type system with respect to the depths of evaluations as defined by the operational big-step semantics. The soundness holds for terminating and non-terminating evaluations.

Type Judgments The typing rules in Figure 4 define a *resource-annotated typing judgment* of the form

$$\Sigma; \Gamma; \{Q_1, \ldots, Q_n\} \overset{M}{\vdash} e : (A, Q')$$

where M is a metric, $n \in \{1, 2\}$, e is an expression, Σ is a resource-annotated signature (see below), $(\Gamma; Q_i)$ is a resource-annotated context for every $i \in \{1, \ldots, n\}$, and (A, Q') is a resource-annotated data type. The intended meaning of this judgment is the following. If there are more than $\Phi(\Gamma; Q_i)$ resource units available for every $i \in \{1, \ldots, n\}$ then this is sufficient to pay for the depth of the evaluation of e under the metric M. In addition, there are more than $\Phi(v{:}(A, Q'))$ resource units left if e evaluates to a value v.

In outermost judgements, we are only interested in the case where $n = 1$ and the judgement is equivalent to the similar judgement for sequential programs [22]. The form in which $n = 2$ is introduced in the type rule E:PAR for parallel bindings and eliminated by multiple applications of the sharing rule E:SHARE (more explanations follow).

The type judgement is affine in the sense that every variable in a context Γ can be used at most once in the expression e. Of course, we have to also deal with expressions in which a variable occurs more than once. To account for multiple variable uses we use the sharing rule T:SHARE that doubles a variable in a context without increasing the potential of the context.

As usual Γ_1, Γ_2 denotes the union of the contexts Γ_1 and Γ_2 provided that $\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \varnothing$. We thus have the implicit side condition $\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \varnothing$ whenever Γ_1, Γ_2 occurs in a typing rule. Especially, writing $\Gamma = x_1{:}A_1, \ldots, x_k{:}A_k$ means that the variables x_i are pairwise distinct.

Programs with Annotated Types *Resource-annotated first-order types* have the form $(A, Q) \to (B, Q')$ for annotated data types (A, Q) and (B, Q'). A *resource-annotated signature* Σ is a finite, partial mapping of function identifiers to *sets of* resource-annotated first-order types. A program with resource-annotated types for the metric M consists of a resource-annotated signature Σ and a family of expressions with variables identifiers $(e_f, y_f)_{f \in \text{dom}(\Sigma)}$ such that $\Sigma; y_f{:}A; Q \overset{M}{\vdash} e_f : (B, Q')$ for every function type $(A, Q) \to (B, Q') \in \Sigma(f)$.

Notations Families that describe type and context annotations are denoted with upper case letters $Q, P, R, \ldots$. We use the convention that the elements of the families are the corresponding lower case letters with corresponding superscripts, i.e., $Q = (q_i)_{i \in I}$, $Q' = (q'_i)_{i \in I}$, and $Q^x = (q^x_i)_{i \in I}$.

Let Q, P, R be annotations with the same index set I. We write $Q \leqslant P$ if $q_i \leqslant p_i$ for every $i \in I$. For $c \in \mathbb{Q}$ we write $Q = Q' + c$ to state that $q_{\vec{0}} = q'_{\vec{0}} + c \geqslant 0$ and $q_i = q'_i$ for $i \neq \vec{0} \in I$. Let $\Gamma = \Gamma_1, \Gamma_2$ be a context, let $i = (i_1, \ldots, i_k) \in \mathcal{I}(\Gamma_1)$ and $j = (j_1, \ldots, j_l) \in \mathcal{I}(\Gamma_2)$. We write (i, j) to denote the index $(i_1, \ldots, i_k, j_1, \ldots, j_l) \in \mathcal{I}(\Gamma)$.

Let Q be an annotation for a context Γ_1, Γ_2. For $j \in \mathcal{I}(\Gamma_2)$ we define the *projection* $\pi^{\Gamma_1}_j(Q)$ of Q to Γ_1 to be the annotation Q' with $q'_i = q_{(i,j)}$. We use the same notation for projections to Γ_2.

Proposition 4. Let $\Gamma, x{:}A; Q$ be an annot. context, $H \vDash V : \Gamma, x{:}A$, and $H \vDash V(x) {\mapsto} a : A$. Then it is true that $\Phi_{V,H}(\Gamma, x{:}A; Q) = \sum_{j \in \mathcal{I}(A)} \Phi_{V,H}(\Gamma; \pi^{\Gamma}_j(Q)) \cdot p_j(a)$.

Additive Shift A key notion in the type system is the *additive shift* that is used to assign potential to typing contexts that result from a pattern match or from the application of a constructor. We first define the additive shift, then illustrate the definition with examples and finally state the soundness of the operation.

Let $\Gamma, y{:}L(A)$ be a context and let $Q = (q_i)_{i \in \mathcal{I}(\Gamma, y:L(A))}$ be a context annotation of degree k. The *additive shift for lists* $\lhd_L(Q)$ of Q is an annotation $\lhd_L(Q) = (q'_i)_{i \in \mathcal{I}(\Gamma, x:A, xs:L(A))}$ of degree k for a context $\Gamma, x{:}A, xs{:}L(A)$ that is defined through

$$q'_{(i,j,\ell)} = \begin{cases} q_{(i,j\,::\,\ell)} + q_{(i,\ell)} & j = 0 \\ q_{(i,j\,::\,\ell)} & j \neq 0 \end{cases}$$

The definition of the additive shift is essential for the type system. It is explained and illustrated with examples in previous articles on amortized resource analysis [18, 22].

Lemmas 5 states the soundness of the shift operation.

Lemma 5. Let $\Gamma, \ell{:}L(A); Q$ be an annotated context, $H \vDash V : \Gamma, \ell{:}L(A)$, $H(\ell) = (v_1, \ell')$ and let $V' = V[x_h \mapsto v_1, x_t \mapsto \ell']$. Then $H \vDash V' : \Gamma, x_h{:}A, x_t{:}L(A)$ and $\Phi_{V,H}(\Gamma, \ell{:}L(A); Q) = \Phi_{V',H}(\Gamma, x_h{:}A, x_t{:}L(A); \lhd_L(Q))$.

This is a consequence of Lemma 3. One takes the linear combination of instances of its second equation and regroups the right hand side according to the base polynomials for the resulting context.

Sharing Let $\Gamma, x_1{:}A, x_2{:}A; Q$ be an annotated context. The *sharing operation* $\curlyvee Q$ defines an annotation for a context of the form $\Gamma, x{:}A$. It is used when the potential is split between multiple occur-

$$\boxed{\Sigma; \Gamma; \{Q_1, \ldots, Q_n\} \vdash^{M} e : (A, Q')}$$

Under signature Σ and metric M, expression e has annotated type (A, Q') in the annotated contexts $\Gamma; Q_1, \ldots, \Gamma; Q_{n-1}$, and $\Gamma; Q_n$.

$$\frac{Q = Q' + M^{\mathsf{var}}}{\Sigma; x{:}A; \{Q\} \vdash^{M} x : (A, Q')}\ (\text{T:VAR}) \qquad \frac{n \in \mathbb{Z} \qquad Q = Q' + M^{\mathsf{const}}}{\Sigma; \cdot; \{Q\} \vdash^{M} n : (\mathsf{int}, Q')}\ (\text{T:CONST}) \qquad \frac{P + M^{\mathsf{app}} = Q \qquad (A, P) \to (A', Q') \in \Sigma(f)}{\Sigma; x{:}A; \{Q\} \vdash^{M} f(x) : (A', Q')}\ (\text{T:APP})$$

$$\frac{Q = Q' + M^{\mathsf{nil}}}{\Sigma; \cdot; \{Q\} \vdash^{M} \mathsf{nil} : (L(A), Q')}\ (\text{T:NIL}) \qquad \frac{Q = \lhd_L(Q') + M^{\mathsf{cons}}}{\Sigma; x_1{:}A, x_2{:}L(A); \{Q\} \vdash^{M} \mathsf{cons}(x_1, x_2) : (L(A), Q')}\ (\text{T:CONS})$$

$$\frac{\Sigma; \Gamma; \{R\} \vdash^{M} e_1 : (B, Q') \qquad R + M^{\mathsf{matN}} = \pi_0^{\Gamma}(Q) \qquad \Sigma; \Gamma, x_1{:}A, x_2{:}L(A); \{P\} \vdash^{M} e_2 : (B, Q') \qquad P + M^{\mathsf{matL}} = \lhd_L(Q)}{\Sigma; \Gamma, x{:}L(A); \{Q\} \vdash^{M} \mathsf{match}\, x\, \mathsf{with} \langle \mathsf{nil} \Rightarrow e_1 \mid \mathsf{cons}(x_1, x_2) \Rightarrow e_2 \rangle : (B, Q')}\ (\text{T:MATL})$$

$$\frac{\Sigma; \Gamma, x_1{:}A_1, x_2{:}A_2; \{P\} \vdash^{M} e : (B, Q') \qquad P + M^{\mathsf{matP}} = Q}{\Sigma; \Gamma, x{:}A; \{Q\} \vdash^{M} \mathsf{match}\, x\, \mathsf{with}\, (x_1, x_2) \Rightarrow e : (B, Q')}\ (\text{T:MATP}) \qquad \frac{Q = Q' + M^{\mathsf{pair}}}{\Sigma; x_1{:}A_1, x_2{:}A_2; \{Q\} \vdash^{M} (x_1, x_2) : (A_1 {*} A_2, Q')}\ (\text{T:PAIR})$$

$$\frac{\Sigma; \Gamma_1, \Gamma_2; R \vdash^{M} e_1 \rightsquigarrow \Gamma_2, x{:}A; R' \qquad \Sigma; \Gamma_2, x{:}A; \{R'\} \vdash^{M} e_2 : (B, Q') \qquad Q = R + M^{\mathsf{let}}}{\Sigma; \Gamma_1, \Gamma_2; \{Q\} \vdash^{M} \mathsf{let}\, x = e_1\, \mathsf{in}\, e_2 : (B, Q')}\ (\text{T:LET})$$

$$\frac{\begin{array}{cc} \Sigma; \Gamma_1, \Gamma_2, \Delta; Q \vdash^{M} e_1 \rightsquigarrow \Gamma_2, \Delta, x_1{:}A_1; Q' & \Sigma; \Gamma_2, \Delta, x_1{:}A_1; Q' \vdash^{\mathsf{cf}} e_2 \rightsquigarrow \Delta, x_1{:}A_1, x_2{:}A_2; R \\ \Sigma; \Gamma_1, \Gamma_2, \Delta; P \vdash^{\mathsf{cf}} e_1 \rightsquigarrow \Gamma_2, \Delta, x_1{:}A_1; P' & \Sigma; \Gamma_2, \Delta, x_1{:}A_1; P' \vdash^{M} e_2 \rightsquigarrow \Delta, x_1{:}A_1, x_2{:}A_2; R \\ \multicolumn{2}{c}{\Sigma; \Delta, x_1{:}A_1, x_2{:}A_2; R \vdash^{M} e : (B, R')} \end{array}}{\Sigma; \Gamma_1, \Gamma_2, \Delta; \{Q + M^{\mathsf{Par}}, P + M^{\mathsf{Par}}\} \vdash^{M} \mathsf{par}\, x_1 = e_1\, \mathsf{and}\, x_2 = e_2\, \mathsf{in}\, e : (B, R')}\ (\text{T:PAR})$$

$$\frac{\Sigma; \Gamma, x_1{:}A, x_2{:}A; \{P_1, \ldots, P_m\} \vdash^{M} e : (B, Q') \qquad \forall i\, \exists j : Q_j = \curlyvee P_i}{\Sigma; \Gamma, x{:}A; \{Q_1, \ldots, Q_n\} \vdash^{M} e[x/x_1, x/x_2] : (B, Q')}\ (\text{T:SHARE})$$

$$\frac{\Sigma; \Gamma; \{P\} \vdash^{M} e : (B, P') \qquad Q \geqslant P + c \qquad Q' \leqslant P' + c}{\Sigma; \Gamma; \{Q\} \vdash^{M} e : (B, Q')}\ (\text{T:WEAK-A}) \qquad \frac{\Sigma; \Gamma; \{P\} \vdash^{M} e : (B, P') \qquad \forall i \in \mathcal{I}(\Gamma):\ p_i = q_{(i,0)}}{\Sigma; \Gamma, x{:}A; \{Q\} \vdash^{M} e : (B, Q')}\ (\text{T:WEAK-C})$$

◆ ◆ ◆

$$\frac{\forall j \in \mathcal{I}(\Delta): \quad j = \vec{0} \implies \Sigma; \Gamma; \pi_j^{\Gamma}(Q) \vdash^{M} e : (A, \pi_j^{x:A}(Q')) \qquad j \neq \vec{0} \implies \Sigma_j; \Gamma; \pi_j^{\Gamma}(Q) \vdash^{\mathsf{cf}} e : (A, \pi_j^{x:A}(Q'))}{\Sigma; \Gamma, \Delta; Q \vdash^{M} e \rightsquigarrow \Delta, x{:}A; Q'}\ (\text{B:BIND})$$

Figure 4: Typing rules for annotated types and the binding rule for multivariate variable binding.

rences of a variable. The following lemma shows that sharing is a linear operation that does not lead to any loss of potential.

Lemma 6. Let A be a data type. Then there are non-negative rational numbers $c_k^{(i,j)}$ for $i, j, k \in \mathcal{I}(A)$ with $\deg(k) \leqslant \deg(i, j)$ such that the following holds. For every context $\Gamma, x_1{:}A, x_2{:}A; Q$ and every H, V with $H \vDash V : \Gamma, x{:}A$ it holds that $\Phi_{V,H}(\Gamma, x{:}A; Q') = \Phi_{V',H}(\Gamma, x_1{:}A, x_2{:}A; Q)$ where $V' = V[x_1, x_2 \mapsto V(x)]$ and $q'_{(\ell,k)} = \sum_{i,j \in \mathcal{I}(A)} c_k^{(i,j)} q_{(\ell,i,j)}$.

The coefficients $c_k^{(i,j)}$ can be computed effectively and are *natural* numbers. For a context $\Gamma, x_1{:}A, x_2{:}A; Q$ we define $\curlyvee Q$ to be the Q' from Lemma 6.

Typing Rules Figure 4 shows the annotated typing rules for expressions. The rules T:WEAK-A, T:SHARE, and T:WEAK-C are structural rules that apply to every expression. The other rules are syntax-driven and there is one rule for every construct of the syntax. In the implementation we incorporated the structural rules in the syntax-driven ones. The rule T:WEAK-A is integrated into T:APP. The weakening rules T:WEAK-A and T:WEAK-C are used to fork and join different branches in the rules T:MATP, T:MAT, and T:PAR.Most of the rules are similar to the rules for multivariate amortized analysis for sequential programs [21, 20]. The main difference it that the rules here operate on annotations that are singleton sets $\{Q\}$ instead of the usual context annotations Q.

T:CONS assigns potential to a lengthened list. The additive shift $\lhd_L(Q')$ transforms the annotation Q' for a list type into an annotation for the context $x_1{:}A, x_2{:}L(A)$. Lemma 5 shows that potential is neither gained nor lost by this operation. The potential Q of the context has to pay for both the potential Q' of the resulting list and the resource cost M^{cons} for list cons.

T:MATL shows how to treat pattern matching of lists. The initial potential defined by the annotation Q of the context $\Gamma, x{:}L(A)$ has to be sufficient to pay the costs of the evaluation of e_1 or e_2 and the potential defined by the annotation Q' of the result type. To type the expression e_1 of the nil case we use the projection $\pi_0^{\Gamma}(Q)$ that results in an annotation for the context Γ. Since the matched list is empty in this case no potential is lost by the discount of the annotations $q_{(i,j)}$ of Q where $j \neq 0$. To type the expression e_2 of the cons case we rely on the shift operation $\lhd_L(Q)$ for lists that results in an annotation for the context $\Gamma, x_1{:}A, x_2{:}L(A)$. Again there is no loss of potential (see Lemma 5). The equalities relate the potential before and after the evaluation of e_1 or e_2, to the potential before the and after the evaluation of the match operation by incorporating the respective resource cost for the matching.

In the binding rules T:LET and T:PAR, we bind the result of the evaluation of an expression e to a variable x. The problem that arises is that the resulting annotated context $\Delta, x{:}A, Q'$ features potential functions whose domain consists of data that is referenced by x as well as data that is referenced by Δ. This potential has to be related to data that is referenced by Δ and the free variables in e.

To express the relations between mixed potentials before and after the evaluation of e, we introduce a new auxiliary binding judgement of the from

$$\Sigma; \Gamma, \Delta; Q \vdash^{M} e \rightsquigarrow \Delta, x{:}A; Q'$$

in the rule B:BIND. The intuitive meaning of the judgement is the following. Assume that e is evaluated in the context Γ, Δ, that $\mathrm{FV}(e) \in \mathrm{dom}(\Gamma)$, and that e evaluates to a value that is bound to the variable x. Then the initial potential $\Phi(\Gamma, \Delta; Q)$ is larger than the cost of evaluating e in the metric M plus the potential of the resulting context $\Phi(\Delta, x{:}A; Q')$. Lemma 7 formalizes this intuition.

Lemma 7. Let $H \models V{:}\Gamma, \Delta$ and $\Sigma; \Gamma, \Delta; Q \overset{M}{\vdash} e \rightsquigarrow \Delta, x{:}A; Q'$.

1. If $V, H \overset{M}{\vdash} e \Downarrow (\ell, H') \mid (w, d)$ then $\Phi_{V,H}(\Gamma, \Delta; Q) \geqslant d + \Phi_{V',H'}(\Delta, x{:}A; Q')$ where $V' = V[x \mapsto \ell]$.
2. If $V, H \overset{M}{\vdash} e \Downarrow \rho \mid (w, d)$ then $d \leqslant \Phi_{V,H}(\Gamma; Q)$.

Formally, Lemma 7 is a consequence of the soundness of the type system (Theorem 4). In the inductive proof of Theorem 4, we use a weaker version of Lemma 7 in which the soundness of the type judgements in Lemma 7 is an additional precondition.

The rule T:PAR for parallel bindings $\mathsf{par}\, x_1 = e_1 \,\mathsf{and}\, x_2 = e_2 \,\mathsf{in}\, e$ is the main novelty in the type system. The idea is that we type the expressions e_1 and e_2 twice using the new binding judgement. In the first group of bindings, we account for the cost of e_1 and derive a context $\Gamma_2, \Delta, x_1{:}A_1; P_1'$ in which the result of the evaluation of e_1 is bound to x_1. This context is then used to bind the result of evaluating e_2 in the context $\Delta, x_1{:}A_1, x_2{:}A_2; R$ without paying for the resource consumption. In the second group of bindings, we also derive the context $\Delta, x_1{:}A_1, x_2{:}A_2; R$ but pay for the cost of evaluating e_2 instead of e_1. The type annotations Q_1 and Q_2 for the initial context $\Gamma = \Gamma_1, \Gamma_2, \Delta$ establish a bound on the depth d of evaluating the whole parallel binding: If the depth of evaluating e_1 is larger than the depth of evaluating e_2 then $\Phi(\Gamma; Q_1) \geqslant d$. Otherwise we have $\Phi(\Gamma; Q_2) \geqslant d$. If the parallel binding evaluates to a value v then we have additionally that $\max(\Phi(\Gamma; Q_1), \Phi(\Gamma; Q_2)) \geqslant d + \Phi(v{:}(B, Q'))$.

It is important that the annotations Q_1 and Q_2 of the initial context $\Gamma_1, \Gamma_2, \Delta$ can defer. The reason is that we have to allow a different sharing of potential in the two groups of bindings. If we would require $Q_1 = Q_2$ then the system would be too restrictive. However, each type derivation has to establish the equality of the two annotations directly after the use of T:PAR by multiple uses of the sharing rule T:SHARE. Note that T:PAR is the only rule that can introduce a non-singleton set $\{Q_1, Q_n\}$ of context annotations.

A useful observation for the implementation is that the rule T:PAR can be simplified for the cost-free metric cf. The two groups of bindings in the premiss of the rule are actually equivalent for cost-free type judgements. As a result, we can just use one group of bindings.

T:SHARE has to be applied to expressions that contain a variable twice (x in the rule). The sharing operation $\curlyvee P$ transfers the annotation P for the context $\Gamma, x_1{:}A, x_2{:}A$ into an annotation Q for the context $\Gamma, x{:}A$ without loss of potential (Lemma 6). This is crucial for the accuracy of the analysis since instances of T:SHARE are quite frequent in typical examples. The remaining rules are affine in the sense that they assume that every variable occurs at most once in the typed expression.

T:SHARE is the only rule whose premiss allows judgements that contain a non-singleton set $\{P_1, \ldots, P_m\}$ of context annotations. It has to be applied to produce a judgement with singleton set $\{Q\}$ before any of the other rules can be applied. The idea is that we always have $n \leqslant m$ for the set $\{Q_1, \ldots, Q_n\}$ and the sharing operation $\curlyvee_i$ is used to unify the different P_i.

Soundness The operational big-step semantics with partial evaluations makes it possible to state and prove a strong soundness result. An annotated type judgment for an expression e establishes a bound on the depth of all evaluations of e in a well-formed environment; regardless of whether these evaluations diverge or fail. (Depending on the metric M there can be diverging evaluations with bounded depth.)

Moreover, the soundness theorem states also a stronger property for terminating evaluations. If an expression e evaluates to a value v in a well-formed environment then the difference between initial and final potential is an upper bound on the depth of the evaluation.

Theorem 4 (Soundness)**.** If $H \models V{:}\Gamma$ and $\Sigma; \Gamma; \mathcal{Q} \vdash e{:}(B, Q')$ then there exists a $Q \in \mathcal{Q}$ such that the following holds.

1. If $V, H \overset{M}{\vdash} e \Downarrow (\ell, H') \mid (w, d)$ then $d \leqslant \Phi_{V,H}(\Gamma; Q) - \Phi_{H'}(\ell{:}(B, Q'))$.
2. If $V, H \overset{M}{\vdash} e \Downarrow \rho \mid (w, d)$ then $d \leqslant \Phi_{V,H}(\Gamma; Q)$.

Theorem 4 is proved by a nested induction on the derivation of the evaluation judgment and the type judgment $\Gamma; Q \vdash e{:}(B, Q')$. The inner induction on the type judgment is needed because of the structural rules. There is one proof for all possible instantiations of the resource constants.

The proof of most rules is very similar to the proof of the rules for multivariate resource analysis for sequential programs [22]. The main novelty is the treatment of parallel evaluation in the rule T:PAR which we described previously.

If the metric M is simple (all constants are 1) then it follows from Theorem 4 that the bounds on the depth also prove the termination of programs.

Corollary 1. Let M be a simple metric. If $H \models V{:}\Gamma$ and $\Sigma; \Gamma; Q \vdash e{:}(A, Q')$ then there are $w \in \mathbb{N}$ and $d \leqslant \Phi_{V,H}(\Gamma; Q)$ such that $V, H \overset{M}{\vdash} e \Downarrow (\ell, H') \mid (w, d)$ for some ℓ and H'.

Type Inference. In principle, type inference consists of four steps. First, we perform a classic type inference for the simple types such as nat array. Second, we fix a maximal degree of the bounds and annotate all types in the derivation of the simple types with variables that correspond to type annotations for resource polynomials of that degree. Third, we generate a set of linear inequalities, which express the relationships between the added annotation variables as specified by the type rules. Forth, we solve the inequalities with an LP solver such as CLP. A solution of the linear program corresponds to a type derivation in which the variables in the type annotations are instantiated according to the solution.

In practice, the type inference is slightly more complex. Most importantly, we have to deal with resource-polymorphic recursion in many examples. This means that we need a type annotation in the recursive call that differs from the annotation in the argument and result types of the function. To infer such types we successively infer type annotations of higher and higher degree. Details can be found in previous work [19]. Moreover, we have to use algorithmic versions of the type rules in the inference in which the non-syntax-directed rules are integrated into the syntax-directed ones [22]. Finally, we use several optimizations to reduce the number of generated constraints.

An concrete example of a type derivation can be found in previous work [22].

6. Nested Data Parallelism

The techniques that we describe in this work for a minimal function language scale to more advanced parallel languages such as Blelloch's NESL [6].

To describe the novel type analysis in this paper, we use a binary binding construct to introduce parallelism. In NESL, parallelism is introduced via built-in functions on sequences as well as parallel sequence comprehension that is similar to Haskell's list comprehension. The depth of all built-in sequence functions such as *append* and *sum* is constant in NESL. Similarly, the depth overhead of the

parallel sequence comprehension is constant too. Of course, it is possible to define equivalent functions in RAML. However, the depth would often be linear since we, for instance, have to sequentially form the resulting list.

Nevertheless, the user definable resource metrics in RAML make it easy to introduce built-in functions and language constructs with customized work and depth. For instance we could implement NESL's *append* like the recursive append in RAML but use a metric inside the function body in which all evaluation steps have depth zero. Then the depth of the evaluation of $\mathsf{append}(x, y)$ is constant and the work is linear in $|x|$.

To demonstrate this ability of our approach, we implemented parallel list comprehensions, NESL's most powerful construct for parallel computations. A list comprehension has the form

$$\{\ e\ :\ x_1 \,\mathsf{in}\, e_1\,; \ldots ;\, x_n \,\mathsf{in}\, e_n \mid e_b\ \}$$

where e is an expression, $e_1, \ldots, e_n$ are expressions of some list type, and e_b is a boolean expression. The semantics is that we bind $x_1, \ldots, x_n$ successively to the elements of the lists $e_1, \ldots, e_n$ and evaluate e_b and e under these bindings. If e_b evaluates to true under a binding then we include the result of e under that binding in the resulting list. In other words, the above list comprehension is equivalent to the Haskell expression $[\ e \mid (x_1, \ldots, x_n) \leftarrow \mathsf{zip}_n\, e_1\ \ldots\ e_n\ ,\ e_b\ ]$.

The *work* of evaluating $\{\ e\ :\ x_1 \,\mathsf{in}\, e_1\,; \ldots ;\, x_n \,\mathsf{in}\, e_n \mid e_b\ \}$ is sum of the cost of evaluating $e_1, \ldots, e_{n-1}$ and e_n plus the sum of the cost of evaluating e_b and e with the successive bindings to the elements of the results of the evaluation of $e_1, \ldots, e_n$. The *depth* of the evaluation is sum of the cost of evaluating $e_1, \ldots, e_{n-1}$ and e_n plus the maximum of the cost of evaluating e_b and e with the successive bindings to the elements of the results of the e_i.

For instance, the following RAML expression computes the list that contains the squares of all negative numbers in the list ℓ.

$$\{ x * x\ :\ x \,\mathsf{in}\, \ell \mid x < 0 \}$$

The work of the evaluation is linear in ℓ and the depth of the evaluation is constant. The second example is an expression that computes all pairs that one can form with the elements of an list ℓ so that the first element is less than the second.

$$\{ \{ (x, y)\ :\ y \,\mathsf{in}\, \ell \mid x < y \} : x \,\mathsf{in}\, \ell \}$$

The work of the evaluation is quadratic in ℓ and the depth is constant.

The automatic cost analysis of a parallel list comprehension is achieved in RAML by a translation of the comprehension into a recursive function in which the depth of certain operations such as pattern matching and list construction is zero. For more examples of parallel list comprehensions refer to the NESL manual [6].

7. Experimental Evaluation

We implemented the developed automatic depth analysis in Resource Aware ML (RAML). The implementation consists mainly of adding the syntactic form for the parallel binding and the parallel list comprehensions together with the treatment in the parser, the interpreter, and the resource-aware type system. RAML is publically available for download and through a user-friendly online interface [1]. On the project web page you also find the source code of all example programs and of RAML itself.

We used the implementation to perform an experimental evaluation of the analysis on typical examples from functional programming. In the compilation of our results we focus on examples that have a different asymptotic worst-case behavior in parallel and sequential evaluation. In many other cases, the worst-case behavior only differs in the constant factors. Also note that many of the classic examples of Blelloch [7]—like quick sort—have a better asymptotic average behavior in parallel evaluation but the same asymptotic worst-case behavior in parallel and sequential cost.

Table 1 contains a representative compilation of our experimental results. For each analyzed function, it shows the function type, the computed bounds on the work and the depth, the run time of the analysis in seconds and the actual asymptotic behavior of the function. The experiments were performed on an iMac with a 3.4 GHz Intel Core i7 and 8 GB memory. The computed bounds are simplified multivariate resource polynomials that are presented to the user by RAML. Note that RAML also outputs the (unsimplified) multivariate resource polynomials. The variables in the computed bounds correspond to the sizes of different parts of the input. The naming convention is that we use the order n, m, x, y, z, u of the variables to name the sizes in a depth-first way: n is the size of the first argument, m is the maximal size of the elements of the first argument, x is the size of the second argument, etc.

All bounds are asymptotically tight if the tight bound is representable by a multivariate resource polynomial. For example, the exponential work bound for fib and the logarithmic bounds for bitonic_sort are not representable as a resource polynomial. Another example is the loose depth bound for dyad_all where we would need the base function $\max_{1 \leqslant i \leqslant n} m_i$ but only have $\sum_{1 \leqslant i \leqslant n} m_i$.

Matrix Operations To study programs that use nested data structures we implemented several matrix operations for matrices that are represented by lists of lists of integers. The implemented operations include, the dyadic product from Section 3 (dyad), transposition of matrices (transpose, see [1]), addition of matrices (m_add, see [1]), and multiplication of matrices (m_mult1 and m_mult2).

Note that we needed to additionally represent the size of the inner lists in the matrix with a natural number n in the functions m_mult2, m_add, and transpose in order to get asymptotically tight bounds for the depth. The reason is that we can only express bounds that are multivariate resource polynomials. So we cannot have bounds such as $8128{\cdot}n$ where n is the maximal length of the inner lists of xs.

To demonstrate the compositionality of the analysis, we have implemented two more involved functions for matrices. The function

$$\mathsf{dyad_all} : L(L(\text{int})) \to L(L(L(\text{int})))$$

computes the dyadic product (using dyad) of all ordered pairs of the inner lists in the argument. The function

$$\mathsf{m_mult_pairs} : L(L(L(\text{int}))) * L(L(L(\text{int}))) \to L(L(L(\text{int})))$$

computes the products $M_1 \cdot M_2$ (using m_mult1) of all pairs of matrices such that M_1 is in the first list of the argument and M_2 is in the second list of the argument.

Sorting Algorithms The sorting algorithms that we implemented include quick sort and bitonic sort for lists of integers (quicksort and bitonic_sort, see [1]).

The analysis computes asymptotically tight quadratic bounds for the work and depth of quick sort. The asymptotically tight bounds for the work and depth of bitonic sort are $O(n \log n)$ and $O(n \log^2 n)$, respectively, and can thus not be expressed by polynomials. However, the analysis computes quadratic and cubic bounds that are asymptotically optimal if we only consider polynomial bounds.

More interesting are sorting algorithms for lists of lists, where the comparisons need linear instead of constant time. In these algorithms we can often perform the comparisons in parallel. For instance, the analysis computes asymptotically tight bounds for quick sort for lists of lists of integers (quicksort_list, see Table 1).

Set Operations We implemented sets as unsorted lists without duplicates. Most list operations such as intersection (Table 1), difference (see [1]), and union (see [1]) have linear depth and quadratic work. The analysis finds these asymptotically tight bounds.

Function Name / Function Type	Computed Depth Bound / Computed Work Bound	Run Time	Asym. Behav.
dyad	$10m + 10n + 3$	0.19 s	$O(n+m)$
$L(\text{int}) * L(\text{int}) \to L(L(\text{int}))$	$10mn + 17n + 3$	0.20 s	$O(nm)$
dyad_all	$1.\bar{6}n^3 - 4n^2 + 10nm + 14.\bar{6}n + 5$	1.66 s	$O(n^2+m)$
$L(L(\text{int})) \to L(L(L(\text{int})))$	$1.\bar{3}n^3 + 5n^2m^2 + 8.5n^2m + 4.5n^2 - 5nm^2 - 8.5nm + 13.1\bar{6}n + 3$	0.96 s	$O(n^3+n^2m^2)$
m_mult1	$15xy + 16x + 10n + 6$	0.37 s	$O(xy)$
$L(L(\text{int})) * L(L(\text{int})) \to L(L(\text{int}))$	$15xyn + 16nm + 18n + 3$	0.36 s	$O(xyn)$
m_mult_pairs	$4n^2 + 15nmx + 10nm + 10n + 3$	3.90 s	O(nm + mx)
$L(L(L(\text{int}))) * L(L(L(\text{int}))) \to L(L(L(\text{int})))$	$7.5n^2m^2x + 7n^2m^2 + n^2mx + 9n^2m + 12.5n^2 - 7.5nm^2x + \ldots$	6.35 s	$O(n^2m^2x)$
m_mult2	$35u + 10y + 15x + 11n + 40$	2.75 s	$O(z+x+n)$
$(L(L(\text{int})) * \text{nat}) * (L(L(\text{int})) * \text{nat}) \to L(L(\text{int}))$	$3.5u^2y + uyz + 14.5uy + 2unx + 13unm + 19un + 31u + 6y + \ldots$	2.99 s	$O(nx(z+y))$
quicksort_list	$12n^2 + 16nm + 12n + 3$	0.67 s	$O(n^2+m)$
$L(L(\text{int})) \to L(L(\text{int}))$	$8n^2m + 15.5n^2 - 8nm + 13.5n + 3$	0.51 s	$O(n^2m)$
intersection	$10m + 12n + 3$	0.49 s	$O(n+m)$
$L(\text{int}) * L(\text{int}) \to L(\text{int})$	$10mn + 19n + 3$	0.28 s	$O(nm)$
product	$8mn + 10m + 14n + 3$	1.05 s	$O(nm)$
$L(\text{int}) * L(\text{int}) \to L(\text{int} * \text{int})$	$18mn + 21n + 3$	0.71 s	$O(nm)$
max_weight	$46n + 44$	0.39 s	$O(n)$
$L(\text{int}) \to \text{int} * L(\text{int})$	$13.5n^2 + 65.5n + 19$	0.30 s	$O(n^2)$
fib	$13n + 4$	0.09 s	$O(n)$
nat * nat → nat	$---$	0.12 s	$O(2^n)$
dyad_comp	13	0.28 s	$O(1)$
$L(\text{int}) * L(\text{int}) \to L(L(\text{int}))$	$6mn + 5n + 2$	0.13 s	$O(nm)$
find	$12m + 29n + 22$	0.38 s	$O(m+n)$
$L(\text{int}) * L(\text{int}) \to L(L(\text{int}))$	$20mn + 18m + 9n + 16$	0.41 s	$O(nm)$

Table 1: The computed depth and work bounds, the actual worst-case time behavior, and the run time of the analysis in seconds. The variables $n, m, x, \ldots$ in the bounds correspond to sizes of the arguments of the functions. For instance, in the case of dyad_all, n is the maximal length of the inner lists and m is the length of the list in the argument. The functions mult and find use parallel list comprehensions. The bounds presented here are simplified multivariate resource polynomials as produced by the analysis.

The function product computes the Cartesian product of two sets. Work and depth of product are both linear and the analysis finds asymptotically tight bounds. However, as it is often the case, the constant factors in the parallel evaluation are much smaller.

Miscellaneous The function max_weight (Table 1) computes the maximal weight of a (connected) sublist of an integer list. The weight of a list is simply the sum of its elements. The work of the algorithm is quadratic but the depth is linear.

Finally, there is a large class of programs that have non-polynomial work but polynomial depth. Since the analysis can only compute polynomial bounds we can only derive bounds on the depth for such programs. A simple example in Table 1 is the function fib that computes the Fibonacci numbers without memoization.

Parallel List Comprehensions The aforementioned examples are all implemented without using parallel list comprehensions. As described in Section 6, parallel list comprehensions have a better asymptotic behavior than semantically-equivalent recursive functions in RAML's current resource metric for evaluation steps.

A simple example is the function dyad_comp which is equivalent to dyad and which is implemented as follows.

```
dyad_comp(xs,ys) = { { x * y : y in ys } : x in xs }
```

As listed in Table 1, the depth of dyad_comp is constant while the depth of dyad is linear. RAML computes tight boundson both work on depth.

A more involved example is the function find that finds a given integer list (needle) in another list (haystack). It returns the starting indices of each occurrence of the needle in the haystack. The algorithm is described by Blelloch [6] and cleverly uses parallel list comprehensions to perform the search in parallel. RAML computes asymptotically tight bounds on the work and depth.

Discussion Our experiments show that the range of the analysis is not reduced when deriving bounds on the depth: The prototype implementation can always infer bounds on the depth of a program if it can infer bounds on the sequential version of the program. The derivation of bounds for parallel programs is also almost as efficient as the derivation of bounds for sequential programs.

We experimentally compared the derived worst-case bounds with the measured work and depth of evaluations with different inputs. In most cases, the derived bounds on the depth are asymptotically tight and the constant factors are close or equal to the optimal ones. As a representative example, Figure 5 presents our experiments for quick sort for lists of lists. The left plot compares the measured depth of manually identified worst-case inputs of different sizes with the inferred depth bound for quicksort_list. The right plot compares the measured work with the inferred work bound. The constant factors in the work bound are optimal and the constant factors in the depth bound are very close to the optimal ones.

However, the analysis of the work is slightly more precise then the analysis of the depth. The reason is that our multivariate resource polynomials cannot express bounds like $q \cdot \binom{n}{2}$ where n is the maximal length of the inner lists of a list. This is a general limitation that also leads to imprecision for bounds on sequential programs but such bounds are more common for depth than for work in practice. Note however that this imprecision can only occur in the analysis of functions with nested data structures in which the resource cost depends on the sizes of the inner data structures.

8. Related Work

Automatic amortized resource analysis was introduced by Hofmann and Jost for a strict first-order functional language [24]. The technique has been applied to higher-order functional programs [28], to derive stack-space bounds for functional programs [10], to func-

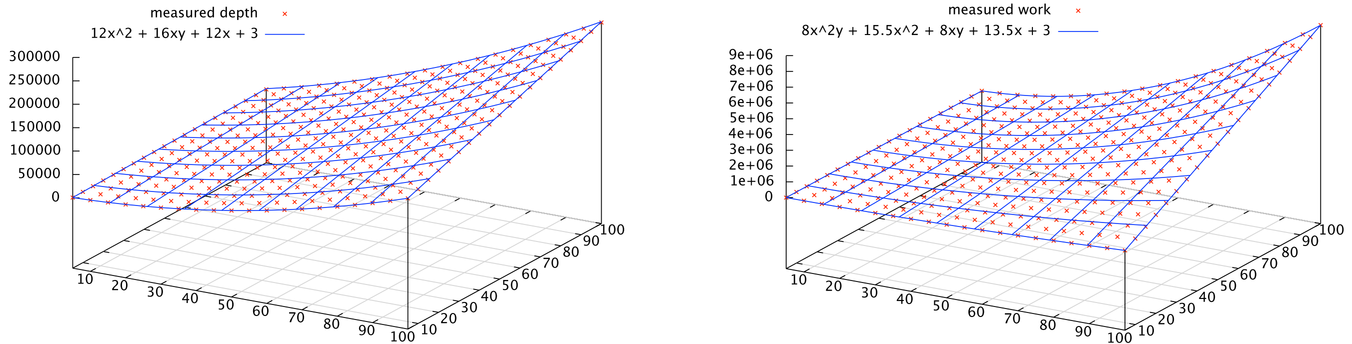

Figure 5: Experimental evaluation of the quality of the bounds for the function quicksort_list. The left plot compares the derived bound (blue lines) on the depth with the measured depths of evaluations with manually identified worst-case input lists (red crosses). The right plot compares the derived bound on the work with the measured work for the worst-case inputs. On the x-axes is the length of the outer list and on the y-axes is maximal length of the inner lists.

tional programs with lazy evaluation [31], to object-oriented programs [25, 26], and to low-level code by integrating it with separation logic [5]. All the aforementioned amortized-analysis–based systems are limited to linear bounds. The polynomial potential functions that we use in this paper were introduced by Hoffmann et al. [18, 21, 22]. In contrast to this work, none of the previous works on amortized analysis considered parallel evaluation. The main technical innovation of this work is the new rule for parallel composition that is not straightforward. The smooth integration of this rule in the existing framework of multivariate amortized resource analysis is one of the main advantages of our work.

Type systems for inferring and verifying cost bounds for sequential programs have been extensively studied. Vasconcelos et al. [35, 34] described an automatic analysis system that is based on sized-types [27] and derives linear bounds for higher-order sequential functional programs. Dal Lago et al. [29, 30] introduced linear dependent types to obtain a complete analysis system for the time complexity of the call-by-name and call-by-value lambda calculus. Crary and Weirich [12] presented a type system for specifying and certifying resource consumption. Danielsson [13] developed a library, based on dependent types and manual cost annotations, that can be used for complexity analyses of functional programs. We are not aware of any type-based analysis systems for parallel evaluation.

Classically, cost analyses are often based on deriving and solving recurrence relations. This approach was pioneered by Wegbreit [36] and has been extensively studied for sequential programs written in imperative languages [2, 4] and functional languages [15, 14].

In comparison, there has been little work done on the analysis of parallel programs. Albert et al. [3] use recurrence relations to derive cost bounds for concurrent object-oriented programs. Their model of concurrent imperative programs that communicate over a shared memory and the corresponding cost measure is however quite different from the depth of functional programs that we study.

The only article on using recurrence relations for deriving bounds on parallel functional programs that we are aware of is a technical report by Zimmermann [37]. The programs that were analyzed in this work are fairly simple and more involved programs such as sorting algorithms seem to be beyond its scope. Additionally, the used technique does not provide the compositionality of amortized resource analysis.

Trinder et al. [33] give a survey of resource analysis techniques for parallel and distributed systems. However, they focus on the usage of analyses for sequential programs to improve the coordination in the parallel systems. Abstract interpretation based approaches to resource analysis [16, 38] are limited to sequential programs.

Finally, there exists research that studies cost models to formally analyze parallel programs. Blelloch and Greiner [7] pioneered the cost measures work and depth that we use in this work. There are more advanced cost models that take into account caches and IO (see, e.g., Blelloch and Harper [8]), However, these works do not provide machine support for deriving or proving static cost bounds.

9. Conclusion

We have introduced the first type-based cost analysis for deriving bounds on the depth of evaluations of parallel function programs. The derived bounds are multivariate resource polynomials that can express a wide range of relations between different parts of the input. As any type system, the analysis is naturally compositional.

The new analysis system has been implemented in Resource Aware ML (RAML) [23]. We have performed a thorough and reproducible experimental evaluation with typical examples from functional programming that shows the practicability of the approach.

An extension of amortized resource analysis to handle non-polynomial bounds such as max and log in a compositional way is an orthogonal research question that we plan to address in the future. Another orthogonal question that we plan to study is the extension of the analysis to additional language features such as higher-order functions, arrays, and user defined data structures. Neither additional language features nor a larger set of potential functions would effect the treatment of parallel bindings in the type system.

The development of our compositional cost analysis for parallel programs is a first step towards a compositional cost analysis system for concurrent programs that dynamically spawn threads.

References

[1] K. Aehlig, M. Hofmann, and J. Hoffmann. RAML Web Site. `http://raml.tcs.ifi.lmu.de`, 2010-2014.

[2] E. Albert, P. Arenas, S. Genaim, G. Puebla, and D. Zanardini. Cost Analysis of Java Bytecode. In *16th Euro. Symp. on Prog. (ESOP'07)*, pages 157–172, 2007.

[3] E. Albert, P. Arenas, S. Genaim, M. Gómez-Zamalloa, and G. Puebla. Cost Analysis of Concurrent OO Programs. In *Prog. Langs. and Systems - 9th Asian Symposium (APLAS'11)*, pages 238–254, 2011.

[4] D. E. Alonso-Blas and S. Genaim. On the limits of the classical approach to cost analysis. In *19th Int. Static Analysis Symposium (SAS'12)*, pages 405–421, 2012.

[5] R. Atkey. Amortised Resource Analysis with Separation Logic. In *19th Euro. Symp. on Prog. (ESOP'10)*, pages 85–103, 2010.

[6] G. E. Blelloch. Nesl: A nested data-parallel language (version 3.1). Technical Report CMU-CS-95-170, CMU, 1995.

[7] G. E. Blelloch and J. Greiner. A Provable Time and Space Efficient Implementation of NESL. In *1st Int. Conf. on Funct. Prog. (ICFP'96)*, pages 213–225, 1996.

[8] G. E. Blelloch and R. Harper. Cache and I/O Efficent Functional Algorithms. In *40th ACM Symp. on Principles Prog. Langs. (POPL'13)*, pages 39–50, 2013.

[9] M. Brockschmidt, F. Emmes, S. Falke, C. Fuhs, and J. Giesl. Alternating Runtime and Size Complexity Analysis of Integer Programs. In *Tools and Alg. for the Constr. and Anal. of Systems - 20th Int. Conf. (TACAS'14)*, pages 140–155, 2014.

[10] B. Campbell. Amortised Memory Analysis using the Depth of Data Structures. In *18th Euro. Symp. on Prog. (ESOP'09)*, pages 190–204, 2009.

[11] A. Charguéraud. Pretty-Big-Step Semantics. In *22nd Euro. Symp. on Prog. (ESOP'13)*, pages 41–60, 2013.

[12] K. Crary and S. Weirich. Resource Bound Certification. In *27th ACM Symp. on Principles of Prog. Langs. (POPL'00)*, pages 184–198, 2000.

[13] N. A. Danielsson. Lightweight Semiformal Time Complexity Analysis for Purely Functional Data Structures. In *35th ACM Symp. on Principles Prog. Langs. (POPL'08)*, pages 133–144, 2008.

[14] N. Danner, J. Paykin, and J. S. Royer. A Static Cost Analysis for a Higher-Order Language. In *7th Workshop on Prog. Languages Meets Prog. Verification (PLPV'13)*, pages 25–34, 2013.

[15] B. Grobauer. Cost Recurrences for DML Programs. In *6th Int. Conf. on Funct. Prog. (ICFP'01)*, pages 253–264, 2001.

[16] S. Gulwani, K. K. Mehra, and T. M. Chilimbi. SPEED: Precise and Efficient Static Estimation of Program Computational Complexity. In *36th ACM Symp. on Principles of Prog. Langs. (POPL'09)*, pages 127–139, 2009.

[17] R. Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, 2012. ISBN 9781107029576.

[18] J. Hoffmann and M. Hofmann. Amortized Resource Analysis with Polynomial Potential. In *19th Euro, Symp. on Prog. (ESOP'10)*, 2010.

[19] J. Hoffmann and M. Hofmann. Amortized Resource Analysis with Polymorphic Recursion and Partial Big-Step Operational Semantics. In *Prog. Langs. and Systems - 8th Asian Symposium (APLAS'10)*, 2010.

[20] J. Hoffmann and Z. Shao. Type-Based Amortized Resource Analysis with Integers and Arrays. In *12th International Symposium on Functional and Logic Programming (FLOPS'14)*, 2014.

[21] J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate Amortized Resource Analysis. In *38th ACM Symp. on Principles of Prog. Langs. (POPL'11)*, 2011.

[22] J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate Amortized Resource Analysis. *ACM Trans. Program. Lang. Syst.*, 2012.

[23] J. Hoffmann, K. Aehlig, and M. Hofmann. Resource Aware ML. In *24rd Int. Conf. on Computer Aided Verification (CAV'12)*, 2012.

[24] M. Hofmann and S. Jost. Static Prediction of Heap Space Usage for First-Order Functional Programs. In *30th ACM Symp. on Principles of Prog. Langs. (POPL'03)*, pages 185–197, 2003.

[25] M. Hofmann and S. Jost. Type-Based Amortised Heap-Space Analysis. In *15th Euro. Symp. on Prog. (ESOP'06)*, pages 22–37, 2006.

[26] M. Hofmann and D. Rodriguez. Automatic Type Inference for Amortised Heap-Space Analysis. In *22nd Euro. Symp. on Prog. (ESOP'13)*, pages 593–613, 2013.

[27] J. Hughes, L. Pareto, and A. Sabry. Proving the Correctness of Reactive Systems Using Sized Types. In *23th ACM Symp. on Principles of Prog. Langs. (POPL'96)*, pages 410–423, 1996.

[28] S. Jost, K. Hammond, H.-W. Loidl, and M. Hofmann. Static Determination of Quantitative Resource Usage for Higher-Order Programs. In *37th ACM Symp. on Principles of Prog. Langs. (POPL'10)*, pages 223–236, 2010.

[29] U. D. Lago and M. Gaboardi. Linear Dependent Types and Relative Completeness. In *26th IEEE Symp. on Logic in Computer Science (LICS'11)*, pages 133–142, 2011.

[30] U. D. Lago and B. Petit. The Geometry of Types. In *40th ACM Symp. on Principles Prog. Langs. (POPL'13)*, pages 167–178, 2013.

[31] H. R. Simões, P. B. Vasconcelos, M. Florido, S. Jost, and K. Hammond. Automatic Amortised Analysis of Dynamic Memory Allocation for Lazy Functional Programs. In *17th Int. Conf. on Funct. Prog. (ICFP'12)*, pages 165–176, 2012.

[32] M. Sinn, F. Zuleger, and H. Veith. A Simple and Scalable Approach to Bound Analysis and Amortized Complexity Analysis. In *Computer Aided Verification - 26th Int. Conf. (CAV'14)*, page 743–759, 2014.

[33] P. W. Trinder, M. I. Cole, K. Hammond, H.-W. Loidl, and G. Michaelson. Resource Analyses for Parallel and Distributed Coordination. *Concurrency and Computation: Practice and Experience*, 25(3):309–348, 2013.

[34] P. Vasconcelos. *Space Cost Analysis Using Sized Types*. PhD thesis, School of Computer Science, University of St Andrews, 2008.

[35] P. B. Vasconcelos and K. Hammond. Inferring Costs for Recursive, Polymorphic and Higher-Order Functional Programs. In *Int. Workshop on Impl. of Funct. Langs. (IFL'03)*, pages 86–101. Springer-Verlag LNCS, 2003.

[36] B. Wegbreit. Mechanical Program Analysis. *Commun. ACM*, 18(9):528–539, 1975.

[37] W. Zimmermann. Automatic Worst Case Complexity Analysis of Parallel Programs. Technical Report TR-90-066, University of California, Berkeley, 1990.

[38] F. Zuleger, M. Sinn, S. Gulwani, and H. Veith. Bound Analysis of Imperative Programs with the Size-change Abstraction. In *18th Int. Static Analysis Symposium (SAS'11)*, 2011.