Accepted Manuscript

Automatically deriving cost models for structured parallel processes using hylomorphisms

David Castro, Kevin Hammond, Susmit Sarkar, Yasir Alguwaifli

Please cite this article as: D. Castro, et al., Automatically deriving cost models for structured parallel processes using hylomorphisms, *Future Generation Computer Systems* (2017), http://dx.doi.org/10.1016/j.future.2017.04.035

- An operational semantics of a queue language for streaming computations.
- An operational semantics of algorithmic skeletons using this queue language.
- Derive cost equations for algorithmic skeletons from the operational semantics.
- Cost equations for algorithmic skeletons are combined with sized types.

Automatically Deriving Cost Models for Structured Parallel Processes Using Hylomorphisms

David Castro, Kevin Hammond, Susmit Sarkar, Yasir Alguwaifli

School of Computer Science, University of St Andrews, St Andrews, UK.

Abstract

Structured parallelism using nested algorithmic skeletons can greatly ease the task of writing parallel software, since common, but hard-to-debug, problems such as *race conditions* are eliminated *by design*. However, choosing the best combination of algorithmic skeletons to yield good parallel speedups for a specific program on a specific parallel architecture is still a difficult problem. This paper uses the unifying notion of *hylomorphisms*, a general recursion pattern, to make it possible to reason about both the functional correctness properties and the extra-functional timing properties of structured parallel programs. We have previously used hylomorphisms to provide a denotational semantics for skeletons, and proved that a given parallel structure for a program satisfies functional correctness. This paper expands on this theme, providing a simple operational semantics for algorithmic skeletons and a cost semantics that can be automatically derived from that operational semantics. We prove that both semantics are sound with respect to our previously defined denotational semantics. This means that we can now automatically and statically choose a provably optimal parallel structure for a given program with respect to a cost model for a (class of) parallel architecture. By deriving an *automatic amortised analysis* from our cost model, we can also accurately predict parallel runtimes and speedups.

1. Introduction

In previous work [9], we have defined a type-based mechanism for reasoning about the safe introduction of parallelism using a structured parallel approach. Our approach allows us to extract parallel program structure as a type. Given a suitable model of a program's execution cost (e.g. in terms of its time performance), we can reason formally about performance improvements for alternative parallelisations, and so select a *provably optimal* parallel implementation. In this paper, we show how to derive appropriate cost models formally from the


Figure 1: Deriving Cost Equations from Operational Semantics.

parallel structure of a program, using a new queue-based operational semantics. This gives a completely formal system for reasoning about the performance of structured parallel programs. We combine related work on *algorithmic skeletons* and *recursion schemes*. Algorithmic skeletons [11] are parametric implementations of common patterns of parallel programming. Using a pattern/skeleton approach, the programmer can design and implement a parallel program in a top-down manner. For example, the programmer could first identify the parallel patterns that occur in a particular piece of software, then select the patterns that potentially lead to the best speedups, and finally select a suitable implementation for those patterns, as a composition of one or more algorithmic skeletons. This composition then exposes the *parallel structure* of the implementation. Developing an equational theory for easily, and automatically, changing the parallel structure of a program has been the subject of research in the skeletons community for the last two decades [42, 3, 1, 8]. Over a similar timeframe, in the functional programming community, a structured form of recursion has been explored, in the form of *patterns of recursion*, or *recursion schemes* [31]. Research on this topic has brought many improvements to equational reasoning in functional languages. One example is the application of the laws and properties of recursion schemes to the well-known *deforestation* optimisation [47, 36].

There is an obvious connection between algorithmic skeletons and recursion schemes: algorithmic skeletons are essentially *higher order functions* that implement some common pattern of parallelism. This has already been exploited a number of times [42, 22, 38, 44]. In [9], we expand on this connection, using the fact that a large number of recursion patterns can be represented as instances of a more general pattern, a *hylomorphism*. The basic idea is to provide a single unifying framework for reasoning both about program transformations and about parallel execution times. This unifying framework provides a type-level abstraction of the *program structure*, given as a combination of the *hylomorphisms* and *algorithmic skeletons* that are used to implement a particular program. This allows us to define a type-based mechanism for reasoning about the safe introduction of parallelism: specific combinations of hylomorphisms can be "replaced" by specific combinations of algorithmic skeletons. We provide strong static guarantees that the resulting program will be functionally equivalent to the original one. Given a suitable *model* for the cost of a parallel implementation (e.g. in terms of execution time), we can then use this type-based approach to reason formally about performance improvements for alternative parallelisations, and so to select a *provably optimal* parallel implementation.

2

1.1. Novel Contributions

In this paper, we show how to formally derive appropriate cost models from the parallel structure of a program, using a new queue-based operational semantics. This gives a completely formal system for reasoning about the performance of structured parallel programs. The main novel contributions of the paper are:

- We define the operational semantics of a *queue language* that is powerful enough to describe the operational semantics of a number of algorithmic skeletons, and small and restrictive enough to facilitate reasoning about correctness and execution times (Sec. 4).

- We define the *operational semantics* of a number of key algorithmic skeletons using this queue language (Sec. 5).

- We derive a set of *cost equations* for a number of algorithmic skeletons from the operational semantics in a systematic way. We combine these cost equations with the notion of *sized types*, and sketch how this process can be automated (Sec. 6).

1.2. Motivating Example

We illustrate our approach using the *image merge* example from [9]. The purpose of the imgMerge function is to mark and then merge pairs of images.

$$\text{imgMerge} \ : \ \text{List}(\text{Img} \times \text{Img}) \ \rightarrow \ \text{List Img}$$
$$\text{imgMerge} \ = \ \text{map} \ (\text{merge} \circ \text{mark})$$

This has many possible, semantically equivalent, parallel implementations.

$$\text{imgMerge}_1 \ = \ \text{farm} \ n \ (\text{fun} \ (\text{merge} \circ \text{mark}))$$
$$\text{imgMerge}_2 \ = \ \text{farm} \ m \ (\text{fun} \ \text{mark}) \ \| \ \text{farm} \ n \ (\text{fun} \ \text{merge})$$
$$\text{imgMerge}_3 \ = \ \text{farm} \ n \ (\text{fun} \ \text{mark}) \ \| \ \text{fun} \ \text{merge}$$
$$\dots$$

Here, farm n is a *task farm* skeleton that replicates its argument n times, fun captures primitive (sequential) functions, $\circ$ is the normal function composition, and $\|$ is a parallel pipeline, i.e. the parallel composition of two skeletons. The structure of each parallel implementation can be lifted into an appropriate type signature, so that

$$\text{imgMerge}_1 \ : \ \text{List}(\text{Img} \times \text{Img}) \ \xrightarrow{\text{FARM}_n(\text{FUN} \ (\text{merge} \circ \text{mark}))} \ \text{List Img}$$

By combining well-known properties of algorithmic skeletons [2] and *hylomorphisms* [31], we can define a *convertibility* relation that relates each of these types for *imgMerge* to any of the other types. Moreover, this same relation also allows the underlying program to be automatically rewritten so that it matches the new type. We can thus automatically select any valid parallel implementation of *imgMerge* simply by changing the type, and without changing its definition. Finally, if we have a cost model for these types, we can now reason about costs, and automatically select the provably best parallel implementation for *imgMerge*, or, indeed, for any appropriately structured parallel program.

3

2. Algorithmic Skeletons

We use *algorithmic skeletons* [11] to represent structured parallel patterns. Skeletons are parameterised templates (*higher-order functions*) that capture the structure of a parallel program. That is, they implement a specific parallel *pattern*, that can be instantiated to produce a specific parallel *algorithm*. In this paper, as in our earlier work [9, 21], we use a "pluggable" approach, where all skeletons are *streaming* entities (that is, they operate over multiple inputs and produce multiple results). Each skeleton takes its inputs from an *input queue*, and produces a *result queue*. This approach allows skeletons to be easily nested or linked together into more complex structures, whose sub-components are connected *via* intermediate queues. In this paper, we will consider four common parallel skeletons, *task farms, parallel pipelines, feedbacks*, and *parallel divide-and-conquer*, plus one basic building block, *structured functions*.

Structured Functions (ꜰᴜɴ f) *lift* basic (sequentially executed) functions so that they can be used as building blocks for our skeletons. These may be either named functions; compositions of one or more structured functions, constructed using the *sequential composition* operator ($\circ$); or recursive functions that have been built using common patterns of recursion. Structured functions expose the underlying structure of a computation, enabling possible parallelisations.

Task farms (farm) apply the same operation to each element of an input stream of tasks, using a fixed number of parallel workers. The input tasks must be independent, and the outputs can be produced in an arbitrary order.

Parallel pipelines ($\|$) compose two other streaming skeletons, in parallel. This can be used to parallelise two or more stages of a computation, e.g. marking and merging in the *image merge* example above.

4

Feedbacks (fb) capture recursion in a streaming computation, operating over some internally nested skeleton with a given predicate e.g. f. This could be used, for example, to repeatedly transform an image using some parallel skeleton until a certain dynamic condition was met.
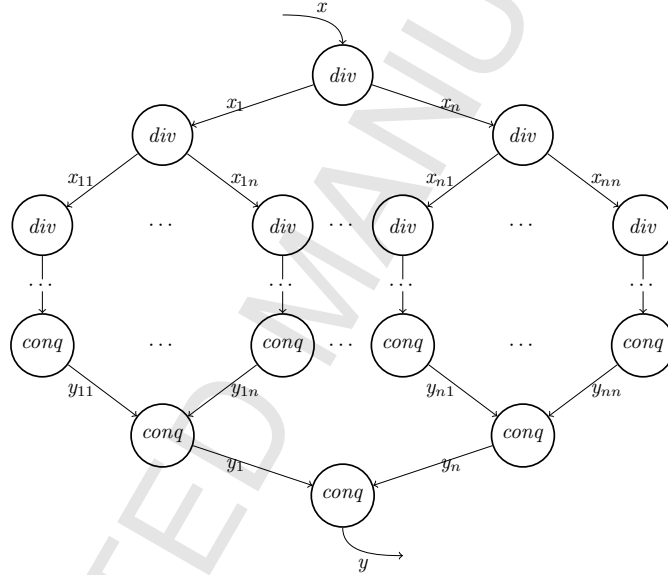
Parallel Divide-and-Conquers (dc) are parallel implementations of classical divide-and-conquer algorithms. Parallelism comes from performing each of the recursive calls in the divide-and-conquer in parallel. In our framework, each instance of a divide-and-conquer skeleton divides an input into n sub-inputs using a $n+1$-ary operation, so that we do not need to assume associativity.

3. A Type-Level Treatment of Parallel Structure

3.1. Functors

We begin by defining some basic concepts from category theory. A *functor* is a structure-preserving mapping between *categories*. We will only require *endofunctors* on CPO, $F : CPO \to CPO$, which will represent type constructors, $F : \texttt{Type} \to \texttt{Type}$. *Bifunctors* are functors that are generalised to multiple arguments. We use them to define polymorphic data types. Our functors are either: i) standard polynomial functors with constant types; ii) the left section of a bifunctor; or iii) a polymorphic type defined as the fixpoint of some bifunctor.

5

Bifunctors are defined using products and sums alone. Functors are defined using a *pointed* notation, with the obvious semantic interpretation. If A and B are type variables, and T is a type, we accept the following definitions:

$$
\begin{array}{rcll}
G\ A\ B & = & T & \text{(bifunctor defined using sums and products)} \\
F\ A & = & T & \text{(functor defined using sums and products)} \\
F\ A & = & G\ T\ A & (G \text{ is a bifunctor}) \\
F\ A & = & \mu(G\ A) & (G \text{ is a bifunctor})
\end{array}
$$

Example 1 (Lists). *Given the bifunctor $L\ A\ B = 1 + A \times B$, the polymorphic* List *data type is defined by the fixpoint of $L\ A$:* List $A = \mu(L\ A)$. *Note that given a base bifunctor $G\ A\ B$, the data type $F\ A = \mu(G\ A)$ is also a functor. The two list constructors are defined in a standard way:*

$$
\begin{array}{rcl rcl}
\texttt{nil} & : & \mathsf{List}\ A & \texttt{cons} & : & A \to \mathsf{List}\ A \to \mathsf{List}\ A \\
\texttt{nil} & = & in_{L\,A}\ (inj_1\ ()) & \texttt{cons}\ x\ l & = & in_{L\,A}\ (inj_2\ (x,l))
\end{array}
$$

3.2. Hylomorphisms

Hylomorphisms are a well known, very general, recursion pattern [31], that can be seen as a generalised divide-and-conquer pattern. Intuitively, $hylo_F\ f\ g$ is a *regular* recursive algorithm (i.e. with no nested or mutual recursion), where g describes how the algorithm divides the input problem into sub-problems, stored in a structure F, and f describes how these results are combined.

$$
\begin{array}{rcl}
hylo_F & : & (F\ B \to B) \to (A \to F\ A) \to A \to B \\
hylo_F\ f\ g & = & f \circ F\ (hylo_F\ f\ g) \circ g
\end{array}
$$

Catamorphisms (folds), *anamorphisms* (unfolds) and *maps* are just special cases of hylomorphisms.

$$
\begin{array}{rcl}
T\ A & = & \mu(F\ A) \\
map_T\ f & = & hylo_{F\,A}\ (in_{F\,B} \circ (F\ f\ id))\ out_{F\,A}, \\
 & & \quad \text{where } A = dom(f) \text{ and } B = codom(f) \\
cata_F\ f & = & hylo_F\ f\ out_F \\
ana_F\ f & = & hylo_F\ in_F\ f
\end{array}
$$

For any recursive type that is defined as the fixpoint of a functor, μF, the standard $in_F : F\mu F \to \mu F$ and $out_F : \mu F \to F\mu F$ capture the isomorphism between $F\mu F$ and μF. Note that, since $out_F \circ in_F = id$, $hylo_F\ f\ g = cata_F\ f \circ ana_F\ g$.

Example 2 (Quicksort). *Assuming a type A, and two functions, $\texttt{leq}$ and $\texttt{gt}$ that filter the elements appropriately, we can implement naïve quicksort as:*

$$
\begin{array}{rcl}
\texttt{qsort} & : & \mathsf{List}\ A \to \mathsf{List}\ A \\
\texttt{qsort}\,\texttt{nil} & = & [] \\
\texttt{qsort}\,(\texttt{cons}\ x\ xs) & = & \texttt{qsort}\,(\texttt{leq}\ x\ xs) \mathbin{+\!\!+} \texttt{cons}\ x(\texttt{qsort}\,(\texttt{gt}\ x\ l))
\end{array}
$$

6

We make the recursive structure explicit by using a tree. The split *function unfolds the arguments into this tree, and the* join *function then flattens it.*

```
split            :   List A → Tree A
split nil        =   empty
split (cons x l) =   node (split (leq x l)) x (split (gt x l))

join             :   Tree A → List A
join empty       =   nil
join (node l x r) =  join l ++ cons x (join r)

qsort            :   List A → List A
qsort            =   join ∘ split
```

We can remove the explicit recursion from these definitions, since split *is a tree anamorphism, and* join *is a tree catamorphism.*

$$
\begin{aligned}
&\texttt{split} &&:\quad \text{List } A \to T\ A\ (\text{List } A)\\
&\texttt{split nil} &&=\quad inj_1\ ()\\
&\texttt{split (cons } x\ l) &&=\quad inj_2\ (x,\ \texttt{leq } x\ l,\ \texttt{gt } x\ l)\\[6pt]
&\texttt{join} &&:\quad T\ A\ (\text{List } A) \to \text{List } A\\
&\texttt{join } (inj_1\ ()) &&=\quad \texttt{nil}\\
&\texttt{join } (inj_2\ (x,l,r)) &&=\quad l \mathbin{+\!\!+} \texttt{cons } x\ r\\[6pt]
&\texttt{qsort} &&:\quad \text{List } A \to \text{List } A\\
&\texttt{qsort} &&=\quad cata_{T\,A}\ \texttt{join} \circ ana_{T\,A}\ \texttt{split}
\end{aligned}
$$

Finally, since we have a composition of a catamorphism and an anamorphism, we can write qsort *as the equivalent hylomorphism.*

$$
\texttt{qsort} = hylo_{T\,A}\ \texttt{join}\ \texttt{split}
$$

3.3. Type System and Semantics

The syntax of expressions, E, is shown below. We distinguish between two levels of expressions: *Structured Expressions (S)* describe primitive (sequential) forms, and *Structured Parallel Processes (P)* introduce specific parallel skeletons.

$$
\begin{aligned}
e \in E &\quad ::= \quad s \quad | \quad \mathsf{par}_T\ p\\
s \in S &\quad ::= \quad f \quad | \quad e_1 \circ e_2 \quad | \quad \mathsf{hylo}_F\ e_1\ e_2\\
p \in P &\quad ::= \quad \mathsf{fun}\ s\ |\ p_1 \parallel p_2\ |\ \mathsf{dc}_{n,F}\ s_1\ s_2\ |\ \mathsf{farm}\ n\ p\ |\ \mathsf{fb}\ p
\end{aligned}
$$

3.3.1. Denotational Semantics

Our denotational semantics is likewise split into two parts: $\mathcal{S}[\![\cdot]\!]$ describes the base semantics, and $[\![\cdot]\!]$ lifts this to a *streaming* form. We use a global environment for atomic function types, ρ, and the corresponding global environment of functions, $\hat{\rho}$:

$$
\rho \;=\; \{f : A \to B,\ \ldots\} \qquad \hat{\rho} \;=\; \{[\![f]\!] \in [\![A \to B]\!],\ \ldots\}
$$

7

$$
\begin{array}{lll}
\mathcal{S}[\![p \; : \; T \; A \to T \; B]\!] & : & [\![A \to B]\!] \\
\mathcal{S}[\![\mathsf{fun} \; f]\!] & = & \hat{\rho}(f) \\
\mathcal{S}[\![p_1 \; \| \; p_2]\!] & = & \mathcal{S}[\![p_2]\!] \circ \mathcal{S}[\![p_1]\!] \\
\mathcal{S}[\![\mathsf{farm} \; n \; p]\!] & = & \mathcal{S}[\![p]\!] \\
\mathcal{S}[\![\mathsf{fb} \; p]\!] & = & iter \; \mathcal{S}[\![p]\!] \\
\mathcal{S}[\![\mathsf{dc}_{n,T,F} \; f \; g]\!] & = & cata_F \; (\hat{\rho}(f)) \circ ana_F \; (\hat{\rho}(g))
\end{array}
$$

$$
\begin{array}{lll}
[\![p \; : \; T \; A \to T \; B]\!] & : & [\![T \; A \to T \; B]\!] \\
[\![p]\!] & = & map_T \; \mathcal{S}[\![p]\!]
\end{array}
$$

Here $\circ$ denotes the usual function composition; *iter* captures the iterative structure of a feedback skeleton (the usual *trampoline* function [9]); and *cata* and *ana* are the usual cata/anamorphisms. Instead of assuming a dynamic termination condition, we require the worker operation, p, of a feedback loop, $\mathsf{fb} \; p$ to return a sum type $A + B$. The feedback skeleton will apply p to a stream of inputs of type A, merging any results of type A with the other inputs, and returning any results of type B as its outputs. Since the *iter*, *cata*, *ana* and *map* operations are all instances of *hylomorphisms*, it follows that there is a single, underpinning theoretical formulation for all our skeleton forms, and it also follows that our structured programs may be easily transformed to introduce/remove parallelism in a provably sound way. We will achieve this using a type system.

3.3.2. The Structure-Annotated Type System

We annotate top-level program types with an abstraction of the *parallel structure* of the program, $\sigma \in \Sigma$. Figure 2 shows our structure-annotated type system. Intuitively, Σ is a "pruned" version of E that retains information about *how* the computation is performed, while removing as many details as possible about *what* is being computed. We once again split its definition into two levels, Σ_{s} for sequential structures, and Σ_{p} for parallel ones.

$$
\begin{array}{llll}
\sigma \; \in \Sigma & ::= & \sigma_{\mathsf{s}} \; | \; \mathrm{PAR}_F \; \sigma_{\mathsf{p}} \\
\sigma_{\mathsf{s}} \in \Sigma_{\mathsf{s}} & ::= & \mathrm{A} \; | \; \sigma \circ \sigma \; | \; \mathrm{HYLO}_F \; \sigma \; \sigma \\
\sigma_{\mathsf{p}} \in \Sigma_{\mathsf{p}} & ::= & \mathrm{FUN} \; \sigma_{\mathsf{s}} \; | & \mathrm{DC}_{n,F} \; \sigma_{\mathsf{s}} \; \sigma_{\mathsf{s}} \\
& | & \sigma_{\mathsf{p}} \| \sigma_{\mathsf{p}} \; | & \mathrm{FARM}_n \; \sigma_{\mathsf{p}} \; | \; \mathrm{FB} \; \sigma_{\mathsf{p}}
\end{array}
$$

3.3.3. Convertibility

Our type system needs to include a non-structural rule that captures the *convertibility relation*, $\equiv$, for Σ.

$$
\frac{\vdash e \; : \; A \xrightarrow{\sigma_1} B \quad \sigma_1 \equiv \sigma_2}{\vdash e \; : \; A \xmapsto{\sigma_2} B}
$$

This allows us to transform between structurally equivalent forms. For example, a *sequential composition* can be transformed into a *parallel pipeline* or *vice-versa*, or a *task farm* can be introduced wherever we are using any other skeleton. $\equiv$ is defined in terms of the relations $\equiv_{\mathsf{s}}$ (for sequential operations) and $\equiv_{\mathsf{p}}$ (for

8

$$\frac{\rho(f) = A \to B}{\vdash f \ : \ A \xrightarrow{\mathrm{A}} B} \qquad \frac{\begin{array}{c}\vdash e_1 \ : \ B \xrightarrow{\sigma_1} C \\ \vdash e_2 \ : \ A \xrightarrow{\sigma_2} B\end{array}}{\vdash e_1 \circ e_2 \ : \ A \xrightarrow{\sigma_1 \circ \sigma_2} C}$$

$$\frac{\begin{array}{c}\vdash e_1 \ : \ F\,B \xrightarrow{\sigma_1} B \\ \vdash e_2 \ : \ A \xrightarrow{\sigma_2} F\,A \quad G = \mathsf{base}\ F\end{array}}{\vdash \mathtt{hylo}_F\ e_1\ e_2 \ : \ A \xrightarrow{\mathrm{HYLO}_G\ \sigma_1\ \sigma_2} B} \qquad \frac{\begin{array}{c}\vdash p \ : \ T\,A \xrightarrow{\sigma} T\,B \\ F = \mathsf{base}\ T\end{array}}{\vdash \mathsf{par}_T\ p \ : \ T\,A \xrightarrow{\mathrm{PAR}_F\ \sigma} T\,B}$$

$$\frac{\vdash s \ : \ A \xrightarrow{\sigma} B}{\vdash \mathsf{fun}\ s \ : \ T\,A \xrightarrow{\mathrm{FUN}\ \sigma} TB}$$

$$\frac{\vdash s_1 \ : \ F\,B \xrightarrow{\sigma_1} B \quad \vdash s_2 \ : \ A \xrightarrow{\sigma_2} F\,A \quad G = \mathsf{base}\ F}{\vdash \mathsf{dc}_{n,F}\ s_1\ s_2 \ : \ T\,A \xrightarrow{\mathrm{DC}_{n,G}\ \sigma_1\ \sigma_2} T\,B}$$

$$\frac{n \ : \ \mathbb{N} \quad \vdash p \ : \ T\,A \xrightarrow{\sigma} T\,B}{\vdash \mathsf{farm}\ n\ p \ : \ T\,A \xrightarrow{\mathrm{FARM}_n\ \sigma} T\,B}$$

$$\frac{\vdash p_1 \ : \ T\,A \xrightarrow{\sigma_1} T\,B \quad \vdash p_2 \ : \ T\,B \xrightarrow{\sigma_2} T\,C}{\vdash p_1 \parallel p_2 \ : \ T\,A \xrightarrow{\sigma_1\ \parallel\ \sigma_2} T\,C}$$

$$\frac{\vdash p \ : \ T\,A \xrightarrow{\sigma} T\,(A+B)}{\vdash \mathsf{fb}\ p \ : \ T\,A \xrightarrow{\mathrm{FB}\ \sigma} T\,B}$$

Figure 2: Structure-Annotated Type System.

parallel skeletons), plus a lifting rule that links the two, PAR-EQUIV.

$$\frac{\sigma_1 \equiv_\mathsf{s} \sigma_2}{\sigma_1 \equiv \sigma_2} \qquad \frac{\sigma_1 \equiv_\mathsf{p} \sigma_2}{\mathrm{PAR}_F\ \sigma_1 \equiv \mathrm{PAR}_F\ \sigma_2}$$

$$\mathrm{PAR}_F\ (\mathrm{FUN}\ \sigma) \equiv \mathrm{MAP}_F\ \sigma \qquad (\text{PAR-EQUIV})$$

In [9], we have defined a number of equivalences, using well-known properties of skeletons (e.g. [2, 7]) plus fundamental hylomorphism laws [31]. For example, a parallel pipeline structure ($\parallel$) is functionally equivalent to a function composition; a task farm FARM can be introduced for any structure; and divide-and-conquer DC and feedback FB skeletons can both be derived from hylomorphisms.

$$\begin{array}{rcll}
\mathrm{FUN}\ \sigma_1 \parallel \mathrm{FUN}\ \sigma_2 & \equiv_\mathsf{p} & \mathrm{FUN}\ (\sigma_2 \circ \sigma_1) & (\text{PIPE-EQUIV}) \\
\mathrm{DC}_{n,F}\ \sigma_1\ \sigma_2 & \equiv_\mathsf{p} & \mathrm{FUN}\ (\mathrm{HYLO}_F\ \sigma_1\ \sigma_2) & (\text{DC-EQUIV}) \\
\mathrm{FARM}_n\ \sigma & \equiv_\mathsf{p} & \sigma & (\text{FARM-EQUIV}) \\
\mathrm{FB}(\mathrm{FUN}\ \sigma) & \equiv_\mathsf{p} & \mathrm{FUN}\ (\mathrm{ITER}\ \sigma) & (\text{FB-EQUIV})
\end{array}$$

These equivalences, plus *reflexivity, symmetry and transitivity*, define an equational theory that allows conversion between different parallel forms, as well as conversion between structured expressions and parallel forms.

9

3.4. Automatic Rewriting using Cost Information

We can use the convertibility relation to define a *rewriting system* that allows us to automatically rewrite functionally equivalent structured parallel forms. We interpret structured expressions and parallel processes as members of families of functionally equivalent expressions E_s, indexed by a representative structured expression s. For all well-typed $s : A \overset{\sigma}{\mapsto} B$, σ is an index of the family defined by s, i.e. $\phi_s(\sigma) \in E_s$. The function $\phi_s : \Sigma \to E_s$ is a partial function whose result is defined for any structure σ that is an index of the family E_s. The type-checking algorithm needs to decide whether rewriting s to the desired structure would result in a member of the family E_s. We achieve this by defining a *confluent* rewriting system on structures, $\Sigma \rightsquigarrow \Sigma$ [9]. Since, for a given s, the family E_s is finite, we can use cost information to find optimal structures. For small enough parallel structures, exhaustive search is feasible. For larger structures, we can equip a family E_s with a strong ordering $\sigma_1, \sigma_2 \in \text{dom}(\phi_s), \sigma_1 \leq \sigma_2$ iff $\text{cost}(\sigma_1) \leq \text{cost}(\sigma_2)$, and use this ordering to optimise the search process, at the cost of not necessarily finding an optimal structure[1] In order to exploit this technique, we must provide a cost model that is strongly connected to our execution semantics. This is the focus of the remainder of this paper.

4. An Operational Semantics for Queues

In order to obtain a sound cost model, we need to first give an operational semantics for each of our skeletons, and derive then cost information from this semantics. We will build our skeleton semantics (Sec. 5) on a small-step trace-based operational semantics for the queues that are used to link these skeletons. Each step in the queue semantics will describe a state transition within a simple parallel process abstraction. We first give a number of basic definitions.

State. A *state* comprises a tuple of three main structures, $\mathcal{W} \times \mathcal{Q} \times \mathcal{S}$:

- an environment of worker definitions, $\mathcal{W}$, which is a mapping from worker identifiers to *worker loop* definitions, i.e. to the code that is run by each worker of the parallel process;

- an environment of *worker states*, $\mathcal{S}$, which represents the instruction that is currently being executed by the worker; and,

- a *queue environment*, $\mathcal{Q}$, which represents the buffers that link the workers.

We will assume that the worker environment, $\mathcal{W}$, is fixed. The rules in our operational semantics therefore have the form:

$$\boxed{(\mathcal{Q}, \mathcal{S}) \overset{\alpha}{\to} (\mathcal{Q}', \mathcal{S}')}$$

---

[1] We conjecture that a family E_s that is equipped with an ordering relation such that $\sigma_1 \leq \sigma_2$ *iff* $\sigma_2 \rightsquigarrow \sigma_1$ forms a lattice, and that we can use this information, together with cost information, to optimise the search for an optimal structure within a family.

They are given in terms of a labelled state transition system. The labels are the actions $\alpha ::= \mathtt{g}^w q \mid \mathtt{e}^w(f, x) \mid \mathtt{p}^w(x, q)$, which respectively *get* an input from a queue, *evaluate* a function f on an input x, or *put* a result on a queue. Actions are performed on specific queues q, and are tagged with the worker, w, that performs the action.

Queue Environments. A queue environment is a mapping from queue identifiers to sequences of values.

$$
\mathcal{Q} \;=\; \left[\begin{array}{ll} q_0 & \mapsto \langle x_0,\ x_1, \ldots \rangle \\ \ldots & \\ q_n & \mapsto \langle \ldots \rangle \end{array} \right]
$$

Queue Operations. We assume the usual basic, thread-safe, *enqueue* and *dequeue* operations.

$$
\begin{array}{ll}
\mathtt{enqueue}\,(\mathcal{Q}[q \mapsto vs], x, q) & \rightarrow \quad \mathcal{Q}[q \mapsto \langle x \mid vs \rangle] \\
\mathtt{dequeue}\,(\mathcal{Q}[q \mapsto \langle vs \mid x \rangle], q) & \rightarrow \quad \mathcal{Q}[q \mapsto vs],\ x
\end{array}
$$

We overload the notation for *enqueue*/*dequeue* operations to also work on sums and products. Enqueuing a product type value into a "product queue" results in a pairwise *enqueue* on each sub-queue. Enqueuing a sum type value into a "sum queue" yields a single *enqueue* operation, on the corresponding queue. We use Q to refer to these *queue structures*, and q to refer to queue identifiers.

$$
Q \quad ::= \quad q \mid q_1 \times \cdots \times q_n \mid q_1 + \cdots + q_n
$$

We interpret the *enqueue* operation on products of queues as a sequence of simple *enqueue* operations:

$$
\begin{array}{l}
\mathtt{enqueue}\,(\mathcal{Q}, q_1 \times \cdots \times q_n, (x_1, \ldots, x_n)) \\
\quad = \quad \mathtt{enqueue}\,(\mathtt{enqueue}\,(\ldots \mathtt{enqueue}\,(\mathcal{Q}, q_1, x_1) \ldots), q_n, x_n)
\end{array}
$$

An *enqueue* on a sum of queues is an *enqueue* on the corresponding queue. For example, if $0 \leq i \leq n$:

$$
\mathtt{enqueue}\,(\mathcal{Q}, q_1 + \cdots + q_n, \mathtt{inj}_1\, x) = \mathtt{enqueue}\,(\mathcal{Q}, q_i, x)
$$

Dequeue operations work similarly. We interpret the *dequeue* operation on products of queues as a sequence of simple *dequeue* operations:

$$
\mathtt{dequeue}\,(\mathcal{Q}, q_1 \times \cdots \times q_n) = (\mathtt{dequeue}\,(\mathcal{Q}, q_1), \ldots, \mathtt{dequeue}\,(\mathcal{Q}, q_1))
$$

A *dequeue* operation on a sum of queues dequeues an element from the first non-empty queue. For example, if $0 \leq i \leq n$, and for all j, $0 \leq j < i$, $\mathcal{Q}[q_j \mapsto \langle \rangle]$, and $\mathcal{Q}[q_i \mapsto \langle vs|x \rangle]$, then:

$$
\mathtt{dequeue}\,(\mathcal{Q}, q_1 + \cdots + q_n) = \mathtt{inj}_i\,(\mathtt{dequeue}\,(\mathcal{Q}, q_i))
$$

This is a purely arbitrary choice. Any alternative ordering (e.g. round-robin) is equally acceptable. In the operational semantics, each action will represent an *enqueue* or *dequeue* operation on a single queue. So dequeuing from a tuple of n queues will result in n *dequeue* actions.

11

Workers. In each iteration, a worker first performs a sequence of *dequeues*, then performs its local computation, f, and finally performs a sequence of *enqueue* operations. We represent this as:

$$\mathcal{W} = \left[\begin{array}{c} \dots \\ w \mapsto \mathtt{worker}\,(Q_i,\ f,\ Q_o) \\ \dots \end{array} \right]$$

where Q_i and Q_o are the input and output queue structures. The example below shows the state of a worker that is in the process of dequeuing from a tuple of queue identifiers, and that has already *dequeue*d n elements from the first n queues:

$$\mathcal{S} = \left[\begin{array}{c} \dots \\ w \mapsto (x_1, \dots, x_n, \mathtt{dequeue}\,(q_{n+1}), \dots, \mathtt{dequeue}\,(q_m)) \\ \dots \end{array} \right]$$

Generally, a worker state can be either a sequence (or sum) of *dequeue*/*enqueue* operations, or an *eval* operation. Let v be a value:

$$
\begin{array}{lll}
\mathtt{st} & ::= \mathtt{inSt} & |\quad \mathtt{outSt} \quad | \quad \mathtt{eval}\,(x) \\
\mathtt{inSt} & ::= v & |\quad \mathtt{dequeue}\,(q) \quad | \quad (\mathtt{inSt}, \dots, \mathtt{inSt}) \\
\mathtt{outSt} & ::= v & |\quad \mathtt{enqueue}\,(q,x) \quad | \quad \mathtt{outSt}; \dots; \mathtt{outSt}
\end{array}
$$

The transition rules for our operational semantics are given in Figure 3. Note that we overload *enqueue*/*dequeue* operations to also deal with sums and products. We also assume that *enqueues* on simple queues are trivial. It is now straightforward to define the operational semantics on a sequence of workers:

$$\frac{(\mathcal{Q}, \mathtt{st}) \xrightarrow{\alpha^w} (\mathcal{Q}', \mathtt{st}')}{(\mathcal{Q}, \mathcal{S}[w \mapsto \mathtt{st}]) \xrightarrow{\alpha^w} (\mathcal{Q}', \mathcal{S}[w \mapsto \mathtt{st}'])}$$

It is easy to see that the only rules that affect the output are those involving communication (*enqueue*/*dequeue*).

Definition 4.1 (Ready and Idle State). *A worker $w \mapsto \mathtt{worker}\,(Q_i, f, Q_o)$ is in a* ready *state if it is at the beginning of its worker loop. A worker is* idle *if it is ready and there are no inputs in its input queue structure Q_i*

Definition 4.2 (Initial and Final State). *A parallel process is* initial *if all its workers are* ready *and all the queues except the input queue are empty. A parallel process is* final *if all its workers are* idle.

$$\frac{\texttt{dequeue}\,(\mathcal{Q}, q_{n+1}) \to (\mathcal{Q}', v)}{(\mathcal{Q}, (v_1, \ldots, v_n, \texttt{dequeue}\,(q_{n+1}), \ldots)) \xrightarrow{\texttt{g}^w q_{n+1}} (\mathcal{Q}', (v_1, \ldots, v_n, v, \ldots))}$$

$$\frac{\texttt{dequeue}\,(\mathcal{Q}, q_{n+1}) \to (\mathcal{Q}', v)}{(\mathcal{Q}, (v_1, \ldots, v_n, \texttt{dequeue}\,(q_{n+1}))) \xrightarrow{\texttt{g}^w q_{n+1}} (\mathcal{Q}', \texttt{eval}\,(v_1, \ldots, v_n, v))}$$

$$\frac{\texttt{enqueue}\,(\mathcal{Q}, q_1, x_1) \to \mathcal{Q}'}{(\mathcal{Q}, \texttt{enqueue}\,(q_1, x_1)\,;\ldots) \xrightarrow{\texttt{p}^w q_1} (\mathcal{Q}', \ldots)}$$

$$\frac{\texttt{enqueue}\,(\mathcal{Q}, q, x) \to \mathcal{Q}' \quad \mathcal{W}[w \mapsto \texttt{worker}\,(Q_i, f, Q_o)]}{(\mathcal{Q}, \texttt{enqueue}\,(q, x)) \xrightarrow{\texttt{p}^w q} (\mathcal{Q}', \texttt{dequeue}\,(Q_i, x))}$$

$$\frac{[\![f]\!](x) \;=\; y \quad \mathcal{W}[w \mapsto \texttt{worker}\,(Q_i, f, Q_o)]}{(\mathcal{Q}, \texttt{eval}\,(x)) \xrightarrow{\texttt{e}^w (x)} (\mathcal{Q}, \texttt{enqueue}\,(Q_o, y))}$$

Figure 3: Transition Rules for queue operations.

5. A Queue-Based Operational Semantics of Algorithmic Skeletons

We can now build the parallel structures that we require.

Parallel Composition of Processes. We define the parallel composition of processes as the union of their queues and workers.

$$(\mathcal{W}_1, \mathcal{Q}_1, \mathcal{S}_1) \parallel (\mathcal{W}_2, \mathcal{Q}_2, \mathcal{S}_2) = (\mathcal{W}_1 \uplus \mathcal{W}_2, \mathcal{Q}_1 \cup \mathcal{Q}_2, \mathcal{S}_1 \uplus \mathcal{S}_2)$$

Worker. We lift sequential functions into workers with input and output queue structures. The new worker is *ready*.

$$\texttt{qfun}(f)(Q_0, Q_1) =$$
$$\mathcal{W} = [w \mapsto \texttt{worker}\,(Q_0, f, Q_1)], \mathcal{Q} = Q_0 \cup Q_1, \mathcal{S} = [w \mapsto \texttt{dequeue}\,(Q_0)]$$

Task farm. A task farm replicates a structure n times:

$$\texttt{qfarm}(n, \mathcal{P})(Q_0, Q_1) = \overbrace{\mathcal{P}(Q_0, Q_1) \parallel \ldots \parallel \mathcal{P}(Q_0, Q_1)}^{n \text{ times}}$$

Parallel pipeline. A parallel pipeline is the parallel composition of two structures, linked by an intermediate queue. Let q be a fresh queue identifier:

$$\texttt{qpipe}(\mathcal{P}_1, \mathcal{P}_2)(Q_0, Q_1) = \mathcal{P}_1(Q_0, q) \parallel \mathcal{P}_2(q, Q_1)$$

13

Feedback loop. A feedback loop is created by simply inspecting the output of a structure, and placing it back in the input queue depending on the result:

$$\texttt{qfeedback}(\mathcal{P})(Q_0, Q_1) = \mathcal{P}(Q_0, Q_0 + Q_1)$$

Parallel Hylomorphism. One of the novelties of this paper lies in describing how to parallelise an arbitrary *hylomorphism* using a divide-and-conquer skeleton. We will first describe a series of simple semantics-preserving transformations for any hylomorphism. The idea is that if the *anamorphism* part of a hylomorphism needs to split an input value into at most n sub-values, then we will create $n + 1$ queues, the first of which will send the corresponding input to the "combine" worker, and the remaining n of which will send the subdivided inputs to the subsequent divide stages. If an input cannot be divided any further, then a synchronisation token, 1, will be sent. Given a functor F described as a combination of sums, products and constants, $F\ A = T$, where A is nested at most n times within a product in T, we define the functor D_F as $D_F\ B = (1 + T[1/A]) \times \overbrace{(B \times \cdots \times B)}^{n \text{ times}}$. If we can convert any hylomorphism to this new structure, then we can use its regular structure to create a regular "divide-and-conquer graph" with the following communication structure:

1. The *divide* worker will communicate a value of type $1 + T[1/A]$ to the corresponding *combine* worker, and values of type B to the subsequent *divide* workers.

   - A value of type $1+T[1/A]$ contains the "non-recursive" part of T, plus a unit to indicate that the input could not be divided any further.

   - A *combine* worker can use a value of this type to decide how to recombine the n values of type B that has been received from the previous stages of the divide-and-conquer skeleton.

2. A *divide* worker takes an element of type $1 + A$. If it is 1, then it transmits $\texttt{inj}_1$ () to all its output queues. This indicates to the subsequent workers that there is no more work to be done. If it receives a value of type A, it divides it, splits the recursive and non-recursive parts of T, and sends the corresponding elements to the output queues.

3. A *combine* worker that receives an element of type 1 from the corresponding *divide* worker will simply discard all values received from the previous level, and send 1 to the *combine* worker of the next level. If it receives an element of type $T[1/A]$, it will need to recombine the corresponding value of type $F\ A$ from the inputs, and apply the *combine* function to it.

We need a way to change a hylomorphism from an F structure to D_F. Note that there are two functions

$$
\begin{aligned}
d &: 1 + F\ A \to D_F\ (1 + A) \\
c &: D_F\ (1 + A) \to 1 + F\ A.
\end{aligned}
$$

14

The d function separates the occurrences of values of type A in some structure $F\ A$ into the corresponding $T[1/A]$ and n-tuple of $1 + A$, and the function c recomposes a structure $F\ A$ from the structure of D_F. If d receives $\mathtt{inj}_1\ ()$, then it just returns a $n+1$ tuple of $\mathtt{inj}_1\ ()$. The c function returns $\mathtt{inj}_1\ ()$ if the first component of the tuple $D_F\ (1 + A)$ is $\mathtt{inj}_1\ ()$. Note that the types F and D_F are not isomorphic, since the function c is *partial*. However, the following properties do hold:

$$c \circ d = \mathtt{id}$$
$$d \circ (\underline{1} + F\ f) = D_F\ (\underline{1} + f) \circ d.$$

Using these properties, we can show that the following condition holds for any functor F:

$$\underline{1} + \mathtt{hylo}_F\ g\ h$$
$$=$$
$$\underline{1} + g \circ F\ (\mathtt{hylo}_F\ g\ h) \circ h$$
$$=$$
$$(\underline{1} + g) \circ (\underline{1} + F\ (\mathtt{hylo}_F\ g\ h)) \circ (\underline{1} + h)$$
$$=$$
$$((\underline{1} + g) \circ c \circ d) \circ (\underline{1} + F\ (\mathtt{hylo}_F\ g\ h)) \circ (\underline{1} + h)$$
$$=$$
$$((\underline{1} + g) \circ c) \circ D_F\ (\underline{1} + \mathtt{hylo}_F\ g\ h) \circ (d \circ (\underline{1} + h))$$
$$=$$
$$\mathtt{hylo}_{D_F}\ ((\underline{1} + g) \circ c)\ (d \circ (\underline{1} + h))$$

Using this, we can convert to a hylomorphism of a different structure. For all $\mathtt{hylo}_F\ g\ h\ :\ A \to B$, given any $x\ :\ B$, then

$$\underline{x} \triangledown \mathtt{id} \circ (\mathtt{hylo}_{D_F}\ ((\underline{1} + g) \circ c)\ (d \circ (\underline{1} + h))) \circ \mathtt{inj}_2$$
$$=$$
$$\underline{x} \triangledown \mathtt{id} \circ (\underline{1} + \mathtt{hylo}_F\ g\ h) \circ \mathtt{inj}_2$$
$$=$$
$$(\underline{x} \circ \underline{1}) \triangledown (\mathtt{id} \circ \mathtt{hylo}_F\ g\ h) \circ \mathtt{inj}_2$$
$$=$$
$$\mathtt{id} \circ \mathtt{hylo}_F\ g\ h$$
$$=$$
$$\mathtt{hylo}_F\ g\ h$$

This shows that the first level of a divide-and-conquer must wrap the input in $\mathtt{inj}_2$, and the last level of a divide-and-conquer must unwrap the result using $\underline{x} \triangledown \mathtt{id}$, for any arbitrary $x\ :\ B$. We define these as:

$$\mathcal{D}_1(F, h) = d \circ \mathtt{inj}_2 \circ h \qquad \mathcal{C}_2(F, g) = (\underline{1} + g) \circ c$$
$$\mathcal{D}_2(F, h) = d \circ (\underline{1} + h) \qquad \mathcal{C}_1(F, g) = (\underline{x} \triangledown \mathtt{id}) \circ (\underline{1} + g) \circ c$$

Although the structure D_F may appear to be complicated, it neatly fits our queue-based model, in that we can create queues that send/receive values of
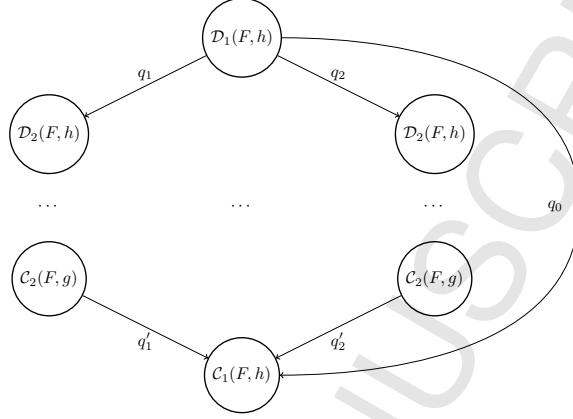
15

Figure 4: Divide-and-conquer skeleton: the q_i are the queues; circles represent the workers.

types $(1 + T[1/A])$ and $1 + A$, and can create workers that split/combine values of these types, as shown in Figure 4. Values of type 1 are simply used to synchronise the different *divide* and *combine* workers, so that the different levels of a divide-and-conquer skeleton can operate on different inputs in parallel. We define this as:

$$\mathtt{qdc}(n, F, g, h)(Q_0, Q_1) = \mathtt{qdc}'(\mathtt{hylo}_F g\ h,\ n,\ F,\ g,\ h,\ \mathcal{C}_1,\ \mathcal{D}_1)(Q_0, Q_1)$$

$$\mathtt{qdc}'(f, 0, F, g, h, c, d)(Q_0, Q_1) = \mathtt{qfun}(f)(Q_0, Q_1)$$
$$\mathtt{qdc}'(f, n, F, g, h, c, d)(Q_0, Q_1) =$$
$$\quad \mathtt{qfun}(d(F, h))(Q_0, q_0 \times q_1 \times \cdots \times q_n)$$
$$\quad \|\ \mathtt{qdc}'(\underline{1} + \mathtt{hylo}_F g\ h, n - 1, F, g, h, \mathcal{C}_2, \mathcal{D}_2)(q_1, q_1')$$
$$\quad \|\ \cdots$$
$$\quad \|\ \mathtt{qdc}'(\underline{1} + \mathtt{hylo}_F g\ h, n - 1, F, g, h, \mathcal{C}_2, \mathcal{D}_2)(q_n, q_n')$$
$$\quad \|\ \mathtt{qfun}(c(F, g))(q_0 \times q_1' \times \cdots \times q_n', Q_1)$$

Soundness. It is obvious from these definitions that the soundness of the operational semantics with respect to the denotational semantics can be reduced to correctly connecting the queues of the algorithmic skeletons.

6. Predicting Parallel Performance

We will now formally derive a set of cost equations from our operational semantics in a systematic way. These cost equations can then be used by our type-system to derive cost-models for the high-level parallel structures. This gives two main benefits: i) the cost models are sound w.r.t. the operational semantics *by construction*; and ii) this provides a way to *automatically* derive cost equations for newly defined parallel structures. Defining this automatic process is out of the scope of this paper, but we will discuss later how this can

16

be realised as an algorithm. In this paper, we are mainly interested in predicting realistic *lower bounds* on the execution times of parallel programs. While the more usual *worst-case* predictions might be useful for safety reasons, they would not provide sufficient information to enable us to choose the best parallel implementation. Moreover, when "initialisation" and "finalisation" times are taken into account, the worst-case parallel execution time will generally be very similar to the sequential case, and so provide very little insight into parallel performance. The *amortised* average case timings that we use here are much more useful for predicting the actual parallel performance of the system.

6.1. Costs and Sizes

The sequential components of our algorithmic skeletons are given as suitably lifted functions. Before we can define their cost models, we must first explain how we derive cost models for hylomorphisms. Vasconcelos [46] showed how to use sized types [24] for developing an accurate space cost analysis. Sized types have also been used for estimating upper bounds of execution times of parallel functional programs [28]. The success on using sized types for cost analysis motivated our approach. We exploit this previous research on *sized types* [46], as well as previous research on *cost equations* [40]. A *cost equation* for a function $f : A \to B$ provides an estimate of the execution time of f given an input of a given size.

$$\mathtt{cost}_f \, : \, \mathtt{size}_A \to \mathtt{time}$$

To obtain the inputs for these cost equations, we use a variant of the usual notion of *sized types* [24]. The only difference is that, since we are interested in amortised costs, we use *average sizes* rather than upper or lower bounds. The notion of *average size* is dependent on each problem, so rather than defining a generic calculation, we require specific definitions for each function and *type constructor*. Suitable definitions include, e.g. the average depth of a tree, the average number of elements in a structure etc.

Type Constructors. These are either constant types C, or recursive types defined as the fixpoint of some base functor μF.

Sizes and Size Constraints. We reuse the idea of *stages* [5]. Essentially, sizes are either size variables or sums of sizes, and size constraints are inequalities on sizes. We write $A^i \mid C$ for a sized type A with size i and constraint C.

Sized types. We require the basic polynomial (bi-) functors to be annotated with size expressions for each alternative (given as a sum-type). For example, the base bifunctor of a binary tree can be annotated as follows:

$$F^{0 \vee 1 + i + j} \, A \, B \; = \; 1 + A \times B^i \times B^j$$

The size expression $0 \vee 1 + i + j$ states that functor F either has size 0, i.e. it contains no elements of type A; or it has size $1 + i + j$, i.e. it has one element of type A, plus the sizes of the elements of type B, one of size i and the other

17

of size j. Note that there are alternative definitions for the size of a functor F, e.g: $F^{1 \vee \max(i,j)} \; A \; B \;=\; 1 + A \times B^i \times B^j$. In a more general sense, a sized-type is a type constructor that is annotated with size information, e.g. $Int^i, List^{j+k} \; A, \ldots$. The $\mathtt{in}_F$ and $\mathtt{out}_F$ functions for a sized base functor of a recursive type must have the following type and constraints:

$$\begin{aligned} \mathtt{in}_F &: \; F^j \mu F \to \mu^i F \mid i = j \\ \mathtt{out}_F &: \; \mu^i F \to F^j \mu F \mid i = j \end{aligned}$$

That is, we require that the sizes of the base functor represent the same information as the sizes of the recursive type. Finally, we will use $|A|$ to denote all the nested size information and size constraints in a type A. This is useful for defining cost models that require access to nested size information.

Deriving Recurrences from Hylomorphisms. We now show how we use sized types to derive recurrences. Essentially, we will define a *cost equation* for each syntactic construct, that takes some cost equation parameters and produces another cost equation. Although the cost of a function is not, in general, compositional, since it may depend on previous computations as well as on other properties of the input data, we will here take a compositional approach under the assumption that only the sizes will affect the execution times. This is a valid way to obtain *amortised* costs. The execution time of each primitive operation is assumed to be a constant that depends on the target architecture. This must be provided as a parameter. The cost of a function composition is the sum of the run-times of each stage, plus some architecture-dependent overhead. The only difficult case involves recursive functions. For recursive functions that are defined as hylomorphisms, we generate recurrence relations, as with Barbosa *et al.* [4]. In a similar sense to Sands' work [40], our cost equations can be thought of as *functions* that we are integrating into our type-system. For example, assuming that we have the following sized-type for *quicksort*:

$$\begin{aligned} \mathtt{merge} &: \; T^{0 \vee 1 + i + j} A \; (\mathsf{List} \; A) \;\to\; \mathsf{List}^k \; A \mid k = 0 \; \vee \; k = i + 1 + j \\ \mathtt{split} &: \; \mathsf{List}^n \; A \;\to\; T^{0 \vee 1 + r + s} A \; (\mathsf{List} \; A) \mid n = 0 \; \vee \; n = 1 + r + s \\[4pt] \mathtt{qs} &: \; \mathsf{List}^i \; A \;\to\; \mathsf{List}^j \; A \mid i = j \\ \mathtt{qs} &= \; \mathtt{hylo}_T \; \mathtt{merge} \; \mathtt{split} \end{aligned}$$

Given suitable cost equations for *split* and *merge*, $\mathtt{cost}_{\mathtt{split}}$, $\mathtt{cost}_{\mathtt{merge}}$, then there are two cases. Either $n = 0$, in which case $t = 0$, so we assume some time $\mathtt{cost}_{\mathtt{split}}(0) + \mathtt{cost}_{\mathtt{merge}}(0)$, or $n = 1 + r + s$, in which case:

$$\mathtt{cost}_{\mathtt{qs}} = \mathtt{cost}_{\mathtt{split}}(1 + r + s) + \mathtt{cost}_{\mathtt{merge}}(r + 1 + s) + \mathtt{cost}_{\mathtt{qs}}(r) + \mathtt{cost}_{\mathtt{qs}}(s)$$

To complete the cost equation, we need now to relate r and s with the input size n. By assuming that these sizes correspond to a balanced computation, we can then automatically calculate an *amortised* cost. Taking more extreme cases

18

into account would, of course, require further programmer input.

$$\mathtt{cost_{qs}}(n) =$$
$$\mathtt{cost_{split}}(1 + r + s) + \mathtt{cost_{merge}}(r + 1 + s) + \mathtt{cost_{qs}}(r) + \mathtt{cost_{qs}}(s)$$
$$\text{where } r = (n - 1)/2 \text{ and } s = (n - 1)/2$$

We simplify this internally to generate the desired recurrence relation:

$$\mathtt{cost_{qs}}(0) = \mathtt{cost_{split}}(0) + \mathtt{cost_{merge}}(0)$$
$$\mathtt{cost_{qs}}(n) = \mathtt{cost_{split}}(n) + \mathtt{cost_{merge}}(n) + 2 \times \mathtt{cost_{qs}}((n - 1)/2)$$

6.2. Costing Traces of Parallel Processes

We now describe a systematic way to derive cost equations. We start with a structure C, with some initial empty $\mathcal{P}_i$ and cost c_i. Taking suitably sound simplifications and approximations of $\mathtt{time}(\alpha_1, \ldots)$, we then derive an "amortised cost equation" $\mathtt{cost_C}$ such that for all input l with sized-type $\langle A \rangle^i$ and trace,

$$C(\mathcal{P}_1, \ldots, \mathcal{P}_n)(q_{in} \mapsto l, q_{out} \mapsto \langle \rangle) \xrightarrow{\alpha_1, \ldots} C(\mathcal{P}_1, \ldots, \mathcal{P}_n)(q_{in} \mapsto \langle \rangle, q_{out} \mapsto l'),$$

then $i \times \mathtt{cost_C}(c_1, \ldots, c_n) |A| \approx \mathtt{time}(\alpha_1, \ldots)$. We differentiate three phases in the execution of a parallel process: an *initialisation phase*, a *steady state*, and a final *flushing phase*. For example, a pipeline of two atomic functions, $w_1 \parallel w_2$, reaches a *steady state* after executing the *initialisation phase* of w_1. At this point, w_2 can run in parallel with w_1:

$$\left(\begin{bmatrix} q_0 & \mapsto \langle x_1, \ x_2, \ \ldots \rangle, \\ q_1 & \mapsto \langle \rangle, \\ q_2 & \mapsto \langle \rangle \end{bmatrix}, \begin{bmatrix} w_1 & \mapsto \mathtt{worker}\,(q_0, f, q_1) \\ w_2 & \mapsto \mathtt{worker}\,(q_1, g, q_2) \end{bmatrix} \right)$$

$$\xrightarrow{\mathtt{g}^{w_1} q_0, \ \mathtt{e}^{w_1} x_1, \ \mathtt{p}^{w_1} q_1}$$

$$\left(\begin{bmatrix} q_0 & \mapsto \langle x_2, \ \ldots \qquad \rangle, \\ q_1 & \mapsto \langle f \ x_1 \rangle, \\ q_2 & \mapsto \langle \rangle \end{bmatrix}, \begin{bmatrix} w_1 & \mapsto \mathtt{worker}\,(q_0, f, q_1) \\ w_2 & \mapsto \mathtt{worker}\,(q_1, g, q_2) \end{bmatrix} \right)$$

We capture these ideas in the definitions below.

Definition 6.1 (Steady State). *A parallel process $\mathcal{P}$ is in a steady state, $\mathtt{steady}(\mathcal{P})$, if for all $w \in \mathcal{P}$, $\neg \mathtt{idle}(w)$.*

Definition 6.2 (Initialisation Phase). *The initialisation phase for a parallel process $\mathtt{initial}(\mathcal{P}_1)$ is the shortest sequence of $a_1, a_2, \ldots, a_n$, such that if $\mathcal{P}_1 \xrightarrow{a_1, a_2, \ldots, a_n} \mathcal{P}_2$, then $\mathtt{steady}(\mathcal{P})$.*

Definition 6.3 (Flushing Phase). *The flushing phase for a parallel process $\mathcal{P}_1$ is the sequence of $a_1, a_2, \ldots, a_n$, such that for all $i \in [1 \ldots n]$ and for all $w \in \mathcal{P}_1$, $a_i \neq \mathtt{dequeue}\,(q_{in}\,)^w$, and if $\mathcal{P}_1 \xrightarrow{a_1, a_2, \ldots, a_n} \mathcal{P}_2$, then $\mathtt{final}(\mathcal{P}_2)$.*

19

$$
\begin{aligned}
&\mathtt{cost}(\text{FUN } \sigma^{n,m}) &&= T_{\text{enqueue}}(n) + \mathtt{cost}\, i + T_{\text{dequeue}}(m) \\
&\quad \text{where } |\sigma| = i \to j \mid C \\[1em]
&\mathtt{cost}(\text{FARM } n\ \sigma^{i,o}) &&= \frac{\mathtt{cost}(\sigma^{i \times n, o \times n})}{n} \\
&\mathtt{cost}(\sigma_1^{i,j} \parallel \sigma_2^{k,l}) &&= \max \left\{ \mathtt{cost}(\sigma_1^{i,j+k}), \mathtt{cost}(\sigma_2^{j+k,l}) \right\} \\[1em]
&\mathtt{cost}(\text{FB } \sigma^{n,m}) &&= \mathtt{if}\,(|\sigma| = 0) \\
&&&\quad \mathtt{then}\ \mathtt{cost}(\sigma^{n+m,m}) \\
&&&\quad \mathtt{else}\ \mathtt{cost}(\sigma^{n+m,m}) + \mathtt{cost}(\text{FB }(\mathtt{resize}(\sigma^{n,m}, |\sigma|))) \\[1em]
&\mathtt{cost}(\text{DC}_{n,F}\ \sigma_1\ \sigma_2) &&= \max \left\{ \begin{array}{l} \mathtt{cost}(\mathcal{D}(F, \sigma_2)), \\ \mathtt{cost}(\text{DC}_{n-1,F}\ \sigma_1\ \sigma_2), \\ \mathtt{cost}(\mathcal{C}(F, \sigma_1)) \end{array} \right\}
\end{aligned}
$$

Figure 5: Cost equations.

Total Cost and Amortised Cost:. Given some initial $\mathcal{P}$ and final $\mathcal{P}'$, where $\mathcal{P} \xrightarrow{\alpha_1,\ldots,\alpha_n} \mathcal{P}'$, the total cost of a parallel computation is $\mathtt{time}(\alpha_1, \ldots, \alpha_n)$. The cost of each action depends the worker environment. Based on [21], we calculate the queue contention on each queue by simply counting the number of workers in which a queue appears in $\mathcal{W}$. Then, the cost of the $\mathtt{g}^w$ and $\mathtt{p}^w$ is adjusted according to the corresponding overhead. If we split the trace into $\mathcal{P} \xrightarrow{\mathtt{init}} \mathcal{P}_1 \xrightarrow{\mathtt{steady}} \mathcal{P}_2 \xrightarrow{\mathtt{flush}} \mathcal{P}'$, where $\mathtt{init} = \alpha_1, \ldots, \alpha_i$, $\mathtt{steady} = \alpha_i, \ldots, \alpha_j$, $\mathtt{flush} = \alpha_j, \ldots, \alpha_n$ such that $\mathtt{initial}(\mathcal{P})$, $\mathtt{steady}(\mathcal{P}_1)$, and $\mathtt{final}(\mathcal{P}')$ then the total time is: $\mathtt{time}(\alpha_1, \ldots, \alpha_n) = \mathtt{time}(\mathtt{init}) + \mathtt{time}(\mathtt{steady}) + \mathtt{time}(\mathtt{flush})$. We can systematically derive cost models for $\mathtt{time}(\mathtt{init})$ and $\mathtt{time}(\mathtt{flush})$ using the method shown below.

6.3. Deriving Cost Equations from the Operational Semantics

We now show how to systematically derive the cost equations in Figure 5 from our operational semantics for skeletons. We base our approach on symbolic execution. The following assumptions are used to automatically derive amortised cost equations for our parallel structures.

1. The queues contain enough elements for each *dequeue* to succeed.

2. Each time an element is dequeued, we will return a size that is extracted from its type.

3. Evaluating a function on a size immediately returns the size of the output type of that function, plus the set of constraints that relate the input and output sizes.

4. An *enqueue* operation always succeeds, and has no effect on the queues.

20

5. Any substructure (e.g. a farm worker) produces a trace that can be safely interleaved with the trace that is produced by other substructures (i.e. without modifying the result of the overall computation), and that has a known cost.

Note that assumption 1 holds only for a `steady` structure, so these cost equations would not be useful for estimating run-times of a computation where `time(init)` and/or `time(flush)` dominate.

Worker cost. A worker, w, computing $f : A^i \rightarrow B^j \mid C$ in environment $\mathcal{W}[w \mapsto \text{worker}\,(q_i, f, q_o)]$ produces the "symbolic trace": $\mathsf{g}^w q_i$, $\mathsf{e}^w |A^i|$, $\mathsf{p}^w q_o$. Assuming there exists some trace $\alpha_1, \ldots, \alpha_n$ that can be safely interleaved with this trace, we want to know the cost of the actions in $\alpha_1, \ldots, \mathsf{g}^w q_i, \ldots,$ $\mathsf{e}^w \, i, \ldots, \mathsf{p}^w q_o, \ldots, \alpha_n$. Here, the cost of $\mathsf{g}^w q_i$ will depend on any other action that is happening in parallel. The only actions that can affect the cost of a queue operation are other queue operations that are acting on the same queue. If n is the number of workers $w_{\text{get}} \in \mathcal{W}$ such that $w_{\text{get}} \neq w_i$ and the number of workers operating on q_o is m, then if we assume that the cost of the *enqueue* and *dequeue* operations is as described by [21], then the cost is:

$$\text{time}(\mathsf{g}^w q_i, \ \mathsf{e}^w \, i, \ \mathsf{p}^w q_o) = T_{\text{enqueue}}(n) + \text{cost}_f(i) + T_{\text{dequeue}}(m)$$

We associate this cost with the corresponding high-level structure, so that it can be used by our type-checking algorithm:

$$\text{cost}(\text{FUN}\sigma) = T_{\text{enqueue}}(n) + \text{cost}(\sigma) + T_{\text{dequeue}}(m)$$

The parameters n, m and $\text{cost}(\sigma)$ can be instantiated from the context, and added to the structure σ by the type-checking algorithm in a straightforward way. We write $\sigma^{n,m}$ for a structure with n contending workers in the input queue, and with m contending workers in the output queue. Note that the real cost, understood as an equation from an input size to an execution time, would be $\text{cost}(\text{FUN}\sigma^{n,m})(T^i \, j) = i \times (T_{\text{enqueue}}(n) + \text{cost}_\sigma(j) + T_{\text{dequeue}}(m))$, if we assume that T contains i elements of size j. Since we can annotate structures, σ, with sizes, we can abuse our notation for extracting sizes of types, $|\sigma|$, and "overload" the meaning of cost to take a structure and yield an amortised cost.

Farm cost. A farm, $\mathcal{C}(Q_0, Q_1) \parallel \cdots \parallel \mathcal{C}(Q_0, Q_1)$, consists of a number of parallel processes that are joined using our parallel composition operator, and that share input and output queues. Since we symbolically evaluate the *enqueue* and *dequeue* operations, we can take n arbitrary traces, one for each farm worker:

$$\mathcal{C}(Q_0, Q_1) \xrightarrow{\alpha_1^1, \alpha_2^1, \cdots} \mathcal{P}', \quad \ldots, \quad \mathcal{C}(Q_0, Q_1) \xrightarrow{\alpha_1^n, \alpha_2^n, \cdots} \mathcal{P}'$$

Each trace has cost $c_i, \ldots, c_n$. To obtain amortised costs, we assume that each of these values is equal to some average cost c. Because we assume that these actions can be safely interleaved, and because the cost of the workers already

considers the queue contention, we simply need to calculate the maximum cost of any worker, assuming that each sub-trace can be performed in parallel:

$$\texttt{time}(\alpha_a^i, \alpha_b^j, \ldots) \;=\; \max \left\{ \begin{array}{l} \texttt{time}(\alpha_1^1, \alpha_2^1, \ldots), \\ \texttt{time}(\ldots), \\ \texttt{time}(\alpha_1^n, \alpha_2^n, \ldots) \end{array} \right\} \;=\; \max\{c_1, \ldots, c_n\} \;=\; c$$

Note that if each $\mathcal{P}$ produces k outputs, then $\texttt{qfarm}(n, \mathcal{P})$ produces $k \times n$ outputs. So, in order to obtain an amortised cost, when we associate it with the high-level structure, we need to divide the total cost by n.

$$\texttt{cost}(\textsc{farm} \; n \; \sigma) \;=\; \frac{\texttt{cost}(\sigma)}{n}$$

Note that, although this cost is obviously correct, the main point is that it was *systematically* derived from the simple queue-based model. It is thus *sound by construction*, and so no longer needs to be a parameter of our type system.

Parallel pipeline. A pipeline, $\mathcal{C}_1(Q_0, q) \parallel \mathcal{C}_2(q, Q_1)$, consists of two parallel processes that are joined using our parallel composition operator, and connected using an intermediate queue q. Again, we take a similar reasoning process:

$$\mathcal{C}_1(Q_0, q) \xrightarrow{\alpha_1^1, \alpha_2^1, \ldots} \mathcal{P}_1 \qquad \mathcal{C}_2(q, Q_1) \xrightarrow{\alpha_1^2, \alpha_2^2, \ldots} \mathcal{P}_2$$

Assume costs c_1, and c_2 for each process:

$$\texttt{time}(\alpha_a^i, \alpha_b^j, \ldots) \;=\; \max \left\{ \begin{array}{l} \texttt{time}(\alpha_1^1, \alpha_2^1, \ldots), \\ \texttt{time}(\alpha_1^2, \alpha_2^2, \ldots) \end{array} \right\} \;=\; \max\{c_1, c_2\}$$

Note that, in order to accurately lift these costs to the high-level structures, we need to consider the size of the output of α_i^1, not the size of the α_j^2 input. We do this by writing $\texttt{resize}(\sigma_2, |\sigma_1|)$, meaning that the input of σ_2 is altered to have the size of the output size in $|\sigma_1|$. We can do this safely since our type-checking algorithm can ensure that sizes meet the necessary constraints. Finally, we associate the cost with the corresponding high-level structure:

$$\texttt{cost}(\sigma_1 \parallel \sigma_2) \;=\; \max\{\texttt{cost}(\sigma_1), \texttt{cost}(\texttt{resize}(\sigma_2, |\sigma_1|))\}$$

Feedback loop. A feedback loop requires us to take into consideration that an element may be written back to the input queue, $\mathcal{C}_1(Q_0, Q_0 + Q_1)$. The structure C must compute some function f, and we require it to have type:

$$F^{n \vee 0} \; A \;=\; A^n + B, \quad \text{s.t.} \quad f : A^i \to F \; A^j \mid j < i.$$

Since we require j to be strictly smaller than i, we can estimate the number of steps that are required until $i = 0$ to be n, *i.e.* the average number of times an element will need to be put back into the input queue. Basically, a trace

$\mathcal{C}_1(Q_0, Q_0 + Q_1) \xrightarrow{\alpha_1^1, \alpha_2^1, \dots} \mathcal{P}_1$ will need to be taken n times to ensure that, on average, at least one element is enqueued into output queue Q_1.

$$\mathcal{C}_1(Q_0, Q_0 + Q_1) \xrightarrow{\alpha_1^1, \alpha_2^1, \dots} \mathcal{P}_1 \xrightarrow{\cdots} \mathcal{P}_n$$

Since this parameter can be estimated from the sized types, we can again assume that a structure σ is parameterised on it, and can use this in our high-level cost models. Again, we need to use `resize` to generate the appropriate cost equation:

$$
\begin{aligned}
\texttt{cost}(\textsc{fb}\ \sigma)\ =\ \texttt{if}\quad &(|\sigma| = 0) \\
&\texttt{then}\ \texttt{cost}(\sigma) \\
&\texttt{else}\ \texttt{cost}(\sigma)\ +\ \texttt{cost}(\textsc{fb}\ (\texttt{resize}(\sigma, \sigma)))
\end{aligned}
$$

Divide-and-Conquer. The divide-and-conquer skeleton requires a little more work, since we first need to transform the structure so that it matches the skeleton. Since this transformation can be done in a fairly standard way, we initially focus on the cost of the divide-and-conquer skeleton:

$$
\begin{aligned}
\texttt{qdc}'(0,\ F,\ g,\ h)(Q_0, Q_1)\ &=\ \texttt{qfun}(\underline{1} + \texttt{hylo}_F g\ h)(Q_0, Q_1) \\
\texttt{qdc}'(n,\ F,\ g,\ h)(Q_0, Q_1)\ &=\ \texttt{qfun}(\mathcal{D}_2(F, h))(Q_0, q_0 \times q_1 \times \cdots \times q_m) \\
&\quad \|\ \texttt{qdc}'(n-1, F, g, h)(q_1, q_1') \\
&\quad \|\ \dots \\
&\quad \|\ \texttt{qdc}'(n-1, F, g, h)(q_m, q_m') \\
&\quad \|\ \texttt{qfun}(\mathcal{C}_2(F, g))(q_0 \times q_1' \times \cdots \times q_m', Q_1)
\end{aligned}
$$

Since we define this skeleton inductively for some "depth" n, we start with the base case (depth 0), which is equivalent to the cost of an atomic function:

$$\texttt{cost}(\textsc{dc}_{0,F}\ \sigma_1\ \sigma_2)\ =\ \texttt{cost}(\textsc{fun}\ (\textsc{hylo}_F\ \sigma_1\ \sigma_2))$$

For the recursive case, we assume $2 + m$ traces, one for each "recursive" case, plus the trace of the divide and the trace of the combine parts. Again, we assume that the traces can be safely interleaved, so we can calculate the cost of the total trace as:

$$
\texttt{time}(\alpha_1, \dots) = \max \left\{
\begin{array}{l}
\texttt{time}(\alpha^{\text{div}}, \dots), \\
\texttt{time}(\alpha^{\text{qdc}}, \dots), \\
\texttt{time}(\dots), \\
\texttt{time}(\alpha^{\text{qdc}}, \dots), \\
\texttt{time}(\alpha^{\text{conq}}, \dots)
\end{array}
\right\} = \max \left\{
\begin{array}{l}
\texttt{time}(\alpha^{\text{div}}, \dots), \\
\texttt{time}(\alpha^{\text{qdc}}, \dots), \\
\texttt{time}(\alpha^{\text{conq}}, \dots)
\end{array}
\right\}
$$

We can then associate this cost equation with the costs of the high-level structures, by simply substituting the cost of the relevant trace. Note that the elements in the queues decrease in size for each level that we descend into the divide-and-conquer structure. Assuming that our structures are annotated with sizes, we need to update the sizes accordingly in the recursive call to the cost of a divide-and-conquer structure, and in the combine part. Although we can once

23

again use `resize`, we assume that the actual sizes have been correctly updated in the DC structure:

$$\mathsf{cost}(\mathrm{DC}_{n,F}\ \sigma_1\ \sigma_2)\ =\ \max \left\{ \begin{array}{l} \mathsf{cost}(\mathcal{D}(F,\sigma_2)), \\ \mathsf{cost}(\mathrm{DC}_{n-1,F}\ \sigma_1\ \sigma_2), \\ \mathsf{cost}(\mathcal{C}(F,\sigma_1)) \end{array} \right\}$$

7. Examples

This section a) illustrates the cost equations that can be derived from some parallel structures, as they are inferred by the type system; and, b) compare the costs that are obtained by cost models for some simple structures with actual runtimes, in order to provide evidence for the feasibility of our approach.

7.1. Type-based Derivation of Cost Equations

7.1.1. Image Merge

Image merge basically composes two functions: `mark` and `merge`. It can be directly parallelised using different combinations of farms and pipelines.

$$\mathrm{IM}\ n\ =\ \mathrm{PAR}_L\ (m_1 \parallel \mathrm{FARM}\ n\ m_2)$$

$$\texttt{imgMerge} : \mathsf{List}^k(\mathsf{Img} \times^{i+j} \mathsf{Img}) \xmapsto{\mathrm{IM}\ n} \mathsf{List}^k(\mathsf{Img}^{\max(i,j)})$$

$$\texttt{imgMerge}\ =\ \mathsf{map}_{\mathsf{List}}\ (\texttt{merge} \circ \texttt{mark})$$

The type system tries the following substitutions for m_1 and m_2 in IM:

$$
\begin{array}{lll}
\delta & = \{m_1 \sim \text{FUN A}, & m_2 \sim \text{FUN A}\} \\
\delta_1 & = \{m_1 \sim \text{FARM } n_1\ (\text{FUN A}), & m_2 \sim \text{FARM } n_2\ (\text{FUN A})\} \\
\delta_2 & = \{m_1 \sim \text{FUN A}, & m_2 \sim \text{FARM } n_2\ (\text{FUN A})\} \\
\delta_3 & = \{m_1 \sim \text{FARM } n_1\ (\text{FUN A}), & m_2 \sim \text{FUN A}\} \\
\delta_4 & = \{m_1 \sim \text{FUN A}, & m_2 \sim \text{FUN A}\}
\end{array}
$$

Note that the substitutions that introduce a farm to m_2 would yield a structure FARM n (FARM n_2) $\_$. In our model, this would not be a problem, since the underlying structure would be equivalent to FARM $(n \times n_2)$ $\_$. We could improve our type-checking implementation by considering the context of a metavariable. This would avoid superfluous rewritings. In our example, we assume that $sz = [d]^{1000}$. This represents the size of 1000 pairs of images of d dimensions. Arithmetic operations on sz are applied to the superscript. The size function of the first stage $|A_{c_1}|$ is the identity, since we are not modifying the images. The parameters for the number of farm workers are fixed to be those with the least cost, given some maximum number of available cores. For δ_1, we determine that $n_1 = 9$, $n = 3$ and $n_2 = 5$. The values of the costs on those sizes, and the overheads of farms and pipelines are given below. In the example cost equation below, $|A_{c_1}|(sz)$ denotes the output size of the first stage, obtained using `resize`, and κ is the overhead of the parallel structure. In the following examples, we

24

omit these numbers and just assume a 24-core machine with similar overheads for farms, pipelines, divide-and-conquer and feedbacks.

$$c_1 \left[(2048 \times 2048, 2048 \times 2048)\right]^n = \times 25.11ms$$
$$c_2 \left[(2048 \times 2048, 2048 \times 2048)\right]^n = \times 45.21ms$$
$$\kappa(9) = 29.66ms \qquad \kappa(15) = 60.93ms$$

$$\text{cost } (\delta_1 \text{IM}(n)) \ sz$$
$$= \max \left\{ c_1 \ (\frac{sz}{n_1}) + \kappa(n_1 + n_2), \ c_2 \ (\frac{|A_{c_1}|(sz)}{n_2}) + \kappa(n_1 + n_2) \right\} = 3145.69ms$$
$$\text{cost } (\delta_2 \text{IM}(n)) \ sz = 25123.81ms$$
$$\text{cost } (\delta_3 \text{IM}(n)) \ sz = 3189.60ms$$
$$\text{cost } (\delta_4 \text{IM}(n)) \ sz = 25123.81ms$$

The structure that results from applying δ_1 is the least cost one, $\delta_1(\text{IM}_2(3))$, with $n_1 = 9$ and $n_2 = 5$.

7.1.2. Quicksort

We will now analyse *quicksort* and show how it can exploit a divide-and-conquer parallel structure. Our cost models show the following estimations, where κ is adjusted to take into account the overheads of a DC structure, again calculated as the cost of the queue contention. In this example, we set the size parameter of our cost model to 1000 lists of 3,000,000 elements:

$$\text{qsorts} \ : \ \text{List}(\text{List } A) \mapsto \text{List}(\text{List } A)$$
$$\text{qsorts} = \text{map}_{\text{List}} \ (\text{hylo}_{F\,A} \ \text{merge} \ \text{div})$$

$$\text{cost } \left(\text{PAR}_L \ (\text{DC}_{n,F} \ A_{c_1} \ A_{c_2})\right) \ sz$$
$$= \max\{c_2 \ (|A_{c_2}|^i sz), \dots$$
$$\quad, \text{cost } \left(\text{HYLO}_F \ A_{c_1} \ A_{c_2}\right) \ (|A_{c_2}|^n sz), \dots$$
$$\quad, \max\{c_1 \ (|A_{c_1}|^i |A_{c_2}|^n sz)\} = 42602.72ms$$
$$\text{cost } \left(\text{PAR}_L \ (\text{FARM}_n \ (\text{FUN } (\text{ANA}_L \ A_{c_2})) \ \| \ (\text{FUN } (\text{CATA}_L \ A_{c_1})))\right) \ sz$$
$$= 27846.13ms$$
$$\text{cost } \left(\text{PAR}_L \ (\text{FARM}_n \ (\text{FUN } (\text{HYLO}_F \ A_{c_1} \ A_{c_2})))\right) \ sz$$
$$= 32179.77ms$$
$$\dots$$

Since the most expensive part of the *quicksort* function is the *divide*, and since flattening a tree is linear, the cost of adding a farm to the *divide* part is less than the cost of using a divide-and-conquer skeleton for this example.

7.1.3. N-Body Simulation

N-Body simulations are widely used in astrophysics. They comprise a simulation of a dynamic system of particles, usually under the influence of physical forces. The Barnes-Hut simulation recursively divides the n bodies, storing them in a `Octree`, or a 8-ary tree. Each node in the tree represents a region of the space, where the topmost node represents the whole space and the eight children

represent the eight octants of the space. The leaves of the tree contain the bodies. We then calculate the cumulative mass and centre of mass of each region of the space. Finally, the algorithm calculates the net force on each particular body by traversing the tree, and updates its velocity and position. This process is repeated for a number of iterations. We will here abstract most of the concrete, well known details of the algorithm, and present its high-level structure, using the following types and functions:

$$
\begin{array}{lll}
C & = & \mathbb{Q} \times \mathbb{Q} \\
F\ A\ B & = & A\ +\ C \times B^8 \\
G\ A & = & F\ \texttt{Body} \\
\texttt{Octree} & = & \mu G \\
\texttt{insert} & : & \texttt{Body} \times \texttt{Octree} \to \texttt{Octree}
\end{array}
$$

Since this algorithm also involves iterating for a fixed number of steps, we define iteration as a hylomorphism, as well as its structure.

$$
\begin{array}{l}
\textsc{loop}\ :\ \Sigma \to \Sigma \\
\textsc{loop}\ \sigma\ =\ \textsc{hylo}_{(+)}\ \big(\textsc{id}\nabla\sigma\big)\ \big((\textsc{a} + (\textsc{a} \triangle (\textsc{a} \circ \textsc{a}))) \circ (\textsc{a} \circ \textsc{a}?)\big)
\end{array}
$$

$$
\begin{array}{l}
\texttt{loop}_A\ :\ (A \overset{m}{\longmapsto} A) \to A \times \mathbb{N} \overset{\textsc{loop}\ m}{\longmapsto} A \\
\texttt{loop}_A\ s\ = \\
\quad \texttt{hylo}_{(A+)}\ (\text{id} \ \triangledown\ s) \\
\qquad\qquad \big((\pi_1 + (\pi_1 \triangle ((-1) \circ \pi_2))) \circ ((== 0) \circ \pi_2)?\big)
\end{array}
$$

This example uses some additional functions: `calcMass` annotates each node with the total mass and centre of mass; `dist` distributes the octree to all the bodies, to allow independent calculations, `calcForce` calculates the force of one body; and `move` updates the velocity and position of the body.

$$
\begin{array}{l}
\texttt{calcMass}\ :\ G\ \texttt{Octree} \to G\ \texttt{Octree} \\
\texttt{dist}\ :\ \texttt{Octree} \times \texttt{List Body} \\
\qquad\qquad \to L\ (\texttt{Octree} \times \texttt{Body})\ (\texttt{Octree} \times \texttt{List Body})
\end{array}
$$

The algorithm is:

$$
\begin{array}{l}
\texttt{nbody}\ :\ \texttt{List Body} \times \mathbb{N} \overset{\textsc{loop}\ \sigma}{\longmapsto} \texttt{List Body} \\
\texttt{nbody}\ =\ \texttt{loop}\ (\texttt{ana}_L\ (L\ (\texttt{move} \circ \texttt{calcForce}) \circ \texttt{dist}) \\
\qquad\qquad \circ((\texttt{cata}_G\ (\texttt{in}_G \circ \texttt{calcMass}) \circ \texttt{cata}_L\ \texttt{insert}) \triangle \text{id}))
\end{array}
$$

Note that we do not allow the fixed structure determined by LOOP to be rewritten. Even if we were to allow this, there is a data dependency in its definition that suggests that no parallelism can be extracted from it. However, the loop body does contain sources of parallelism, which our type system is able to exploit. In particular, the structure of the loop body is:

$$
\sigma = \textsc{ana}_L\big(L\big(\textsc{a} \circ \textsc{a}\big) \circ \textsc{a}\big) \circ \big(\textsc{cata}_G(\textsc{in} \circ \textsc{a}) \circ \textsc{cata}_L\ \textsc{a}\big) \triangle \textsc{id}
$$

The normalised structure reveals more opportunities for parallelisation:

$$
\sigma = \textsc{map}_L\textsc{a} \circ \textsc{map}_L\ \textsc{a} \circ \textsc{ana}_L\ \textsc{a} \circ \big(\textsc{cata}_G(\textsc{in} \circ \textsc{a}) \circ \textsc{cata}_L\ \textsc{a}\big) \triangle \textsc{id}
$$

After normalisation, this structure is equivalent to:

$$\sigma = \text{PAR}_L \left(\text{FUN}\left(\text{A} \circ \text{A}\right)\right) \circ \_$$

The structure makes it clear that there are many possibilities for parallelism using farms and pipelines. As before, parallelism can be introduced semi-automatically using our cost models. For example, setting the input size to 20,000 bodies:

$$\sigma \; = \text{PAR}_L \left(\text{FARM } n \; \_ \parallel \_\right) \circ \_$$
$$\sigma' = \text{PAR}_L \left(\texttt{min cost}\left(\_ \parallel \_\right)\right) \circ \_$$

$$\texttt{cost}\left(\text{FUN } \text{A}_{c_1} \parallel \text{FUN } \text{A}_{c_2}\right) sz \;\; = \;\; 310525.67ms$$
$$\texttt{cost}\left(\text{FARM}_6\left(\text{FUN } \text{A}_{c_1}\right) \parallel \left(\text{FUN } \text{A}_{c_2}\right)\right) sz \;\; = \;\; 55755.43ms$$
$$\texttt{cost}\left(\text{FUN } \text{A}_{c_1} \parallel \text{FARM}_1\left(\text{FUN } \text{A}_{c_2}\right)\right) sz \;\; = \;\; 310525.67ms$$
$$\texttt{cost}\left(\text{FARM}_{20}\left(\text{FUN } \text{A}_{c_1}\right) \parallel \text{FARM}_4\left(\text{FUN } \text{A}_{c_2}\right)\right) sz \;\; = \;\; 15730.46ms$$

7.1.4. Iterative Convolution

Image convolution is also widely used in image processing applications. We assume the type Img of images, the type Kern of kernels, the functor $F\ A\ B = A + B \times B \times B \times B$, and the *split*, *combine*, *kern* and *finished* functions. The split function splits an image into 4 sub-images with overlapping borders, as required for the kernel. The combine function concatenates the sub-images in the corresponding positions. The kern function applies a kernel to an image. Finally, the finished function tests whether an image has the desired properties, in which case the computation terminates. We can represent image convolution on a list of input images as follows:

$$\text{conv} : \text{Kern} \to (\text{List Img} \overset{\sigma}{\mapsto} \text{List Img})$$
$$\text{conv } k =$$
$$\text{map}_{\text{List}} \left(\text{iter}_{\text{Img}} \left(\texttt{finished?} \circ \text{hylo}_F \left(\texttt{combine} \circ F \left(\texttt{kern } k\right)\right)\right.$$
$$\left.\left(\texttt{split } k\right)\right))$$

The structure of conv is equivalent to a feedback loop, which exposes many opportunities for parallelism. Again, we assume a suitable cost model. Our estimates are given for 1000 images, of size 2048×2048.

$$
\begin{aligned}
\sigma \;=\;& \text{PAR}_L \left(\text{FB} \left(\text{DC}_{n,L,F} \left(\text{A} \circ F \text{ A} \right) \text{A} \parallel \_ \right) \right) \\
=\;& \text{PAR}_L \left(\text{FB} \left(\text{ FARM } n \;\_ \parallel \; \_ \parallel \_ \right) \right) \\
=\;& \text{min cost} \left(\text{PAR}_L \left(\text{FB} \left(\_ \parallel \; \_ \right) \right) \right) \\
=\;& \ldots
\end{aligned}
$$

$$
\begin{aligned}
& \text{cost} \left(\text{PAR}_L \left(\text{FB} \left(\text{A}_{c_1} \parallel \text{A}_{c_2} \right) \right) \right) \, sz \;=\; \\
& \quad \sum_{1 \le i, |\text{A}_{c_1} \parallel \text{A}_{c_2}|^i sz > 0} \text{cost} \left(\text{A}_{c_1} \parallel \text{A}_{c_2} \right) \left(|\text{A}_{c_1} \parallel \text{A}_{c_2}|^i sz \right) \\
& \quad =\; 20923.02ms \\
& \text{cost} \left(\text{PAR}_L \left(\text{FB} \left(\text{FARM}_4 \left(\text{FUN } \text{A}_{c_1} \right) \parallel \left(\text{FUN } \text{A}_{c_2} \right) \right) \right) \right) \, sz \\
& \quad =\; 6649.55ms \\
& \text{cost} \left(\text{PAR}_L \left(\text{FB} \left(\text{FUN } \text{A}_{c_1} \parallel \text{FARM}_1 \left(\text{FUN } \text{A}_{c_2} \right) \right) \right) \right) \, sz \\
& \quad =\; 20923.02ms \\
& \text{cost} \left(\text{PAR}_L \left(\text{FB} \left(\text{FARM}_{14} \left(\text{FUN } \text{A}_{c_1} \right) \parallel \text{FARM}_4 \left(\text{FUN } \text{A}_{c_2} \right) \right) \right) \right) \, sz \\
& \quad =\; 2694.30ms \\
& \ldots
\end{aligned}
$$

Collectively, these four examples have demonstrated the use of our techniques for all the parallel structures that we consider in this paper. In the next section, we provide evidence to validate the cost models that we have derived from our operational semantics against actual parallel executions.

7.2. Actual vs Predicted Speedups

In order to validate our results, we have compared predicted *versus* actual speedups for a number of example parallel programs. Our results not only validate our cost equations, but also confirm our previous experimental results [21]. We have taken the results of type-checking a number of examples, and implemented the corresponding structures in C, following our model, using the *pthreads*-based queue implementation that we described in [21]. In Section 6.3, we showed that most of the cost equations can be simplified to the expected ones, apart from the cost of a FUN σ structure. Basically, the idea is that the overhead of the parallel structures can be safely pushed down into its workers. We show here how that overhead cost affects the actual runtimes, comparing this with the predicted speedups for some of our parallel structures.

We use two different real multicore machines: *titanic*, a 800MHz 24-core, AMD Opteron 6176, running Centos Linux 2.6.18;and *lovelace*, a 1.4GHz 64-core, AMD Opteron 6376, running GNU/Linux 2.6.32.All the speedups shown here have been calculated as the mean of ten runs. Figure 6 shows the real *vs*. predicted speedups for task farms, using a simple *matrix multiplication* example; Figure 7 shows the corresponding speedups for the *image merge* example, parallelised as a farm of a pipeline; and Figure 8 shows the corresponding speedups for the *image convolution* example. The latter example was originally described using a similar structured expression to that for *image merge*, i.e. as the composition of two functions, read and process. However, the cost models predict that a different parallel structure is optimal: a pipeline of two farms. Finally, Figures 9 and 10 compare several different possible parallel structures
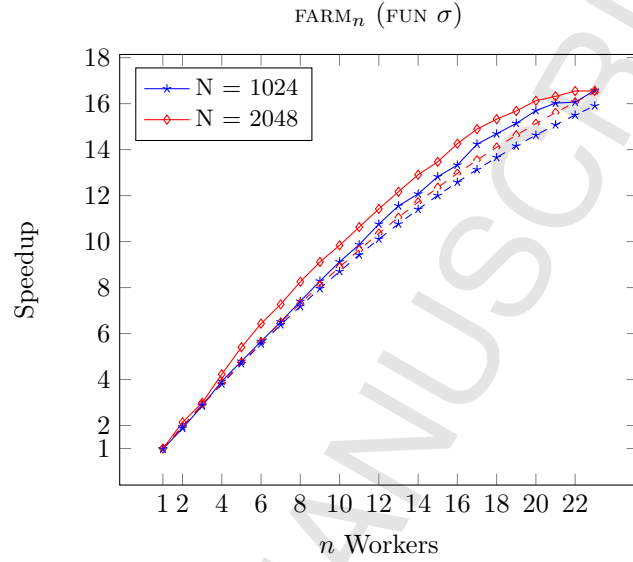
28

FARM$_n$ (FUN σ)

Figure 6: Speedup (solid lines) vs prediction (dashed lines). Matrix Multiplication of matrices of sizes N×N (*titanic*).

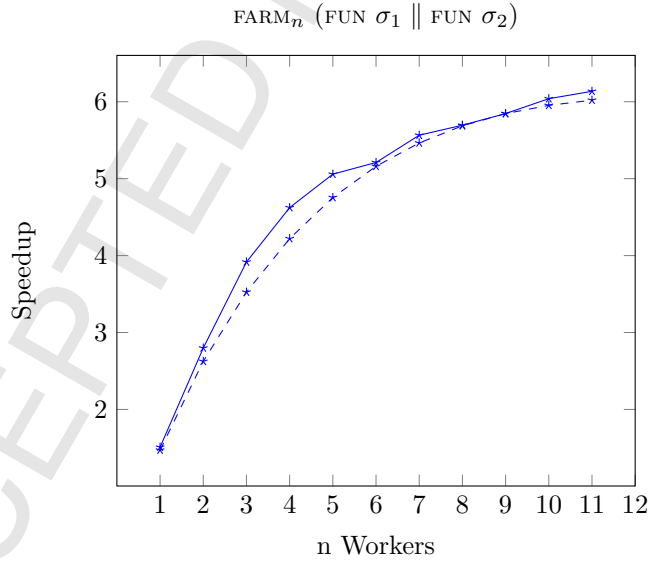FARM$_n$ (FUN σ_1 ‖ FUN σ_2)

Figure 7: Speedup (solid lines) vs prediction (dashed lines). Image Merge, 500 input tasks (*titanic*).

over farms and pipelines for the *image convolution* example. These examples collectively verify the experimental results that we previously showed in [21], and show that the cost models defined in this paper are able to correctly capture
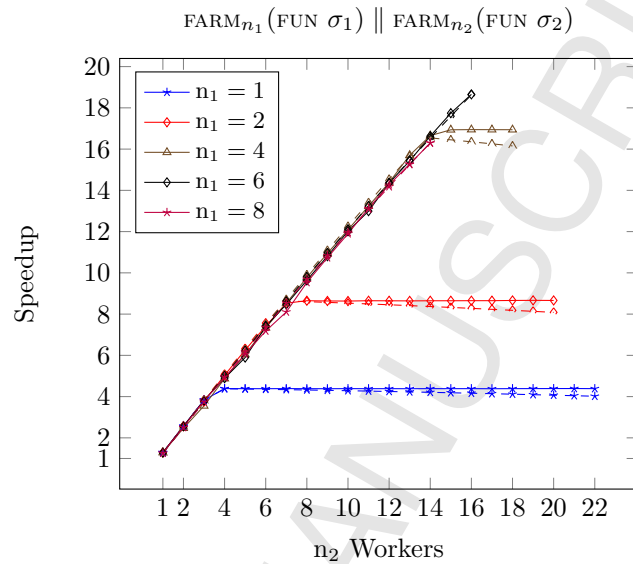
Figure 8: Speedup (solid lines) vs prediction (dashed lines). Image Convolution, 500 input tasks (*titanic*).
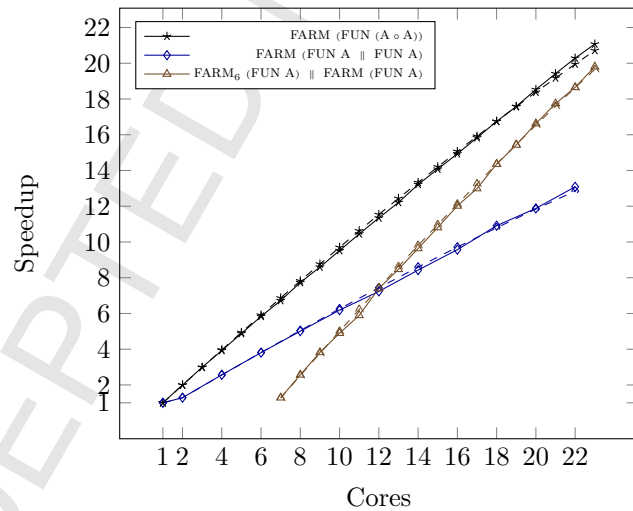
Figure 9: Speedup (solid lines) vs predicted (dashed lines). Different Parallel Structures for Image Convolution, 500 Images 1024 * 1024: *titanic*

queue contentions. We are thus able to predict a safe, tight upper bound on execution times.
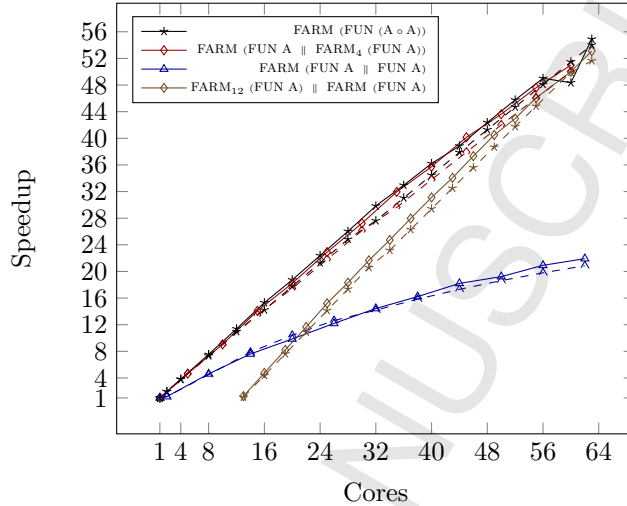
Figure 10: Speedup (solid lines) vs predicted (dashed lines). Different Parallel Structures for Image Convolution, 500 Images 1024 * 1024: *lovelace*.

8. Related Work

There have been a few previous treatments of parallelism using types, but these deal only with sizes and productivity. One line of work uses *sized types* [24] to incorporate a notion of the sizes of streaming data into a type system. This has been extended to cover a small number of skeletons in the Eden dialect of Haskell [37]. More recently, López et al [29] have presented a type-based methodology, inspired by session types, for verifying that parallel MPI programs follow some given protocol. This approach provides a scalable methodology to statically guarantee that a program satisfies a number of interesting safety properties, such as the absence of deadlocks. Both approaches show that types are indeed useful to prove a number of interesting properties of parallel programs, e.g. termination and productivity. In contrast, our work focuses on the different, but equally important, properties of semantic equivalence and cost.

The expressive power of hylomorphisms for parallel programming was first explored by Fischer and Gorlatch [12], who showed that a programming language based on catamorphisms and anamorphisms is *Turing-universal*. The idea of using hylomorphisms for parallel programming also appears in Morihata's work [33]. Morihata explores a theory for developing parallelisation theorems based on the *third homomorphism theorem* and *shortcut fusion*, and generalises it to hylomorphisms. In contrast, our work directly exploits the properties of hylomorphisms, in order to choose a suitable parallel skeleton implementation for hylomorphisms. The two lines of work are therefore orthogonal, and we can potentially benefit from Morihata's results in our future work.

Deriving parallel implementations from small, simple specifications has been widely studied. The third homomorphism theorem, list homomorphisms, and

31

the Bird-Meertens Formalism are amongst the many techniques that have been explored [18, 22, 23, 26, 30, 32, 34, 39, 42, 41]. The third homomorphism theorem states that if a function can be written both as a left fold and a right fold, then it can also be evaluated in a divide-and-conquer manner [16]. This theorem has been widely used for parallelism [10, 14, 15, 17, 27, 33, 35]. The majority of this work enables suitable automation and derivation of efficient parallel implementations. Our work differs in that we allow part of the parallel structure to be chosen in an automated way. This adds flexibility, enabling a parallel implementation to be changed quickly and easily by changing only a single type annotation. One possible extension of our work is to include some automatic transformations derived from the third homomorphism theorem. By parameterising our type system over some cost function on parallel structures, we smoothly integrate the introduction of parallelism with the ability to reason about the run-time behaviour of the parallel program. Skillicorn and Cai [43] have previously shown the utility of such an integration of a cost calculus with derivational software development, illustrating the approach for the Bird-Meertens theory of lists. In this paper, we take this approach one step further by using a more general equational theory based on hylomorphisms. Moreover, our type-based approach introduces new benefits, by providing a mechanism for specifying new parallel structures whose denotational semantics can be described as a composition of hylomorphisms.

Specific algorithmic skeletons are frequently associated with timing cost models [20]. However, these are almost exclusively obtained through measurement or approximation rather than being systematically derived from an operational semantics, as here. Much of the work on developing cost models for parallel execution has focused on data parallelism. The Parallel Random-Access Machine (PRAM) execution model [13] acts as a theory of complexity for parallel algorithms on idealised shared-memory SIMD machines. In the basic PRAM execution model, basic computations and shared-memory accesses are both assumed to take unit time. Unfortunately, PRAM costs underestimate actual machine execution costs, but in an unpredictable way [43]. The Bulk Synchronous Parallel (BSP) model [45] extends the PRAM model in a more realistic way, introducing a synchronising communication step after each set of computation steps. Lisper [19] has investigated the use of a Bulk Synchronous Parallel skeleton for determining worst-case execution times, but only informally, and not in the context of real processor models. Skillicorn and Cai have likewise developed a high-level cost calculus for data parallel computations [43], based on the *shape* of data structures, and known properties of primitive parallel operations, but have not based this on a strong machine-level semantics. Blelloch and Greiner have demonstrated provable time and space bounds for nested data parallel computations in NESL [6]. None of these models has the generality of the one given here. By linking the model given in this paper with the machine-level operational semantics for queueing that we have previously produced [21], we are now able to give a complete operational model for parallel algorithmic skeletons that accounts for all levels from the memory model upwards, including dealing with real memory effects on *x86* architectures.

32

In a structured parallelism setting, Janjic *et al.* [25] define a high-level DSL, the Refactoring Pattern Language (RPL), that aims to represent the parallel structure of a program, and capture its execution time. This DSL is a powerful tool, since it allows suitable parallelisations to be found for a given program, and then to apply them to a real C++ program. There are a number of differences between the approach that is described here and RPL. First, our type-based approach does not need to realise parallelisations as refactoring rules: parallel structures are tied to programs in a systematic way by each syntactic construct. The advantage of our approach is that we can use type information to automatically generate parallel code at compile-time. The corresponding disadvantage is lack of flexibility: the RPL approach can be combined with a number of refactorings that can take into account the user input in a more interactive way. The second important difference is that we use hylomorphisms as an unifying construct. This enables us to use the rich theory of hylomorphisms for parallelism. Moreover, we base our approach on a decision procedure that is derived from an equational theory that is both sound and complete w.r.t. the rules of the underlying equational theory, and we use a standard type unification algorithm to instantiate parallel structures from sequential code. Finally, our parallel structures and cost models are not built-in, but are derived in a systematic way from an underlying cost model and operational semantics.

Finally, Steuwer *et al.* [44] generate high-performance OpenCL code from a high-level specification by applying a simple set of rewrite rules, and using Monte-Carlo search to traverse the corresponding search space to find an implementation. Our approach provides a way to narrow down this search space, while using cost models to automate the rest of the search. It is, in a sense, more general, since we allow our parallel structures to be easily extended. However, we could benefit from exploiting Steuwer *et al.*'s work in GPU-specific rewriting rules and skeletons.

9. Conclusions

This paper has introduced new formally-based cost models for common algorithmic skeletons. Unlike previous approaches, these cost models are derived from a formal operational semantics for these skeletons, and consider both queue contention and scheduling. Using our approach, we can be sure that our cost models are sound w.r.t. the operational semantics *by construction* and we can *automatically* derive cost equations for newly defined parallel structures. A key aspect of our approach is the use of *hylomorphisms*, combinations of *catamorphisms* and *anamorphisms* (or *fold/unfold* operations). We have used hylomorphisms to capture some common patterns of parallelism: task farms, pipelines, divide-and-conquers and feedbacks. As shown in our examples, this single construct is surprisingly powerful, providing a system of canonical representations that is easy to understand and that is also easy to transform.

We have shown how our cost models can be used to drive type-level decisions about parallelisation, supporting a type-level reasoning system that we recently introduced [9]. Collectively, this allows the construction of provably optimal

parallel solutions, that need to take into account only basic, easily obtained, information about the costs of executing sequential operations. The use of a type-level approach avoids the usual separation of analysis from program, allowing structures and costs to be directly and accurately associated with the program. In particular, all transformations are performed internally by the type checker and we ensure the preservation of the underlying functional behaviour *simply by construction*. Moreover, our concrete results show that our cost models are good predictors for actual parallel performance.

A number of obvious extensions can be made to this work. Firstly, there are a few forms of parallel pattern that we have not yet considered: *map (reduce)* and *fold* are clearly instances of hylomorphisms, but *stencil* and *bulk synchronous parallel* patterns, for example, may require deeper thought. Secondly, we could investigate more complex and detailed cost models that take into account even more details of a parallel system. Finally, we believe that hylomorphisms and type structures can form the basis for a good, new parallel programming methodology. We intend to explore this by building our ideas into a new parallel library for Haskell that will enable the automatic introduction of parallelism into normal Haskell programs, under type control.

Acknowledgements

References

[1] M. Aldinucci. Automatic program transformation: The Meta tool for skeleton-based languages. *Constructive Methods for Parallel Programming, Advances in Computation: Theory and Practice*, pages 59–78, 2002.

[2] M. Aldinucci and M. Danelutto. Stream parallel skeleton optimization. In *Proc. PDCS '99: IASTED Intl. Conf. on Parallel and Distributed Computing and Systems*, pages 955–962, Nov. 1999.

[3] M. Aldinucci, M. Coppola, and M. Danelutto. Rewriting skeleton programs: How to evaluate the data-parallel stream-parallel tradeoff. In *Proc. International Workshop on Constructive Methods for Parallel Programming. Technical Report MIP-9805. University of Passau. Passau*, 1998.

[4] M. Barbosa, A. Cunha, and J. S. Pinto. Recursion Patterns and Time-analysis. *ACM SIGPLAN Notices*, 40(5):45–54, 2005.

[5] G. Barthe, B. Grégoire, and C. Riba. Type-Based Termination with Sized Products. In *International Workshop on Computer Science Logic*, pages 493–507. Springer, 2008.

[6] G. E. Blelloch. Nesl: A nested data-parallel language.(version 3.1). Technical report, DTIC Document, 1995.

[7] C. Brown, K. Hammond, M. Danelutto, and P. Kilpatrick. A language-independent parallel refactoring framework. In *Proceedings of the Fifth Workshop on Refactoring Tools*, pages 54–58. ACM, 2012.

[8] C. Brown, M. Danelutto, K. Hammond, P. Kilpatrick, and A. Elliott. Cost-directed refactoring for parallel Erlang programs. *International Journal of Parallel Programming*, 42(4):1–19, 2013.

[9] D. Castro, K. Hammond, and S. Sarkar. Farms, Pipes, Streams and Reforestation: Reasoning About Structured Parallel Processes Using Types and Hylomorphisms. In *Proc. ICFP 2016: Intl. Conf. on Functional Programming*, pages 4–17, 2016.

[10] Y.-Y. Chi and S.-C. Mu. Constructing List Homomorphisms from Proofs. In *Proc. APLAS '11: Asian Symposium on Programming Languages & Systems*, pages 74–88. 2011.

[11] M. I. Cole. *Algorithmic Skeletons: Structured Management of Parallel Computation*. Pitman, London, 1989.

[12] J. Fischer and S. Gorlatch. Turing Universality of Recursive Patterns for Parallel Programming. *Parallel Processing Letters*, 12(02):229–246, 2002.

[13] S. Fortune and J. Wyllie. Parallelism in Random Access Machines. In *Proc. STOC '78: 10th Symp. on Theory of Computing*, pages 114–118, 1978.

[14] A. Geser and S. Gorlatch. Parallelizing Functional Programs by Generalization. *Journal of Functional Programming (JFP)*, 9(06):649–673, 1999.

[15] J. Gibbons. Computing Downwards Accumulations on Trees Quickly. *Theoretical Computer Science*, 169(1):67–80, 1996.

[16] J. Gibbons. The Third Homomorphism Theorem. *Journal of Functional Programming (JFP)*, 6(4):657–665, 1996.

[17] S. Gorlatch. Extracting and Implementing List Homomorphisms in Parallel Program Development. *Sci. of Computer Programming*, 33(1):1 – 27, 1999.

[18] S. Gorlatch and C. Lengauer. Parallelization of Divide-and-Conquer in the Bird-Meertens Formalism. *Formal Aspects of Comp.*, 7(6):663–682, 1995.

[19] A. Gustavsson, J. Gustafsson, and B. Lisper. Toward static timing analysis of parallel software. In *OASIcs-OpenAccess Series in Informatics*, volume 23. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2012.

[20] M. M. Hamdan. A Survey of Cost Models for Algorithmic Skeletons. Technical report, Technical report, Heriot-Watt University, Dept. of Computers and Electrical Engineering, 1999.

[21] K. Hammond, C. Brown, and S. Sarkar. Timing Properties and Correctness for Structured Parallel Programs on x86-64 Multicores. In *Proc. FOPARA 2015: Foundational and Practical Aspects of Resource Analysis, Revised Selected Papers, Springer LNCS 9964*, pages 101–125, 2016.

[22] Z. Hu, H. Iwasaki, and M. Takechi. Formal Derivation of Efficient Parallel Programs by Construction of List Homomorphisms. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 19(3):444–461, 1997.

[23] Z. Hu, M. Takeichi, and W.-N. Chin. Parallelization in Calculational Forms. In *Proc. POPL '98: 25th ACM Symposium on Principles of Programming Languages*, pages 316–328, 1998.

[24] J. Hughes, L. Pareto, and A. Sabry. Proving the Correctness of Reactive Systems Using Sized Types. In *Proc. POPL '96: 23rd ACM Symposium on Principles of Programming Languages*, pages 410–423, 1996.

[25] V. Janjic, C. Brown, K. Mackenzie, K. Hammond, M. Danelutto, M. Aldinucci, and J. D. Garcia. RPL: A Domain-Specific Language for Designing and Implementing Parallel C++ Applications. In *Proc. PDP 2016: Euromicro International Conf. on Parallel, Distributed, and Network-Based Processing*, pages 288–295. IEEE, 2016.

[26] G. Keller and M. Chakravarty. Flattening Trees. In *Proc. Euro-Par '98: European Conference on Parallelism*, pages 709–719. 1998.

[27] Y. Liu, Z. Hu, and K. Matsuzaki. Towards Systematic Parallel Programming over Mapreduce. In *Proc. Euro-Par 2011: European Conference on Parallelism*, pages 39–50. 2011.

[28] H.-W. Loidl and K. Hammond. A Sized Time System for a Parallel Functional Language. In *Proc. Glasgow Workshop on Functional Programming, Ullapool, Scotland*, 1996.

[29] H. A. López, E. R. B. Marques, F. Martins, N. Ng, C. Santos, V. T. Vasconcelos, and N. Yoshida. Protocol-based Verification of Message-passing Parallel Programs. In *Proc. OOPSLA 2015: Intl. Conf. on Object-Oriented Programming, Systems, Languages, and Applications*, pages 280–298, 2015.

[30] K. Matsuzaki, Z. Hu, and M. Takeichi. Towards Automatic Parallelization of Tree Reductions in Dynamic Programming. In *Proc. SPAA 2006: Symp. on Parallelism in Algorithms and Architecture*, pages 39–48, 2006.

[31] E. Meijer, M. Fokkinga, and R. Paterson. Functional Programming with Bananas, Lenses, Envelopes and Barbed Wire. In *Proc. ICFP '91: ACM Conf. on Functional Programming*, pages 124–144, 1991.

[32] J. Misra. Powerlist: A Structure for Parallel Recursion. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 16(6):1737–1767, 1994.

[33] A. Morihata. A Short Cut to Parallelization Theorems. In *Proc. ICFP 2013: 18th ACM Conf. on Functional Programming*, pages 245–256, 2013.

[34] A. Morihata and K. Matsuzaki. Automatic Parallelization of Recursive Functions using Quantifier Elimination. In *Proc. FLOPS '10: Functional and Logic Programming*, pages 321–336. 2010.

[35] K. Morita, A. Morihata, K. Matsuzaki, Z. Hu, and M. Takeichi. Automatic Inversion Generates Divide-and-Conquer Parallel Programs. In *Proc. PLDI '07: ACM Conf. on Prog. Lang. Design and Impl.*, pages 146–155, 2007.

[36] Y. Onoue, Z. Hu, M. Takeichi, and H. Iwasaki. A Calculational Fusion System HYLO. In *Proc. Intl. Workshop on Alg. Lang. and Calculi*, pages 76–106, 1997.

[37] R. Peña and C. Segura. Sized Types for Typing Eden Skeletons. In *Proc. IFL '01: Intl. Symp. on Impl. of Functional Langs.*, pages 1–17, 2001.

[38] F. A. Rabhi and S. Gorlatch. *Patterns and Skeletons for Parallel and Distributed Computing*. Springer, 2003.

[39] J. H. Reif. *Synthesis of Parallel Algorithms*. Morgan Kaufmann, 1993.

[40] D. Sands. A Naïve Time Analysis and its Theory of Cost Equivalence. *Journal of Logic and Computation*, 5(4):495–541, 1995.

[41] D. B. Skillicorn. Models for Practical Parallel Computation. *International Journal of Parallel Programming*, 20(2):133–158, 1991.

[42] D. B. Skillicorn. *The Bird-Meertens Formalism as a Parallel Model*. Springer, 1993.

[43] D. B. Skillicorn and W. Cai. A Cost Calculus for Parallel Functional Programming. *J. Parallel Distrib. Comput.*, 28(1):65–83, 1995.

[44] M. Steuwer, C. Fensch, S. Lindley, and C. Dubach. Generating Performance Portable Code Using Rewrite Rules. In *Proc ICFP 2015: 20th ACM Conf. on Functional Prog. Lang. and Comp. Arch.*, pages 205–217, 2015.

[45] L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.

[46] P. B. Vasconcelos. *Space Cost Analysis using Sized Types*. PhD thesis, University of St Andrews, 2008.

[47] P. Wadler. Deforestation: Transforming Programs to Eliminate Trees. *Theoretical Computer Science*, 73(2):231–248, 1990.

*Biographies (Text)

Prof. Kevin Hammond (Principal Investigator) leads the Programming Languages group in the School of Computer Science. He has over 30 years of research experience, having worked extensively in the field of advanced programming language design and implementation, with a focus on understanding and reasoning about extra-functional properties. His work concentrates on declarative language designs, including that of the standard non-strict functional language Haskell [34], where he served on the international design committee, worked on the dominant compiler, GHC [33], and developed the GUM parallel Haskell implementation of GHC with colleagues at Glasgow and elsewhere [36]. In 2011, GHC was awarded the ACM SigPlan prize for significant programming language software. He has published over 100 books, book chapters, journal papers and other refereed publications focusing on functional programming, domain-specific programming languages, type systems, real-time systems, cost issues, adaptive run-time environments, lightweight concurrency, high-level programming language design, parallel programming and performance monitoring/visualisation. He is the recently-appointed Director of SICSA, the pan-Scottish pooling agreement for Informatics and Computer Science, which groups the 14 Scottish universities. He is a Senior Member of the ACM, a founder member of IFIP WG 2.11 (Generative Programming), and an honorary Professor at Heriot-Watt University. He sits on the board of the Scottish government funded Data Lab Innnovation Centre. He was awarded a personal award for his services to the parallel programming community at the HiPEAC 2016 conference.

*Biographies (Photograph)
Click here to download high resolution image