

The Third Homomorphism Theorem on Trees

Downward & Upward Lead to Divide-and-Conquer

Akimasa Morihata

University of Tokyo
JSPS Research Fellow

morihata@ipl.t.u-tokyo.ac.jp

Kiminori Matsuzaki

University of Tokyo

kmatsu@ipl.t.u-tokyo.ac.jp

Zhenjiang Hu

National Institute of
Informatics

hu@nii.ac.jp

Masato Takeichi

University of Tokyo

takeichi@mist.i.u-tokyo.ac.jp

Abstract

Parallel programs on lists have been intensively studied. It is well known that associativity provides a good characterization for divide-and-conquer parallel programs. In particular, *the third homomorphism theorem* is not only useful for systematic development of parallel programs on lists, but it is also suitable for automatic parallelization. The theorem states that if two sequential programs iterate the same list leftward and rightward, respectively, and compute the same value, then there exists a divide-and-conquer parallel program that computes the same value as the sequential programs.

While there have been many studies on lists, few have been done for characterizing and developing of parallel programs on trees. Naive divide-and-conquer programs, which divide a tree at the root and compute independent subtrees in parallel, take time that is proportional to the height of the input tree and have poor scalability with respect to the number of processors when the input tree is ill-balanced.

In this paper, we develop a method for systematically constructing scalable divide-and-conquer parallel programs on trees, in which two sequential programs lead to a scalable divide-and-conquer parallel program. We focus on paths instead of trees so as to utilize rich results on lists and demonstrate that associativity provides good characterization for scalable divide-and-conquer parallel programs on trees. Moreover, we generalize the third homomorphism theorem from lists to trees. We demonstrate the effectiveness of our method with various examples. Our results, being generalizations of known results for lists, are generic in the sense that they work well for all polynomial data structures.

Categories and Subject Descriptors D.1.3 [Concurrent Programming]: Parallel programming; I.2.2 [Automatic Programming]: Program transformation; D.1.1 [Applicative (Functional) Programming]

General Terms Algorithm, Theory

Keywords Divide-and-conquer, Huet's zippers, Polynomial data structures, The third homomorphism theorem

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL'09, January 18–24, 2009, Savannah, Georgia, USA.
Copyright © 2009 ACM 978-1-60558-379-2/09/01...\$5.00

1. Introduction

*What are little boys made of?
Snips and snails, and puppy-dogs' tails
That's what little boys are made of!
What are little girls made of?
Sugar and spice and all things nice
That's what little girls are made of!
(an old nursery rhyme)*

What are parallel programs on lists made of? Consider summing up the elements in a list $[a_1, a_2, a_3, a_4, a_5, a_6, a_7, a_8]$ as an example. It is easy to derive sequential algorithms; both the rightward summation

$$((((((a_1 + a_2) + a_3) + a_4) + a_5) + a_6) + a_7) + a_8$$

and the leftward summation

$$a_1 + (a_2 + (a_3 + (a_4 + (a_5 + (a_6 + (a_7 + a_8))))))$$

are sequential algorithms. Can we derive a parallel algorithm for summation? Yes, the *divide-and-conquer* summation

$$((a_1 + a_2) + (a_3 + a_4)) + ((a_5 + a_6) + (a_7 + a_8))$$

is a parallel algorithm, which divides the list at its center and computes each part in parallel. What is the key to such divide-and-conquer parallel algorithms? Compare the sequential algorithms with the divide-and-conquer algorithm. The only difference is the structure of the parentheses, and the *associative law* of $+$, namely $a + (b + c) = (a + b) + c$, enables us to rearrange the parentheses.

This observation, i.e., an associative operator provides a divide-and-conquer parallel computation, is formalized as the notion of *list homomorphisms* (Bird 1987). Function h is a list homomorphism if there exists associative operator \odot such that

$$h(x \uparrow y) = h x \odot h y$$

holds, where operator \uparrow denotes the concatenation of two lists. What is nice about list homomorphisms is the scalability¹: given a list of length n , list homomorphisms yield linear speedup up to $O(n/\log n)$ processors. Scalability is one of the most important properties in parallel programming. If a parallel program has good scalability, it shows great speedup that other optimization techniques can barely achieve. If scalability is poor, parallelization is nothing but anti-optimization, because parallelization requires overheads such as communication and synchronization.

Then, what are list homomorphisms made of? It may be surprising that parallel programming on lists is made of writing two

¹In this paper, we consider scalability with respect to numbers of processors (?).




Figure 1. Aggressively dividing binary tree

sequential programs, which proved as *the third homomorphism theorem* (Gibbons 1996). The theorem states that if two sequential programs iterate the same list leftward and rightward, respectively, and compute the same value, then there exists a list homomorphism that computes the same value as these sequential programs. In other words, two sequential programs lead to a scalable divide-and-conquer parallel program. The third homomorphism theorem is useful for developing parallel programs, and automatic parallelization methods have been proposed based on the theorem (Geser and Gorlatch 1999; Morita et al. 2007).

In summary, scalable parallel programs on lists are made of “sugar and spice and all things nice”: associative operators and list homomorphisms, which are obtained from two sequential programs.

Then, what are parallel programs on trees made of? For example, let us consider summing up all values in a tree. Someone may think of a divide-and-conquer parallel algorithm raised by subtree structures, i.e., computing independent subtrees in parallel. However, such a naive parallel algorithm is generally not scalable. Its speedup is limited by the height of the input tree, and thus, it has miserable scalability if the input tree is ill-balanced. To obtain better scalability, we need to introduce more aggressive divisions, as outlined in Figure 1. In this case, aggressive divisions yield a scalable parallel algorithm, which computes the summation in logarithmic time on the size of the tree with a sufficient number of processors.

Although we have successfully constructed a scalable divide-and-conquer parallel summation algorithm based on aggressive divisions, it is nontrivial to generalize the algorithm. What algebraic properties will enable us to compute each part in parallel? How can we obtain a parallel program from sequential ones? The nonlinear structure of trees makes it hard to develop parallel programs. Therefore, it has been considered that scalable parallel programs on trees were made of “snips and snails, and puppy-dogs’ tails”.

In this paper, we explain that scalable parallel programs on trees are in fact made of “sugar and spice and all things nice”—even though they are a bit spicier than those on lists. We focus on the similarity between lists and paths, formalize scalable parallel programs on trees as path-based computations, and demonstrate that associative computations on paths lead to scalable parallel programs. Moreover, we prove that the following proposition holds, which is a tree version of “the third homomorphism theorem” that enables us to derive a scalable parallel program on trees from two sequential programs.

If two sequential programs iterate the same tree downward and upward, respectively, and compute the same value, then there exists a scalable divide-and-conquer parallel program that computes the same value as these sequential programs.

This paper makes three main contributions:

- *Formalization of path-based computations on trees:* We introduce path-based computations, which include downward and upward computations. To express path-based computations, we borrow the notion of *Huet’s zippers* (Huet 1997; McBride 2001). Spotlighting paths is the key to developing parallel programs, because it makes theories on lists applicable to trees.
- *Characterization of scalable parallel computations on trees:* We propose an algebraic characterization of scalable parallel programs on trees. The main idea is to consider divide-and-conquer algorithms on one-hole contexts instead of those on trees. We can evaluate programs that belong to our characterization efficiently in parallel: given a tree of n nodes, they finish in $O(n/p + \log p)$ steps of primitive operations on an exclusive-read/exclusive-write parallel random access machine with p processors. It is worth noting that this computational complexity implies they have good scalability in the sense that they show linear speedup up to $O(n/\log n)$ processors.
- *“The third homomorphism theorem” on trees:* We prove “the third homomorphism theorem” on trees. It states that if a function is both downward and upward, then it can be evaluated efficiently in parallel. The theorem is useful for systematically developing parallel programs. We will demonstrate its effectiveness with various examples.

Our results are *generic* in the sense that they work well for all *polynomial data structures*, which can capture a large class of algebraic data structures on functional languages; besides, they are generalizations of known results on lists.

The rest of this paper is organized as follows. After the preliminaries in Section 2, we explain our ideas and results on node-valued binary trees in Sections 3 and 4. We develop our theory in Section 3 and present examples in Section 4. After that, we generalize our results to generic trees, namely polynomial data structures, in Section 6. We discuss related and future works in Section 7.

2. The Third Homomorphism Theorem

2.1 Basic definitions

In this paper, we basically borrow notations of functional programming language Haskell (Peyton Jones 2003). The parentheses for function applications may be omitted. Note that applications for functions have higher priority than those for operators, thus $f x \oplus y$ means $(f x) \oplus y$. Operator \circ denotes a function composition, and its definition is $(f \circ g)(x) = f(g(x))$. Underscore $_$ is used to express “don’t care” pattern. $x :: X$ means the type of x is X .

A list is denoted by brackets split by commas. The list concatenation operator is denoted by $++$. Note that $++$ is associative. $[A]$ denotes the type of lists whose elements are in A .

The disjoint sum of two sets A and B is denoted by *Either* $A B$.

data *Either* $a b = \text{Left } a \mid \text{Right } b$

We will write L and R instead of *Left* and *Right* for shorthand.

We will use several standard functions in Haskell, and their definitions are given in Figure 2.

2.2 Right Inverses

Right inverses, which are generalizations of inverse functions, are useful for developing divide-and-conquer parallel programs (Gibbons 1996; Morita et al. 2007).

Definition 2.1 (right inverse). For function $f :: A \rightarrow B$, a *right inverse* of f , denoted by f° , is a function satisfying the following equation.

$$f \circ f^\circ \circ f = f \quad \square$$

$$\begin{aligned}
id\ x &= x \\
fst\ (a, b) &= a \\
map\ f\ [] &= [] \\
map\ f\ ([a] ++ x) &= [f\ a] ++ map\ f\ x \\
foldr\ (\oplus)\ e\ [] &= e \\
foldr\ (\oplus)\ e\ ([a] ++ x) &= a \oplus (foldr\ (\oplus)\ e\ x) \\
foldl\ (\otimes)\ e\ [] &= e \\
foldl\ (\otimes)\ e\ (x ++ [a]) &= (foldl\ (\otimes)\ e\ x) \otimes a
\end{aligned}$$

Figure 2. Definitions of standard functions

Two things are worth noting. First, as a right inverse exists for any function, it is unnecessary to worry about its existence. Second, a right inverse of a function is generally not unique, and f° denotes one of the right inverses of f .

2.3 List Homomorphisms and The Third Homomorphism Theorem

List homomorphisms are an expressive computation pattern for scalable divide-and-conquer parallel programs on lists.

Definition 2.2 (list homomorphism (Bird 1987)). Function $h :: [A] \rightarrow B$ is said to be a *list homomorphism* if there exists function $\phi :: A \rightarrow B$ and associative operator $(\odot) :: B \rightarrow B \rightarrow B$ such that

$$\begin{aligned}
h\ [] &= \iota_\odot \\
h\ [a] &= \phi\ a \\
h\ (x ++ y) &= h\ x \odot h\ y
\end{aligned}$$

hold, where ι_\odot is the unit of \odot . Here, we write $h = hom\ (\odot)\ \phi$.² \square

It is worth noting that associative operator \odot characterizes a list homomorphism. The associativity of \odot guarantees that the result of computation is not affected by where to divide the list. List homomorphisms are useful for developing parallel programs on lists (Bird 1987; Cole 1994, 1995; Gibbons 1996; Hu et al. 1997a).

The third homomorphism theorem demonstrates a necessary and sufficient condition of existence of a list homomorphism.

Theorem 2.3 (the third homomorphism theorem (Gibbons 1996)). Function h is a list homomorphism if and only if there exist two operators \oplus and \otimes such that the following equations hold.

$$\begin{aligned}
h\ ([a] ++ x) &= a \oplus h\ x \\
h\ (x ++ [a]) &= h\ x \otimes a
\end{aligned}
\quad \square$$

The third homomorphism theorem states that if we can compute a function in both leftward and rightward manners, then there exists a divide-and-conquer parallel algorithm to evaluate the function. What the theorem indicates is not only the existence of parallel programs but also a way of systematically developing parallel programs. The following lemma plays a central role in parallelization.

Lemma 2.4 (Gibbons (1996); Morita et al. (2007)). Assume that the following equations hold for function h .

$$\begin{aligned}
h\ ([a] ++ x) &= a \oplus h\ x \\
h\ (x ++ [a]) &= h\ x \otimes a
\end{aligned}$$

Then, $h = hom\ (\odot)\ \phi$ holds, where \odot and ϕ are defined as follows.

$$\begin{aligned}
\phi\ a &= h\ [a] \\
a \odot b &= h\ (h^\circ\ a ++ h^\circ\ b)
\end{aligned}
\quad \square$$

Lemma 2.4 states that we can derive a parallel program from two sequential programs through a right inverse of h .

²We usually neglect ι_\odot because parallel computations on empty lists are useless. If ι_\odot is necessary, we can prepare it by introducing a special value that behaves as the unit.




Figure 3. A zipper structure, which expresses a path from the root to the black leaf

3. “The Third Homomorphism Theorem” on Binary Trees

Here, let us consider node-valued binary trees.

$$\begin{aligned}
\mathbf{data}\ Tree &= Node\ Int\ Tree\ Tree \\
&| Leaf
\end{aligned}$$

The goal is to formalize and prove “the third homomorphism theorem” on node-valued binary trees, which consists of the three notions of downward computations, upward computations, and scalable divide-and-conquer computations on trees. To formalize these notions, we will focus on paths. Given a path, the downward and upward computations are computations that are accomplished while walking the path in downward and upward manners, respectively. Scalable divide-and-conquer computations on trees are computations that split a path in the middle and compute each substructure in parallel. To specify a path on a tree, we will borrow *Huet’s zippers* (Huet 1997).

3.1 Zippers on Binary Trees

A zipper is a list whose elements are contexts that are left after a walk. Based on walking downward from the root of a tree, we construct a zipper as follows: when we go down-right from a node, we add its left child to the zipper; when we go down-left, we add the right child to the zipper. For example, Figure 3 shows the correspondence between a zipper and a walk from the root to the black leaf.

The zipper structures for node-valued binary trees can be specified in the following type, where components of the *Either* type correspond to the left and the right child.

$$\mathbf{type}\ Zipper = [Either\ (Int, Tree)\ (Int, Tree)]$$

In this paper, a walk usually ends at a leaf. Then, a zipper stores the whole tree. Function $z2t$ below restores a tree from a zipper.

$$\begin{aligned}
z2t\ [] &= Leaf \\
z2t\ ([L\ (n, l)] ++ r) &= Node\ n\ l\ (z2t\ r) \\
z2t\ ([R\ (n, r)] ++ l) &= Node\ n\ (z2t\ l)\ r
\end{aligned}$$

When a walk ends at a leaf, a zipper corresponds to a tree. When a walk ends at an internal node, a zipper corresponds not to a tree but to a one-hole context of a tree. Look at Figure 3 again. On one hand, the zipper represents the tree with its path from the root to the black leaf. On the other hand, the zipper also represents the one-hole context in which the black circle represents the hole. We will formalize scalable divide-and-conquer parallel programs on trees based on the second viewpoint, viewing a zipper as a one-hole context. However, we usually regard a zipper as a tree and call the hole the *terminal leaf* because a walk usually ends at a leaf.

We would like to summarize the correspondences among a zipper, a tree (or a one-hole context), and a path. A zipper corresponds to a tree with a terminal leaf (a one-hole context) or a path from the root to the terminal leaf. An initial segment of a zipper corresponds to a one-hole context containing the root or a path from the root to

a node. A tail segment of a zipper corresponds to a subtree containing the terminal leaf (a one-hole context whose hole is the same as the hole of the original one) or a path from a node to the terminal leaf.

3.2 Downward and Upward Computations on Binary Trees

Next, let us formalize downward and upward computations.

Consider function $sumTree$ below as an example, which sums up all values in a tree.

$$\begin{aligned} sumTree \text{ Leaf} &= 0 \\ sumTree (\text{Node } n \ l \ r) &= n + sumTree \ l + sumTree \ r \end{aligned}$$

First, we would like to give its downward version. Since an initial segment of a zipper corresponds to a path from the root to an internal node, function $sumTree_{\downarrow}$ below performs its computation downward from the root to the terminal leaf.

$$\begin{aligned} sumTree_{\downarrow} \square &= 0 \\ sumTree_{\downarrow} (x \# [L(n, l)]) &= sumTree_{\downarrow} x + n + sumTree \ l \\ sumTree_{\downarrow} (x \# [R(n, r)]) &= sumTree_{\downarrow} x + n + sumTree \ r \end{aligned}$$

Note that for computing summations for trees in the zipper, we use the function $sumTree$.

Similarly, we can give its upward version. Since a tail segment of a zipper corresponds to a path from an internal node to the terminal leaf, function $sumTree_{\uparrow}$ below traverses a tree from its terminal leaf to its root.

$$\begin{aligned} sumTree_{\uparrow} \square &= 0 \\ sumTree_{\uparrow} ([L(n, l)] \# r) &= n + sumTree \ l + sumTree_{\uparrow} \ r \\ sumTree_{\uparrow} ([R(n, r)] \# l) &= n + sumTree_{\uparrow} \ l + sumTree \ r \end{aligned}$$

Here, $sumTree_{\uparrow}$ and $sumTree_{\downarrow}$ are equivalent to $sumTree$ in the sense that $sumTree_{\uparrow} = sumTree_{\downarrow} = sumTree \circ z2t$ holds. However, computations on a path may require more information than those on trees. To formalize the correspondence between computations on paths and those on trees, we introduce a notion of *path-based computations*.

Definition 3.1 (path-based computation on binary trees). Function $h' :: Zipper \rightarrow B$ is said to be a path-based computation of $h :: Tree \rightarrow A$ if there exists function $\psi :: B \rightarrow A$ such that the following equation holds.

$$\psi \circ h' = h \circ z2t \quad \square$$

This equation means that h' simulates the computation of h and the result is extracted by ψ . Note that $z2t$ is a path-based computation of any function; however, this is useless in practice because no significant computations are managed on paths. In other words, it is important to specify an appropriate path-based computation.

Now, let us introduce downward and upward computations.

Definition 3.2 (downward computations on binary trees). Function $h' :: Zipper \rightarrow B$, which is a path-based computation of $h :: Tree \rightarrow A$, is said to be *downward* if there exists operator $(\otimes) :: B \rightarrow \text{Either } (Int, A) \ (Int, A) \rightarrow B$ such that the following equations hold.

$$\begin{aligned} h' (x \# [L(n, t)]) &= h' x \otimes L(n, h t) \\ h' (x \# [R(n, t)]) &= h' x \otimes R(n, h t) \end{aligned} \quad \square$$

Definition 3.3 (upward computations on binary trees). Function $h' :: Zipper \rightarrow B$, which is a path-based computation of $h :: Tree \rightarrow A$, is said to be *upward* if there exists operator $(\oplus) :: \text{Either } (Int, A) \ (Int, A) \rightarrow B \rightarrow B$ such that the following equations hold.

$$\begin{aligned} h' ([L(n, t)] \# x) &= L(n, h t) \oplus h' x \\ h' ([R(n, t)] \# x) &= R(n, h t) \oplus h' x \end{aligned} \quad \square$$




Figure 4. Recursive division on one-hole context: at each step, the one-hole context is divided at the node represented by a concentric circle.

3.3 Parallel Computations on Binary Trees

Next, let us consider scalable divide-and-conquer parallel computations on trees. For scalable divide-and-conquer computations, we would like to divide a tree at an arbitrary place and compute each part in parallel. However, as seen in Figure 1, splitting a tree does not yield two trees of the original type: one is a tree of the original type, but the other is a one-hole context. Therefore, it is difficult to formalize recursive division on trees.

Our idea is to consider recursive division on one-hole contexts, instead of that on trees. We divide a path from the root to the hole, as Figure 4. At each step, we select a node on a path from the root to the hole, divide the one-hole context into three one-hole contexts, i.e., the upper part, the lower part, and a complete subtree with the node. Apparently, we can divide the upper part and the lower part once again; we can divide the complete subtree by taking off an arbitrary leaf and obtaining a one-hole context. Then, we can accomplish recursive division on a one-hole context.

Now, let us characterize scalable divide-and-conquer parallel computations on node-valued binary trees. For parallel computation on trees, we require three operations: an operation (say \odot) that merges the results of two one-hole contexts, an operation (say ϕ) that takes the result of a complete tree and yields the result of the complete tree with its parent, and an operation (say ψ) that computes a result of a complete tree from the result of a one-hole context. Since a one-hole context corresponds to a zipper, computation on a one-hole context can be specified by path-based computations.

Definition 3.4 (decomposition on binary trees). A *decomposition* of function $h :: Tree \rightarrow A$ is triple (ϕ, \odot, ψ) that consists of associative operator $\odot :: B \rightarrow B \rightarrow B$ and two functions $\phi :: \text{Either } (Int, A) \ (Int, A) \rightarrow B$ and $\psi :: B \rightarrow A$ such that

$$\begin{aligned} \psi \circ h' &= h \circ z2t \\ h' \square &= \iota_{\odot} \\ h' [L(n, t)] &= \phi(L(n, h t)) \\ h' [R(n, t)] &= \phi(R(n, h t)) \\ h' (x \# y) &= h' x \odot h' y \end{aligned}$$

hold, where ι_{\odot} is the unit of \odot . In this case, h is said to be *decomposable*. \square

It is worth noting that the associativity of \odot is necessary to guarantee that the result of computation is not affected by where to divide the tree. It is also worth noting that a decomposition can be seen as a list homomorphism on zippers: Define function ϕ' as $\phi'(L(n, t)) = \phi(L(n, h t))$ and $\phi'(R(n, t)) = \phi(R(n, h t))$; then, $h' = \text{hom}(\odot) \phi'$ holds.

Actually decomposable functions can be efficiently evaluated in parallel, which is a consequence of theories of *parallel tree contraction* (Miller and Reif 1985; Cole and Vishkin 1988; Gibbons and Rytter 1989; Abrahamson et al. 1989; Reif 1993).

Theorem 3.5. If (ϕ, \odot, ψ) is a decomposition of function h and all of ϕ , \odot , and ψ are constant-time computations, then h can be evaluated for a tree of n nodes in $O(n/p + \log p)$ time on an exclusive-read/exclusive-write parallel random access machine with p processors.

Proof. It is not difficult to confirm that h satisfies the premise of Theorem 3.1 in (Abrahamson et al. 1989). \square

Complexity implies good scalability of decomposable functions, because we can obtain linear speedup up to $O(n/\log n)$ processors.

The function $sumTree$ is decomposable as the following equations show.

$$\begin{aligned} sumPara &= sumTree \circ z2t \\ sumPara [] &= 0 \\ sumPara [L(n, l)] &= n + sumTree l \\ sumPara [R(n, r)] &= n + sumTree r \\ sumPara (x ++ y) &= sumPara x + sumPara y \end{aligned}$$

In other words, $sumTree$ has a decomposition $(\phi, +, id)$, where ϕ is defined by $\phi(L(n, v)) = n + v$ and $\phi(R(n, v)) = n + v$. It is not difficult to derive a parallel program for $sumTree$ because of the associativity of $+$. However, it is generally difficult to derive an associative operator that provides a decomposition.

3.4 “The Third Homomorphism Theorem” on Binary Trees

Finally, let us introduce “the third homomorphism theorem” on node-valued binary trees, which demonstrates a necessary and sufficient condition of existence of a decomposition.

Theorem 3.6 (“the third homomorphism theorem” on binary trees). Function h is decomposable if and only if there exists a path-based computation of h that is both downward and upward.

Proof. We will prove a stronger theorem later (Theorem 6.7). \square

The statement of Theorem 3.6 is similar to the third homomorphism theorem on lists. The observation underlying the theorem is that associative operators provide parallel programs not only on lists but also on trees. In addition, the following lemma is useful for developing parallel programs.

Lemma 3.7. Assume that function h' , which is a path-based computation of h so that $\psi \circ h' = h \circ z2t$, is both downward and upward; then, there exists a decomposition (ϕ, \odot, ψ) of h such that the following equations hold.

$$\begin{aligned} \phi(L(v, h t)) &= h' [L(v, t)] \\ \phi(R(v, h t)) &= h' [R(v, t)] \\ a \odot b &= h' (h'^{\circ} a ++ h'^{\circ} b) \end{aligned}$$

Proof. We will prove a stronger lemma later (Lemma 6.6). \square

Combining Theorem 3.6 and Lemma 3.7, we can derive decompositions of functions. Recall $sumTree$. A path-based computation of $sumTree$ (say st) is both downward and upward, because of $st = sumTree_{\downarrow} = sumTree_{\uparrow}$. Therefore, Theorem 3.6 proves that there is a decomposition of $sumTree$. Lemma 3.7 shows a way of obtaining a decomposition (ϕ, \odot, ψ) . Here, $\psi = id$, because $sumTree \circ z2t = id \circ st$ holds. Obtaining function ϕ is easy as the following calculation shows, where C is either L or R.

$$\begin{aligned} \phi(C(n, sumTree t)) &= \{ \text{Lemma 3.7} \} \\ &= \{ \text{definition of } st \} \\ &= n + sumTree t \end{aligned}$$

Thus $\phi(C(n, r)) = n + r$. The last is associative operator \odot . Lemma 3.7 states that a right inverse of st enables us to derive the operator. It is not difficult to find a right inverse.

$$st^{\circ} s = [L(s, Leaf)]$$

Function st° above is certainly a right inverse of st , because $st(st^{\circ} s) = st[L(s, Leaf)] = s$ holds. Now we can obtain a definition of \odot as follows.

$$\begin{aligned} a \odot b &= \{ \text{Lemma 3.7} \} \\ &= st(st^{\circ} a ++ st^{\circ} b) \\ &= \{ \text{definition of } st^{\circ} \} \\ &= st[L(a, Leaf), L(b, Leaf)] \\ &= \{ \text{definition of } st \} \\ &= a + b \end{aligned}$$

We have obtained a decomposition of $sumTree$, which is exactly the same as the one we previously showed.

4. Examples

In this section, we demonstrate how to develop scalable divide-and-conquer parallel programs. Our development consists of two steps. First, we seek an appropriate path-based computation that is both downward and upward. After that, we obtain a decomposition that brings scalable parallelism.

4.1 Maximum Path Weight

Let us consider a small optimization problem as an initial example to compute the maximum weight of paths from the root to a leaf. For simplicity, we will assume that the value of each node is non-negative. The following sequential program solves the problem.

$$\begin{aligned} maxPath Leaf &= 0 \\ maxPath (Node n l r) &= n + max(maxPath l) (maxPath r) \end{aligned}$$

Our objective here is to develop a parallel program to solve the problem. First, we try to obtain a downward definition and think of the following program.

$$\begin{aligned} maxPath_{\downarrow} [] &= 0 \\ maxPath_{\downarrow} (x ++ [L(n, l)]) &= \max(maxPath_{\downarrow} x) (pathWeight x + n + maxPath l) \\ maxPath_{\downarrow} (x ++ [R(n, r)]) &= \max(maxPath_{\downarrow} x) (pathWeight x + n + maxPath r) \\ pathWeight [] &= 0 \\ pathWeight (x ++ [L(n, l)]) &= pathWeight x + n \\ pathWeight (x ++ [R(n, r)]) &= pathWeight x + n \end{aligned}$$

Note that $maxPath_{\downarrow}$ is not downward, because it uses auxiliary function $pathWeight$ that computes the weight of the path from the root to the terminal leaf. The tupling transformation (Fokkinga 1989; Chin 1993; Hu et al. 1997b) is helpful in dealing with such situations. Consider function $maxPath'_{\downarrow}$, which computes the weight of the path to the terminal leaf together with the maximum path weight of the tree, namely $maxPath'_{\downarrow} x = (maxPath_{\downarrow} x, pathWeight x)$. Apparently $maxPath'_{\downarrow}$ is a path-based computation of $maxPath$; in addition, it is downward.

$$\begin{aligned} maxPath'_{\downarrow} [] &= (0, 0) \\ maxPath'_{\downarrow} (x ++ [L(n, l)]) &= \mathbf{let} (m, w) = maxPath'_{\downarrow} x \\ &\quad \mathbf{in} (max m (w + n + maxPath l), w + n) \\ maxPath'_{\downarrow} (x ++ [R(n, r)]) &= \mathbf{let} (m, w) = maxPath'_{\downarrow} x \\ &\quad \mathbf{in} (max m (w + n + maxPath r), w + n) \end{aligned}$$

Therefore, $maxPath'_{\downarrow}$ seems an appropriate path-based computation for $maxPath$, and we would like to derive its upward defini-


```

leftOdd ◦ z2t = ψ ◦ lo
ψ * = *
ψ (L v) = v
ψ (R v) = v
lo [] = *
lo [L (n, t)] = case leftOdd t of
    * → if odd n then L n else *
    v → L v
lo [R (n, t)] = if odd n then R n
    else case leftOdd t of * → *
    v → R v
lo (z1 ++ z2) = case (lo z1, lo z2) of
    (L v, _) → L v
    (R v, *) → R v
    (_, r) → r

```

Figure 6. Divide-and-conquer parallel program for *leftOdd*

This problem is similar to the maximum-path-weight problem, and we can specify downward and upward definitions in a similar way.

```

height↓ [] = (1, 1)
height↓ (x ++ [L _ l]) = let (h, d) = height↓ x
    in (max h (d + height l), d + 1)
height↓ (x ++ [R _ r]) = let (h, d) = height↓ x
    in (max h (d + height r), d + 1)
height↑ [] = (1, 1)
height↑ ([L _ l] ++ x) = let (h, d) = height↑ x
    in (1 + max h (height l), d + 1)
height↑ ([R _ r] ++ x) = let (h, d) = height↑ x
    in (1 + max h (height r), d + 1)

```

Function $height_{\downarrow} = height_{\uparrow}$ (say *ht*) computes the height of a tree in its first result, and its second result retains the depth of the terminal leaf. Function *ht* is a path-based computation of *height*, because $fst \circ ht = height \circ z2t$ holds.

Let us parallelize it. The only nontrivial part is that to obtain an associative operator from its right inverse. In this case, different from the previous examples, we should define right inverse ht° in a recursive manner because $ht^{\circ}(h, d)$ yields a tree of height *h*. Therefore, it first seems difficult to simplify the definition of \odot , even though the naive definition $a \odot b = ht(ht^{\circ} a ++ ht^{\circ} b)$ is inefficient. Actually, that simplification is not too difficult. Let us look at Figure 7, which outlines the tree $ht^{\circ}(h_1, d_1) ++ ht^{\circ}(h_2, d_2)$. The left and right trees correspond to $ht^{\circ}(h_1, d_1)$ and $ht^{\circ}(h_2, d_2)$, and the curved arrow corresponds to the concatenation operation on zippers. Now it is easy to see that the height of this tree is $\max h_1 (d_1 + h_2)$ and the depth of the terminal leaf is $d_1 + d_2$. In short, the following gives a definition of \odot .

$$(h_1, d_1) \odot (h_2, d_2) = (\max h_1 (d_1 + h_2), d_1 + d_2)$$

Then, we obtain the following parallel program for *height*.

```

height ◦ z2t = fst ◦ ht
ht [] = (1, 1)
ht [L (n, t)] = (1 + height t, 2)
ht [R (n, t)] = (1 + height t, 2)
ht (z1 ++ z2) = let (h1, d1) = ht z1
    (h2, d2) = ht z2
    in (max h1 (d1 + h2), d1 + d2)

```

We have considered how to merge the results of substructures with the abstraction in Figure 7. The most significant thing is that Theorem 3.6 guarantees the correctness of the merging operation obtained from the abstraction. Theorem 3.6 proves that the results of *ht*, namely the height of the tree and the depth of the terminal




Figure 7. Outline of $ht^{\circ}(h_1, d_1) ++ ht^{\circ}(h_2, d_2)$: curved arrowed line denotes plugging operation, which corresponds to concatenation of two zippers.

leaf, are sufficient for merging the results of two parts; thus, we can derive correct merging operation no matter what shape the trees are we image for ht° .

5. “The Third Homomorphism Theorem” on Polynomial Data Structures

This section discusses our generalization of the method so that it can deal with all *polynomial data structures*, which can capture a large class of algebraic data structures on functional languages.

5.1 Polynomial Functors

To describe our results, we use a few notions from category theory. A category consists of a set of objects and a set of morphisms. A functor (say *F*) is a morphism of categories, and it respects both identity and composition.

$$\begin{aligned} F id &= id \\ F(f \circ g) &= Ff \circ Fg \end{aligned}$$

We will consider the category **Set**, where an object is a set and morphisms from object *A* to object *B* are all total functions from *A* to *B*. A functor gives a mapping between functions together with sets.

A functor is said to be polynomial if it is made up from identity functor *I*, constant functors such as $!A$ where *A* is an object, and bifunctors $+$ and \times . Let two operators for morphisms \times and $+$ be $(f \times g)(x, y) = (f x, g y)$, $(f + g)(1, x) = (1, f x)$, and $(f + g)(2, y) = (2, g y)$. Then, the definitions of polynomial functors are given as follows, where *F* and *G* are polynomial functors, *A* and *B* are objects, and *f* is a morphism.

$$\begin{aligned} (!A)B &= A \\ (!A)f &= id \\ !A &= A \\ !f &= f \\ (F \times G)A &= FA \times GA \\ (F \times G)f &= Ff \times Gf \\ (F + G)A &= (\{1\} \times FA) \cup (\{2\} \times GA) \\ (F + G)f &= Ff + Gf \end{aligned}$$

The least fixed point of functor *F*, denoted by μF , is the smallest set that satisfies $F(\mu F) = \mu F$. The least fixed point exists for each polynomial functor. It is well known that the least fixed points of polynomial functors provide a good characterization of a large class of algebraic data structures (Backhouse et al. 1998), and these are called *polynomial data structures*. For example, lists having elements of type *A* can be recognized as the least fixed point of functor *L* below, where $!1$ denotes a singleton set.

$$L = !1 + !A \times L$$

Node-valued trees having values of type *A* can also be recognized as the least fixed point of functor *T* below.

$$T = !1 + !A \times L \times L$$

5.2 Zippers for Polynomial Data Structures

Next, let us define zippers for polynomial data structures. We will follow McBride (2001) who showed a systematic derivation of zippers based on derivatives of functors. The derivative of polynomial functor F , denoted by ∂F , is a polynomial functor defined as follows, where \bullet denotes a distinguishable element.

$$\begin{aligned}\partial(!A) &= !\emptyset \\ \partial(1) &= !\{\bullet\} \\ \partial(F \times G) &= F \times \partial G + \partial F \times G \\ \partial(F + G) &= \partial F + \partial G\end{aligned}$$

Functor ∂F corresponds to a one-hole context structure of F . The zipper structure of μF , denoted by Z_F , is recognized as $Z_F = [\partial F(\mu F)]$.

To convert zippers to trees, we use an operator (\triangleleft_F) :: $Z_F \rightarrow \mu F \rightarrow \mu F$, which is defined by

$$z \triangleleft_F t = \text{foldr} (\triangleleft_F) t z$$

where (\triangleleft_F) :: $\partial F(\mu F) \rightarrow \mu F \rightarrow \mu F$ is the plugging-in operator (McBride 2001) defined as follows.

$$\begin{aligned}a \triangleleft_{!A} t &= a \\ \bullet \triangleleft_1 t &= t \\ (1, (a, b)) \triangleleft_{F \times G} t &= (a, b \triangleleft_G t) \\ (2, (a, b)) \triangleleft_{F \times G} t &= (a \triangleleft_F t, b) \\ (1, a) \triangleleft_{F+G} t &= (1, a \triangleleft_F t) \\ (2, b) \triangleleft_{F+G} t &= (2, b \triangleleft_G t)\end{aligned}$$

We will omit the subscript of \triangleleft_F when it is apparent from its context.

Note that operator \triangleleft takes an additional tree to convert a zipper to a tree, because a zipper corresponds to a one-hole context. As we want to minimize the differences between zippers and trees, we will force the difference between zippers and trees to be leaves. A set of leaves of μF , denoted by $leaves_{\mu F}$, is formalized as follows.

$$\begin{aligned}leaves_{\mu F} &= \llbracket F \rrbracket_{leaves} \\ \llbracket !A \rrbracket_{leaves} &= A \\ \llbracket 1 \rrbracket_{leaves} &= \emptyset \\ \llbracket F \times G \rrbracket_{leaves} &= \llbracket F \rrbracket_{leaves} \times \llbracket G \rrbracket_{leaves} \\ \llbracket F + G \rrbracket_{leaves} &= (\{1\} \times \llbracket F \rrbracket_{leaves}) \cup (\{2\} \times \llbracket G \rrbracket_{leaves})\end{aligned}$$

5.3 Sequential and Parallel Computations on Zippers

Now that the notion of zippers has been clarified, it is not difficult to provide definitions for downward and upward computations.

Definition 5.1 (path-based computation). Function $h' :: Z_F \rightarrow B$ is said to be a *path-based computation* of function $h :: \mu F \rightarrow A$ if there exists operator $\ominus :: B \rightarrow \mu F \rightarrow A$ such that the following equation holds for all $t \in leaves_{\mu F}$.

$$h(z \triangleleft t) = h' z \ominus t \quad \square$$

Definition 5.2 (downward computation). Function $h' :: Z_F \rightarrow B$, which is a path-based computation of function $h :: \mu F \rightarrow A$, is said to be *downward* if there exist operator $(\otimes) :: B \rightarrow \partial F A \rightarrow B$ and value $e :: B$ such that the following equation holds.

$$h' z = \text{foldl} (\otimes) e (\text{map} (\partial F h) z) \quad \square$$

Definition 5.3 (upward computation). Function $h' :: Z_F \rightarrow B$, which is a path-based computation of function $h :: \mu F \rightarrow A$, is said to be *upward* if there exist operator $(\oplus) :: \partial F A \rightarrow B \rightarrow B$ and value $e :: B$ such that the following equations holds.

$$h' z = \text{foldr} (\oplus) e (\text{map} (\partial F h) z) \quad \square$$

We characterize scalable divide-and-conquer parallel programs by recursive division on one-hole contexts as the same as the case of node-valued binary trees.

Definition 5.4 (decomposition). A *decomposition* of function $h :: \mu F \rightarrow A$ is tuple (ϕ, \odot, \ominus) that consists of function $\phi :: \partial F A \rightarrow B$, associative operator $\odot :: B \rightarrow B \rightarrow B$, and operator $\ominus :: B \rightarrow \mu F \rightarrow A$ such that the following equation holds for all $t \in leaves_{\mu F}$.

$$h(z \triangleleft t) = \text{hom} (\odot) (\phi \circ \partial F h) z \ominus t \quad \square$$

One of the most difficult issues is to find a scalable divide-and-conquer parallel algorithm to evaluate such functions. We (Mori-hata and Matsuzaki 2008) devised an algorithm by generalizing the tree contraction algorithm on binary trees given by Abrahamson et al. (1989). We have not discussed this in detail here, because it is beyond the scope of this paper.

Theorem 5.5. If (ϕ, \odot, \ominus) is a decomposition of function h and all of ϕ , \odot , and \ominus are constant-time computations, then h can be evaluated for a tree of n nodes in $O(n/p + \log p)$ time on an exclusive-read/exclusive-write parallel random access machine with p processors.

Proof. It is not difficult to confirm that h satisfies the premise of Theorem 9 in (Mori-hata and Matsuzaki 2008). \square

5.4 “The Third Homomorphism Theorem” on Polynomial Data Structures

We have shown that list homomorphisms characterize scalable divide-and-conquer parallel programs on polynomial data structures. Therefore, we can utilize parallelization methods on lists for free. For instance, “the third homomorphism theorem” on trees is a direct consequence of that on lists.

Lemma 5.6. Assume that function h' , which is a path-based computation of h so that $h(z \triangleleft t) = h' z \ominus t$ holds for all $t \in leaves_{\mu F}$, is both downward and upward; then, there exists a decomposition (ϕ, \odot, \ominus) of h such that $\phi(\partial F h a) = h'[a]$ and $a \odot b = h'(h'^{\circ} a \# h'^{\circ} b)$ hold. \square

Proof. Since h' is both leftward and rightward, \odot is associative and $h' = \text{hom} (\odot) (\phi \circ \partial F h)$ holds from Lemma 2.4. Therefore, (ϕ, \odot, \ominus) forms a decomposition of h . \square

Theorem 5.7 (“the third homomorphism theorem” on polynomial data structures). Function $h :: \mu F \rightarrow A$ is decomposable if and only if there exist operators $(\oplus) :: \partial F A \rightarrow B \rightarrow B$, $(\otimes) :: B \rightarrow \partial F A \rightarrow B$, and $(\ominus) :: B \rightarrow \mu F \rightarrow A$, such that the following equations hold for all $t \in leaves_{\mu F}$.

$$\begin{aligned}h(z \triangleleft t) &= \text{foldr} (\oplus) e (\text{map} (\partial F h) z) \ominus t \\ &= \text{foldl} (\otimes) e (\text{map} (\partial F h) z) \ominus t\end{aligned}$$

Proof. The “if” part is a direct consequence of Lemma 6.6, and the “only if” part is straightforward. \square

6. Discussion

We have developed a method for systematically constructing scalable divide-and-conquer parallel programs on polynomial data structures. We have focused on paths and expressed paths by zippers to utilize the known theories on lists. We have proposed a characterization of scalable divide-and-conquer parallel programs on polynomial data structures, proved “the third homomorphism theorem” on polynomial data structures, and demonstrated the effectiveness of our method with examples.

Our motivation was to derive “the third homomorphism theorem” on trees. The theorem, which was first introduced by Gibbons (1996), is useful for developing parallel programs, and automatic parallelization methods have been proposed (Geser and Gorchatch 1999; Morita et al. 2007) based on the theorem. Our theorem is an

exact generalization of the original third homomorphism theorem, and our results are built on list homomorphisms. Therefore, existing automatic parallelization methods should be applicable.

One of the aims of this paper is to explain ideas for constructing scalable parallel programs on trees through the development of our theory. Parallel tree contraction, which was first proposed by Miller and Reif (1985), is known to be a useful framework for developing scalable parallel programs on trees, and many computations have been implemented on it (Cole and Vishkin 1988; Gibbons and Rytter 1989; Abrahamson et al. 1989; Skillicorn 1996; Gibbons et al. 1994; ?; Matsuzaki 2007); however, parallel tree contraction is hard to use, because it requires a set of operations that satisfy a certain condition. In fact, the requirements for decomposable functions is equivalent to the sufficient condition for parallel tree contraction discussed in (?), and “the third homomorphism theorem” on trees brings sets of operations that satisfy the condition. Moreover, our results can deal with polynomial structures, even though most of the existing studies have only considered binary trees.

The third homomorphism theorem requires two sequential programs, while conventional parallelization methods generate a parallel program from a sequential program. Even though this requirement may have its shortcomings, it is arguably true. There is generally little hope of obtaining a parallel program from a sequential program, because parallel programs are more complicated than sequential ones. In other words, extra information is necessary to develop parallel programs from sequential ones. What the third homomorphism theorem provides is a systematic way of revealing such extra information.

We have developed our methods on polynomial data structures, which correspond to trees of bounded degree. Regular data structures are a generalization of polynomial data structures, and include trees of unbounded degree. Since McBride (2001) demonstrated a systematic derivation of zipper structures for regular data structures, there are no problems with formalizing path-based computations, downward computations, or upward computations on regular data structures. However, it is difficult to achieve scalability on regular data structures. We have required the operation to merge the results of siblings in $O(1)$ time. As this requirement is not realistic on regular data structures, it is necessary to parallelize this merging operation. In summary, it would be interesting to investigate “the third homomorphism theorem” on regular data structures.

Acknowledgments

The authors are grateful to Nao Hirokawa, who proposed using Huet’s zippers for formalizing computations on paths. The authors are also grateful to anonymous referees of POPL 2009 for their helpful comments. The first author is supported by Grant-in-Aid for JSPS research fellows 20 · 2411.

References

- Karl R. Abrahamson, N. Dadoun, David G. Kirkpatrick, and Teresa M. Przytycka. A simple parallel tree contraction algorithm. *Journal of Algorithms*, 10(2):287–302, 1989.
- Roland Carl Backhouse, Patrik Jansson, Johan Jeuring, and Lambert G. L. T. Meertens. Generic programming: An introduction. In *Advanced Functional Programming*, pages 28–115, 1998.
- Richard S. Bird. An introduction to the theory of lists. In M. Broy, editor, *Logic of Programming and Calculi of Discrete Design*, pages 3–42. Springer-Verlag, 1987. NATO ASI Series F Volume 36. Also available as Technical Monograph PRG-56, from the Programming Research Group, Oxford University.
- Wei-Ngan Chin. Towards an automated tupling strategy. In *Proceedings of the ACM SIGPLAN symposium on Partial evaluation and semantics-based program manipulation, PEPM’93*, pages 119–132, New York, NY, USA, 1993. ACM Press.
- Murray Cole. Parallel programming, list homomorphisms and the maximum segment sum problem. In Gerhard R. Joubert, Denis Trystram, Frans J. Peters, and David J. Evans, editors, *Parallel Computing: Trends and Applications, PARCO 1993, Grenoble, France*, pages 489–492. Elsevier, 1994.
- Murray Cole. Parallel programming with list homomorphisms. *Parallel Processing Letters*, 5:191–203, 1995.
- Richard Cole and Uzi Vishkin. The accelerated centroid decomposition technique for optimal parallel tree evaluation in logarithmic time. *Algorithmica*, 3:329–346, 1988.
- Maarten M. Fokkinga. Tupling and mutomorphisms. In *Squiggolist*, volume 1(4), 1989.
- Alfons Geser and Sergei Gorlatch. Parallelizing functional programs by generalization. *Journal of Functional Programming*, 9(6):649–673, 1999.
- Alan Gibbons and Wojciech Rytter. Optimal parallel algorithm for dynamic expression evaluation and context-free recognition. *Information and Computation*, 81(1):32–45, 1989.
- Jeremy Gibbons. The third homomorphism theorem. *Journal of Functional Programming*, 6(4):657–665, 1996.
- Jeremy Gibbons, Wentong Cai, and David B. Skillicorn. Efficient parallel algorithms for tree accumulations. *Science of Computer Programming*, 23(1):1–18, 1994.
- Zhenjiang Hu, Hideya Iwasaki, and Masato Takeichi. Formal derivation of efficient parallel programs by construction of list homomorphisms. *ACM Transactions on Programming Languages and Systems*, 19(3):444–461, 1997a.
- Zhenjiang Hu, Hideya Iwasaki, Masato Takeichi, and Akihiko Takano. Tupling calculation eliminates multiple data traversals. In *Proceedings of the 2nd ACM SIGPLAN International Conference on Functional Programming, ICFP’97, Amsterdam, The Netherlands*, pages 164–175. ACM Press, 1997b.
- G erard P. Huet. The zipper. *Journal of Functional Programming*, 7(5):549–554, 1997.
- Kiminori Matsuzaki. *Parallel Programming with Tree Skeletons*. PhD thesis, Graduate School of Information Science and Technology, The University of Tokyo, 2007.
- Kiminori Matsuzaki and Zhenjiang Hu. Implementation of tree skeletons on distributed-memory parallel computers. *Technical Report METR 2006-65*, Department of Mathematical Informatics, University of Tokyo, December 2006.
- Conor McBride. The derivative of a regular type is its type of one-hole contexts. Unpublished manuscript, 2001.
- Gary L. Miller and John H. Reif. Parallel tree contraction and its application. In *26th Annual Symposium on Foundations of Computer Science, 21-23 October 1985, Portland, Oregon, USA*, pages 478–489. IEEE, 1985.
- Akimasa Morihata and Kiminori Matsuzaki. A tree contraction algorithm on non-binary trees. *Technical Report METR 2008-27*, Department of Mathematical Informatics, University of Tokyo, June 2008.
- Kazutaka Morita, Akimasa Morihata, Kiminori Matsuzaki, Zhenjiang Hu, and Masato Takeichi. Automatic inversion generates divide-and-conquer parallel programs. In Jeanne Ferrante and Kathryn S. McKinley, editors, *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation, San Diego, California, USA, June 10-13, 2007*, pages 146–155, 2007.
- Simon Peyton Jones, editor. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- John H. Reif, editor. *Synthesis of Parallel Algorithms*. Morgan Kaufmann Publishers, 1993.
- David B. Skillicorn. Parallel implementation of tree skeletons. *Journal of Parallel and Distributed Computing*, 39(2):115–125, 1996.