

Translation of tree-processing programs into stream-processing programs based on ordered linear type¹

KOICHI KODAMA

Tokyo Institute of Technology

KOHEI SUENAGA

University of Tokyo

(e-mail: kohei@y1.is.s.u.tokyo.ac.jp)

NAOKI KOBAYASHI

Tohoku University

Abstract

There are two ways to write a program for manipulating tree-structured data such as XML documents: One is to write a tree-processing program focusing on the logical structure of the data and the other is to write a stream-processing program focusing on the physical structure. While tree-processing programs are easier to write than stream-processing programs, tree-processing programs are less efficient in memory usage since they use trees as intermediate data. Our aim is to establish a method for automatically translating a tree-processing program to a stream-processing one in order to take the best of both worlds. We first define a programming language for processing binary trees and a type system based on ordered linear type, and show that every well-typed program can be translated to an equivalent stream-processing program. We then extend the language and the type system to deal with XML documents. We have implemented an XML stream processor generator based on our algorithm, and obtained promising experimental results.

1 Introduction

There are two ways to write a program for manipulating tree-structured data such as XML documents (Bray *et al.* 2000): One is to write a tree-processing program focusing on the logical structure of the data and the other is to write a stream-processing program focusing on the physical structure. For example, as for XML processing, DOM (Document Object Model) API and programming language XDuce (Hosoya & Pierce 2003) are used for tree-processing programs, while SAX (Simple API for XML) is for stream-processing programs.

Figure 1 illustrates what tree-processing and stream-processing programs look like for the case of binary trees. The tree-processing program f takes a binary tree t as an

¹This article is a revised and extended version of the paper presented at the Second ASIAN Symposium on Programming Languages and Systems (APLAS 2004).

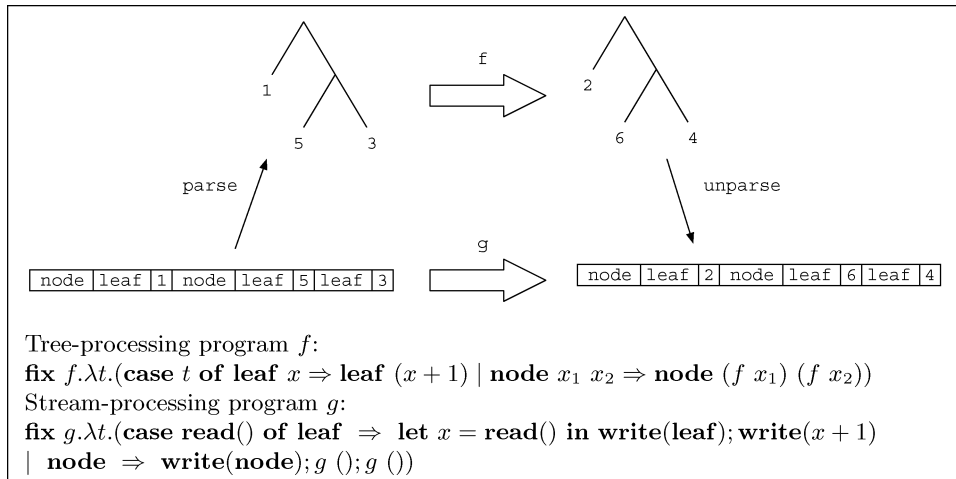


Fig. 1. Tree-processing and stream-processing programs.

input, and performs case analysis on t . If t is a leaf, it increments the value of the leaf. If t is a branch, f recursively processes the left and right subtrees. If actual tree data are represented as a sequence of tokens (as is often the case for XML documents), f must be combined with the function *parse* for parsing the input sequence, and the function *unparse* for unparsing the result tree into the output sequence, as shown in the figure. The stream-processing program g directly reads/writes data from/to streams. It reads an element from the input stream using the **read** primitive and performs case analysis on the element. If the input is the **leaf** tag, g outputs **leaf** to the output stream with the **write** primitive, reads another element, adds 1 to it, and outputs it. If the input is the **node** tag, g outputs **node** to the output stream and recursively calls the function g twice with the argument $()$.

Both of the approaches explained above have advantages and disadvantages. Tree-processing programs are written on the basis of the logical structure of data, so that it is easier to write, read, and manipulate (e.g., apply program transformation like deforestation (Wadler 1988)) than stream-processing programs. On the other hand, stream-processing programs have their own advantage in that intermediate tree structures are not needed, so that they often run faster than the corresponding tree-processing programs if input/output trees are physically represented as streams, as in the case of XML.

The goal of this article is to achieve the best of both approaches, by allowing a programmer to write a tree-processing program and automatically translating the program to an equivalent stream-processing program. To clarify the essence, we first use a λ -calculus with primitives on binary trees, and show how the translation works. We then extend the language for XML processing.

The key observation is that (1) stream processing is most effective *when trees are traversed and constructed from left to right in the depth-first preorder*, and (2) in that case, we can obtain from the tree-processing program the corresponding stream-processing program simply by replacing case analysis on an input tree with case

Tree-processing program:
fix $f.\lambda t.(\text{case } t \text{ of leaf } x \Rightarrow \text{leaf } x \mid \text{node } x_1 x_2 \Rightarrow \text{node } (f x_2) (f x_1))$

Fig. 2. A program that swaps children of every node.

analysis on input tokens, and replacing tree constructions with stream outputs. In fact, the stream-processing program in Figure 1, which satisfies the above criterion, is obtained from the tree-processing program in that way.

To check that a program satisfies the criterion, we use the idea of ordered linear types (Polakow 2001; Petersen *et al.* 2003). Ordered linear types, which are linear types (Baker 1992; Turner *et al.* 1995) extended with order constraints, describe not only how often but also *in which order* data are used. Our type system designed on the basis of the ordered linear types guarantees that a well-typed program traverses and constructs trees from left to right and in the depth-first preorder. Thus, every well-typed program can be translated to an equivalent stream-processing program. The tree-processing program f in Figure 1 is well-typed in our type system, so that it can automatically be translated to the stream-processing program g . On the other hand, the program in Figure 2 is not well-typed in our type system since it accesses the right subtree of an input before accessing the left subtree. In fact, we would obtain a wrong stream-processing program if we simply applied the above-mentioned translation to the program in Figure 2.

Our contributions can be summarized as follows:

- Formalization of the type system and the translation algorithm mentioned above for an extension of λ -calculus with binary trees.
- Proof of correctness of the translation algorithm.
- Extension of the type system and the translation algorithm for XML documents.
- Implementation and experiments. We implemented the type-checking algorithm and the translation algorithm for XML documents. With that implementation, we performed experiments using four micro-benchmarks in XSLTMark (DataPower Technology 2001) and one our own micro-benchmark.

The rest of the article is organized as follows: To clarify the essence, we first focus on a minimal calculus that deals with only binary trees in Sections 2–4. In Section 2, we define the source language and the target language of the translation. We define an ordered linear type system of the source language in Section 3. As mentioned above, the type system guarantees that every node of input trees is accessed exactly once in left-to-right, depth-first preorder. Section 4 presents a translation algorithm, proves its correctness, and discusses the improvement gained by the translation. The minimal calculus is not so expressive; especially, one can only write a program that does not at all store input/output trees on memory. Section 5 describes an extension that allows selective buffering of input/output trees. Section 6 shows how our framework can be applied to XML documents. Section 7 reports implementation and experiment. After discussing related work in Section 8, we conclude in Section 9.

Terms, values, and evaluation contexts:	
M (terms)	$::= i \mid \lambda x.M \mid x \mid M_1 M_2 \mid M_1 + M_2 \mid \mathbf{fix} f.M$ $\mid \mathbf{leaf} M \mid \mathbf{node} M_1 M_2$ $\mid (\mathbf{case} M \mathbf{of leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$
V (tree values)	$::= \mathbf{leaf} i \mid \mathbf{node} V_1 V_2$
W (values)	$::= i \mid \lambda x.M \mid V$
E_s (evaluation contexts)	$::= [] \mid E_s M \mid (\lambda x.M) E_s \mid E_s + M \mid i + E_s$ $\mid \mathbf{leaf} E_s \mid \mathbf{node} E_s M \mid \mathbf{node} V E_s$ $\mid (\mathbf{case} E_s \mathbf{of leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$
Reduction rules:	
	$E_s[i_1 + i_2] \rightarrow E_s[\mathit{plus}(i_1, i_2)]$ (ES-PLUS)
	$E_s[(\lambda x.M)W] \rightarrow E_s[[W/x]M]$ (ES-APP)
	$E_s[\mathbf{fix} f.M] \rightarrow E_s[[\mathbf{fix} f.M/f]M]$ (ES-FIX)
	$E_s[\mathbf{case leaf} i \mathbf{of leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2] \rightarrow E_s[[i/x]M_1]$ (ES-CASE1)
	$E_s[\mathbf{case node} V_1 V_2 \mathbf{of leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2] \rightarrow$ $E_s[[V_1/x_1, V_2/x_2]M_2]$ (ES-CASE2)

Fig. 3. The syntax, evaluation context, and reduction rules of the source language. $\mathit{plus}(i_1, i_2)$ is the sum of i_1 and i_2 .

2 Language definitions

We define the source and target languages in this section. The source language is a call-by-value functional language with primitives for manipulating binary trees. The target language is a call-by-value, impure functional language that uses imperative streams for input and output.

2.1 Source language

The syntax and operational semantics of the source language is summarized in Figure 3.

The meta-variables x and i range over the sets of variables and integers, respectively. The meta-variable W ranges over the set of values, which consists of integers i , λ -abstractions $\lambda x.M$, and binary trees V . A binary tree V is either a leaf labeled with an integer or a tree with two children. ($\mathbf{case} M \mathbf{of leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2$) performs case analysis on a tree. If M is a leaf, x is bound to its label and M_1 is evaluated. Otherwise, x_1 and x_2 are bound to the left and right children, respectively, and M_2 is evaluated. $\mathbf{fix} f.M$ is a recursive function that satisfies $f = M$. Bound and free variables are defined as usual. We assume that α -conversion is implicitly applied so that bound variables are always different from each other and free variables.

We write $\mathbf{let} x = M_1 \mathbf{in} M_2$ for $(\lambda x.M_2) M_1$. Especially, if M_2 contains no free occurrence of x , we write $M_1; M_2$ for it.

2.2 Target language

The syntax and operational semantics of the target language are summarized in Figure 4. A stream, represented by the meta-variable S , is a sequence consisting of

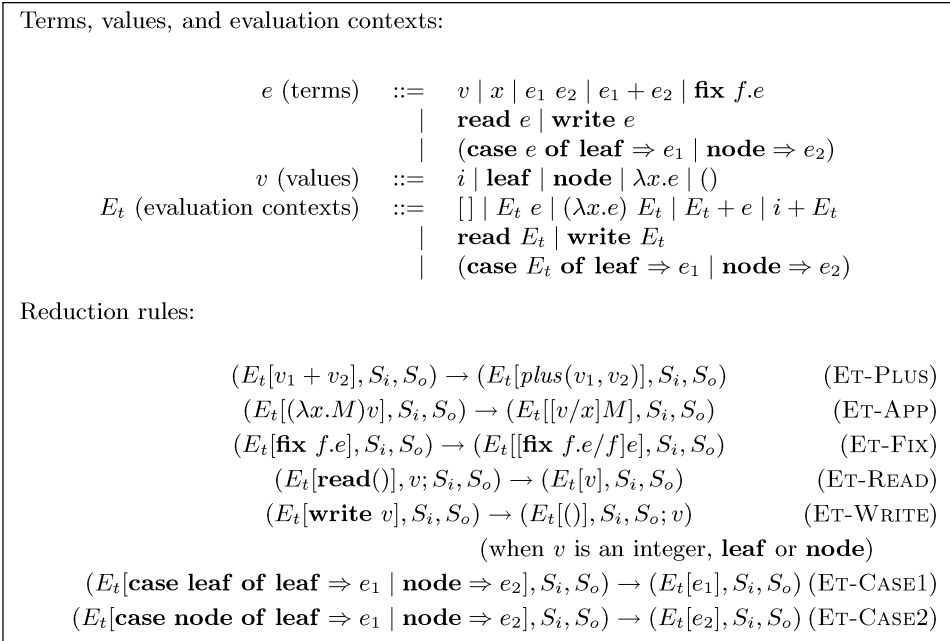


Fig. 4. The reduction rules of the target language.




Fig. 5. Programs that calculate the sum of leaf elements of a binary tree.

leaf, **node**, and integers. We write \emptyset for the empty sequence and write $S_1; S_2$ for the concatenation of the sequences S_1 and S_2 .

A primitive **read** is taken for reading a token (**leaf**, **node**, or an integer) from the input stream; **write** is a primitive for writing a value to the output stream. The term $(\mathbf{case} e \text{ of } \mathbf{leaf} \Rightarrow e_1 \mid \mathbf{node} \Rightarrow e_2)$ performs a case analysis on the value of e . If e evaluates to **leaf**, e_1 is evaluated, and if e evaluates to **node**, e_2 is evaluated. $\mathbf{fix} f.e$ is a recursive function that satisfies $f = e$. Bound and free variables are defined as usual.

We write $\mathbf{let} x = e_1 \text{ in } e_2$ for $(\lambda x.e_2) e_1$. Especially, if e_2 does not contain x as a free variable, we write $e_1; e_2$ for it.

Figure 5 shows programs that take a tree as an input and calculate the sum of leaf elements. The source program takes a tree t as an argument of the function, and performs a case analysis on t . If t is a leaf, the program binds x to the element and returns it. If t is a branch node, the program recursively applies f to the left and right children, and returns the sum of the results. The target program reads a tree (as a sequence of tokens) from the input stream, performs a case analysis on tokens, and returns the sum of leaf elements. Here, we assume that the input

stream represents a valid tree. If the input stream is in a wrong format (e.g., when the stream is **node**;1;2), the execution gets stuck.

3 Type system

In this section, we present a type system of the source language, which guarantees that a well-typed program reads every node of an input tree exactly once from left to right in the left-to-right, depth-first preorder. Thanks to this guarantee, any well-typed program can be translated to an equivalent, stream-processing program without changing the structure of the program, as shown in the next section. To enforce the correct access order on input trees, we use ordered linear types (Polakow 2001; Petersen *et al.* 2003).

3.1 Type and type environment

Definition 3.1 (Type)

The set of types, ranged over by τ , is defined by:

$$\begin{aligned} \tau \text{ (type)} &::= \mathbf{Int} \mid \mathbf{Tree}^d \mid \tau_1 \rightarrow \tau_2 \\ d \text{ (mode)} &::= - \mid +. \end{aligned}$$

where **Int** is the type of integers. For a technical reason, we distinguish between input trees and output trees by types. We write **Tree**⁻ for the type of input trees and write **Tree**⁺ for the type of output trees. $\tau_1 \rightarrow \tau_2$ is the type of functions from τ_1 to τ_2 .

We introduce two kinds of type environments for our type system: ordered linear type environments and (non-ordered) type environments.

Definition 3.2 (Ordered linear type environment)

An *ordered linear type environment*, ranged over by the meta-variable Δ , is a sequence of the form $x_1 : \mathbf{Tree}^-, \dots, x_n : \mathbf{Tree}^-$, where x_1, \dots, x_n are different from each other. We write Δ_1, Δ_2 for the concatenation of Δ_1 and Δ_2 .

An ordered linear type environment $x_1 : \mathbf{Tree}^-, \dots, x_n : \mathbf{Tree}^-$ specifies not only that x_1, \dots, x_n are bound to trees but also that each of x_1, \dots, x_n must be accessed exactly once in this order and that each of the trees bound to x_1, \dots, x_n must be accessed in the left-to-right, depth-first preorder.

Definition 3.3 (Non-ordered type environment)

A *(non-ordered) type environment*, ranged over by a meta-variable Γ , is a set of the form $\{x_1 : \tau_1, \dots, x_n : \tau_n\}$, where x_1, \dots, x_n are different from each other and $\{\tau_1, \dots, \tau_n\}$ does not contain **Tree**^d.

Note that a non-ordered type environment must not contain variables of tree types. **Tree**⁻ is excluded since input trees must be accessed in the specific order. **Tree**⁺ is excluded in order to forbid output trees from being bound to variables. For example, we exclude a program like **let** $x_1 = t_1$ **in** **let** $x_2 = t_2$ **in** **node** $x_1 x_2$ when t_1 and t_2 have type **Tree**⁺. This restriction is convenient for ensuring that trees

are constructed in the specific (from left to right and in the depth-first pre-) order. Note that a non-ordered type environment can contain types that use tree types like $\mathbf{Tree}^- \rightarrow \mathbf{Tree}^+$. See the first example in Section 3.2.

3.2 Type judgment

A type judgement is of the form $\Gamma \mid \Delta \vdash M : \tau$, where Γ is a non-ordered type environment, Δ is an ordered linear type environment and Γ and Δ do not share the same variables. The judgment means, “If M evaluates to a value under an environment described by Γ and Δ , the value has type τ and the variables in Δ are accessed in the order specified by Δ .” For example, if $\Gamma = \{f : \mathbf{Tree}^- \rightarrow \mathbf{Tree}^+\}$ and $\Delta = x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-$,

$$\Gamma \mid \Delta \vdash \mathbf{node} (f \ x_1) (f \ x_2) : \mathbf{Tree}^+$$

holds, while

$$\Gamma \mid \Delta \vdash \mathbf{node} (f \ x_2) (f \ x_1) : \mathbf{Tree}^+$$

does not. The latter program violates the restriction specified by Δ such that x_1 and x_2 must be accessed in this order.

$\Gamma \mid \Delta \vdash M : \tau$ is the least relation that is closed under the rules in Figure 6. Note that in the typing rules, there is an implicit assumption that the conclusion must be a well-formed judgment in order for a rule to be applied. Therefore, in T-APP, T-PLUS, and T-NODE, it is implicitly assumed that the variables in Δ_1 and Δ_2 are disjoint.

T-VAR1, T-VAR2, and T-INT are the rules for variables and integer constants. As in ordinary linear type systems, these rules prohibit variables that do not allow a term to occur in the ordered linear type environment. (In other words, weakening is not allowed on an ordered linear type environment.) This restriction is necessary to guarantee that each variable in an ordered linear type environment is accessed exactly once.

T-ABS1 and T-ABS2 are rules for λ -abstraction. Note that the ordered type environments of the conclusions of these rules must be empty. This restriction prevents input trees from being stored in function closures. This makes it easy to enforce the access order on input trees. For example, without this restriction, the function

$$\lambda t. \mathbf{let} \ g = \lambda f. (f \ t) \ \mathbf{in} \ (g \ \mathit{sumtree}) + (g \ \mathit{sumtree})$$

would be well-typed, where *sumtree* is the function given in Figure 5. However, when a tree is passed to this function, its nodes are accessed twice because the function g is called twice. The program above is actually rejected by our type system since the closure $\lambda f. (f \ t)$ is not well-typed because of the restriction of T-ABS2.²

² We can relax the restriction by controlling usage of not only trees but also functions, as in the resource usage analysis (Igarashi & Kobayashi 2002). The resulting type system would, however, become very complex.

$\Gamma \mid x : \mathbf{Tree}^- \vdash x : \mathbf{Tree}^-$	(T-VAR1)
$\Gamma, x : \tau \mid \emptyset \vdash x : \tau$	(T-VAR2)
$\Gamma \mid \emptyset \vdash i : \mathbf{Int}$	(T-INT)
$\frac{\Gamma \mid x : \mathbf{Tree}^- \vdash M : \tau}{\Gamma \mid \emptyset \vdash \lambda x.M : \mathbf{Tree}^- \rightarrow \tau}$	(T-ABS1)
$\frac{\Gamma, x : \tau_1 \mid \emptyset \vdash M : \tau_2}{\Gamma \mid \emptyset \vdash \lambda x.M : \tau_1 \rightarrow \tau_2}$	(T-ABS2)
$\frac{\Gamma \mid \Delta_1 \vdash M_1 : \tau_2 \rightarrow \tau_1 \quad \Gamma \mid \Delta_2 \vdash M_2 : \tau_2}{\Gamma \mid \Delta_1, \Delta_2 \vdash M_1 M_2 : \tau_1}$	(T-APP)
$\frac{\Gamma \mid \Delta_1 \vdash M_1 : \mathbf{Int} \quad \Gamma \mid \Delta_2 \vdash M_2 : \mathbf{Int}}{\Gamma \mid \Delta_1, \Delta_2 \vdash M_1 + M_2 : \mathbf{Int}}$	(T-PLUS)
$\frac{\Gamma, f : \tau_1 \rightarrow \tau_2 \mid \emptyset \vdash M : \tau_1 \rightarrow \tau_2}{\Gamma \mid \emptyset \vdash \mathbf{fix} f.M : \tau_1 \rightarrow \tau_2}$	(T-FIX)
$\frac{\Gamma \mid \Delta \vdash M : \mathbf{Int}}{\Gamma \mid \Delta \vdash \mathbf{leaf} M : \mathbf{Tree}^+}$	(T-LEAF)
$\frac{\Gamma \mid \Delta_1 \vdash M_1 : \mathbf{Tree}^+ \quad \Gamma \mid \Delta_2 \vdash M_2 : \mathbf{Tree}^+}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{node} M_1 M_2 : \mathbf{Tree}^+}$	(T-NODE)
$\frac{\Gamma \mid \Delta_1 \vdash M : \mathbf{Tree}^- \quad \Gamma, x : \mathbf{Int} \mid \Delta_2 \vdash M_1 : \tau \quad \Gamma \mid x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-, \Delta_2 \vdash M_2 : \tau}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{case} M \mathbf{of} \mathbf{leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2 : \tau}$	(T-CASE)

Fig. 6. Typing rules.

T-APP is the rule for function application. The ordered linear type environments of M_1 and M_2 , Δ_1 and Δ_2 , respectively, are concatenated in this order because when $M_1 M_2$ is evaluated, (1) M_1 is first evaluated, (2) M_2 is then evaluated, and (3) M_1 is finally applied to M_2 . In the first step, the variables in Δ_1 are accessed in the order specified by Δ_1 . In the second and third steps, the variables in Δ_2 are accessed in the order specified by Δ_2 . On the other hand, because there is no restriction on usage of the variables in a non-ordered type environment, the same type environment (Γ) is used for both subterms.

T-LEAF and T-NODE are rules for tree construction. We concatenate the ordered type environments of M_1 and M_2 , Δ_1 and Δ_2 , in this order as we did in T-APP. Also as in the rule T-APP, Δ_1 and Δ_2 cannot share the same variables in T-NODE.

T-CASE is the rule for case expressions. If M matches **node** $x_1 x_2$, subtrees x_1 and x_2 have to be accessed in this order after that. This restriction is expressed by $x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-, \Delta_2$, the ordered linear type environment of M_2 .

T-FIX is the rule for recursion. Note that the ordered type environment must be empty as in T-ABS2.

The program in Figure 1 is typed as shown in Figure 7. On the other hand, the program in Figure 2 is ill-typed: $\Gamma \mid x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^- \vdash \mathbf{node} (f x_2) (f x_1) : \mathbf{Tree}^+$

$$\begin{array}{c}
 \vdots \qquad \qquad \qquad \vdots \\
 \frac{\Gamma \mid t : \mathbf{Tree}^- \vdash t : \mathbf{Tree}^- \quad \Gamma' \mid \emptyset \vdash M : \mathbf{Tree}^+ \quad \Gamma \mid \Delta \vdash \mathbf{node} (f x_1) (f x_2) : \mathbf{Tree}^+}{\Gamma \mid t : \mathbf{Tree}^- \vdash \mathbf{case} t \text{ of leaf } x \Rightarrow M \mid \mathbf{node} x_1 x_2 \Rightarrow \mathbf{node} (f x_1) (f x_2) : \mathbf{Tree}^+} \\
 \frac{\Gamma \mid \emptyset \vdash \lambda t. \mathbf{case} t \text{ of leaf } x \Rightarrow M \mid \mathbf{node} x_1 x_2 \Rightarrow \mathbf{node} (f x_1) (f x_2) : \mathbf{Tree}^- \rightarrow \mathbf{Tree}^+}{\emptyset \mid \emptyset \vdash \mathbf{fix} f. \lambda t. \mathbf{case} t \text{ of leaf } x \Rightarrow M \mid \mathbf{node} x_1 x_2 \Rightarrow \mathbf{node} (f x_1) (f x_2) : \mathbf{Tree}^- \rightarrow \mathbf{Tree}^+}
 \end{array}$$

Fig. 7. An example of typing derivation. $\Gamma = \{f : \mathbf{Tree}^- \rightarrow \mathbf{Tree}^+\}$, $M = \mathbf{leaf} (x + 1)$, $\Gamma' = \{f : \mathbf{Tree}^- \rightarrow \mathbf{Tree}^+, x : \mathbf{Int}\}$, and $\Delta = x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-$.

must hold for the program to be well-typed, but it cannot be derived by using T-NODE.

3.3 Examples of programs

Figure 8 shows more examples of source programs. The first and second programs (or the catamorphism; Meijer *et al.* 1991) apply the same operation on every node of the input tree. (The return value of the function *tree.fold* cannot, however, be a tree because a return value of *tree.fold* is passed to *g*. Note that functions cannot take a value of type \mathbf{Tree}^+ as argument because binding variables to \mathbf{Tree}^+ values is forbidden in our type system, as we mentioned in Section 3.1.) The third program returns a tree obtained by incrementing all the leaf values of odd depth by 1. In this way, one can also write functions that process nodes in a non-uniform manner. The fourth program is an example of a well-typed program in which variables of type \mathbf{Tree}^- (y_1, y_2 and x_2) are not bound to sibling nodes in the input tree.

The fifth program in Figure 8 takes a tree as an input and returns the right subtree. Because of the restriction imposed by the type system, the program uses subfunctions *copy_tree* and *skip_tree* for explicitly copying and skipping trees.³ (See Section 9 for a method for automatically inserting these functions.)

The last program in Figure 8 takes a tree as an input, and returns a tree consisting of a copy of the input tree and the number of leaves of the input. Here, we assume that the source language has been extended with ML-style reference cells. The type system can be easily extended: The only restriction is that input/output trees cannot be stored in reference cells.

Remark 3.1

The reader may think that our type system is too restrictive. The following two functions, which are obtained by modifying the first and third programs in Figure 8,

³ Because of the restriction that λ -abstractions cannot contain variables of type \mathbf{Tree}^d , we need to introduce **let** expression as a primitive and extend typing rules with the following rule:

$$\frac{\Gamma \mid \Delta_1 \vdash M_1 : \tau' \quad \Gamma \cup \{x : \tau'\} \mid \Delta_2 \vdash M_2 : \tau \quad \tau' \neq \mathbf{Tree}^d}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{let} x = M_1 \mathbf{in} M_2 : \tau} \quad (\text{T-LET})$$

```

fix tree_map.λf.λt.
  case t of
    leaf x ⇒ leaf (f x)
    | node x1 x2 ⇒ node (tree_map f x1) (tree_map f x2)

fix tree_fold.λf.λg.λt.
  case t of
    leaf n ⇒ (f n)
    | node t1 t2 ⇒ (g (tree_fold f g t1) (tree_fold f g t2))

fix inc_alt.
  let inc_alt_even =
    λt.case t of
      Leaf x ⇒ leaf x + 1
      | node x1 x2 ⇒ node (inc_alt x1) (inc_alt x2)
  in
    λt.case t of
      Leaf x ⇒ leaf x
      | node x1 x2 ⇒ node (inc_alt_even x1) (inc_alt_even x2)

fix inc_left.λt.
  case t of
    leaf x ⇒ leaf x
    | node x1 x2 ⇒
      node
        (case x1 of leaf y ⇒ leaf (y + 1)
         | node y1 y2 ⇒ node (inc_left y1) (inc_left y2))
        (inc_left x2)

let copy_tree =
  fix copy_tree.λt.
    case t of
      leaf x ⇒ leaf x
      | node x1 x2 ⇒ node (copy_tree x1) (copy_tree x2) in
let skip_tree =
  fix skip_tree.λt.
    case t of
      leaf x ⇒ 0
      | node x1 x2 ⇒ (skip_tree x1); (copy_tree x2) in
  λt.(case t of leaf x ⇒ leaf x | node x1 x2 ⇒ (skip_tree x1); (copy_tree x2))

let copy_and_count_aux = fix f.λr.λt.
  case t of
    leaf x ⇒ (r :=!r + 1; leaf x)
    | node x1 x2 ⇒ node (f r x1) (f r x2) in
let copy_and_count =
  λt.let x = ref 0 in node (copy_and_count_aux x t) (leaf !x)

```

Fig. 8. Program examples.

highlight the main limitations of our type system.

```

fix tree_map'.λt.λf.
  case t of
    leaf x ⇒ leaf (f x)
  | node x1 x2 ⇒ node (tree_map' x1 f) (tree_map' x2 f))

fix tree_map''.λf.λt.
  case t of
    leaf x ⇒ leaf (f x)
  | node x1 x2 ⇒ let t1 = tree_map'' f x1 in node t1 (tree_map'' f x2)

```

They are operationally equivalent to the first program in Figure 8, but are *not* well-typed. The subexpression $\lambda f.\mathbf{case} t \mathbf{of} \dots$ of the first program above violates the rule T-ABS2, which says that a variable of type \mathbf{Tree}^- must not occur free in a function. The subexpression $\mathbf{let} t_1 = \mathit{tree_map}'' f x_1 \mathbf{in} \dots$ violates the condition that no variable can be bound to a tree of type \mathbf{Tree}^+ .

It would be possible to remove the limitations to some extent by introducing a more complex type system.⁴ We do not do so, however. We believe that in many cases, the above problems can be avoided by adding a simple pre-processing phase; The first problem can often be avoided by uncurrying transformation, and the second problem can often be avoided by inlining arguments of the **node** constructor. Thus, it is unclear how much benefit is obtained in practice by introducing a more complex type system. Moreover, the simplicity of the type system is beneficial for various extensions of the framework, especially for automatic insertion of buffering primitives discussed elsewhere (Suenaga *et al.* 2005).

3.4 Type-checking algorithm

We sketch a type-checking algorithm that takes a type-annotated term M and a non-ordered type environment Γ as input, and outputs τ and Δ such that $\Gamma \mid \Delta \vdash M : \tau$, or reports failure if such τ or Δ does not exist. Here, by a type-annotated term, we mean a term whose bound variables are annotated with types.

The rules in Figure 6 can be interpreted as such as an algorithm. Consider Γ and M in each rule as input and Δ and τ in each rule as output. Then, read each rule from bottom to top. For example, T-APP can be read as follows. For input $(\Gamma, M_1 M_2)$, do the following:

1. Execute the algorithm with input (Γ, M_1) and obtain the result $(\Delta_1, \tau_2 \rightarrow \tau_1)$.
2. Execute the algorithm with input (Γ, M_2) and obtain the result (Δ_2, τ'_2) .
3. Check (Δ_1, Δ_2) is well-formed (i.e., does not share the same variables) and $\tau_2 = \tau'_2$.
4. Output $((\Delta_1, \Delta_2), \tau_1)$.

⁴ For the second problem, we also have to modify the translation algorithm discussed in the next section.

$$\begin{aligned}
\mathcal{A}(x) &= x \\
\mathcal{A}(i) &= i \\
\mathcal{A}(\lambda x.M) &= \lambda x.\mathcal{A}(M) \\
\mathcal{A}(M_1 M_2) &= \mathcal{A}(M_1) \mathcal{A}(M_2) \\
\mathcal{A}(M_1 + M_2) &= \mathcal{A}(M_1) + \mathcal{A}(M_2) \\
\mathcal{A}(\mathbf{fix} f.M) &= \mathbf{fix} f.\mathcal{A}(M) \\
\mathcal{A}(\mathbf{leaf} M) &= \mathbf{write}(\mathbf{leaf}); \mathbf{write}(\mathcal{A}(M)) \\
\mathcal{A}(\mathbf{node} M_1 M_2) &= \mathbf{write}(\mathbf{node}); \mathcal{A}(M_1); \mathcal{A}(M_2) \\
\mathcal{A}(\mathbf{case} M \mathbf{of} \mathbf{leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2) &= \\
&\quad \mathbf{case} \mathcal{A}(M); \mathbf{read}() \mathbf{of} \mathbf{leaf} \Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} \mathcal{A}(M_1) \\
&\quad \mid \mathbf{node} \Rightarrow [()/x_1, ()/x_2]\mathcal{A}(M_2)
\end{aligned}$$

Fig. 9. Translation algorithm.

4 Translation algorithm

Definition 4.1 (Translation algorithm)

The translation function \mathcal{A} from a source program into the target program is defined in Figure 9.

\mathcal{A} maps a source program to a target program, preserving the structure of the source program and replacing operations on trees with operations on streams. For example, the tree construction primitives (**leaf** M and **node** $M_1 M_2$) are translated into writing operations to the output stream. The pattern matching primitive (**case** $M \mathbf{of} \mathbf{leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2$) is translated into the case analysis on the head of the input stream. As we have observed in Section 1, this simple translation works because the type system guarantees that well-typed programs access an input tree in left-to-right, depth-first preorder.

In the rest of this section, we prove the correctness of the translation algorithm \mathcal{A} and discuss the efficiency of the translated program.

4.1 Correctness of translation algorithm

The correctness of the translation algorithm \mathcal{A} is stated as follows:

Definition 4.2

A function $\llbracket \cdot \rrbracket$ from the set of trees to the set of streams is defined by:

$$\begin{aligned}
\llbracket \mathbf{leaf} i \rrbracket &= \mathbf{leaf}; i \\
\llbracket \mathbf{node} V_1 V_2 \rrbracket &= \mathbf{node}; \llbracket V_1 \rrbracket; \llbracket V_2 \rrbracket.
\end{aligned}$$

Theorem 4.1 (Correctness of translation)

If $\emptyset \mid \emptyset \vdash M : \mathbf{Tree}^- \rightarrow \tau$ and τ is **Int** or **Tree**⁺, the following properties hold for any tree value V :

- (i) $M V \rightarrow^* i$ if and only if $(\mathcal{A}(M)(), \llbracket V \rrbracket, \emptyset) \rightarrow^* (i, \emptyset, \emptyset)$;
- (ii) $M V \rightarrow^* V'$ if and only if $(\mathcal{A}(M)(), \llbracket V \rrbracket, \emptyset) \rightarrow^* ((, \emptyset, \llbracket V' \rrbracket))$.

The above theorem means that a source program and the corresponding target program evaluate to the same value. The clause (i) is for the case where the result is an integer and (ii) is for the case where the result is a tree.

$$\begin{array}{l}
 U ::= x \mid i \mid \lambda x.M \\
 (E_s[i_1 + i_2], \delta) \rightarrow (E_s[\text{plus}(i_1, i_2)], \delta) \quad (\text{ES2-PLUS}) \\
 (E_s[(\lambda x.M)U], \delta) \rightarrow (E_s[[U/x]M], \delta) \quad (\text{ES2-APP}) \\
 (E_s[\text{case } y \text{ of leaf } x \Rightarrow M_1 \mid \text{node } x_1 x_2 \Rightarrow M_2], (y \mapsto \text{leaf } i, \delta)) \rightarrow \\
 \quad (E_s[[i/x]M_1], \delta) \quad (\text{ES2-CASE1}) \\
 (E_s[\text{case } y \text{ of leaf } x \Rightarrow M_1 \mid \text{node } x_1 x_2 \Rightarrow M_2], (y \mapsto \text{node } V_1 V_2, \delta)) \rightarrow \\
 \quad (E_s[M_2], (x_1 \mapsto V_1, x_2 \mapsto V_2, \delta)) \quad (\text{ES2-CASE2}) \\
 (E_s[\text{fix } f.M], \delta) \rightarrow (E_s[[\text{fix } f.M/f]M], \delta) \quad (\text{ES2-FIX})
 \end{array}$$

Fig. 10. The new reduction semantics of the source language.

We give an outline of the proof of Theorem 4.1 in the following. The basic idea of the proof is to show a correspondence between reduction steps of a source program and those of the target program. However, the reduction semantics given in Section 2 is not convenient for showing the correspondence because the target language semantics imposes a restriction on access to stream input/outputs, while the source language semantics in Section 2 does not. So, we introduce another semantics of the source language that has the corresponding order restriction on environments holding input/output trees, and prove: (1) for well-typed programs, the new semantics is equivalent to the one in Section 2 (Corollary 4.3 below), and (2) evaluation of the source program based on the new semantics agrees with evaluation of the corresponding target program (Theorem 4.4 below). First, we define the new operational semantics of the source language. The semantics takes the access order of input trees into account.

Definition 4.3 (Ordered environment)

An ordered environment is a sequence of the form $x_1 \mapsto V_1, \dots, x_n \mapsto V_n$, where x_1, \dots, x_n are distinct from each other.

We use a meta-variable δ to represent an ordered environment. Given an ordered environment $x_1 \mapsto V_1, \dots, x_n \mapsto V_n$, a program must access variables x_1, \dots, x_n in this order.

Definition 4.4 (New reduction semantics)

The reduction relation $(M, \delta) \rightarrow (M', \delta')$ is the least relation that satisfies the rules in Figure 10.

The meta-variable U in Figure 10 ranges over the set of variables (bound to tree values) and non-tree values.

The differences between the new reduction semantics above and the original one in Section 2 are (1) input trees are substituted in the original semantics while they are held in ordered environments in the new semantics (compare ES-CASE2 with ES2-CASE2) and (2) input trees must be accessed in the order specified by δ in the new semantics (note that variable y that is being referred to must be at the head of the ordered environment in ES2-CASE1 and ES2-CASE2). Thus, evaluation based on the new semantics can differ from the one in Section 2 only when the latter

one succeeds while the former one gets stuck because of the restriction on access to input trees. As the following theorem (Theorem 4.2) states, this cannot happen if the program is well-typed, so that both semantics are equivalent for well-typed programs (Corollary 4.3).

Definition 4.5

A function $\langle\langle\cdot\rangle\rangle$ from the set of ordered environments to the set of ordered linear type environments is defined by:

$$\begin{aligned} \langle\langle\emptyset\rangle\rangle &= \emptyset \\ \langle\langle x \mapsto V, \delta \rangle\rangle &= x : \mathbf{Tree}^-, \langle\langle\delta\rangle\rangle \end{aligned}$$

Theorem 4.2

Suppose $\emptyset \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$. Then the following conditions hold:

- M is a value or a variable, or $(M, \delta) \rightarrow (M', \delta')$ holds for some M' and δ' .
- If $(M, \delta) \rightarrow (M', \delta')$ holds, then $\emptyset \mid \langle\langle\delta'\rangle\rangle \vdash M' : \tau$.

Proof

See Appendix 9. □

Corollary 4.3

If $\emptyset \mid \emptyset \vdash M : \mathbf{Tree}^- \rightarrow \tau$ and if $\tau \in \{\mathbf{Int}, \mathbf{Tree}^+\}$, $MV \rightarrow^* W$ if and only if $(Mx, x \mapsto V) \rightarrow^* (W, \emptyset)$ for any tree value V .

Proof

See Appendix 9. □

The following theorem states that the evaluation of a source program under the new rules agrees with the evaluation of the target program.

Theorem 4.4

If $\emptyset \mid \emptyset \vdash M : \mathbf{Tree}^- \rightarrow \tau$ and $\tau \in \{\mathbf{Int}, \mathbf{Tree}^+\}$ holds, the following statements hold for any tree value V .

- (i) $(Mx, x \mapsto V) \rightarrow^* (i, \emptyset)$ holds if and only if $(\mathcal{A}(M)(\cdot), \llbracket V \rrbracket, \emptyset) \rightarrow^* (i, \emptyset, \emptyset)$
- (ii) If $(Mx, x \mapsto V) \rightarrow^* (V', \emptyset)$ holds if and only if $(\mathcal{A}(M)(\cdot), \llbracket V \rrbracket, \emptyset) \rightarrow^* ((\cdot), \emptyset, \llbracket V' \rrbracket)$

We hereafter give an outline of the proof of Theorem 4.4. Figure 11 illustrates the idea of the proof (for the case where the result is a tree). The relation \sim (defined later in Definition 4.7) in the diagram expresses the correspondence between an evaluation state of a source program (M, δ) and a state of a target program (e, S_i, S_o) . We show that the target program $\mathcal{A}(M)$ can always be reduced to a state corresponding to the initial state of the source program M (Lemma 4.5 below) and that reductions and the correspondence relation commute (Lemma 4.6). These imply that the whole diagram in Figure 11 commutes, that is, the second statement of Theorem 4.4 holds.

To define the correspondence $(M, \delta) \sim (e, S_i, S_o)$ between states, we use the following function $\langle\cdot\rangle$, which maps an ordered environment to the corresponding stream.




Fig. 11. Evaluation of a source and the target program.

Definition 4.6

A function $\langle \cdot \rangle$ from the set of ordered environments to the set of streams is defined by:

$$\langle \emptyset \rangle = \emptyset$$

$$\langle x \mapsto V, \delta \rangle = [V]; \langle \delta \rangle$$

Definition 4.7 (Correspondence between states)

The relations $(M, \delta) \sim (e, S_i, S_o)$ and $M \sim_\gamma (e, S_o)$ are the least relations closed under the rules in Figure 12.

In the figure, the meta-variable γ denotes a set of variables. $FV(M)$ is the set of free variables in M . $\mathcal{A}_\gamma(M)$ is the term obtained from $\mathcal{A}(M)$ by replacing every occurrence of variables in γ with $()$. The meta-variable I represents the term that is being reduced. Note that any term M can be written as $E_s[I]$ if it is reducible.

In the relation $(M, \delta) \sim (e, S_i, S_o)$, e represents the rest of computation, S_i is the input stream, and S_o is the already generated output stream. For example, $(\mathbf{node}(\mathbf{leaf} \ 1)(\mathbf{leaf} \ (2 + 3)), \emptyset)$ corresponds to $(2 + 3, \emptyset, \mathbf{node}; \mathbf{leaf}; 1; \mathbf{leaf})$.

We explain some of the rules in Figure 12 below.

- C-TREE: A source program V represents a state where the tree V has been constructed. Thus, it corresponds to $((), [V])$, where there is nothing to be computed and V has been written to the output stream.
- C-NODE1: A source program $\mathbf{node} \ E_s[I] \ M$ represents a state where the left subtree is being computed. Thus, the rest computation of the target program is $(e; \mathcal{A}_\gamma(M))$, where e is the rest computation in $E_s[I]$ and $\mathcal{A}_\gamma(M)$ represents the computation for constructing the right subtree. The corresponding output stream is $\mathbf{node}; S_o$ because \mathbf{node} represents the root of the tree being constructed, and S_o represents the part of the left subtree that has been already constructed.

Lemmas 4.5 and 4.6 below imply that the whole diagram in Figure 12 commutes, which completes the proof of Theorem 4.4.

Lemma 4.5

Suppose $\emptyset \mid \langle \langle \delta \rangle \rangle \vdash M : \tau$. Then, there exist e and S_o that satisfy

- $(M, \delta) \sim (e, \langle \delta \rangle, S_o)$
- $(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S_o)$

Proof

See Appendix 9. □

$I ::=$	$(\lambda x.M)U \mid i_1 + i_2 \mid (\text{case } y \text{ of leaf } x \Rightarrow M_1 \mid \text{node } x_1 x_2 \Rightarrow M_2) \mid \text{fix } f.M$	
	$\frac{M \sim_{FV(M)} (e, S_o) \quad S_i = \langle \delta \rangle}{(M, \delta) \sim (e, S_i, S_o)}$	
	$\frac{}{U \sim_\gamma (\mathcal{A}_\gamma(U), \emptyset)}$	(C-VALUE)
	$\frac{}{V \sim_\gamma ((), \llbracket V \rrbracket)}$	(C-TREE)
	$\frac{}{I \sim_\gamma (\mathcal{A}_\gamma(I), \emptyset)}$	(C-INST)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{E_s[I] M \sim_\gamma (e \mathcal{A}_\gamma(M), S_o)}$	(C-APP1)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{(\lambda x.M) E_s[I] \sim_\gamma ((\lambda x.\mathcal{A}_\gamma(M)) e, S_o)}$	(C-APP2)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{E_s[I] + M \sim_\gamma (e + \mathcal{A}_\gamma(M), S_o)}$	(C-PLUS1)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{i + E_s[I] \sim_\gamma (i + e, S_o)}$	(C-PLUS2)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{\text{leaf } E_s[I] \sim_\gamma (\text{write}(e), \text{leaf}; S_o)}$	(C-LEAF)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{\text{node } E_s[I] M \sim_\gamma (e; \mathcal{A}_\gamma(M), \text{node}; S_o)}$	(C-NODE1)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{\text{node } V E_s[I] \sim_\gamma (e, \text{node}; \llbracket V \rrbracket; S_o)}$	(C-NODE2)
	$\frac{E_s[I] \sim_\gamma (e, S_o)}{\text{case } E_s[I] \text{ of leaf } x \Rightarrow M_1 \mid \text{node } x_1 x_2 \Rightarrow M_2 \sim_\gamma$	(C-CASE)
	$(\text{case } e; \text{read } () \text{ of leaf } \Rightarrow \text{let } x = \text{read } () \text{ in } \mathcal{A}_\gamma(M_1), S_o$	
	$\text{node } \Rightarrow [()/x_1, ()/x_2]\mathcal{A}_\gamma(M_2)$	

Fig. 12. Correspondence between run-time states of source and target programs.

Lemma 4.6

If $\emptyset \mid \langle \langle \delta \rangle \rangle \vdash M : \tau$ and $(M, \delta) \sim (e, S_i, S_o)$, the following conditions hold:

- $S_i = \langle \delta \rangle$.
- If $(M, \delta) \rightarrow (M', \delta')$, then there exist e' and S'_i and S'_o that satisfy $(e, \langle \delta \rangle, S_o) \rightarrow^+ (e', \langle \delta' \rangle, S'_o)$ and $(M', \delta') \sim (e', S'_i, S'_o)$.
- If $(e, \langle \delta \rangle, S_o)$ is reducible, there exist M' and δ' that satisfy $(M, \delta) \rightarrow (M', \delta')$.

Proof

The first condition is obvious. For the proof of the other conditions, see Appendix 9. \square

4.2 Efficiency of translated programs

Let M be a source program of type $\mathbf{Tree}^- \rightarrow \mathbf{Tree}^+$. We argue below that the target program $\mathcal{A}(M)$ runs more efficiently than the source program $\text{unparse} \circ M \circ \text{parse}$,

where *parse* is a function that parses the input stream and returns a binary tree, and *unparse* is a function that takes a binary tree as an input and writes it to the output stream. Note that the fact that the target program is a stream-processing program does not necessarily imply that it is more efficient than the source program. In fact, if the translation \mathcal{A} were defined by $\mathcal{A}(M) = \text{unparse} \circ M \circ \text{parse}$, obviously there would be no improvement.

The target program being more efficient follows from the fact that the translation function \mathcal{A} preserves the structure of the source program, with only replacing tree constructions with stream outputs, and case analyses on trees with stream inputs and case analyses on input tokens. More precisely, by inspecting the proof of soundness of the translation, we can observe⁵:

- When a closure is allocated in the execution of M (so that the heap space is consumed), the corresponding closure is allocated in the corresponding reduction step of $\mathcal{A}(M)$, and vice versa.
- When a function is called in the execution of M (so that the stack space is consumed), the corresponding function is called in the corresponding reduction step of $\mathcal{A}(M)$, and vice versa.
- When a case analysis on an input tree is performed in the execution of M , a token is read from the input stream and a case analysis on the token is performed in the corresponding reduction step of $\mathcal{A}(M)$.
- When a tree is constructed in the execution of M , the corresponding sequence of tokens is written on the output stream in the corresponding reduction steps of $\mathcal{A}(M)$.

By the observation above, we can conclude:

- The memory space allocated by $\mathcal{A}(M)$ is less than the one allocated by $\text{unparse} \circ M \circ \text{parse}$ by the amount of the space for storing the intermediate trees output by *parse* and M (except for an implementation-dependent constant factor).
- The number of computation steps for running $\mathcal{A}(M)$ is the same as the one for running $\text{unparse} \circ M \circ \text{parse}$ (up to an implementation-dependent constant factor).

Thus, our translation is effective especially when the space for evaluating M is much smaller than the space for storing input and output trees.

5 Constructs for storing trees on memory

By adding primitives for constructing and destructing trees on memory, we can allow programmers to selectively buffer input/output trees at the cost of efficiency of target programs. Let us extend the syntax of the source and target languages as follows:

⁵ To completely formalize these observations, we need to define another operational semantics that makes both the heap and the stack explicit.

$$\begin{aligned}
M &::= \dots \mid \mathbf{mleaf} \ M \mid \mathbf{mnode} \ M_1 \ M_2 \\
&\quad \mid (\mathbf{mcase} \ M \ \mathbf{of} \ \mathbf{mleaf} \ x \Rightarrow M_1 \mid \mathbf{mnode} \ x_1 \ x_2 \Rightarrow M_2) \\
e &::= \dots \mid \mathbf{mleaf} \ e \mid \mathbf{mnode} \ e_1 \ e_2 \\
&\quad \mid (\mathbf{mcase} \ e \ \mathbf{of} \ \mathbf{mleaf} \ x \Rightarrow e_1 \mid \mathbf{mnode} \ x_1 \ x_2 \Rightarrow e_2).
\end{aligned}$$

Here, $\mathbf{mleaf} \ M$ and $\mathbf{mnode} \ M_1 \ M_2$ are constructors of trees on memory and $\mathbf{mcase} \dots$ is a destructor.

We also add type \mathbf{Tree}^ω , the type of trees stored on memory. The type system imposes no restriction on access order between variables of type \mathbf{Tree}^ω like type \mathbf{Int} (so \mathbf{Tree}^ω is put in the ordinary type environment, not in the ordered linear type environment). The translation algorithm \mathcal{A} simply translates a source program, preserving the structure:

$$\begin{aligned}
\mathcal{A}(\mathbf{mleaf} \ M) &= \mathbf{mleaf} \ \mathcal{A}(M) \\
\mathcal{A}(\mathbf{mnode} \ M_1 \ M_2) &= \mathbf{mnode} \ \mathcal{A}(M_1) \ \mathcal{A}(M_2) \\
&\dots
\end{aligned}$$

These extensions are summarized in Figure 13.

With these primitives, a function $strm_to_mem$, which copies a tree from the input stream to memory, and mem_to_strm , which writes a tree on memory to the output stream, can be defined as shown in Figure 14.

Using the functions above, one can write a program that selectively buffers only a part of the input tree, while the type system guarantees that the selective buffering is correctly performed. For example, the program in Figure 15, which swaps children of nodes whose depth is more than n , only buffers the nodes whose depth is more than n . In that example, we assume that we have booleans, comparison operators, and $\mathbf{if} \ M \ \mathbf{then} \ M_1 \ \mathbf{else} \ M_2$ in our source and target language syntax. The typing rule for $\mathbf{if} - \mathbf{then} - \mathbf{else}$ expressions is as follows:

$$\frac{\Gamma \mid \Delta_1 \vdash M : \mathbf{bool} \quad \Gamma \mid \Delta_2 \vdash M_1 : \tau \quad \Gamma \mid \Delta_2 \vdash M_2 : \tau}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{if} \ M \ \mathbf{then} \ M_1 \ \mathbf{else} \ M_2 : \tau} \quad (\text{T-IF})$$

The proof of Theorem 4.1 can be easily adapted for the extended language.

One drawback of the framework described in this section is that programmers need to write tree buffering operations. We report elsewhere an algorithm that automatically inserts $strm_to_mem$ and mem_to_strm (Suenaga et al. 2005).

6 Extension for dealing with XML

We extend our framework to deal with XML documents in this section. Note that we do not consider schemata for XML in this article (i.e., XML documents are considered as untyped labeled unranked trees.) The use of schemata in our framework is left for future work.

The difference between binary trees and (untyped) XML documents is that the latter ones (i) are unranked trees and (ii) contain end tags that mark the end of sequences in the stream format. The first point can be captured as the difference between the type $\mathbf{binarytree}$ and the type $\mathbf{xmltree}$ in the following ML-style type declarations.

M	$::=$	$i \mid \lambda x.M \mid x \mid M_1 M_2 \mid M_1 + M_2$ $\mathbf{leaf} M \mid \mathbf{node} M_1 M_2 \mid \mathbf{mleaf} M \mid \mathbf{mnode} M_1 M_2$ $(\mathbf{case} M \mathbf{of} \mathbf{leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$ $(\mathbf{mcase} M \mathbf{of} \mathbf{mleaf} x \Rightarrow M_1 \mid \mathbf{mnode} x_1 x_2 \Rightarrow M_2)$ $\mathbf{fix} f.M$
e	$::=$	$v \mid x \mid e_1 e_2 \mid e_1 + e_2$ $\mathbf{read} e \mid \mathbf{write} e$ $\mathbf{mleaf} M \mid \mathbf{mnode} M_1 M_2$ $(\mathbf{mcase} M \mathbf{of} \mathbf{mleaf} x \Rightarrow M_1 \mid \mathbf{mnode} x_1 x_2 \Rightarrow M_2)$ $(\mathbf{case} e \mathbf{of} \mathbf{leaf} \Rightarrow e_1 \mid \mathbf{node} \Rightarrow e_2)$ $\mathbf{fix} f.e$
		$\frac{\Gamma \mid \Delta \vdash M : \mathbf{Int}}{\Gamma \mid \Delta \vdash \mathbf{mleaf} M : \mathbf{Tree}^\omega} \quad (\text{T-MLEAF})$
		$\frac{\Gamma \mid \Delta_1 \vdash M_1 : \mathbf{Tree}^\omega \quad \Gamma \mid \Delta_2 \vdash M_2 : \mathbf{Tree}^\omega}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{mnode} M_1 M_2 : \mathbf{Tree}^\omega} \quad (\text{T-MNODE})$
		$\frac{\Gamma \mid \Delta_1 \vdash M : \mathbf{Tree}^\omega \quad \Gamma, x : \mathbf{Int} \mid \Delta_2 \vdash M_1 : \tau \quad \Gamma, x_1 : \mathbf{Tree}^\omega, x_2 : \mathbf{Tree}^\omega \mid \Delta_2 \vdash M_2 : \tau}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{mcase} M \mathbf{of} \mathbf{mleaf} x \Rightarrow M_1 \mid \mathbf{mnode} x_1 x_2 \Rightarrow M_2 : \tau} \quad (\text{T-MCASE})$
		$\begin{aligned} \mathcal{A}(\mathbf{mleaf} M) &= \mathbf{mleaf} \mathcal{A}(M) \\ \mathcal{A}(\mathbf{mnode} M_1 M_2) &= \mathbf{mnode} \mathcal{A}(M_1) \mathcal{A}(M_2) \\ \mathcal{A}(\mathbf{mcase} M \mathbf{of} \mathbf{mleaf} x \Rightarrow M_1 \mid \mathbf{mnode} x_1 x_2 \Rightarrow M_2) &= \\ &\mathbf{mcase} \mathcal{A}(M) \mathbf{of} \mathbf{mleaf} x \Rightarrow \mathcal{A}(M_1) \mid \mathbf{mnode} x_1 x_2 \Rightarrow \mathcal{A}(M_2) \end{aligned}$

Fig. 13. Extended languages, type system, and translation algorithm.

$\mathbf{fix} \mathit{strm_to_mem}.$ $\lambda t. \mathbf{case} t \mathbf{of} \mathbf{leaf} x \Rightarrow \mathbf{mleaf} x$ $\mid \mathbf{node} x_1 x_2 \Rightarrow \mathbf{mnode} (\mathit{strm_to_mem} x_1) (\mathit{strm_to_mem} x_2) \mathbf{fix} \mathit{mem_to_strm}.$ $\lambda t. \mathbf{mcase} t \mathbf{of} \mathbf{mleaf} x \Rightarrow \mathbf{leaf} x$ $\mid \mathbf{mnode} x_1 x_2 \Rightarrow \mathbf{node} (\mathit{mem_to_strm} x_1) (\mathit{mem_to_strm} x_2)$
--

 Fig. 14. Definition of *strm_to_mem* and *mem_to_strm*.

$\mathbf{let} \mathit{mswap} =$ $\mathbf{fix} f. \lambda t. \mathbf{mcase} t \mathbf{of} \mathbf{mleaf} x \Rightarrow \mathbf{leaf} x$ $\mid \mathbf{mnode} x_1 x_2 \Rightarrow \mathbf{node} (f x_2) (f x_1) \mathbf{in}$ $\mathbf{fix} \mathit{swap_deep}. \lambda n. \lambda t.$ $\quad \mathbf{if} n = 0 \mathbf{then} \mathit{mswap} (\mathit{strm_to_mem} t)$ $\quad \mathbf{else}$ $\quad \quad \mathbf{case} t \mathbf{of}$ $\quad \quad \quad \mathbf{leaf} x \Rightarrow \mathbf{leaf} x$ $\quad \quad \mid \mathbf{node} x_1 x_2 \Rightarrow \mathbf{node} (\mathit{swap_deep} (n - 1) x_1) (\mathit{swap_deep} (n - 1) x_2)$

 Fig. 15. A program that swaps children of nodes whose depth is more than n .

Remove **leaf**, **node**, and **case** expressions from the definition of M and add the following constructs.

$$M ::= \dots$$

	leaf M node $M_l M_a M$
	nil cons (M_1, M_2)
	caseElem M of leaf $x \Rightarrow M_1$ node ($l, attr, tl$) $\Rightarrow M_2$
	caseElem M of nil $\Rightarrow M_1$ cons (x, xs) $\Rightarrow M_2$

Fig. 16. Change to the syntax of the source language to deal with XML documents.

```
type binarytree = Leaf of int | Node of tree * tree;
```

```
type xmltree = Leaf of pcdat
  | Node of label * attribute * treelist
and treelist = Nil | Cons of xmltree * treelist;
```

While the type `binarytree` represents binary trees, `xmltree` represents unranked trees. On the basis of the difference between these types, we can replace the **case**-construct of the source language with the following two pattern-matching primitives.

```
caseElem ... of Leaf x -> ...
  | Node(l, attr, tl) -> ...
```

```
caseSeq ... of Nil -> ...
  | Cons(x, xs) -> ...
```

The above changes to the syntax are summarized in Figure 16.

We have not incorporated regular expression patterns (Hosoya & Pierce 2003) into our framework because type-checking regular pattern matching requires document type information. This makes our current framework unsatisfactory as a high-level programming language for XML processing. We plan to address this by designing a translation algorithm from a language with regular pattern matching into our language.

Typing rules can also be naturally extended. For example, the typing rules for the **leaf** and **node** constructs are:

$$\frac{\Gamma \mid \Delta \vdash M : \text{pcdata}}{\Gamma \mid \Delta \vdash \text{leaf } M : \text{xmltree}^+}$$

$$\frac{\Gamma \mid \Delta_1 \vdash M_l : \text{label} \quad \Gamma \mid \Delta_2 \vdash M_a : \text{attribute} \quad \Gamma \mid \Delta_3 \vdash M_3 : \text{treelist}^+}{\Gamma \mid \Delta_1, \Delta_2, \Delta_3 \vdash \text{node}(M_l, M_a, M)}$$

and the typing rule for the **caseSeq** construct is:

$$\frac{\Gamma \mid \Delta_1 \vdash tl : \text{treelist}^- \quad \Gamma \mid \Delta_2 \vdash M_1 : \tau \quad \Gamma \mid x : \text{xmltree}^-, xl : \text{treelist}^-, \Delta_2 \vdash M_2 : \tau}{\Gamma \mid \Delta_1, \Delta_2 \vdash \text{caseSeq } tl \text{ of nil } \Rightarrow M_1 \mid \text{cons}(x, xl) \Rightarrow M_2 : \tau}$$

$$\begin{array}{c}
 \frac{\Gamma \mid \Delta \vdash M : \mathbf{pcdata}}{\Gamma \mid \Delta \vdash \mathbf{leaf} M : \mathbf{xmltree}^+} \quad (\text{T-LEAF}) \\
 \frac{\Gamma \mid \Delta_1 \vdash M_l : \mathbf{label} \quad \Gamma \mid \Delta_2 \vdash M_a : \mathbf{attribute} \quad \Gamma \mid \Delta_3 \vdash M_3 : \mathbf{treelist}^+}{\Gamma \mid \Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{node}(M_l, M_a, M)} \quad (\text{T-NODE}) \\
 \frac{\Gamma \mid \Delta_1 \vdash M : \mathbf{xmltree} \quad \Gamma, x : \mathbf{pcdata} \mid \Delta_2 \vdash M_2 : \tau \quad \Gamma, l : \mathbf{label}, \mathbf{attr} : \mathbf{attribute} \mid tl : \mathbf{treelist}, \Delta_2 \vdash M_1 : \tau}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{caseElem} M \text{ of } \mathbf{node}(l, \mathbf{attr}, tl) \Rightarrow M_1 \mid \mathbf{leaf}(x) \Rightarrow M_2 : \tau} \quad (\text{T-CASEELEM}) \\
 \frac{\Gamma \mid \Delta_1 \vdash tl : \mathbf{treelist} \quad \Gamma \mid \Delta_2 \vdash M_1 : \tau \quad \Gamma \mid x : \mathbf{xmltree}, xl : \mathbf{treelist}, \Delta_2 \vdash M_2 : \tau}{\Gamma \mid \Delta_1, \Delta_2 \vdash \mathbf{caseSeq} tl \text{ of } \mathbf{nil} \Rightarrow M_1 \mid \mathbf{cons}(x, xl) \Rightarrow M_2 : \tau} \quad (\text{T-CASESEQ})
 \end{array}$$

Fig. 17. Typing rules for XML extension.

$$\begin{array}{l}
 \mathcal{A}(\mathbf{caseElem} M \text{ of } \mathbf{node}(l, \mathbf{attr}, tl) \Rightarrow M_1 \mid \mathbf{leaf}(x) \Rightarrow M_2) = \\
 \quad \mathbf{case} \mathbf{readtoken}() \text{ of} \\
 \quad \quad \mathbf{start_tag}(l, \mathbf{attr}) \Rightarrow \mathbf{push}(l); [()/tl]\mathcal{A}(M_1) \\
 \quad \quad \mid \mathbf{pcdata}(x) \Rightarrow \mathcal{A}(M_2) \\
 \quad \quad \mid - \Rightarrow \mathbf{raise} \mathit{IllFormedException} (* \mathbf{end_tag} \text{ or } \mathbf{end \text{ of } file} *) \\
 \\
 \mathcal{A}(\mathbf{caseSeq} M \text{ of } \mathbf{nil} \Rightarrow M_1 \mid \mathbf{cons}(x, xs) \Rightarrow M_2) = \\
 \quad \mathbf{case} \mathbf{peektoken}() \text{ of} \\
 \quad \quad \mathbf{end_tag}(l) \Rightarrow \mathbf{let} l' = \mathbf{pop}() \text{ in} \\
 \quad \quad \quad \mathbf{if} l = l' \text{ then } \mathbf{readtoken}(); \mathcal{A}(M_1) \\
 \quad \quad \quad \mathbf{else} \mathbf{raise} \mathit{IllFormedException} \\
 \quad \quad \mid - \Rightarrow [()/x, ()/xs]\mathcal{A}(M_2)
 \end{array}$$

Fig. 18. Definition of \mathcal{A} for XML-processing constructs.

The restriction on the access order is expressed by $x : \mathbf{xmltree}^-$, $xl : \mathbf{treelist}^-$, Δ_2 as in T-NODE. Figure 17 shows the typing rules for newly added constructs.

Figure 18 shows the definition of \mathcal{A} for XML-processing constructs. $\mathbf{readtoken}()$ reads an element from the input stream where the element is $\mathbf{start_tag}(l, \mathbf{attr})$, $\mathbf{end_tag}(l)$, $\mathbf{pcdata}(x)$, or $\mathbf{end_of_file}$. The primitive $\mathbf{peektoken}()$ peeks the top of the input stream without deleting the element. Values of type $\mathbf{attribute}$ are stored on memory-like values of type \mathbf{pcdata} , because attributes can occur in an arbitrary order (e.g., $\langle a \ b='foo' \ c='baa' \rangle$ and $\langle a \ c='baa' \ b='foo' \rangle$ have the same meaning). Thus, type $\mathbf{attribute}$ is put in the ordinary type environment, not in the ordered linear type environment, in the rule T-CASEELEM.

In target programs, a stack is used to check well-formedness of input documents. When a start tag is read, the tag is pushed on the stack. When an end tag is read, it is compared with the tag stored on the top of the stack.

```

...
let rec process_row = fun /t/ ->
  match t with
  /id/::/tl/ -> skip_tree id;
  match tl with /firstname/::/tl/ ->
    let firstname = get_firstname firstname in
    match tl with /lastname/::/tl/ ->
      let lastname = get_lastname lastname in
      match tl with /street/::/tl/ -> skip_tree street;
      match tl with /city/::/tl/ -> skip_tree city;
      match tl with /state/::/tl/ -> skip_tree state;
      match tl with /zip/::/tl/ -> skip_tree zip;
      match tl with [] ->
        <row>@{[
          <firstname>@{[PC firstname];
          <lastname>@{[PC lastname]
        ]
...

```

Fig. 19. An X-P program that extracts first-name and last-name fields from an input document.

Note that the pattern **nil** in the source language is translated to the pattern for closing tags.

Remark 6.1

The data representation and pattern-matching constructs for XML are essentially the same as those used in *general-purpose* programming languages that provide no primitive support for XML processing (except that there is an additional restriction on access order). Programs written with such data representation and pattern matching are certainly more awkward than those written in XML-centric programming languages (Benzaken *et al.* 2003; Hosoya & Pierce 2003), which provide a number of useful features such as regular expression pattern matching. (For example, compare the program in Figure 19 with a program using regular expression pattern matching.) We, however, believe that writing a program in our language is still much easier than writing it using a low-level stream-processing language (like the target language of our framework).

Programming can be made much easier by introduction of additional layers of transformations. For example, we have already implemented a translator that takes an ordinary XML-processing program as an input, and automatically inserts buffering primitives, so that the output is well-typed in our type system (Suenaga *et al.* 2005; Sato 2007). The result of the transformation is then given as an input of the transformation described in this article and transformed into a stream-processing program. Thus, the source language discussed in this section actually serves as an *intermediate* language of the entire transformation framework. We also plan to add an additional layer of transformation for compiling regular expression pattern matching.

<pre> Input: <table> <row> <id>0001</id> <firstname>Al</firstname> <lastname>Aranow</lastname> <street>1 Any St.</street> <city>Anytown</city> <state>AL</state> <zip>22000</zip> </row> ... </table> </pre>	<pre> Output: <table> <row> <firstname>Al</firstname> <lastname>Aranow</lastname> </row> ... </table> </pre>
--	--

Fig. 20. A sample input and the intended output for programs in Figures 19 and 21.

7 Implementation and experiments

We have implemented a translator X-P based on the type system and the translation algorithm presented. X-P can deal with both tree buffering in Section 5 and XML processing in Section 6.

Figure 19 shows a part of a sample X-P program. The program extracts `firstname` and `lastname` fields from input data, by filtering out `id`, `street`, `city`, `state`, and `zip` fields. An example of input and output is shown in Figure 20.

The current implementation of X-P translator adopts ML-like syntax. Here is an overview of the syntax of X-P language.

- a tree expressed with an XML document

```

<row sex='male'><firstname>Kohei</firstname>
      <lastname>Suenaga</lastname></row>

```

is written like

```

<row>@{sex = 'male'}[
  <firstname>@{}[PC 'Kohei'];
  <lastname>@{}[PC 'Suenaga']
]

```

- **case M of nil** $\Rightarrow M_1 \mid \mathbf{cons}(x, y) \Rightarrow M_2$ is written like `match M with [] -> M1 | /x/::/y/ -> M2`. We mark bound variables of type **Tree**⁻ with surrounding slashes. These marks are used in the type-checking algorithm described in Section 3.4.
- In the current implementation, we allow only two patterns for matching lists: `[]` and `x::y`.

Remark 7.1

The program in Figure 19 is unnecessarily involved because of the limited pattern-matching constructs and the need for automatic insertion of the `skip_tree` instructions. By using the usual ML-like pattern-matching constructs and the translator

```

let rec _process_row29 = (fun _t30 ->
match _t30;peek_token () with
StartTag(_, _) | PCDATA _ ->
  skip_tree (); match ();peek_token () with
  StartTag(_, _) | PCDATA _ ->
    let _firstname35 = _get_firstname7 () in
    match ();peek_token () with
    StartTag(_, _) | PCDATA _ ->
      let _lastname38 = _get_lastname8 () in
      match ();peek_token () with
      StartTag(_, _) | PCDATA _ ->
        skip_tree (); match ();peek_token () with
        StartTag(_, _) | PCDATA _ ->
          skip_tree (); match ();peek_token () with
          StartTag(_, _) | PCDATA _ ->
            skip_tree (); match ();peek_token () with
            StartTag(_, _) | PCDATA _ ->
              skip_tree (); match ();peek_token () with
              EndTag __65 ->
                read_token ();
                assert (pop ()) = PCDATA(__65);
                write (StartTag("row", []));
                write (StartTag("first", []));
                write (PCDATA(_firstname35));
                write (EndTag("first"));
                write (StartTag("last", []));
                write (PCDATA(_lastname38));
                write (EndTag("last"));
                write (EndTag("row"))

```

Fig. 21. The result of translation of the program in Figure 19.

for automatic insertion of buffering (recall Remark 6.1), the program in Figure 19 can be actually simplified to:

```

let rec process_row = fun /t/ ->
match _::/firstname/::t1 with
  let firstname = get_firstname firstname in
  match t1 with /lastname/::_ ->
    let lastname = get_lastname lastname in
    <row>@{}[
      <firstname>@{}[PC firstname];
      <lastname>@{}[PC lastname]
    ]

```

Figure 21 shows the result of translation. Based on the framework so far, the system first type checks the source program and then transforms it into a stream-processing program written in Objective Caml (OCaml).

We measured the maximum heap size during the execution for four micro-benchmarks in XSLTMark (DataPower Technology 2001) (dbonerow, dbtail,


```

<site>
...
<people>
  <person id='person0'><name>Kohei Suenaga</name> ... </person>
  ...
</people>
...
<closed_auctions>
  <closed_auction> ... <buyer person='person0' /> ... </closed_auction>
  ...
</closed_auctions>
</site>

```

Fig. 22. Sample input of benchmarks Q_1 and Q_8 .

avts, and stringsort), one in our own micro-benchmark (evensort), and two in XMark (Schmidt *et al.* 2001) (Q_1 and Q_8). We also measured the execution time of benchmarks Q_1 and Q_8 .

Benchmarks *dbonerow*, *dbtail*, *avts*, *stringsort*, and *evensort* use XML documents of the form shown in Figure 20 as input. The following is the description of those benchmarks:

- *dbonerow*: Scans the whole input and extracts the row record whose *id* field is equal to 0432 in the form of an HTML document.
- *dbtail*: Scans the whole input and extracts *firstname* field and *lastname* field of each row record (A program presented in part in Figure 19).
- *avts*: Scans the whole input and converts each row record to an address tag whose attributes have the information the row record has.
- *stringsort*: Sorts all the row records in input documents by values of *firstname* field.
- *evensort*: Sorts row records whose *id* field is an even number by values of *firstname* field.

Benchmarks Q_1 and Q_8 use auction data as input of the form shown in Figure 22. The followings are the description of each benchmark:

- Q_1 : Extracts the name of the person whose ID is *person0*.
- Q_8 : Extracts the names of persons and the number of items bought by each person from the join of *people* and *closed_auctions*.

For each benchmark, we prepared (1) an OCaml program generated by our implementation from a manually written tree-processing program in our language and (2) a manually written OCaml tree-processing program. All the OCaml programs were compiled by *ocamlpt* 3.08.4. The experiments were performed on Intel(R) CPU 1.06 GHz with 1 GB memory. We executed each of *dbonerow*, *dbtail*, *avts*, *stringsort*, and *evensort* with 12 input documents, each of which contains 10, 50, 100, 250, 500, 750, 1,000, 2,000, 4,000, 6,000, 8,000, and 10,000 records. Benchmarks Q_1 and Q_8 were executed with 5, 10, 50, and 100 MB documents.

Figure 23 shows the memory consumption of *dbonerow*, *dbtail*, *avts*, *stringsort*, and *evensort*. In these graphs, the horizontal axes show the number of row records.

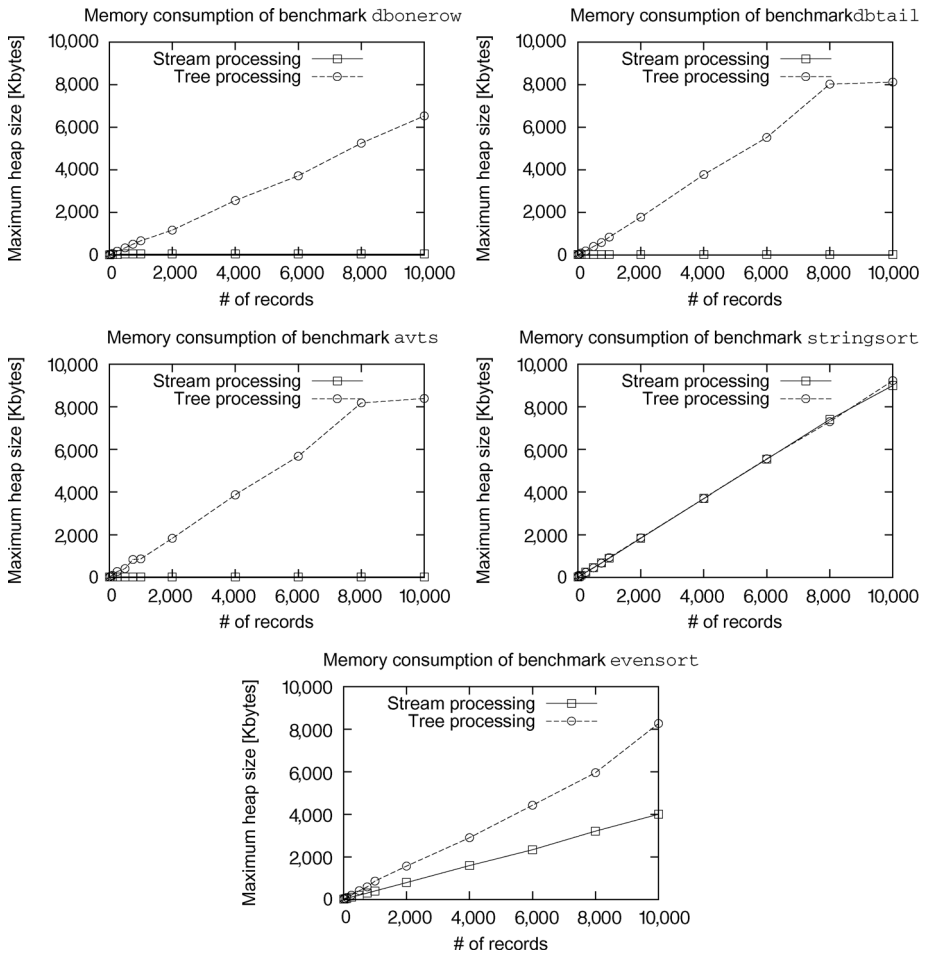


Fig. 23. Memory consumption of benchmarks dbonerow, dbtail, avts, stringsort, and evensort.




Fig. 24. Memory consumption of benchmarks Q_1 and Q_8 .

Figure 24 shows the memory consumption of Q_1 and Q_8 . In these graphs, the horizontal axes show the size of input documents. In both figures, the vertical axes are the maximum heap size during the execution (measured in kilobytes.)

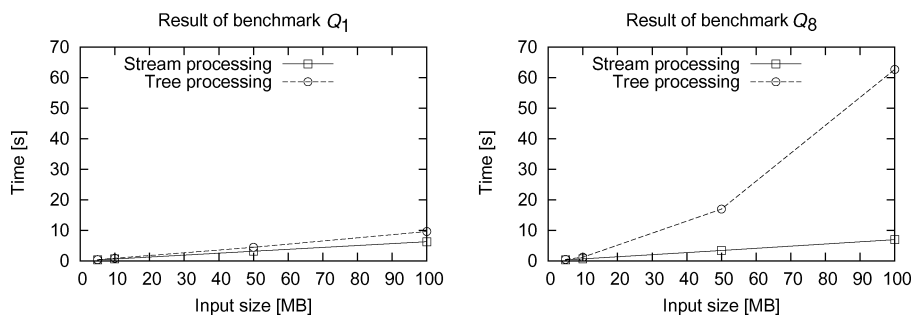


Fig. 25. Time consumption of benchmarks Q_1 and Q_8 .

In each benchmark, memory consumption of tree-processing programs is proportional to the number of the size of input. In `dbonerow`, `dbtail`, `avts`, and Q_1 benchmarks, in which all the tree variables have ordered linear type, the heap memory consumption of stream-processing programs is constant w.r.t. the number of records. (On the other hand, the stack memory consumption is linear in the depth of the input document. The input documents used in our experiments, however, have the constant depth, so that the stack memory consumption is also constant.) In `stringsort` benchmark, which requires all the records to be stored on memory in processing, so that all the tree variables in the program have buffered tree type, the program consumes as much memory as the tree-processing version does. Benchmark `evensort` requires records with even `id` field values to be stored on memory. As expected, the memory consumption of the stream-processing version is almost a half of that of the tree-processing version. In benchmark Q_8 , which performs a join operation and hence requires a part of input to be on heap, the memory consumption of the stream-processing program is about 3% of that of tree-processing one. For example, for 100-MB input, memory consumption is 5,576 KB in the stream-processing program, while 179,232 KB in the tree-processing program. Note that we cannot simply compare this result with that in Koch and Scherzinger (2003), because our tree-processing program is optimized manually to produce a memory-efficient stream-processing program,⁶ while that Koch and Scherzinger (2003) automatically derives stream-processing programs from XMark queries.

Figure 25 shows the execution time of benchmark Q_1 and Q_8 . In both benchmarks, the stream-processing program is faster than the tree-processing program.

8 Related work

Attribute grammars are widely used for specifying XML processing. Nakano and Nishimura (2001) and Nakano (2004) proposed a method for translating tree-processing attribute grammars into stream-processing ones by descriptive composition (Ganzinger & Giegerich 1984) with parsing and unparsing grammars.

⁶ Specifically, we took advantage of the following facts in translating Q_8 into our language: (1) `people` comes before `closed.auctions` (2) in `people` tree, only the name and the ID of each person are needed, and thus should be buffered.

```

N → node N1 N2
    N1.inh = f1 N.inh; N2.inh = f2 N.inh N1.syn N1.inh
    N.syn = f3 N.inh N1.syn N1.inh N2.syn N2.inh
N → leaf i
    N.syn = f4 N.inh i

fix f.λinh.λt.case t of
  leaf x ⇒ f4 inh x
  node x1 x2 ⇒ let N1.inh = f1 inh in
                  let N1.syn = f N1.inh x1 in
                  let N2.inh = f2 N.inh N1.syn N1.inh in
                  let N2.syn = f N2.inh x2 in
                    f3 N.inh N1.syn N1.inh N2.syn N2.inh

```

Fig. 26. L-attributed grammar over binary trees and corresponding programs.

Quasi-SSUR condition in Nakano and Nishimura (2001) and *single use requirement* in Nakano (2004), which force attributes of non-terminal symbols to be used at most once, seems to correspond to our linearity restriction on variables of tree types, but there seems to be no restriction that corresponds to our order restriction. As a result, their method can deal with programs (written as attribute grammars) that violate the order restriction of our type system, although in that case, generated stream-processing programs store parts of trees in memory, so that the translation may not improve the efficiency. On the other hand, an advantage of our method is that programs are easier to read and write since one can write programs as ordinary functional programs except for the restriction imposed by the type system, rather than as attribute grammars. Another advantage of our method is that we can deal with source programs that involve side effects (e.g., programs that print the value of every leaf) while that seems difficult in their method based on attribute grammars (since the order is important for side effects).

Nakano and Mu (2006) proposed a method of deriving a stream-processing program from a tree-processing program written as a macro forest transducer (MFT). They fuse parsing and unparsing MFTs with a given tree-processing program, and derive an equivalent stream processor. A main difference between their method and ours seems to be that stream processors derived in their method tend to buffer outputs, while those in our method (used with the automatic insertion algorithm of buffering primitives, presented in Suenaga *et al.* (2005), tend to buffer inputs. For example, their method translates a program that reverses children of every node (a program presented in Figure 2) into a stream processor that buffers output, while our framework buffers input and performs computation on the buffered input. Another difference is that their stream processors may perform redundant computation due to their evaluation strategy. For example, consider a program that transform a tree of the form **node** M (**leaf** i) into $f_1 M$ if i is an even number and into $f_2 M$ if i is an odd number. To avoid buffering the input tree, their stream processor evaluates *both* $f_1 M$ and $f_2 M$, while ours buffers M and evaluates *either* $f_1 M$ or $f_2 M$.

The class of well-typed programs in our language seems to be closely related to the class of L-attributed grammars (Aho *et al.* 1986). The upper half of Figure 26

shows the general form of L-attributed grammar over binary trees. In the grammar, *syn* represents a synthesized attribute and *inh* represents an inherited attribute. Also, f_1, \dots, f_4 are functions that calculate values of attributes. For example, the inherited attribute of N_2 is calculated from the inherited attributes of N and N_1 and the synthesized attribute of N_1 with the function f_2 . (Note that $N_2.inh$ depends only on $N.inh, N_1.syn,$ and $N_1.inh$ from the definition of L-attributed grammar.) That L-attributed grammar can be expressed as a program is shown in the lower half of Figure 26. If output trees are not used in attributes, the program is well-typed. Conversely, any program that is well-typed in our language seems to be definable as an L-attribute grammar. The corresponding attribute grammar may, however, be awkward, since one has to encode control information into attributes.

Koch and Scherzinger (2003) proposed XML Stream Attribute Grammars (XSAGs), which are extended regular tree grammars (Neven 2005) with several primitives for manipulating the output stream. By limiting the grammars to L-attributed ones, they ensure that any XSAGs are evaluated as stream-processing programs that read the input stream from head to tail with bounded memory consumption, and thus programmers need not care about the input stream. However, their framework does not ensure that output is well-formed. We guess well-typed XML processing programs in our framework are at least as expressive as XSAGs, which output well-formed documents by the following observation. According to Koch and Scherzinger (2003), expressive power of XSAGs is equivalent to that of XML-DPDTs, deterministic pushdown transducers with an auxiliary stack for checking well-formedness of input documents. XML-DPDTs are essentially XML-processing programs that (1) read well-formed input documents in left-to-right, depth-first preorder and (2) generate an output in document order, which are also the properties of well-typed programs of our framework. Thus, the expressive power of XSAGs is equivalent to that of our well-typed programs.

There are other studies on translation of tree-processing programs into stream-processing programs. Some of them (Green *et al.* 2001; Avila-Campillo *et al.* 2002; Olteanu *et al.* 2002; Gupta & Suciuc 2003; Bar-Yossef *et al.* 2005) deal with XPath expressions (Berglund *et al.* 2003; Scardina & Fernandez 2003) and others (Ludäscher *et al.* 2002) deal with XQuery (Boag *et al.* 2003; Chamberlin *et al.* 2003). These translations are more aggressive than ours in the sense that the structure of source programs is changed so that input trees can be processed in one pass. On the other hand, their target languages (XPath and XQuery languages) are restricted in the sense that they do not contain functions and side effects.⁷

There are many studies on program transformation (Wadler 1988; Meijer *et al.* 1991) for eliminating intermediate data structures of functional programs, known as deforestation or fusion. Although the goal of our translation is also to remove intermediate data structures from $unparse \circ f \circ parse$, the previous methods are not directly applicable since they do not guarantee that transformed programs access inputs in a stream-processing manner. In fact, *swap* in Figure 2, which

⁷ Although XQuery can deal with user-defined functions, the fragment dealt with in Ludäscher *et al.* (2002) excludes user-defined functions.

violates the access order, can be expressed as a treeless program (Wadler 1988) or a catamorphism (Meijer *et al.* 1991), but the result of deforestation is not an expected stream-processing program.

Actually, there are many similarities between the restriction of treeless program (Wadler 1988) and that of our type system. In treeless programs, (1) variables have to occur only once and (2) only variables can be passed to functions. (1) corresponds to the linearity restriction of our type system; (2) is the restriction for prohibiting trees generated in programs to be passed to functions, which corresponds to the restriction that functions cannot take values of type **Tree**⁺ in our type system. The main differences are:

- Our type system additionally imposes a restriction on the access order. This is required to guarantee that translated programs read input streams sequentially.
- We restrict programs with a type system, while the restriction on treeless programs is syntactic. Our type-based approach enables us to deal with higher-order functions. The type-based approach is also useful for automatic inference of selective buffering of trees, as discussed in Section 9.

The type system used in this article is based on the ordered linear logic proposed by Polakow (2001). He proposed a logic programming language Olli, logical framework OLF, and ordered λ -calculus based on the logic. There are many similarities between our typing rules and his derivation rules for the ordered linear logic. For example, our type judgment $\Gamma \mid \Delta \vdash M : \tau$ corresponds to the judgment $\Gamma ; ; \Delta \vdash A$ of ordered linear logic. The rule T-ABS1 corresponds to a combination of the rules for an ordered linear implication and the modality (!). However, we cannot use ordered linear logic directly since it would make our type system unsound.

Petersen *et al.* (2003) used ordered linear types to guarantee correctness of memory allocation and data layout. While they used an ordered linear type environment to express a spatial order, we used it to express a temporal order.

As stated in Sections 3 and 5, one can write tree-processing programs that selectively skip and/or buffer trees by using *skip_tree*, *copy_tree*, *strm_to_mem*, and *mem_to_strm*. However, inserting those functions manually is sometimes tedious. To solve this problem, we have developed a type-directed source-to-source translation for automatically inserting these functions (Suenaga *et al.* 2005).

9 Conclusion

We have proposed a type system on the basis of ordered linear types to enable translation of tree-processing programs into stream-processing programs, and proved the correctness of the translation.

Since our translation algorithm preserves the structure of source programs, the translation works in the presence of side effects other than stream inputs/outputs. Our framework can be easily extended to deal with multiple input trees as well.

In addition to application to XML processing, our translation framework may also be useful for optimization of distributed programs that process and communicate complex data structures. Serialization/unserialization of data corresponds to

unparsing/parsing in Figure 1, so that our translation framework can be used for eliminating intermediate data structures and processing serialized data directly.

Acknowledgments

We thank Keisuke Nakano for his comment on this research. We also thank the anonymous referees for their fruitful comments.

Appendix A

Proof of Theorem 4.2

We prepare the following lemma to prove the theorem.

Lemma A.1 (type substitution)

If $\Gamma \cup \{x : \tau'\} \mid \Delta \vdash M : \tau$ and $\Gamma \mid \emptyset \vdash N : \tau'$ hold, $\Gamma \mid \Delta \vdash [N/x]M : \tau$.

Proof

Straightforward. \square

Proof of Theorem 4.2. The first condition can be easily proved by induction on the derivation tree of $\Gamma \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$. Here, we only show the proof of the second condition.

From the assumption $(M, \delta) \rightarrow (M', \delta')$, there exist E_s and I that satisfy $M = E_s[I]$. We use structural induction on E_s .

- Case $E_s = []$.

Case analysis on I .

- Case $I = (\lambda x.N) U$.

First, suppose $U = i$ or $U = \lambda y.N'$ for some i or y and N' . Then, $\delta = \emptyset$ and $\Gamma \mid \emptyset \vdash U : \tau'$ and $\Gamma \cup \{x : \tau'\} \mid \emptyset \vdash N : \tau$ follow from the assumption $\Gamma \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$ and inversion of T-APP and T-ABS2. $M' = [U/x]N$ and $\delta' = \delta$ (and thus $\delta' = \emptyset$) follow from Es2-APP. Thus, $\Gamma \mid \langle\langle\delta'\rangle\rangle \vdash M' : \tau$ follows from Lemma A.1 as required.

Next, suppose $U = y$ for some y . Then, $M' = [y/x]N$ and $\delta' = \delta$ follow from Es2-APP. $\langle\langle\delta\rangle\rangle = y : \mathbf{Tree}^-$ and $\Gamma \mid x : \mathbf{Tree}^- \vdash N : \tau$ follow from the assumption $\Gamma \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$ and inversion of T-APP and T-ABS1. Thus, as can be easily seen, $\Gamma \mid y : \mathbf{Tree}^- \vdash [y/x]N : \tau$. Thus, $\Gamma \mid \langle\langle\delta'\rangle\rangle \vdash M' : \tau$ follows as required.

- Case $I = (\mathbf{case} \ y \ \mathbf{of} \ \mathbf{leaf} \ x \Rightarrow M_1 \mid \mathbf{node} \ x_1 \ x_2 \Rightarrow M_2)$ with $\delta = (y \mapsto \mathbf{leaf} \ i, \delta'')$.

$\Gamma \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$ must have been derived by using T-CASE. Thus, we have $\Gamma \cup \{x : \mathbf{Int}\} \mid \langle\langle\delta''\rangle\rangle \vdash M_1 : \tau$. Because $\Gamma \mid \emptyset \vdash i : \mathbf{Int}$, we have $\Gamma \mid \langle\langle\delta''\rangle\rangle \vdash [i/x]M_1 : \tau$ from Lemma A.1. Because $M' = [i/x]M_1$ and $\delta' = \delta''$ follow from Es2-CASE1, we have $\Gamma \mid \langle\langle\delta'\rangle\rangle \vdash M' : \tau$ as required.

- Case $I = (\mathbf{case} \ y \ \mathbf{of} \ \mathbf{leaf} \ x \Rightarrow M_1 \mid \mathbf{node} \ x_1 \ x_2 \Rightarrow M_2)$ with $\delta = (y \mapsto (\mathbf{node} \ V_1 \ V_2), \delta'')$.

$\Gamma \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$ must have been derived by using T-CASE. Thus, we have $\Gamma \mid x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-, \langle\langle\delta''\rangle\rangle \vdash M_2 : \tau$. Because $M' = M_2$ and $\delta' = x_1 \mapsto V_1, x_2 \mapsto V_2, \delta''$ follow from Es2-CASE2, we have $\Gamma \mid \langle\langle\delta'\rangle\rangle \vdash M' : \tau$ as required.

- Case $E_s = (\mathbf{case} E'_s \text{ of leaf } x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$.

$\Gamma \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$ must have been derived by using T-CASE. Thus, we have $\Gamma \mid \langle\langle\delta_1\rangle\rangle \vdash E'_s[I] : \mathbf{Tree}^-$ and $\Gamma \cup \{x : \mathbf{Int}\} \mid \langle\langle\delta_2\rangle\rangle \vdash M_1 : \tau$ and $\Gamma \mid x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-, \langle\langle\delta_2\rangle\rangle \vdash M_2 : \tau$ and $\delta = \delta_1, \delta_2$ for some δ_1 and δ_2 . By the induction hypothesis, there exist δ'_1 and M'' that satisfy $\Gamma \mid \langle\langle\delta'_1\rangle\rangle \vdash M'' : \mathbf{Tree}^-$ and $(E'_s[I], \delta_1) \rightarrow (M'', \delta'_1)$. Because $M' = (\mathbf{case} M'' \text{ of leaf } x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$ and $\delta' = \delta'_1, \delta_2$, we have $\Gamma \mid \langle\langle\delta'\rangle\rangle \vdash M' : \tau$ from T-CASE as required.

Appendix B

Proof of Corollary 4.3

The following lemma guarantees that the reducibility in the original source language semantics and that in the revised semantics are equivalent for well-typed programs.

Lemma B.1

If $\emptyset \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$ and if $\tau \in \{\mathbf{Int}, \mathbf{Tree}^+\}$, then $(M, \delta) \rightarrow^* (M', \delta')$ if and only if $\{\!\{\delta\}\!\}M \rightarrow \{\!\{\delta'\}\!\}M'$, where $\{\!\{\cdot\}\!\}$ is the function from ordered environments to substitutions defined by

$$\begin{aligned} \{\!\{\emptyset\}\!\} &= (\text{identity function}) \\ \{\!\{x \mapsto V, \delta\}\!\} &= [V/x] \circ \{\!\{\delta\}\!\}. \end{aligned}$$

Proof

(\implies) is obvious. We only show (\impliedby) by induction on the length of $\{\!\{\delta\}\!\}M \rightarrow^* \{\!\{\delta'\}\!\}M'$. The case where $M = M'$ and $\delta = \delta'$ is obvious. Suppose that $\{\!\{\delta\}\!\}M \rightarrow \{\!\{\delta''\}\!\}M'' \rightarrow^* \{\!\{\delta'\}\!\}M'$. From Theorem 4.2, we have $(M, \delta) \rightarrow (M''', \delta''')$ and $\emptyset \mid \langle\langle\delta'''\rangle\rangle \vdash M''' : \tau$ for some M''' and δ''' . From the (\implies) direction of this lemma, and that the first semantics is deterministic, we have $M''' = M''$ and $\delta''' = \delta''$. Then, we have $(M'', \delta'') \rightarrow^* (M', \delta')$ from the induction hypothesis. \square

Lemma B.2

If $\emptyset \mid \Delta \vdash W : \tau$, then $\Delta = \emptyset$.

Proof

Case analysis on the last rule that derives $\emptyset \mid \Delta \vdash W : \tau$. \square

Proof of Corollary 4.3. From $\emptyset \mid \emptyset \vdash M : \mathbf{Tree}^- \rightarrow \tau$, we have $\emptyset \mid x : \mathbf{Tree}^- \vdash M x : \tau$. Thus, $(M x, x \mapsto V) \rightarrow^* (W, \delta)$ if and only if $M V \rightarrow^* \{\!\{\delta\}\!\}W$. (Note that $\{\!\{x \mapsto V\}\!\}(M x) = M V$.) It suffices to show that $\delta = \emptyset$. From Theorem 4.2, we have $\emptyset \mid \langle\langle\delta\rangle\rangle \vdash W : \tau$. Thus, we have $\delta = \emptyset$ from Lemma B.2.

Appendix C

Proof of Lemma 4.5

For the proof of Lemma 4.5, we prepare the following lemma.

Lemma C.1

If $x \notin \gamma$, then $\mathcal{A}_\gamma([M_1/x]M_2) = [\mathcal{A}_\gamma(M_1)/x]\mathcal{A}_\gamma(M_2)$.

Proof

Induction on the structure of M_2 . \square

Proof of Lemma 4.5. We prove $M \sim_{\mathbf{FV}(M)} (e, S_\delta)$. We hereafter write γ for $\mathbf{FV}(M)$ and S for S_δ .

First, suppose that M is not reducible. Then, $M = U$ or $M = V$.

- Case $M = U$.

Let $e = \mathcal{A}_\gamma(M)$ and $S = \emptyset$. Then, $M \sim_\gamma (e, S)$ follows from C-VALUE. $(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$ is obvious.

- Case $M = V$.

Let $e = ()$ and $S = \llbracket V \rrbracket$. Then, $M \sim_\gamma (e, S)$ follows from C-TREE.

$(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$ follows from the structural induction on V below:

— Case $V = \mathbf{leaf} \ i$.

In this case, $\mathcal{A}_\gamma(M) = (\mathbf{write}(\mathbf{leaf}); \mathbf{write}(i))$. Thus,

$(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$ holds.

— Case $V = \mathbf{node} \ V_1 \ V_2$.

In this case, $\mathcal{A}_\gamma(M) = (\mathbf{write}(\mathbf{node}); \mathcal{A}_\gamma(V_1); \mathcal{A}_\gamma(V_2))$ and

$S = \mathbf{node}; \llbracket V_1 \rrbracket; \llbracket V_2 \rrbracket$. Because $\emptyset \mid \emptyset \vdash V_1 : \mathbf{Tree}^+$ and $\emptyset \mid \emptyset \vdash V_2 : \mathbf{Tree}^+$, we have

$(\mathcal{A}_\gamma(V_1), \emptyset, \emptyset) \rightarrow^* ((), \emptyset, \llbracket V_1 \rrbracket)$ and $(\mathcal{A}_\gamma(V_2), \emptyset, \emptyset) \rightarrow^* ((), \emptyset, \llbracket V_2 \rrbracket)$ from the induction hypothesis. Thus, $(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$. (Note that $\delta = \emptyset$ because $\emptyset \mid \emptyset \vdash V : \mathbf{Tree}^+$.)

Next, suppose that M is reducible. Then, there exist E_s and I such that $M = E_s[I]$. We use structural induction on E_s . We show only important cases.

- Case $E_s = []$.

In this case, $M = I$. Let $e = \mathcal{A}_\gamma(M)$ and $S = \emptyset$. Then, $M \sim_\gamma (e, S)$ follows from C-INST. $(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$ is obvious.

- Case $E_s = \mathbf{node} \ E'_s \ M'$.

In this case, $M = \mathbf{node} \ E'_s[I] \ M'$. We have $\emptyset \mid \langle \langle \delta_1 \rangle \rangle \vdash E'_s[I] : \mathbf{Tree}^+$ and $\emptyset \mid \langle \langle \delta_2 \rangle \rangle \vdash M' : \mathbf{Tree}^+$ and $\delta = \delta_1, \delta_2$ for some δ_1 and δ_2 from the assumption $\emptyset \mid \langle \langle \delta \rangle \rangle \vdash M : \tau$. Thus, $E'_s[I] \sim_\gamma (e', S')$ and $(\mathcal{A}_\gamma(E'_s[I]), \langle \delta_1 \rangle, \emptyset) \rightarrow^* (e', \langle \delta_1 \rangle, S')$ follow for some e' and S' from the induction hypothesis. Let e be e' ; $\mathcal{A}_\gamma(M')$ and S be $\mathbf{node}; S'$. Then, $M \sim_\gamma (e, S)$ follows from C-NODE1. Because $\mathcal{A}_\gamma(M) = \mathbf{write}(\mathbf{node}); \mathcal{A}_\gamma(E'_s[I]); \mathcal{A}_\gamma(M')$, we have $(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$.

- Case $E_s = \mathbf{node} \ V \ E'_s$.

In this case, $M = \mathbf{node} \ V \ E'_s[I]$. We have $\emptyset \mid \langle \langle \delta \rangle \rangle \vdash E'_s[I] : \mathbf{Tree}^+$ from the assumption $\emptyset \mid \langle \langle \delta \rangle \rangle \vdash M : \tau$. Thus, $E'_s[I] \sim_\gamma (e', S')$ and $(\mathcal{A}_\gamma(E'_s[I]), \langle \delta \rangle, \emptyset) \rightarrow^* (e', \langle \delta \rangle, S')$ follow for some e' and S' from the induction hypothesis. Let e be e' and S be $\mathbf{node}; \llbracket V \rrbracket; S'$. Then, $M \sim_\gamma (e, S)$ follows from C-NODE2. Because $\mathcal{A}_\gamma(M) = \mathbf{write}(\mathbf{node});$

$\mathcal{A}_\gamma(V); \mathcal{A}_\gamma(E'_s[I])$, we have $(\mathcal{A}_\gamma(M), \langle \delta \rangle, \emptyset) \rightarrow^* (e, \langle \delta \rangle, S)$.

- Case $E_s = (\mathbf{case} E'_s \text{ of leaf } x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$.

In this case, $M = (\mathbf{case} E'_s[I] \text{ of leaf } x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$. We have $\emptyset \mid \langle\langle\delta_1\rangle\rangle \vdash E'_s[I] : \mathbf{Tree}^-$ and $x : \mathbf{Int} \mid \langle\langle\delta_2\rangle\rangle \vdash M_1 : \tau$ and $\emptyset \mid x_1 : \mathbf{Tree}^-, x_2 : \mathbf{Tree}^-, \langle\langle\delta_2\rangle\rangle \vdash M_1 : \tau$ and $\delta = \delta_1, \delta_2$ for some δ_1 and δ_2 from the assumption $\emptyset \mid \langle\langle\delta\rangle\rangle \vdash M : \tau$. Thus, $E'_s[I] \sim_\gamma (e', S')$ and $(\mathcal{A}_\gamma(E'_s[I]), \langle\delta_1\rangle, \emptyset) \rightarrow^* (e', \langle\delta_1\rangle, S')$ follows for some e' and S' from the induction hypothesis. Let e be

$$\begin{aligned} \mathbf{case} e'; \mathbf{read}() \quad \mathbf{leaf} &\Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} \mathcal{A}_\gamma(M_1) \\ \mathbf{node} &\Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{aligned}$$

and S be S' . Then, $M \sim_\gamma (e, S)$ follows from C-CASE. Because $\mathcal{A}_\gamma(M) = \mathbf{case} \mathcal{A}_\gamma(E'_s[I]); \mathbf{read}() \text{ of leaf } \Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} \mathcal{A}_\gamma(M_1) \mid \mathbf{node} \Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2)$. $(\mathcal{A}_\gamma(M), \langle\delta\rangle, \emptyset) \rightarrow^* (e, \langle\delta\rangle, S)$ holds.

Appendix D

Proof of Lemma 4.6

We prove Lemma 4.6 in this section.

Proof

The second property follows immediately from the definition of $M \sim_\gamma (e, S)$. (If M is irreducible, then $M \sim_\gamma (e, S)$ must follow either from C-VALUE or from C-TREE, which implies that e is irreducible too.)

We prove the first property below. We prove $M' \sim_{\mathbf{FV}(M)} (e', S'_o)$. We hereafter write γ for $\mathbf{FV}(M)$ and S' for S'_o .

Suppose $(M, \delta) \rightarrow (M', \delta')$. Then, $M = E_s[I]$ for some E_s and I . We use structural induction on E_s . We show only important cases.

- Case $E_s = []$. Case analysis on I . We show only important cases.

— Case $I = \mathbf{case} y \text{ of leaf } x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2$ with $\delta = (y \mapsto \mathbf{leaf} i, \delta_1)$.

$(M, \delta) \rightarrow (M', \delta')$ must have been derived by using ES2-CASE1. So, it must be the case that $M' = [i/x]M_1$ and $\delta' = \delta_1$. $M \sim_\gamma (e, S)$ implies

$$e = \mathcal{A}_\gamma(I) = \begin{aligned} \mathbf{case} () ; \mathbf{read}() \quad \mathbf{leaf} &\Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} \mathcal{A}_\gamma(M_1) \\ \mathbf{node} &\Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{aligned}$$

and $S = \emptyset$. $(e, \mathbf{leaf}; i; \langle\delta_1\rangle, \emptyset) \rightarrow^+ (\mathcal{A}_\gamma(M'), \langle\delta_1\rangle, \emptyset)$ follows from Lemma C.1.

From Lemma 4.5, there exist e'' and S'' that satisfy $(\mathcal{A}_\gamma(M'), \langle\delta_1\rangle, \emptyset) \rightarrow (e'', \langle\delta_1\rangle, S'')$ and $M' \sim_\gamma (e'', S'')$. Let e' be e'' and S' be S'' . Then, we have $(e, \langle\delta\rangle, S) \rightarrow^+ (e', \langle\delta'\rangle, S')$ and $M' \sim_\gamma (e', S')$ as required.

— Case $I = \mathbf{case} y \text{ of leaf } x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2$ with $\delta = (y \mapsto \mathbf{node} V_1 V_2, \delta_1)$.

$(M, \delta) \rightarrow (M', \delta')$ must have been derived by using ES2-CASE2. So, it must be the case that $M' = M_2$ and $\delta' = x_1 \mapsto V_1, x_2 \mapsto V_2, \delta_1$. $M \sim_\gamma (e, S)$ implies

$$e = \mathcal{A}_\gamma(I) = \begin{aligned} \mathbf{case} () ; \mathbf{read}() \quad \mathbf{leaf} &\Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} \mathcal{A}_\gamma(M_1) \\ \mathbf{node} &\Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{aligned}$$

and $S = \emptyset$. As can be easily seen, $[(\)/x_1, (\)/x_2].\mathcal{A}_\gamma(M_2) = \mathcal{A}_{\gamma \cup \{x_1, x_2\}}(M_2)$. Thus, $(e, \mathbf{node}; \llbracket V_1 \rrbracket; \llbracket V_2 \rrbracket; \langle \delta_1 \rangle, \emptyset) \rightarrow^+ (\mathcal{A}_{\gamma \cup \{x_1, x_2\}}(M_2), \llbracket V_1 \rrbracket; \llbracket V_2 \rrbracket; \langle \delta_1 \rangle, \emptyset)$. From Lemma 4.5, there exist e'' and S'' that satisfy $(\mathcal{A}_{\gamma \cup \{x_1, x_2\}}(M_2), \llbracket V_1 \rrbracket; \llbracket V_2 \rrbracket; \langle \delta_1 \rangle, \emptyset) \rightarrow^* (e'', \llbracket V_1 \rrbracket; \llbracket V_2 \rrbracket; \langle \delta_1 \rangle, S'')$ and $M_2 \sim_\gamma (e'', S'')$. Let $e' = e''$ and $S' = S''$. Then, we have $(e, \langle \delta \rangle, S) \rightarrow^+ (e', \langle \delta' \rangle, S')$ and $M' \sim_\gamma (e', S')$ as required.

- Case $E_s = \mathbf{leaf} E_1$.

There exists M'_1 that satisfies $M' = \mathbf{leaf} M'_1$ and $(E_1[I], \delta) \rightarrow (M'_1, \delta')$. $M \sim_\gamma (e, S)$ must have been derived from C-LEAF. Thus, there exist e_1 and S_1 that satisfies $e = \mathbf{write}(e_1)$ and $E_1[I] \sim_\gamma (e_1, S_1)$ and $S = \mathbf{leaf}; S_1$. Because M is well-typed, $E_1[I]$ is also well-typed. Thus, from the induction hypothesis, we have $(e_1, \langle \delta \rangle, S_1) \rightarrow^+ (e'_1, \langle \delta' \rangle, S'_1)$ and $M'_1 \sim_\gamma (e'_1, S'_1)$ for some e'_1 and S'_1 .

First, suppose that M'_1 is reducible. Let e' be $\mathbf{write}(e'_1)$ and S' be $\mathbf{leaf}; S_1$. Then, $M' \sim_\gamma (e', S')$ follows from C-LEAF as required.

Next, suppose that M'_1 is a value. Because M is well-typed, M'_1 is an integer (let the integer be i'_1) and $M'_1 \sim_\gamma (e'_1, S'_1)$ must have been derived from C-VALUE. Thus, $e'_1 = i'_1$ and $S'_1 = \emptyset$. Let e' be $(\)$ and S' be $\mathbf{leaf}; i'_1$. Then, $M' \sim_\gamma (e', S')$ follows from C-TREE and $(\mathbf{write}(e'_1), \langle \delta' \rangle, \mathbf{leaf}; S'_1) \rightarrow (e', \langle \delta' \rangle, S')$ holds.

- Case $E_s = \mathbf{node} E_1 M_2$.

There exists M'_1 that satisfies $M' = \mathbf{node} M'_1 M_2$ and $(E_1[I], \delta) \rightarrow (M'_1, \delta')$. $M \sim_\gamma (e, S)$ must have been derived from C-NODE1. Thus, there exist e_1 and S_1 that satisfies $e = e_1; \mathcal{A}_\gamma(M_2)$ and $E_1[I] \sim_\gamma (e_1, S_1)$ and $S = \mathbf{node}; S_1$. Because we assume that M is well-typed, $E_1[I]$ is also well-typed. Thus, from the induction hypothesis, we have $(e_1; \mathcal{A}_\gamma(M_2), \langle \delta \rangle, \mathbf{node}; S_1) \rightarrow^+ (e'_1; \mathcal{A}_\gamma(M_2), \langle \delta' \rangle, \mathbf{node}; S'_1)$ and $M'_1 \sim_\gamma (e'_1, S'_1)$ for some e'_1 and S'_1 .

First, suppose that M'_1 is reducible. Let e' be $e'_1; \mathcal{A}_\gamma(M_2)$ and S' be $\mathbf{node}; S'_1$. Then, $M' \sim_\gamma (e', S')$ follows from C-NODE1 as required.

Next, suppose that M'_1 is a value (let the value be V'_1). $M'_1 \sim_\gamma (e'_1, S_1)$ must have been derived from C-TREE. Thus, $e'_1 = (\)$ and $S_1 = \llbracket V'_1 \rrbracket$. Thus, $(e, \langle \delta \rangle, S) \rightarrow^+ (\mathcal{A}_\gamma(M_2), \langle \delta' \rangle, \mathbf{node}; \llbracket V'_1 \rrbracket)$. From Lemma 4.5, there exist e_2 and S_2 that satisfy $M_2 \sim_\gamma (e_2, S_2)$ and $(\mathcal{A}_\gamma(M_2), \langle \delta' \rangle, \mathbf{node}; \llbracket V'_1 \rrbracket) \rightarrow^* (e_2, \langle \delta' \rangle, \mathbf{node}; \llbracket V'_1 \rrbracket; S_2)$. Let e' be e_2 and S' be $\mathbf{node}; \llbracket V'_1 \rrbracket; S_2$. Then, $M' \sim_\gamma (e', S')$ follows from C-NODE2 and $(e, \langle \delta \rangle, S) \rightarrow^+ (e', \langle \delta' \rangle, S')$ holds.

- Case $E_s = \mathbf{node} V_2 E_1$.

There exists M'_1 that satisfies $M' = \mathbf{node} V_2 M'_1$ and $(E_1[I], \delta) \rightarrow (M'_1, \delta')$. $M \sim_\gamma (e, S)$ must have been derived from C-NODE2. Thus, there exist e_1 and S_1 that satisfy $e = e_1$ and $E_1[I] \sim_\gamma (e_1, S_1)$ and $S = \mathbf{node}; \llbracket V_2 \rrbracket; S_1$. Because M is well-typed, $E_1[I]$ is also well-typed. Thus, from the induction hypothesis, we have $(e_1, \langle \delta \rangle, \mathbf{node}; \llbracket V_2 \rrbracket; S_1) \rightarrow^+ (e'_1, \langle \delta' \rangle, \mathbf{node}; \llbracket V_2 \rrbracket; S'_1)$ and $M'_1 \sim_\gamma (e'_1, S'_1)$ for some e'_1 and S'_1 .

First, suppose that M'_1 is reducible. Let e' be e'_1 and S' be $\mathbf{node}; \llbracket V_2 \rrbracket; S'_1$. Then, $M' \sim_\gamma (e', S')$ follows from C-NODE1 as required.

Next, suppose that M'_1 is a value (let the value be V'_1). $M'_1 \sim_\gamma (e'_1, S_1)$ must have been derived from C-TREE. Thus, $e'_1 = ()$ and $S_1 = \llbracket V'_1 \rrbracket$, and thus $(e, \langle \delta \rangle, S) \rightarrow^+ ((), \langle \delta' \rangle, \mathbf{node}; \llbracket V_2 \rrbracket; \llbracket V'_1 \rrbracket)$. Let e' be $()$ and S' be $\mathbf{node}; \llbracket V_2 \rrbracket; \llbracket V'_1 \rrbracket$. Then, $M' \sim_\gamma (e', S')$ follows from C-TREE as required.

- Case $E_s = (\mathbf{case} E_1 \mathbf{of leaf} x \Rightarrow M_1 \mid \mathbf{node} x_1 x_2 \Rightarrow M_2)$.

There exists M'_1 that satisfies

$$M' = \begin{array}{l} \mathbf{case} M'_1 \quad \mathbf{of} \quad \mathbf{leaf} x \Rightarrow M_1 \\ \quad \quad \quad \mid \quad \mathbf{node} x_1 x_2 \Rightarrow M_2 \end{array}$$

and $(E_1[I], \delta) \rightarrow (M'_1, \delta')$. $M \sim_\gamma (e, S)$ must have been derived from C-CASE. Thus, there exists e_1 that satisfies

$$e = \begin{array}{l} \mathbf{case} e_1; \mathbf{read}() \quad \mathbf{of} \quad \mathbf{leaf} \Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} M_1 \\ \quad \quad \quad \mid \quad \mathbf{node} \Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{array}$$

and $E_1[I] \sim_\gamma (e_1, S)$. Because M is well-typed, $E_1[I]$ is also well-typed. Thus, from the induction hypothesis, we have

$$\left(\begin{array}{l} \mathbf{case} e_1; \mathbf{read}() \quad \mathbf{of} \quad \mathbf{leaf} \Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} M_1 \\ \quad \quad \quad \mid \quad \mathbf{node} \Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{array} \right), \langle \delta \rangle, S \rightarrow^+ \left(\begin{array}{l} \mathbf{case} e'_1; \mathbf{read}() \quad \mathbf{of} \quad \mathbf{leaf} \Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} M_1 \\ \quad \quad \quad \mid \quad \mathbf{node} \Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{array} \right), \langle \delta' \rangle, S''$$

and $M'_1 \sim_\gamma (e'_1, S'')$ for some e'_1 and S'' .

First, suppose that M'_1 is reducible. Let e' be

$$\begin{array}{l} \mathbf{case} e'_1; \mathbf{read}() \quad \mathbf{of} \quad \mathbf{leaf} \Rightarrow \mathbf{let} x = \mathbf{read}() \mathbf{in} M_1 \\ \quad \quad \quad \mid \quad \mathbf{node} \Rightarrow [()/x_1, ()/x_2].\mathcal{A}_\gamma(M_2) \end{array}$$

and S' be S'' . Then, $M' \sim_\gamma (e', S')$ follows from C-CASE.

Next, suppose that M'_1 is not reducible. Because M'_1 is a tree-typed term, M'_1 is a variable (let it be y'_1). Because $M'_1 \sim_\gamma (e'_1, S'')$ must have been derived from C-VALUE, $e'_1 = \mathcal{A}_\gamma(y'_1)$. Thus, $(e, \langle \delta \rangle, S) \rightarrow^+ (\mathcal{A}_\gamma(M'), \langle \delta' \rangle, S'')$. From Lemma 4.5, there exist e'' and S_1 that satisfy $M' \sim_\gamma (e'', S_1)$ and $(\mathcal{A}_\gamma(M'), \langle \delta' \rangle, S'') \rightarrow^* (e'', \langle \delta' \rangle, S_1)$. Let e' be e'' and S' be S_1 . Then, $(e, \langle \delta \rangle, S) \rightarrow^+ (e', \langle \delta' \rangle, S')$ and $M' \sim_\gamma (e', S')$ hold as required.

□

Appendix E

Proof of Theorem 4.4

This section provides proof of Theorem 4.4.

Proof

First of all, note that $\emptyset \mid x : \mathbf{Tree}^- \vdash M x : \tau$ follows an assumption $\emptyset \mid \emptyset \vdash M : \mathbf{Tree}^- \rightarrow \tau$ and $\emptyset \mid x : \mathbf{Tree}^- \vdash x : \mathbf{Tree}^-$.

We prove only (ii) hereafter; (i) can be proved in the same way.

Assume $((M\ x), x \mapsto V) \rightarrow^* (V', \emptyset)$. Because $\emptyset \mid x : \mathbf{Tree}^- \vdash (M\ x) : \tau$ holds, there exist e, S_i , and S_o such that $((M\ x), x \mapsto V) \sim (e, S_i, S_o)$ and $(\mathcal{A}(M)(), S_i, \emptyset) \rightarrow^* (e, S_i, S_o)$ from Lemma 4.5.⁸ From the definition of \sim , $S_i = \llbracket V \rrbracket$. Because of Theorem 4.2 and Lemma 4.6, there exists a sequence of reduction $(e, \llbracket V \rrbracket, S_o) \rightarrow^* (e', \emptyset, S'_o)$ that satisfies $(V', \emptyset) \sim (e', \emptyset, S'_o)$. From the definition of \sim , $e' = ()$ and $S'_o = \llbracket V' \rrbracket$. Thus, $(\mathcal{A}(M)(), \llbracket V \rrbracket, \emptyset) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$ holds.

Next, assume $(\mathcal{A}(M)(), \llbracket V \rrbracket, \emptyset) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$. As we stated above, there exist e, S_i , and S_o such that $((M\ x), x \mapsto V) \sim (e, \llbracket V \rrbracket, S_o)$ and $(\mathcal{A}(M)(), \llbracket V \rrbracket, \emptyset) \rightarrow^* (e, \llbracket V \rrbracket, S_o)$. Because applicable reduction rule can be uniquely determined at each step of reduction, $(e, \llbracket V \rrbracket, S_o) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$ holds.

In the following, we prove “if $\emptyset \mid \langle \delta \rangle \vdash M' : \mathbf{Tree}^+$ and $(e, \langle \delta \rangle, S'_o) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$ and $(M', \delta) \sim (e, S'_i, S'_o)$ hold, $(M', \delta) \rightarrow^* (V', \emptyset)$ holds.” We use mathematical induction on the number of reduction step of $(e, \langle \delta \rangle, S'_o) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$. With this fact, by letting M' be $M\ x$ and δ be $x \mapsto V$, $((M\ x), x \mapsto V) \rightarrow^* (V', \emptyset)$ holds because $\emptyset \mid x : \mathbf{Tree}^- \vdash (M\ x) : \mathbf{Tree}^+$ follows $((M\ x), x \mapsto V) \rightarrow^* (V, \emptyset)$ and Theorem 4.2.

- In the case of $n = 0$, $M' = V'$ and $\delta = \emptyset$ hold because $e = (), \langle \delta \rangle = \emptyset, S'_o = \llbracket V' \rrbracket$ and $(M', \emptyset) \sim (e, \emptyset, S'_o)$ hold. Thus, $(M', \delta) \rightarrow^* (V', \emptyset)$ holds.
- In the case of $n \geq 1$, there exist e', S_i , and S''_o that satisfy

$$(e, \langle \delta \rangle, S'_o) \rightarrow (e', S_i, S''_o) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$$

From Lemma 4.6, there exist M'', δ' , and S'''_o that satisfy

- $(M', \delta) \rightarrow (M'', \delta')$
- $(M'', \delta') \sim (e'', \langle \delta' \rangle, S'''_o)$
- $(e, \langle \delta \rangle, S'_o) \rightarrow^+ (e'', \langle \delta' \rangle, S'''_o)$

Since the reduction is deterministic, $(e'', \langle \delta' \rangle, S'''_o) \rightarrow^* ((), \emptyset, \llbracket V' \rrbracket)$ holds. From the induction hypothesis, $(M'', \delta') \rightarrow^* (V', \emptyset)$. Thus, $(M', \delta) \rightarrow^* (V', \emptyset)$ holds.

□

References

- Aho, A. V., Sethi, R. & Ullman, J. D. (1986) *Compilers*. Boston, MA: Addison-Wesley.
- Avila-Campillo, I., Green, T. J., Gupta, A., Onizuka, M., Raven, D. & Suciu, D. (2002) October XMLTK: An XML toolkit for scalable XML stream processing. In *Proceedings of PLAN-X*.
- Baker, H. G. (1992) Lively Linear Lisp—Look Ma, No Garbage! *ACM SIGPLAN notices*, 27(8), 89–98.
- Bar-Yossef, Z., Fontoura, M., & Josifovski, V. (2005) Buffering in query evaluation over XML streams. In *Proceedings of the 24th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS2005)* pp. 216–227.
- Benzaken, V., Castagna, G. & Frisch, A. (2003 August) CDuce: An XML-centric general-purpose language. In *Proceedings of 8th ACM SIGPLAN International Conference on Functional Programming (ICFP 2003)*. New York: ACM Press, pp. 51–63.

⁸ Because $\emptyset \mid \emptyset \vdash M : \mathbf{Tree}^- \rightarrow \tau$ holds, $\mathbf{FV}(M) = \emptyset$. Thus, $\mathcal{A}_{\mathbf{FV}(M) \cup \{x\}}(M\ x) = \mathcal{A}(M)()$.

- Berglund, A., Boag, S., Chamberlin, D., Fernández, M. F., Kay, M., Robie, J. & Simeon, J. (2003 November). *XML path language (XPath) 2.0*. World Wide Web Consortium. Available at: <http://www.w3.org/TR/xpath20/>. Accessed 29 November 2007.
- Boag, S., Chamberlin, D., Fernández, M. F., Florescu, D., Robie, J. & Siméon, J. (2003 November) *XQuery 1.0: An XML Query Language*. World Wide Web Consortium. <http://www.w3.org/TR/xquery/>. Accessed 29 November 2007.
- Bray, T., Paoli, J., Sperberg-McQueen, C. M. & Maler, E. (2000 October). *Extensible Markup Language (XML) 1.0, 2nd ed.* Tech. rept. World Wide Web Consortium. Available at: <http://www.w3.org/TR/REC-xml>. Accessed 29 November 2007.
- Chamberlin, D., Frankhauser, P., Marchiori, M. & Robie, J. (2003 November) *XML query (XQuery) requirements*. World Wide Web Consortium. Available at: <http://www.w3.org/TR/xquery-requirements/>. Accessed 29 November 2007.
- DataPower Technology, Inc. (2001 March) *XSLTMark*. Available at: <http://www.datapower.com/xmldev/xsltmark.html>. Accessed 2 November 2005.
- Ganzinger, H. & Giegerich, R. (1984) Attribute coupled grammars. In *Proceedings of the ACM SIGPLAN '84 Symposium on Compiler Construction*. New York: ACM Press, pp. 157–170.
- Green, T., Onizuka, M. & Suciu, D. (2001) *Processing XML Streams With Deterministic Automata and Stream Indexes*. Tech. rept. Seattle, WA: University of Washington.
- Gupta, A. K. & Suciu, D. (2003 June) Stream processing of XPath queries with predicates. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*. New York: ACM Press, pp. 419–30.
- Hosoya, H. & Pierce, B. C. (2003). XDuce: A typed XML processing language. *ACM Trans. Internet Technol. (TOIT)* 3(2), 117–148.
- Igarashi, A. & Kobayashi, N. (2002) Resource usage analysis. In *Proceedings of ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. New York: ACM Press, pp. 331–342.
- Koch, C. & Scherzinger, S. (2003 September). Attribute grammars for scalable query processing on XML streams. In *Proceedings of the 9th International Conference on Data Base Programming Language (DBPL)*. New York: Springer-Verlag, pp. 233–56.
- Ludäscher, B., Mukhopadhyay, P. & Papakonstantinou, Y. (2002) A Transducer-based XML query processor. *Proceedings of the 28th International Conference on Very Large Data Bases*. Franklin County, Ohio: VLDB Endowment Inc, pp. 227–38.
- Meijer, E., Fokkinga, M. & Paterson, R. (1991) Functional programming with bananas, lenses, envelopes and barbed wire. In *Proceedings of the 5th ACM Conference on Functional Programming Languages and Computer Architecture*. New York: Springer-Verlag, pp. 124–144.
- Nakano, K. (2004). *Composing Stack-Attributed Tree Transducers*. Tech. rept. METR–2004–01. Department of Mathematical Informatics, University of Tokyo, Japan.
- Nakano, K. & Mu, S.-C. (2006) A pushdown machine for recursive XML processing. In *Proceedings of the 4th ASIAN Symposium on Programming Languages and Systems (APLAS2006)*. New York: Springer-Verlag, pp. 340–56.
- Nakano, K. & Nishimura, S. (2001). Deriving event-based document transformers from tree-based specifications. *Electron. Notes Theor. Comput. Sci.* 44(2), 181–205.
- Neven, F. (2005) Attribute grammars for unranked trees as a query language for structured documents. *J. Comput. Syst. Sci.* 70(2), 221–257.
- Olteanu, D., Meuss, H., Furche, T. & François. (2002) XPath: Looking forward. In *Proceedings of the EDBT Workshop on XML Data Management (XMLDM)*. pp. 109–127.

- Petersen, L., Harper, R., Crary, K. & Pfenning, F. (2003) A type theory for memory allocation and data layout. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. New York: ACM, pp. 172–84.
- Polakow, J. (2001 June) *Ordered Linear Logic and Applications*. Ph.D. thesis. Carnegie Mellon University. Available as Technical Report CMU-CS-01-152.
- Sato, S. (2007 March). *Automatic Insertion of Buffering Primitives for Generating XML Stream Processor* (in Japanese). M.Phil. thesis. Japan: Tohoku University.
- Scardina, M. & Fernandez, M. (2003 August) *XPath Requirements Version 2.0*. World Wide Web Consortium. Available at: <http://www.w3.org/TR/xpath20req/>. Accessed 29 November 2007.
- Schmidt, A. R., Waas, F., Kersten, M. L., Florescu, D., Manolescu, I., Carey, M. J. & Busse, R. (2001 April) *The XML Benchmark Project*. Tech. rept. Centrum voor Wiskunde en Informatica.
- Suenaga, K., Kobayashi, N. & Yonezawa, A. (2005 September) Extension of type-based approach to generation of stream-processing programs by automatic insertion of buffering primitives. In *International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR 2005)*. New York: Springer-Verlag, pp. 98–114.
- Turner, D. N., Wadler, P. & Mossin, C. (1995) Once upon A type. In *Proceedings of Functional Programming Languages and Computer Architecture*. New York: Springer-Verlag, pp. 1–11.
- Wadler, P. (1988) Deforestation: Transforming programs to eliminate trees. In *ESOP '88. European Symposium on Programming, Nancy, France, Lecture Notes in Computer Science, vol. 300*. Berlin: Springer-Verlag, pp. 344–358.