Caching Intermediate Results for Program Improvement

Yanhong A. Liu* Tim Teitelbaum*
Department of Computer Science, Cornell University, Ithaca, NY 14853
Email: {yanhong, tt}@cs.cornell.edu

Abstract

A systematic approach is given for symbolically caching intermediate results useful for deriving incremental programs from non-incremental programs. Our method can be applied straightforwardly to provide a systematic approach to program improvement via caching.

1 Introduction

Incremental programs take advantage of repeated computations on inputs that differ only slightly from one another, making use of the old output in computing a new output rather than computing from scratch. Methods of incremental computation have widespread application, e.g., optimizing compilers [2, 9, 11], transformational programming [30, 33, 43], interactive editing systems [4, 39], *etc.*

Deriving incremental programs. Given a program f and an input change $\oplus$, a program f' that computes the result of $f(x \oplus y)$ efficiently by making use of the value of $f(x)$ is called an incremental version of f under $\oplus$.

Liu and Teitelbaum [27] give a systematic transformational approach for deriving an incremental program f' from a given program f and an input change $\oplus$. The basic idea is to identify in the computation of $f(x \oplus y)$ those subcomputations that are also performed in the computation of $f(x)$ and whose values can be retrieved from the cached result r of $f(x)$. The computation of $f(x \oplus y)$ is transformed symbolically to avoid re-performing these subcomputations by replacing them with corresponding retrievals. This efficient way of computing $f(x \oplus y)$ is captured in the definition of $f'(x, y, r)$.

Caching intermediate results. The above approach has a limitation. The derived program $f'(x, y, r)$ avoids only subcomputations that are performed by $f(x)$ *and* whose values can be retrieved from the cached result r of $f(x)$. There may be subcomputations of $f'(x, y, r)$ that are also performed by $f(x)$ *but* whose values can not be retrieved from r. If the values of these computations are also cached and returned, then after the input change $\oplus$, we can avoid recomputing them by retrieving their values from the extended return value of the old computation. We call the values of these computations *intermediate results useful for computing f incrementally under $\oplus$.*

Examples where intermediate results are needed for incremental computation include incremental parsing [16] and incremental attribute evaluation [24, 40, 50]. An incremental parser may cache, in addition to the derived parse tree, the $LR(0)$ state corresponding to each shift and reduction. An attribute evaluator may only return some designated synthesized attribute of the root [18], but the corresponding incremental attribute evaluator may cache the whole attributed tree.

Program improvement via caching. Deriving incremental programs and caching intermediate results provide a principled approach for program improvement using caching. In essence, every program computes by fixed point iteration, which is why loop optimizations are so important. A straightforward idea then is to compute the result of each iteration incrementally using the stored result of the previous iteration, which is why strength reduction [3] and related techniques [32] are crucial for performance. Observe that, most of the time, not only the result, but also the intermediate results computed in one iteration can be useful for efficiently computing the result of the next iteration. Thus, these intermediate results need to be identified, used, and maintained as well.

We can regard a loop body as a program f, and a loop increment as a change $\oplus$. Then the goal of computing loops incrementally corresponds to transforming f into f', an incremental version of f under $\oplus$, and the problem of identifying intermediate results that can be used from iteration to iteration corresponds to deciding which intermediate results are useful for computing f incrementally under $\oplus$. Once these intermediate results have been determined and the program f has been extended to a program $\hat{f}$ that returns them as well, a program $\hat{f}'$ can be derived that incrementally computes these intermediate results as well as the value of f.

This paper. We present a clean three-stage method, called *cache-and-prune*, for caching intermediate results useful for computing f incrementally under $\oplus$. The basic idea is to (I) extend the program f to a program $\bar{f}$ that returns

*The authors gratefully acknowledge the support of the Office of Naval Research under contract No. N00014-92-J-1973.

****
Appears in Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation, La Jolla, California, June 21-23, 1995.

all intermediate results, (II) incrementalize the program $\bar{f}$ under $\oplus$ to obtain an incremental version $\bar{f}'$, and (III) using the dependencies in $\bar{f}'$, prune the extended program $\bar{f}$ to a program $\hat{f}$ that returns only the useful intermediate results. Additionally, we also prune the program $\bar{f}'$ to directly obtain a program $\hat{f}'$ that incrementally maintains only the useful intermediate results. The contributions of this paper are as follows.

1. We give a systematic approach for caching intermediates results useful for computing f incrementally under $\oplus$, and for constructing a corresponding program that incrementally maintains these intermediate results. Previous work on this relies on a fixed set of rules [3, 32], applies only to programs with certain properties or schemas [5, 10, 34, 35], or requires program annotations [14, 19, 44].

2. Our cache-and-prune method consists of three independent stages, and thus is modular. It has certain nice properties. Stage I gives us maximality by providing all the intermediate results possibly used by Stage II. Stage II uses these intermediate results for the exclusive purpose of incrementalization. Stage III gives us a kind of minimality by preserving only the intermediate results actually used by Stage II. Therefore, the whole method is optimal with respect to the incremental techniques of Stage II (for which we use [27]). Stages I and III are simple, clean, and fully-automatable.

3. We develop in Stage III a backward dependency analysis that uses domain projections to specify sufficient information, which is a natural application of the techniques previously used for other analyses [22, 46]. Our projections specify specific components of compound values, rather than just heads or tails of list values, and thus provide more accurate information. The technique may also be used to assist general program optimizations in context, like tuple elimination [45].

4. Our result can be applied straightforwardly and systematically to program improvement via caching. The classical example of the Fibonacci function, which can be improved dramatically by various caching techniques, is shown in Section 7. A comprehensive comparison with work in program improvement via caching is given in Section 8.

This paper is organized as follows: Section 2 defines the problem of caching intermediate results. Section 3 outlines the cache-and-prune method and its correctness. Sections 4, 5, and 6 describe caching, incrementalization, and pruning, respectively. Several examples are given in Section 7. Finally, we discuss related work and conclude in Section 8. For more details of some transformations, examples, and related work, see the extended version of this paper [25].

2 Defining the problem

We use a simple first-order functional programming language. The expressions of our language are given by the following grammar:

$e ::= v$	variable
$\mid\ c(e_1, ..., e_n)$	constructor application
$\mid\ p(e_1, ..., e_n)$	primitive function application
$\mid\ f(e_1, ..., e_n)$	function application
$\mid\ \textbf{if}\ e_1\ \textbf{then}\ e_2\ \textbf{else}\ e_3$	conditional expression
$\mid\ \textbf{let}\ v = e_1\ \textbf{in}\ e_2$	binding expression

Each constructor c, primitive function p, and user-defined function f has a fixed arity. A program is a set F of mutually recursive function definitions of the form

$$f(v_1, ..., v_n) = e \tag{1}$$

and a function f_0 that is to be evaluated with some input $x = \langle x_1, ..., x_n \rangle$. The semantics of the language is strict. Figure 1 gives some example definitions.

An input change $\oplus$ to a function f_0 combines an old input $x = \langle x_1, ..., x_n \rangle$ and a change $y = \langle y_1, ..., y_m \rangle$ to form a new input $x' = \langle x'_1, ..., x'_n \rangle = x \oplus y$, where each x'_i is some function of x_j's and y_k's. For example, an input change $\oplus_1$ to the function foo or fib in Figure 1 can be defined by $x' = x \oplus_1 y = x + 1$, and an input change $\oplus_2$ to $sort$ can be $x' = x \oplus_2 y = cons(y, x)$.

We need cost models to discuss efficient computation and program improvement. In this paper, we use an asymptotic time model, and write

$$t(f(v_1, .., v_n)) \tag{2}$$

to denote the asymptotic time of computing $f(v_1, ..., v_n)$. Since only asymptotic time is of concern, it is sufficient to consider only the values of function applications as candidate intermediate results to be cached. Of course, caching intermediate results takes extra space, which reflects the well-known principle of trading space for speed. We assume that we have unlimited space to be used for achieving the least asymptotic time possible. The pruning saves time as well as space for computing and maintaining intermediate results that are not useful for incremental computation. There are standard constructions for mechanical time analysis [41, 48], but automatic space analysis and the trade-off between time and space are problems open for study.

Given a program f_0 and an input change $\oplus$, we can use the approach in [27] to derive a program f'_0, an incremental version of f_0 under $\oplus$, such that, if $f_0(x) = r$, then whenever $f_0(x \oplus y)$ returns a value, $f'_0(x, y, r)$ returns the same value and is asymptotically at least as fast.[1] Instead of trivially defining $f'_0(x, y, r)$ to be $f_0(x \oplus y)$, we attempt to make $f'_0(x, y, r)$ as efficient as possible by having it use the cached result r of $f_0(x)$ as much as possible. A number of examples like outer product and selection sort are given in [27]. For the function foo in Figure 1 and input change $x \oplus_1 y = x + 1$, the function foo' given in Figure 2 can be derived. Unfortunately, computing $foo'(x, r)$ is not much faster than computing $foo(x + 1)$ from scratch.

This foo example illustrates that there are cases where we can compute the value of $f_0(x \oplus y)$ more quickly by caching and using, in addition to the value of $f_0(x)$, the intermediate results computed in $f_0(x)$. For example, the value of $foo(x - 1) + foo(x - 2)$, which could be used in computing $foo'(x, r)$, is also computed by $foo(x)$ but can not be retrieved from r. Thus, we can cache this intermediate value and use it in computing the value of $foo(x + 1)$ faster.

We need consistent notations for the mechanical transformation that caches intermediate results. We use $<>$ to denote a tuple constructed by the transformation that bundles intermediate results with the original return value, with $1st$ returning the first element, which is always the original value, and rst returning a tuple of the remaining elements,

[1] While $f_0(x)$ abbreviates $f_0(x_1, ..., x_n)$, and $f_0(x \oplus y)$ abbreviates $f_0(\langle x_1, ..., x_n \rangle \oplus \langle y_1, ..., y_m \rangle)$, $f'_0(x, y, r)$ abbreviates $f'_0(x_1, ..., x_n, y_1, ..., y_m, r)$. Note that some of the parameters of f'_0 may be dead and eliminated [27].

$foo(x)$: sum three preceding "foo" numbers

$$
\begin{aligned}
foo(x) &= \textbf{if } x \le 2 \textbf{ then } 1 \\
& \textbf{ else } boo(x) + foo(x-3) \\
boo(x) &= foo(x-1) + foo(x-2)
\end{aligned}
$$

$fib(x)$: compute the x-th Fibonacci number

$$
\begin{aligned}
fib(x) &= \textbf{if } x \le 1 \textbf{ then } 1 \\
& \textbf{ else } fib(x-1) + fib(x-2)
\end{aligned}
$$

$sort(x)$: sort a list x using merge sort

$$
\begin{aligned}
sort(x) &= \textbf{if } null(x) \textbf{ then } nil \\
& \textbf{ else if } null(cdr(x)) \textbf{ then } x \\
& \textbf{ else } merge(sort(odd(x)), sort(even(x))) \\
odd(x) &= \textbf{if } null(x) \textbf{ then } nil \textbf{ else } cons(car(x), even(cdr(x))) \\
even(x) &= \textbf{if } null(x) \textbf{ then } nil \textbf{ else } odd(cdr(x)) \\
merge(x,y) &= \textbf{if } null(x) \textbf{ then } y \\
& \textbf{ else if } null(y) \textbf{ then } x \\
& \textbf{ else if } car(x) \le car(y) \textbf{ then } \\
&\phantom{= \textbf{ else if }} cons(car(x), merge(cdr(x), y)) \\
& \textbf{ else } cons(car(y), merge(x, cdr(y)))
\end{aligned}
$$

Figure 1: Example function definitions

which are the corresponding intermediate results. We use nth to get the nth element of such a tuple, and we use an infix operation @ to concatenate two such tuples.

For typographical convenience, we shall always use x to refer to the previous input to f_0, r the cached result of $f_0(x)$, y the change parameter to the input x, x' the new input $x \oplus y$, and f_0' an incremental version of f_0 under $\oplus$. We use $\bar{f}_0$ to refer to the extended function that returns all intermediate results of f_0, $\bar{r}$ the cached result of $\bar{f}_0(x)$, and $\bar{f}_0'$ an incremental version of $\bar{f}_0$ under $\oplus$. Similarly, we use $\hat{f}_0$ to refer to the pruned function that returns only the intermediate results of f_0 useful for incremental computation, $\hat{r}$ the cached result of $\hat{f}_0(x)$, and $\hat{f}_0'$ a function that incrementally maintains only the useful intermediate results.

We use the function foo in Figure 1 as a running example. At the end, we obtain the functions $\widehat{foo}$, $\widehat{boo}$, and $\widehat{foo}'$ as shown in Figure 2. Rather than computing foo or $\widehat{foo}$ from scratch using $O(3^n)$ time, $\widehat{foo}'$ computes incrementally using only $O(1)$ time.

3 Approach

The cache-and-prune method consists of three stages.

Stage I constructs a program $\bar{f}_0$, an extended version of f_0, such that $\bar{f}_0(x)$ returns the values of all function calls made in computing $f_0(x)$. Basically, $\bar{f}_0(x)$ returns a tuple containing both the intermediate results and the value of $f_0(x)$, such that

$$1st(\bar{f}_0(x)) = f_0(x) \quad \text{and} \quad t(\bar{f}_0(x)) \le t(f_0(x)). \tag{3}$$

Stage II derives a function $\bar{f}_0'$, an incremental version of $\bar{f}_0$ under $\oplus$, using the approach in [27], such that if $\bar{f}_0(x) = \bar{r}$, then we have if $\bar{f}_0(x \oplus y) = \bar{r}'$, then

$$\bar{f}_0'(x, y, \bar{r}) = \bar{r}' \quad \text{and} \quad t(\bar{f}_0'(x, y, \bar{r})) \le t(\bar{f}_0(x \oplus y)) \tag{4}$$

and thus, together with (3), we have

$$1st(\bar{f}_0'(x, y, \bar{r})) = 1st(\bar{f}_0(x \oplus y)) = f_0(x \oplus y). \tag{5}$$

Stage III generates a function $\hat{f}_0$, a pruned version of $\bar{f}_0$, such that $\hat{f}_0(x)$ returns $\Pi(\bar{r})$, where $\bar{r}$ is the return value of $\bar{f}_0(x)$, and $\Pi(\bar{r})$ projects out the first and other components of $\bar{r}$ on which $1st(\bar{f}_0'(x, y, \bar{r}))$ *transitively* depends. The dependency is transitive in the sense that if $1st(\bar{f}_0'(x, y, \bar{r}))$ depends on $\Pi_1(\bar{r})$, and $\Pi_1(\bar{f}_0'(x, y, \bar{r}))$ depends on $\Pi_2(\bar{r})$, then $1st(\bar{f}_0'(x, y, \bar{r}))$ depends also on $\Pi_2(\bar{r})$. This transitivity is caused by the need to *maintain* intermediate results corresponding to those that are *used* for computing $1st(\bar{f}_0'(x, y, \bar{r}))$. In other words, this stage eliminates those intermediate results cached in $\bar{r}$ that are not transitively needed in incrementally computing $1st(\bar{f}_0'(x, y, \bar{r}))$, the value of $f_0(x \oplus y)$.[2] In particular, if $f_0(x) = r$, then

$$1st(\hat{f}_0(x)) = r \quad \text{and} \quad t(\hat{f}_0(x)) \le t(f_0(x)). \tag{6}$$

Additionally, we obtain a function $\hat{f}_0'$, a pruned version of $\bar{f}_0'$, such that if $\bar{f}_0'(x, y, \bar{r})$ returns $\bar{r}'$, then $\hat{f}_0'(x, y, \hat{r})$, where $\hat{r}$ is $\Pi(\bar{r})$ as above, returns $\Pi(\bar{r}')$. This pruning is possible because $\Pi(\bar{r}')$ depends only on $\Pi(\bar{r})$, which can be easily shown using the transitivity above. With the relationship between $\hat{f}_0$ and $\bar{f}_0$, together with (3) and (4), we can prove that if $f_0(x) = r$, then we have if $\hat{f}_0(x) = \hat{r}$ and $f_0(x \oplus y) = r'$, then

$$\hat{f}_0'(x, y, \hat{r}) = \hat{f}_0(x \oplus y) \quad \text{and} \quad t(\hat{f}_0'(x, y, \hat{r})) \le t(f_0(x \oplus y)) \tag{7}$$

and thus, together with (6), we have

$$1st(\hat{f}_0'(x, y, \hat{r})) = 1st(\hat{f}_0(x \oplus y)) = r'. \tag{8}$$

Thus, $\hat{f}_0'(x, y, \hat{r})$ incrementally computes the desired output and the corresponding intermediate results and is asymptotically at least as fast as computing the desired output from scratch. Therefore, we do not have to conduct a derivation on $\hat{f}_0$ and $\oplus$ to obtain such an incremental function.

At the end, putting (6), (7), and (8) together, we have if $f_0(x) = r$, then

$$1st(\hat{f}_0(x)) = r \quad \text{and} \quad t(\hat{f}_0(x)) \le t(f_0(x)) \tag{9}$$

and if $f_0(x \oplus y) = r'$ and $\hat{f}_0(x) = \hat{r}$, then

$$\begin{aligned} 1st(\hat{f}_0'(x, y, \hat{r})) = r', \quad \hat{f}_0'(x, y, \hat{r}) = \hat{f}_0(x \oplus y), \\ \text{and} \quad t(\hat{f}_0'(x, y, \hat{r})) \le t(f_0(x \oplus y)). \end{aligned} \tag{10}$$

[2] Note that this is different from the partial dead code elimination in [20], where partial dead code refers to code that is dead on some but not all computation paths.

If $foo(x)$ returns r, then $foo'(x,r)$ computes $foo(x+1)$.
For x of length n, $foo'(x,r)$ takes time $O(3^n)$;
$foo(x+1)$ takes time $O(3^n)$.

$foo(x) = 1st(\widehat{foo}(x))$.
For x of length n, $\widehat{foo}(x)$ takes time $O(3^n)$;
$foo(x)$ takes time $O(3^n)$.

If $\widehat{foo}(x)$ returns $\hat{r}$, then $\widehat{foo}'(x,\hat{r})$ computes $\widehat{foo}(x+1)$.
For x of length n, $\widehat{foo}'(x,\hat{r})$ takes time $O(1)$;
$\widehat{foo}(x+1)$ takes time $O(3^n)$.

$$
\begin{array}{lcl}
foo'(x,r) & = & \textbf{if } x \le 1 \textbf{ then } 1 \\
 & & \textbf{else if } x = 2 \textbf{ then } 3 \\
 & & \textbf{else } r + foo(x-1) + foo(x-2) \\
\widehat{foo}(x) & = & \textbf{if } x \le 2 \textbf{ then } < 1 > \\
 & & \textbf{else let } v_1 = \widehat{boo}(x) \textbf{ in} \\
 & & \quad < 1st(v_1) + foo(x-3),\ v_1 > \\
\widehat{boo}(x) & = & \textbf{let } v_1 = foo(x-1) \textbf{ in} \\
 & & < v_1 + foo(x-2),\ < v_1 >> \\
\widehat{foo}'(x,\hat{r}) & = & \textbf{if } x \le 1 \textbf{ then } < 1 > \\
 & & \textbf{else if } x = 2 \textbf{ then } < 3, < 2, < 1 >>> \\
 & & \textbf{else } < 1st(\hat{r}) + 1st(2nd(\hat{r})), \\
 & & \qquad < 1st(\hat{r}) + 1st(2nd(2nd(\hat{r}))),\ < 1st(\hat{r}) >>>
\end{array}
$$

Figure 2: Resulting "foo" function definitions

i.e, the functions $\hat{f}_0$ and $\hat{f}_0{}'$ preserve the semantics and compute asymptotically at least as fast. Note, however, that $\hat{f}_0(x)$ may terminate more often than $f_0(x)$ and $\hat{f}_0{}'(x,y,\hat{r})$ may terminate more often than $f_0(x \oplus y)$ due to the transformations used in Stages II and III.

4 Stage I: Caching all intermediate results

Stage I transforms the program for f_0 to embed all intermediate results in the final return value.

We first perform a local and structure-preserving transformation called *extension*. For each function definition $f(v_1, ..., v_n) = e$, we construct a function definition

$$\bar{f}(v_1, ..., v_n) = \mathcal{E}xt[\![e]\!] \tag{11}$$

where $\mathcal{E}xt[\![e]\!]$ extends an expression e to return the values of all function calls made in computing e, i.e., it considers subexpressions of e in applicative and left-to-right order, introduces bindings that name the results of function calls, builds up tuples of these values together with the values of the original subexpressions, and passes these values from subcomputations to enclosing computations.

The definition of $\mathcal{E}xt$ is given in Figure 3. We assume that each binding introduced uses a fresh variable name. For transforming a conditional expression, the transformation $\mathcal{G}us[\![e]\!]$ generates a tuple of $\_$'s of length equal to the number of the function applications in e, where $\_$ is a dummy constant that just occupies a spot. The lengths of the tuples generated by $\mathcal{G}us[\![e_2]\!]$ and $\mathcal{G}us[\![e_3]\!]$ can easily be determined statically. Actually, they are just the lengths of $rst(v_2)$ and $rst(v_3)$, respectively. This mechanism assures that the extended function returns a uniform tuple no matter what the value of the Boolean expression is, which makes the pruning stage simpler.

$\bar{f}(v_1, ..., v_n)$ and $f(v_1, ..., v_n)$ perform essentially the same computation, and thus take the same asymptotic time. In particular, they have the same termination behavior, and, if they terminate,

$$1st(\bar{f}(v_1, ..., v_n)) = f(v_1, ..., v_n). \tag{12}$$

We then perform administrative simplifications to the resulting function definitions: simplifying operations on tuples that are built up for passing intermediate results, unwinding binding expressions that become unnecessary as a result of simplifying their subexpressions, and lifting bindings out of enclosing expressions.

The result of this stage is a set of extended function definitions that embed the values of all function calls in the return values.

For the functions foo and boo in Figure 1, after the extension transformation and the administrative simplifications, we obtain the functions $\overline{foo}$ and $\overline{boo}$ as follows:

$$
\begin{array}{lcl}
\overline{foo}(x) & = & \textbf{if } x \le 2 \textbf{ then } < 1, \_, \_ > \\
 & & \textbf{else let } u_1 = \overline{boo}(x) \textbf{ in} \\
 & & \quad \textbf{let } u_2 = \overline{foo}(x-3) \textbf{ in} \\
 & & \quad < 1st(u_1) + 1st(u_2),\ u_1,\ u_2 > \\
\overline{boo}(x) & = & \textbf{let } u_1 = \overline{foo}(x-1) \textbf{ in} \\
 & & \textbf{let } u_2 = \overline{foo}(x-2) \textbf{ in} \\
 & & < 1st(u_1) + 1st(u_2),\ u_1,\ u_2 >
\end{array}
\tag{13}
$$

An obvious optimization can be incorporated into the extension transformation, i.e, we introduce bindings only for subexpressions that contain function applications. Thus, there would be fewer tuple operations for passing intermediate results and fewer bindings to be unwound or lifted, leaving less work for the administrative simplifications.

We could also make certain improvements to caching all intermediate results. First, we can omit caching values of function applications that are embedded in the values of their enclosing applications, since these omitted values can be retrieved from the results of the enclosing computations. Also, we can omit making an extended version for a function that does not contain function applications, and reference its return value directly rather than having to refer to the first component of the extended version.

The transformation $\mathcal{E}xt$ is similar to the construction of call-by-value complete recursive programs by Cartwright [6]. However, a call-by-value computation sequence returned by such a program is a flat list of all intermediate results, while our extended function returns a computation tree, a structure that mirrors the hierarchy of function calls. The transformations in this stage also mimic the CPS transformations in some aspects [36, 23]: sequencing subexpressions, naming intermediate results, passing the collected information,

$$
\begin{array}{lcl}
\mathcal{E}xt[\![v]\!] & = & <v> \\
\mathcal{E}xt[\![g(e_1,...,e_n)]\!] \quad \text{where } g \text{ is } c \text{ or } p & = & \textbf{let } v_1 = \mathcal{E}xt[\![e_1]\!] \textbf{ in } ... \textbf{ let } v_n = \mathcal{E}xt[\![e_n]\!] \textbf{ in} \\
 & & <g(1st(v_1),...,1st(v_n))> @rst(v_1)@...@rst(v_n) \\
\mathcal{E}xt[\![f(e_1,...,e_n)]\!] & = & \textbf{let } v_1 = \mathcal{E}xt[\![e_1]\!] \textbf{ in } ... \textbf{ let } v_n = \mathcal{E}xt[\![e_n]\!] \textbf{ in} \\
 & & \textbf{let } v = \bar{f}(1st(v_1),...,1st(v_n)) \textbf{ in} \\
 & & <1st(v)> @rst(v_1)@...@rst(v_n)@<v> \\
\mathcal{E}xt[\![\textbf{if } e_1 \textbf{ then } e_2 \textbf{ else } e_3]\!] & = & \textbf{let } v_1 = \mathcal{E}xt[\![e_1]\!] \textbf{ in} \\
 & & \textbf{if } 1st(v_1) \textbf{ then let } v_2 = \mathcal{E}xt[\![e_2]\!] \textbf{ in} \\
 & & \qquad <1st(v_2)> @rst(v_1)@rst(v_2)@\mathcal{G}us[\![e_3]\!] \\
 & & \quad \textbf{else let } v_3 = \mathcal{E}xt[\![e_3]\!] \textbf{ in} \\
 & & \qquad <1st(v_3)> @rst(v_1)@\mathcal{G}us[\![e_2]\!]@rst(v_3) \\
\mathcal{E}xt[\![\textbf{let } v = e_1 \textbf{ in } e_2]\!] & = & \textbf{let } v_1 = \mathcal{E}xt[\![e_1]\!] \textbf{ in} \\
 & & \textbf{let } v = 1st(v_1) \textbf{ in let } v_2 = \mathcal{E}xt[\![e_2]\!] \textbf{ in} \\
 & & \qquad <1st(v_2)> @rst(v_1)@rst(v_2)
\end{array}
$$

Figure 3: Definition of $\mathcal{E}xt$

and performing administrative reductions on the resulting program. However, they are simpler than the CPS transformations since the collected intermediate results are passed directly to the return values, rather than to continuation functions.

5 Stage II: Incrementalization

Stage II derives a function $\bar{f}_0'$, an incremental version of $\bar{f}_0$ under $\oplus$. Basically, one may identify subcomputations in the expanded $\bar{f}_0(x \oplus y)$ whose values can be retrieved from the cached result $\bar{r}$ of $\bar{f}_0(x)$, replace them by corresponding retrievals, and capture the resulting way of computing $\bar{f}_0(x \oplus y)$ in the incremental version $\bar{f}_0'(x, y, \bar{r})$. Such a derivation method is given in [27], and, depending on the power one expects from the derivation, the method can be made semi-automatic or fully-automatic.

This stage is not the subject of this paper. We give only the result of incrementalization using [27], namely, a function $\overline{foo}'$, an incremental version of $\overline{foo}$ in (13) under $\oplus_1$, such that, if $\overline{foo}(x) = \bar{r}$, then $\overline{foo}'(x, \bar{r}) = \overline{foo}(x+1)$.

$$
\begin{array}{l}
\overline{foo}'(x,\bar{r}) \\
\;= \textbf{if } x \le 1 \textbf{ then } <1,\_,\_> \\
\quad\;\; \textbf{else if } x = 2 \textbf{ then} \\
\qquad\quad <3,\ <2,<1,\_,\_>,<1,\_,\_>>,\ <1,\_,\_>> \\
\quad\;\; \textbf{else } <1st(\bar{r})+1st(2nd(\bar{r})), \\
\qquad\quad <1st(\bar{r})+1st(2nd(2nd(\bar{r}))),\ \bar{r},\ 2nd(2nd(\bar{r}))>, \\
\qquad\quad 3rd(2nd(\bar{r}))>
\end{array}
\tag{14}
$$

Clearly, $\overline{foo}'(x, \bar{r})$ computes $\overline{foo}(x+1)$ in only $O(1)$ time.

6 Stage III: Pruning

Stage III mainly analyzes the function $\bar{f}_0'$ to determine the components of $\bar{r}$ that $1st(\bar{f}_0'(x, y, \bar{r}))$ transitively depends on, as discussed in Section 3. We first depict the transitive dependencies and address a cost issue. Then we give an algorithm that computes the needed components based on a dependency analysis using domain projections [42, 12]. With this result, we prune the function $\bar{f}_0$ to return only the intermediate results that are useful for computing $1st(\bar{f}_0'(x, y, \bar{r}))$. In addition, we prune the function $\bar{f}_0'$ to incrementally maintain only the useful intermediate results.

6.1 Maintaining intermediate results: transitive dependency and cost

Transitive dependency. Figure 4 illustrates the transitive dependencies for the example foo under change $\oplus_1$. By definitions of foo and boo and associativity of '+', we have

$$
\begin{aligned}
foo(x+1) &= boo(x+1) + foo(x-2) \\
&= \big(foo(x)+foo(x-1)\big) + foo(x-2) \\
&= foo(x) + \big(foo(x-1)+foo(x-2)\big) \\
&= foo(x) + boo(x)
\end{aligned}
$$

Thus, to compute the value v_1' of $foo(x+1)$, $\overline{foo}'$ uses the value v_1 of $foo(x)$ and the intermediate result v_2 of $boo(x)$ returned by $\overline{foo}(x)$. Therefore, the corresponding value v_1' of $foo(x+1)$ and the intermediate result v_2' of $boo(x+1)$ need to be maintained. The value v_1' of $foo(x+1)$ has just been considered. To compute the intermediate result v_2' of $boo(x+1)$, $\overline{foo}'$ uses the value v_1 of $foo(x)$ and the intermediate result v_3 of $foo(x-1)$ returned by $\overline{foo}(x)$. Therefore, the corresponding value v_1' of $foo(x+1)$ and the intermediate result v_3' of $foo(x)$ also need to be maintained. Again, the value v_1' of $foo(x+1)$ has just been considered. To compute the value v_3' of $foo(x)$, $\overline{foo}'$ just uses the value v_1 of $foo(x)$ returned by $\overline{foo}(x)$.

Thus, to summarize, the value v_1' of $foo(x+1)$ transitively depends on the components of intermediate results corresponding to v_1, v_2, and v_3, which are maintained as $v_1' = v_1 + v_2$, $v_2' = v_1 + v_3$, and $v_3' = v_1$, respectively. Other components of intermediate results are not needed and therefore do not need to be computed or maintained; they can be be pruned out.

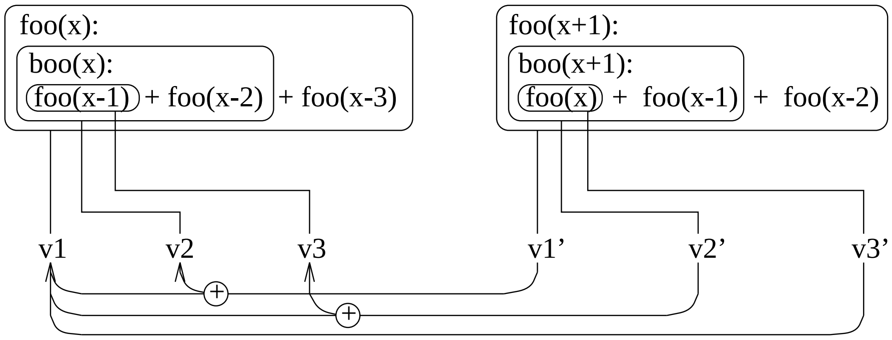

Figure 4: Dependencies for the function foo

Cost. Is it always true that the time spent in maintaining intermediate results will not surpass that saved by using them?

First, we consider the problem in general. Given a way of computing a function f, let g be a function that computes some intermediate results of f in the way f does, and let $\hat{f} = \langle f, g \rangle$. Suppose $f'(x, y, r)$ computes $f(x \oplus y)$ given $r = f(x)$, and $\hat{f}'(x, y, r)$ computes $\hat{f}(x \oplus y)$ given $\hat{r} = \hat{f}(x)$. Then in general, it is not true that $t(\hat{f}'(x, y, \hat{r})) \leq t(f'(x, y, r))$. This is mainly because f, and thus g, could be arbitrary.

But, consider the particular functions f' and $\hat{f}'$ derived using the derivation approach. Suppose we compute $\hat{f}(x \oplus y)$ using the cached result $\hat{r}$ of $\hat{f}(x)$, and suppose computing $1st(\hat{f}(x \oplus y))$, i.e., $f(x \oplus y)$, uses any common subcomputation, say $g(x)$, of $f(x \oplus y)$ and $f(x)$, and the value of $g(x)$ can be retrieved from $\hat{r}$ but not r. Then, on the one hand, the value of $g(x \oplus y)$ needs to be maintained by $\hat{f}'(x, y, \hat{r})$; on the other hand, if we compute $f(x \oplus y)$ using only the cached result r of $f(x)$, then $f'(x, y, r)$ has the cost of recomputing $g(x)$.

Now, for intermediate results of f like the value of g above, if (a) the size of y is bounded, (b) when the size of y is bounded, the time of computing $x \oplus y$ is bounded, and (c) g is at most *linear-power exponential* time, i.e., g is polynomial time or exponential time but with linear exponent, then we have

$$t(\hat{f}'(x, y, \hat{r})) \leq t(f'(x, y, r)). \tag{15}$$

It is easy to see that the three conditions are true for all practical and feasible incremental applications, and therefore, we assume that they are satisfied. To prove (15), we notice that

$$\begin{array}{ll} t(\hat{f}'(x, y, \hat{r})) & \\ \leq t(f'(x, y, r)) + t(x \oplus y) + t(g(x')) & \text{by def. of } \hat{f}' \text{ and derivation} \\ \leq t(f'(x, y, r)) + t(g(x)) & \text{by conditions on } y, \oplus, \text{and } g \\ \leq t(f'(x, y, r)) + t(f'(x, y, r)) & \text{by } g \text{ being subcomp. of } f' \\ \leq t(f'(x, y, r)) & \text{by definition of } t \end{array}$$

We conclude that, with the conditions above, using and maintaining intermediate results is always asymptotically at least as fast. Therefore, in order to achieve as fast incremental computation as possible, we should compute the closure of the transitive dependencies for maintaining intermediate results.

6.2 Dependency analysis using projections

We first describe our use of projections to represent components of the tuple values constructed in Stage I and manipulated by Stage II. Then, we give a backward dependency analysis that determines which components of $\bar{r}$ are needed for computing certain components of $\bar{f_0}'(x, y, \bar{r})$. Finally, we present an algorithm that computes the closure of the transitive dependencies for maintaining intermediate results.

Projections. Our domain of interest D contains $\perp$, indicating a computation diverges, values d returned by functions in the original program for f_0, and constructed tuples $< d_1, ..., d_n >$, where each d_i is (recursively) an element of D (other than $\perp$). The length of a constructed tuple is statically bounded, but the depth of tuple nesting may not be bounded, since it is dynamically determined. Intuitively, any components of a constructed tuple value can be replaced by the dummy constant $\_$, introduced in Stage I, if we do not care about the values of those components. If a subcomputation involves $\_$, then the result of that subcomputation is $\_$, but the result of the parent computation need not be $\_$. For any value d in domain D, $\perp \sqsubseteq d$. For two values d_1 and d_2 other than $\perp$'s in D, $d_1 \sqsubseteq d_2$ iff

$$\begin{array}{l} d_1 = \_, \quad d_1 = d_2, \quad \text{or} \\ d_1 = < d_{11}, ..., d_{1n} >, \; d_2 = < d_{21}, ..., d_{2n} >, \\ \qquad \text{and } d_{1i} \sqsubseteq d_{2i} \text{ for } i = 1..n. \end{array}$$

A projection over the domain D is a function $\Pi : D \to D$ such that $\Pi(\Pi(d)) = \Pi(d) \sqsubseteq d$ for any $d \in D$. Three important projections are ID, ABS, and BOT. ID is the identity function $ID(d) = d$. ABS is the function $ABS(d) = \_$ for any $d \neq \perp$. BOT is the function $BOT(d) = \perp$.

A non-bottom projection Π of interest here can be represented as a set of selection functions π, each of which is a sequence of 1^{st}, 2^{nd}, ..., n^{th}. The null sequence is denoted ϵ. Intuitively, if Π contains a sequence $i_k{}^{th} i_{k-1}{}^{th} ... i_1{}^{th}$, then the i_kth element of the i_{k-1}th element of the $\cdots$ of the i_1th element of Π's argument is selected, and if Π contains ϵ, then all components of Π's argument are selected. A projection Π replaces those components of its argument that are not selected with the constant $\_$. For example $\{1^{st}\}$, $\{1^{st}, 1^{st}2^{nd}\}$, and $\{1^{st}1^{st}2^{nd}, \epsilon\}$ are projections, and

$$\begin{array}{l} \{1^{st}, 1^{st}2^{nd}\}(< d_1, << d_{211}, d_{212} >, d_{22} >>) \\ \qquad = \; < d_1, << d_{211}, d_{212} >, \_ >> \end{array}$$

For convenience of presentation, we use $\Pi_{(i)}$ to denote the set $\{\pi \mid \pi\, i^{th} \in \Pi\}$, i.e., $\Pi_{(i)}$ is the part of Π that considers the ith component. With the set representation, a projection $\Pi = ID$ iff $\epsilon \in \Pi$ or $\Pi_{(i)} = ID$ for $i = 1..n$ for arguments of Π of length n. A projection $\Pi = ABS$ iff $\Pi = \emptyset$. For $\Pi \notin \{ID, ABS\}$, $\Pi(< d_1, ..., d_n >) = < \Pi_{(1)}(d_1), ..., \Pi_{(n)}(d_n) >$. For any two projections Π_1 and Π_2 other than BOT's, $\Pi_1 \sqsubseteq \Pi_2$ iff

$$\begin{array}{l} \Pi_1 = ABS, \quad \Pi_2 = ID, \quad \text{or} \\ \Pi_{1(i)} \sqsubseteq \Pi_{2(i)} \text{ for } i = 1..n \\ \qquad\qquad \text{for arguments of } \Pi_1 \text{ and } \Pi_2 \text{ of length } n. \end{array}$$

For any projection Π, $BOT \sqsubseteq \Pi$.

Dependency analysis. To compute which components of $\bar{r}$ are needed for computing certain components of $\bar{f_0}'(x, y, \bar{r})$, we apply a backward dependency analysis to the program for $\bar{f_0}'$.

Following the style of [46], for each function f of n parameters, and each i from 1 to n, we define f^i to be a *dependency transformer* that takes a projection that is applied to the result of f and returns a projection that is *sufficient* to be applied to the ith parameter. The *sufficiency condition* that f^i must satisfy is: if $\Pi_i = f^i\Pi$ then

$$\Pi(f(v_1, ..., v_i, ..., v_n)) \sqsubseteq f(v_1, ..., \Pi_i(v_i), ..., v_n) \qquad (16)$$

Similarly, we define e^v to be a dependency transformer that takes a projection that is applied to e and returns a projection that is sufficient to be applied to every instance of v in e. A similar sufficiency condition must be satisfied: if $\Pi' = e^v\Pi$ then

$$\Pi(e) \sqsubseteq e[\Pi'(v)/v] \qquad (17)$$

For a function f whose definition is $f(v_1, ..., v_n) = e$, we define $f^i\Pi = e^{v_i}\Pi$. The definition of e^v may in turn refer to f^i, thus the definitions may be mutually recursive. We define

$$e^v BOT = BOT \qquad \text{and} \qquad e^v ABS = ABS. \qquad (18)$$

For $\Pi \neq BOT, ABS$, we give the definition of $e^v\Pi$ in Figure 5. We can easily show that each rule guarantees sufficient information. Thus, the sufficiency conditions are satisfied by recursion induction.

Let $i_{\bar{r}}$ be the index of $\bar{r}$ in the parameters of $\bar{f_0}'$. With the above definitions, we know that $\bar{f_0}'^{i_{\bar{r}}}\Pi$ computes how much of $\bar{r}$ is needed when Π of $\bar{f_0}'(x, y, \bar{r})$ is needed.

To compute $f^i\Pi$ for some f^i and Π, if the definition of f^i does not involve recursion, then we can compute directly using the definition. If the definition of f^i involves recursion, then the argument projections and resulting projections of some dependency transformers may contain selection functions of unbounded depth. To approximate the result, we restrict the selection functions of the projections to be of bounded depth d, namely, if a projection contains a selection function $i_k{}^{th} i_{k-1}^{th} ... i_1{}^{th}$ but $k > d$, then we truncate it to $i_d{}^{th} i_{d-1}^{th} ... i_1{}^{th}$. A simple choice for the depth bound would be 1. A more prudent choice could be the length of the longest cycle that contains f in the call graph. This limits the domain of projections to be finite. Now, to solve the recursive definitions of these dependency transformers, we just compute the limits of the ascending chains by starting at $f^i\Pi = ABS$ for every f^i and Π and iterating using the definitions. This always terminates since the ascending chains are finite.

Our backward dependency analysis uses domain projections to specify sufficient information, which is natural and thus simple. Other uses of projections include the strictness analysis by Wadler and Hughes [46], where necessary information needs to be specified and thus accounts for some complications, and the binding time analysis by Launchbury [21], which is a forward analysis and is proved equivalent to strictness analysis [22]. The necessity interpretation by Jones and Le Métayer [17] is in the same spirit of our analysis, where their notion of necessity patterns correspond to our notion of projections. While necessity patterns specify heads and tails of list values, our projections specify specific components of tuple values and thus provide more accurate information.

Computing the closure of the transitive dependencies. To compute the components of $\bar{r}$ on which $1st(\bar{f_0}'(x, y, \bar{r}))$ transitively depends, we start with Π being $\{1^{st}\}$ and compute the smallest projection Π of $\bar{r}$ on which $\Pi(\bar{f_0}'(x, y, \bar{r}))$ depends, i.e., the smallest projection Π such that

$$\{1^{st}\} \sqsubseteq \Pi \ \text{ and } \ \Pi(\bar{f_0}'(x, y, \bar{r})) \sqsubseteq \bar{f_0}'(x, y, \Pi(\bar{r})). \qquad (19)$$

Of course, the projection $\Pi = ID$ is always a solution to (19). But our goal is to make Π as small as possible, and thus to avoid as much unnecessary caching as possible.

Since $\bar{f_0}'^{i_{\bar{r}}}\Pi$ computes the components of $\bar{r}$ on which $\Pi(\bar{f_0}'(x, y, \bar{r}))$ depends, we define

$$\begin{array}{rcl} \Pi^{(0)} & = & \{1^{st}\} \\ \Pi^{(i+1)} & = & \Pi^{(i)} \ \cup \ \bar{f_0}'^{i_{\bar{r}}}\Pi^{(i)} \end{array} \qquad (20)$$

and compute the least fixed point of Π. In other words, Π is the least projection that satisfies $\{1^{st}\} \sqsubseteq \Pi$ and $\bar{f_0}'^{i_{\bar{r}}}\Pi \sqsubseteq \Pi$. We call this projection the *closure projection*. Note that the above computation always terminates since $\bar{f_0}'^{i_{\bar{r}}}\Pi^{(i)}$ terminates and returns only sets of selection functions of bounded depth.

The time complexity of the closure computation depends on the required size of the projection domain and the complexity of the dependency analysis. Suppose d is the maximum depth of selection functions we consider, and l is the maximum length of the constructed tuples, i.e., the largest number of function applications in a function definition in the program for f_0. Then the maximum number c of disjoint components in these projections is at most l^d, which characterizes the maximum size of the projection domain.

We estimate the complexity of the dependency analysis in the simplest manner. Consider the program for $\bar{f_0}'$. Let n be the number of function definitions, and a be the maximum number of parameters in any of these definitions. Then there are at most na dependency transformers. Since an argument projection may contain any of c components, there are at most 2^c argument projections to each transformer. Thus, the number of projections $f^i\Pi$ to be computed is at most $na2^c$. Now, let s_f be the maximum number of transformers used in a transformer definition, i.e., the number of function applications in a function definition. Being careful, we can recompute each $f^i\Pi$ only when any computed projections used by $f^i\Pi$ changes, where each can change at most c times. Thus, the total number of computations of $f^i\Pi$ using its immediate definition is at most $na2^c cs_f$. Each such computation takes at most sc time, where s is the maximum

$$
\begin{array}{lcll}
v^v\,\Pi & = & \Pi & \\
u^v\,\Pi & = & ABS & \text{if } v \neq u \\
(<e_1,...,e_n>)^v\,\Pi & = & e_1^v\,\Pi_{(1)} \;\cup\; ... \;\cup\; e_n^v\,\Pi_{(n)} & \\
(ith(e))^v\,\Pi & = & e^v\{\pi\, i^{th} \mid \pi \in \Pi\} & \\
(g(e_1,...,e_n))^v\,\Pi & = & e_1^v\; ID \;\cup ... \cup\; e_n^v ID & \text{if } g \text{ is } c \text{ or } p \text{ but not } <> \text{ or } ith \\
(f(e_1,...,e_n))^v\,\Pi & = & e_1^v(f^1\,\Pi) \;\cup\; ... \;\cup\; e_n^v(f^n\,\Pi) & \\
(\textbf{if } e_1 \textbf{ then } e_2 \textbf{ else } e_3)^v\,\Pi & = & e_1^v\; ID \;\cup e_2^v\,\Pi \;\cup\; e_3^v\,\Pi & \\
(\textbf{let } u = e_1 \textbf{ in } e_2)^v\,\Pi & = & e_1^v(e_2^u\,\Pi) \;\cup\; e_2^v\,\Pi &
\end{array}
$$

Figure 5: Definition of $e^v\,\Pi$ for $\Pi \neq BOT, ABS$

size of a function definition, i.e., the number of subexpressions in the defining expression, and c is the time needed to compute operations, such as union, on two projections. Therefore, the total time is at most $na2^c c^2 s s_f$.

If we limit depth of selection functions to be independent of the number of function definitions, then a, c, s, and s_f are all constant factors determined by the size of a function definition. Thus the total time is linear in the number of function definitions, although the constant factors could be very big.

Now that the above estimate includes the computations of all $f^i\Pi$, computing the dependency closure takes at most c projection unions, each taking at most c time. Thus, the total time of closure computation can be no worse than the above bound.

Example. Applying the dependency analysis to the function $\overline{foo}'$ in (14), we get

$$
\begin{array}{l}
\overline{foo}'^2\,\Pi \\
\;=\; (x \leq 1)^{\bar r}\; ID \;\cup (<1,\_,\_>)^{\bar r}\Pi \;\cup \\
\qquad (x = 2)^{\bar r}\; ID \;\cup \\
\qquad\quad (<3,<2,<1,\_,\_>,<1,\_,\_>>,<1,\_,\_>>)^{\bar r}\Pi \;\cup \\
\qquad (<1st(\bar r)+1st(2nd(\bar r)), \\
\qquad\quad <1st(\bar r)+1st(2nd(2nd(\bar r))),\; \bar r,\; 2nd(2nd(\bar r))>, \\
\qquad\quad 3rd(2nd(\bar r))>)^{\bar r}\Pi \\
\;=\; (1st(\bar r)+1st(2nd(\bar r)))^{\bar r}\Pi_{(1)} \;\cup \\
\qquad (1st(\bar r)+1st(2nd(2nd(\bar r))))^{\bar r}\Pi_{(2)(1)} \;\cup \\
\qquad\quad (\bar r)^{\bar r}\Pi_{(2)(2)} \;\cup\; (2nd(2nd(\bar r)))^{\bar r}\Pi_{(2)(3)} \;\cup \\
\qquad (3rd(2nd(\bar r)))^{\bar r}\Pi_{(3)}
\end{array}
$$

For this example, since the definition of $\overline{foo}'^2$ is not recursive, we can compute $\overline{foo}'^2\,\Pi$ for a given Π directly without iteration and approximation. For example,

$$
\begin{array}{lcl}
\overline{foo}'^2\{1^{st}\} & = & \{1^{st}, 1^{st}2^{nd}\} \\
\overline{foo}'^2\{1^{st}2^{nd}\} & = & \{1^{st}, 1^{st}2^{nd}2^{nd}\} \\
\overline{foo}'^2\{1^{st}2^{nd}2^{nd}\} & = & \{1^{st}\}
\end{array}
$$

which illustrates the dependencies depicted in Figure 4. Now, we compute the closure projection:

$$
\begin{array}{lclcl}
\Pi^{(1)} & = & \Pi^{(0)} \cup \overline{foo}'^2\,\Pi^{(0)} & = & \{1^{st}, 1^{st}2^{nd}\} \\
\Pi^{(2)} & = & \Pi^{(1)} \cup \overline{foo}'^2\,\Pi^{(1)} & = & \{1^{st}, 1^{st}2^{nd}, 1^{st}2^{nd}2^{nd}\} \\
\Pi^{(3)} & = & \Pi^{(2)} \cup \overline{foo}'^2\,\Pi^{(2)} & = & \{1^{st}, 1^{st}2^{nd}, 1^{st}2^{nd}2^{nd}\}
\end{array}
$$

We obtain the projection $\{1^{st}, 1^{st}2^{nd}, 1^{st}2^{nd}2^{nd}\}$.

6.3 Pruning under the closure projection

We have obtained a closure projection Π such that $1^{st} \in \Pi$ and $\Pi(\bar f_0'(x,y,\bar r)) \sqsubseteq \bar f_0'(x,y,\Pi(\bar r))$. Now, we prune the extended function $\bar f_0$ to get a function $\hat f_0$ such that $\Pi(\bar f_0(x)) \sqsubseteq \hat f_0(x)$, and prune the incremental function $\bar f_0'$ to get a function $\hat f_0'$ such that $\Pi(\bar f_0'(x,y,\bar r)) \sqsubseteq \hat f_0'(x,y,\Pi(\bar r))$. Setting $\hat f_0$ to be $\bar f_0$ and $\hat f_0'$ to be $\bar f_0'$ would always work, but we want to make $\hat f_0(x)$ as close to $\Pi(\bar f_0(x))$, and $\hat f_0'(x,y,\Pi(\bar r))$ as close to $\Pi(\bar f_0'(x,y,\bar r))$ as possible, and thereby avoid caching and maintaining unnecessary intermediate results as much as possible.

To do this, for each expression e that defines a function $f(v_1,...v_n)$, we associate a projection with each subexpression of e indicating how much of the subexpression is needed assuming Π of $\bar f_0$ (respectively $\bar f_0'$) is needed. The definition and computation of the associated projections can be done in a fashion similar to the dependency analysis. For the program for $\bar f_0'$ and the closure projection Π, the final projection computed associated with each variable will be the same as computed for the variable using dependency analysis.

When the computation reaches the limit of the ascending chain of projections, subexpressions associated with *ID* are left unchanged in the resulting function, and subexpressions associated with *ABS* are replaced by $\_$. If a variable whose value is a constructed tuple is associated with a projection Π other than *ID* or *ABS*, then we construct a tuple with the components selected by Π filled with the corresponding selections and the rest filled with $\_$. For example, if a variable v is associated with a projection $\{1^{st}, 1^{st}2^{nd}\}$, and v represents a tuple of length three whose second component is a tuple of length two, then v is replaced by $<1st(v), <1st(2nd(v)), \_>, \_>$.

As the result of such replacements, we have $\Pi(\bar f_0(x)) \sqsubseteq \hat f_0(x)$, but not $\hat f_0(x) = \Pi(\bar f_0(x))$ as anticipated in Section 3.

Nevertheless, the resulting $\hat{f}_0$ is still good enough to guarantee (9). We can just project $\Pi(\bar{r})$ out of the return value of $\hat{f}_0(x)$. But we do have $\hat{f}_0'(x, y, \Pi(\bar{r})) = \Pi(\bar{f}_0'(x, y, \bar{r}))$. Thus, assuming $\hat{r} = \Pi(\bar{r})$, we have (10). As a matter of fact, we intend to use the function $\hat{f}_0$ only once to get the initial value, and then use the function $\hat{f}_0'$ repeatedly to compute all successive values. Recall that $\hat{f}_0'$ incrementally computes the desired output and the corresponding intermediate results, as shown in (10).

Consider the functions $\overline{foo}$ and $\overline{boo}$ in (13) and $\overline{foo}'$ in (14). After the replacements as above, we obtain

$$\begin{array}{ll}
\widehat{foo}_1(x) = & \textbf{if } x \le 2 \textbf{ then } <1, _, _> \\
& \textbf{else let } u_1 = \widehat{boo}_1(x) \textbf{ in} \\
& \quad \textbf{let } u_2 = \widehat{foo}_1(x-3) \textbf{ in} \\
& \quad <1st(u_1)+1st(u_2), \\
& \quad\ <1st(u_1),\ <1st(2nd(u_1)), _, _>,\ _>, \\
& \quad\ _> \\
\widehat{boo}_1(x) = & \textbf{let } u_1 = \widehat{foo}_1(x-1) \textbf{ in} \\
& \textbf{let } u_2 = \widehat{foo}_1(x-2) \textbf{ in} \\
& <1st(u_1)+1st(u_2),\ <1st(u_1), _, _>,\ _>
\end{array}$$

and

$$\begin{array}{l}
\widehat{foo}_1'(x, \hat{r}_1) \\
= \ \textbf{if } x \le 1 \textbf{ then } <1, _, _> \\
\quad \textbf{else if } x = 2 \textbf{ then} \\
\qquad <3,\ <2, <1, _, _>, <_, _, _>>,\ <_, _, _>> \\
\quad \textbf{else } <1st(\hat{r}_1)+1st(2nd(\hat{r}_1)), \\
\qquad <1st(\hat{r}_1)+1st(2nd(2nd(\hat{r}_1))), <1st(\hat{r}_1), _, _>, _>, \\
\qquad _>
\end{array}$$

such that, if $\widehat{foo}_1(x) = \hat{r}_1$, then $\widehat{foo}_1'(x, \hat{r}_1) = \widehat{foo}_1(x+1)$.

A number of simplifications can be made to the resulting functions: (a) unfolding a **let** expression if a binding variable occurs at most once in the body due to some replacements by _'s, (b) combining unnecessarily split components resulting from some replacements for variables whose values are constructed tuples, (c) lifting common selection computations to avoid unnecessarily computing a compound value and using only part of it, and (d) replacing occurrences of $1st(\hat{f}(e_1, ..., e_n))$ by occurrences of $f(e_1, ..., e_n)$.

Furthermore, we can eliminate _ components. But we must be careful if such a component precedes a non-_ component in a tuple. In particular, if k of the components preceding a component i are eliminated from a tuple, we must replace all uses of the selector ith for the tuple with $(i-k)$th. This elimination needs to be done consistently for $\hat{f}_0$ and $\hat{f}_0'$.

These simplifications and eliminations can be fully automated. At the end, we obtain the final functions $\widehat{foo}$, $\widehat{boo}$, and $\widehat{foo}'$ given in Figure 2.

7 Examples

Fibonacci function. We can apply the cache-and-prune method to the definition of the Fibonacci function fib given in Figure 1 and input change $\oplus_1$. The derivation is similar but much simpler than in the case of foo, yielding

$$\begin{array}{ll}
\widehat{fib}(x) = & \textbf{if } x \le 1 \textbf{ then } <1> \\
& \textbf{else let } u_1 = fib(x-1) \textbf{ in} \\
& \qquad <u_1 + fib(x-2),\ u_1>
\end{array} \tag{21}$$

$$\begin{array}{ll}
\widehat{fib}'(x, \hat{r}) = & \textbf{if } x \le 0 \textbf{ then } <1> \\
& \textbf{else if } x = 1 \textbf{ then } <2, 1> \\
& \textbf{else } <1st(\hat{r}) + 2nd(\hat{r}),\ 1st(\hat{r})>
\end{array} \tag{22}$$

Clearly, $\widehat{fib}'(x, \hat{r})$ takes only $O(1)$ time. Note that $fib(x) = 1st(\widehat{fib}(x))$ and, if $\widehat{fib}(x) = \hat{r}$, then $\widehat{fib}'(x, \hat{r}) = \widehat{fib}(x+1)$. Using the definition of $\widehat{fib}'$ above in this last equation, we obtain a new definition for $\widehat{fib}$:

$$\begin{array}{ll}
\widehat{fib}(x+1) = & \textbf{if } x \le 0 \textbf{ then } <1> \\
& \textbf{else if } x = 1 \textbf{ then } <2, 1> \\
& \textbf{else let } \hat{r} = \widehat{fib}(x) \textbf{ in} \\
& \qquad <1st(\hat{r})+2nd(\hat{r}),\ 1st(\hat{r})>
\end{array}$$

Letting $v = x + 1$, we get

$$\begin{array}{ll}
\widehat{fib}(v) = & \textbf{if } v \le 1 \textbf{ then } <1> \\
& \textbf{else if } v = 2 \textbf{ then } <2, 1> \\
& \textbf{else let } \hat{r} = \widehat{fib}(v-1) \textbf{ in} \\
& \qquad <1st(\hat{r})+2nd(\hat{r}),\ 1st(\hat{r})>
\end{array} \tag{23}$$

Finally, we define $fib(v) = 1st(\widehat{fib}(v))$ using the definition of $\widehat{fib}$ in (23). Clearly, this computes the Fibonacci function in linear time, as desired.

Merge sort. The definition of the merge sort function $sort$ is as given in Figure 1. We consider the input change $x \oplus_2 y = cons(y, x)$.

Stage I. We cache all intermediate results of $sort$ and obtain

$$\begin{array}{ll}
\overline{sort}(x) = & \textbf{if } null(x) \textbf{ then} \\
& \quad <nil, _, _, _, _, _> \\
& \textbf{else if } null(cdr(x)) \textbf{ then} \\
& \quad <x, _, _, _, _, _> \\
& \textbf{else let } v_{11} = \overline{odd}(x) \textbf{ in} \\
& \quad \textbf{let } u_1 = \overline{sort}(1st(v_{11})) \textbf{ in} \\
& \quad \textbf{let } v_{21} = \overline{even}(x) \textbf{ in} \\
& \quad \textbf{let } u_2 = \overline{sort}(1st(v_{21})) \textbf{ in} \\
& \quad \textbf{let } v = \overline{merge}(1st(u_1), 1st(u_2)) \textbf{ in} \\
& \quad <1st(v), v_{11}, u_1, v_{21}, u_2, v>
\end{array} \tag{24}$$

where the definitions of $\overline{odd}$, $\overline{even}$, and $\overline{merge}$ are straightforward, become irrelevant, and are thus omitted here.

Stage II. We incrementalize $\overline{sort}$ under $\oplus_2$ and obtain

$$\begin{array}{l}
\overline{sort}'(y, \bar{r}) \\
= \ \textbf{if } null(1st(\bar{r})) \textbf{ then} \\
\qquad <cons(y, nil), _, _, _, _, _> \\
\quad \textbf{else if } null(cdr(1st(\bar{r}))) \textbf{ then} \\
\qquad \textbf{let } v_{11} = <cons(y, nil),\ <nil,\ < nil >>> \textbf{ in} \\
\qquad \textbf{let } v_{21} = <1st(\bar{r}),\ <1st(\bar{r}),\ < nil >>> \textbf{ in} \\
\qquad \textbf{let } v = \overline{merge}(cons(y, nil), 1st(\bar{r})) \textbf{ in} \\
\qquad <1st(v), v_{11}, <cons(y, nil)>, v_{21}, <1st(\bar{r})>, v> \\
\quad \textbf{else let } v_1 = 4th(\bar{r}) \textbf{ in} \\
\qquad \textbf{let } v_{11} = <cons(y, 1st(v_1)), v_1> \textbf{ in} \\
\qquad \textbf{let } u_1 = \overline{sort}'(y, 5th(\bar{r})) \textbf{ in} \\
\qquad \textbf{let } v_2 = 2nd(\bar{r}) \textbf{ in} \\
\qquad \textbf{let } v_{21} = <1st(v_2), v_2> \textbf{ in} \\
\qquad \textbf{let } u_2 = 3rd(\bar{r}) \textbf{ in} \\
\qquad \textbf{let } v = \overline{merge}(1st(u_1), 1st(u_2)) \textbf{ in} \\
\qquad <1st(v), v_{11}, u_1, v_{21}, u_2, v>
\end{array} \tag{25}$$

Stage III. First, using dependency analysis, for $\Pi \neq ABS$, we have

$$
\begin{array}{l}
\overline{sort}'^{2}\,\Pi = \\
\quad (null(1st(\bar{r})))^{\bar{r}}\ ID\ \cup\ ABS\ \cup \\
\quad (null(cdr(1st(\bar{r}))))^{\bar{r}}\ ID\ \cup\ (1st(\bar{r}))^{\bar{r}}(\overline{merge}^{2}\{1^{st}\})\ \cup \\
\quad (4th(\bar{r}))^{\bar{r}}((1st(v_1))^{v_1}\Pi_{(2)(1)}\ \cup\ \Pi_{(2)(2)})\ \cup \\
\quad (5th(\bar{r}))^{\bar{r}}(\overline{sort}'^{2}((1st(u_1))^{u_1}(\overline{merge}^{1}((1st(v))^{v}\Pi_{(1)} \cup \Pi_{(6)})) \cup \Pi_{(3)}))\ \cup \\
\quad (2nd(\bar{r}))^{\bar{r}}((1st(v_2))^{v_2}\Pi_{(4)(1)}\ \cup\ \Pi_{(4)(2)})\ \cup \\
\quad (3rd(\bar{r}))^{\bar{r}}((1st(u_2))^{u_2}(\overline{merge}^{2}((1st(v))^{v}\Pi_{(1)}\ \cup\ \Pi_{(6)}))\ \cup\ \Pi_{(5)})
\end{array}
$$

which is recursively defined, and can be simplified for $1^{st} \in \Pi$, yielding

$$
\begin{array}{lll}
\overline{sort}'^{2(0)}\,\Pi & = & ABS \\
\overline{sort}'^{2(i+1)}\,\Pi & = & \{1^{st}\}\ \cup \\
 & & (4th(\bar{r}))^{\bar{r}}((1st(v_1))^{v_1}\Pi_{(2)(1)}\ \cup\ \Pi_{(2)(2)})\ \cup \\
 & & (5th(\bar{r}))^{\bar{r}}(\overline{sort}'^{2(i)}(\{1^{st}\}\ \cup\ \Pi_{(3)}))\ \cup \\
 & & (2nd(\bar{r}))^{\bar{r}}((1st(v_2))^{v_2}\Pi_{(4)(1)}\ \cup\ \Pi_{(4)(2)})\ \cup \\
 & & (3rd(\bar{r}))^{\bar{r}}(\{1^{st}\}\ \cup\ \Pi_{(5)})
\end{array}
$$

Limiting the depth of selection functions to 1, we compute and obtain

$$
\begin{array}{l}
\Pi^{(0)} = \{1^{st}\} \\
\Pi^{(1)} = \{1^{st}, 3^{rd}, 5^{th}\} \\
\Pi^{(2)} = \{1^{st}, 3^{rd}, 5^{th}\}
\end{array}
$$

Thus, we get the closure projection $\{1^{st}, 3^{rd}, 5^{th}\}$. Then, we prune the functions $\overline{sort}$ and $\overline{sort}'$ and obtain

$$
\begin{array}{l}
\widehat{sort}_1(x) = \textbf{if } null(x) \textbf{ then} \\
\qquad\qquad < nil, \_, \_, \_, \_, \_ > \\
\qquad\quad \textbf{else if } null(cdr(x)) \textbf{ then} \\
\qquad\qquad < x, \_, \_, \_, \_, \_ > \\
\qquad\quad \textbf{else let } u_1 = \widehat{sort}(odd(x)) \textbf{ in} \\
\qquad\qquad \textbf{let } u_2 = \widehat{sort}(even(x)) \textbf{ in} \\
\qquad\qquad < merge(1st(u_1), 1st(u_2)), \_, u_1, \_, u_2, \_ >
\end{array}
$$

$$
\begin{array}{l}
\widehat{sort}'_1(y, \hat{r}_1) = \textbf{if } null(1st(\hat{r}_1)) \textbf{ then} \\
\qquad\qquad < cons(y, nil), \_, \_, \_, \_, \_ > \\
\qquad\quad \textbf{else if } null(cdr(1st(\hat{r}_1))) \textbf{ then} \\
\qquad\qquad < merge(cons(y, nil), 1st(\hat{r}_1)), \\
\qquad\qquad\quad \_, < cons(y, nil) >, \_, < 1st(\hat{r}_1) >, \_ > \\
\qquad\quad \textbf{else let } u_1 = \widehat{sort}'(y, 5th(\hat{r}_1)) \textbf{ in} \\
\qquad\qquad \textbf{let } u_2 = 3rd(\hat{r}_1) \textbf{ in} \\
\qquad\qquad < merge(1st(u_1), 1st(u_2)), \_, u_1, \_, u_2, \_ >
\end{array}
$$

Finally, we eliminate $\_$ components, adjust the indexing, and obtain

$$
\begin{array}{l}
\widehat{sort}(x) = \textbf{if } null(x) \textbf{ then } < nil > \\
\qquad\quad \textbf{else if } null(cdr(x)) \textbf{ then } < x > \\
\qquad\quad \textbf{else let } u_1 = \widehat{sort}(odd(x)) \textbf{ in} \\
\qquad\qquad \textbf{let } u_2 = \widehat{sort}(even(x)) \textbf{ in} \\
\qquad\qquad < merge(1st(u_1), 1st(u_2)),\ u_1,\ u_2 >
\end{array}
\tag{26}
$$

$$
\begin{array}{l}
\widehat{sort}'(y, \hat{r}) = \textbf{if } null(1st(\hat{r})) \textbf{ then } < cons(y, nil) > \\
\qquad\quad \textbf{else if } null(cdr(1st(\hat{r}))) \textbf{ then} \\
\qquad\qquad < merge(cons(y, nil), 1st(\hat{r})), \\
\qquad\qquad\quad < cons(y, nil) >,\ < 1st(\hat{r}) >> \\
\qquad\quad \textbf{else let } u_1 = \widehat{sort}'(y, 3rd(\hat{r})) \textbf{ in} \\
\qquad\qquad \textbf{let } u_2 = 2nd(\hat{r}) \textbf{ in} \\
\qquad\qquad < merge(1st(u_1), 1st(u_2)),\ u_1,\ u_2 >
\end{array}
\tag{27}
$$

For x of length n, merge sort $sort$ takes $O(n \log n)$ time. Incremental merge sort $\widehat{sort}'$ takes only $O(n)$ time, although it uses $O(n \log n)$ space to store the intermediate results of the previous sort.

Attribute evaluation. Given an attribute grammar, a set of recursive functions can be constructed to evaluate the attribute values for any derivation tree of the grammar [18]. Basically, each function evaluates a synthesized attribute of a non-terminal, and the value of a synthesized attribute of the root symbol is the final return value of interest. Thus, for the given grammar, the set of recursive functions takes a derivation tree of the grammar as input, and returns the value of a synthesized attribute at the root as output.

We consider subtree replacement as the input change, given by a new subtree and a path from the root of the whole tree to the root of the subtree to be replaced.

First, caching all intermediate results leads to a set of extended recursive functions that returns an attributed tree instead of just the value of an synthesized attribute at the root. Then, incrementalizing the set of extended functions under a subtree replacement is just composing a new attributed tree from the old one, evaluating only attributes whose values are affected by the subtree replacement, yielding a set of incremental recursive functions.

Suppose a given batch attribute evaluation program evaluates each attribute only once, then the derived incremental program computes in $O(|\text{PATH}|+|\text{AFFECTED}|)$ time, where PATH is the path from the root of the whole tree to the root of the new subtree, and AFFECTED is the set of attributes whose values are different in the new tree than in the old after the subtree replacement [40].

8 Related work and conclusion

Caching has been the basis of many techniques for developing efficient programs and optimizing programs. Bird [5] and Cohen [10] provide nice overviews. Most of the techniques fall into one of the following three classes.

In the first class, a global cache separate from a subject program is employed to record values of subcomputations that may be needed later, and certain strategies are chosen for using and managing the cache [28]. Work in recent years includes [15, 29, 37]. Two trends seem obvious: studying *specialized* cache strategies for classes of problems [38], and adding *annotations* or certain *specifications* to subject programs [14, 19, 44]. A drawback of these methods is that they are based on dynamic methods, which are fundamentally interpretive and are hard to be simultaneously both general and powerful.

The above drawback can be overcome by transforming subject programs to integrate caching into the transformed programs. In particular, some techniques apply transformations based on special properties and schemas of subject programs, and they form the second class. Typical examples of these techniques are dynamic programming [1], schemas of redundancies [10], and tupling [7, 8, 34, 35]. A drawback of these techniques is their lack of generality.

The third class analyzes and transforms programs under general principles. Often, a set of rules are derived from such principles and are used to transform programs. Examples are the conventional strength reduction technique

[3] and finite differencing technique [32] used in the APTS system [30, 31]. Seeking more flexibility and broader applicability, KIDS [43] and CIP [33] propose certain high-level strategies, but leave the choice of which intermediate results to maintain to manual decisions. Recently, certain principles that can directly guide program transformations have been proposed [13, 47], but implementations based on these principles employ heavy inference engines that are computationally exorbitant.

Our approach to the problem of program improvement via caching is a principled approach that integrates caching in the transformed programs. The approach is a crucial complement to any incremental computation technique for achieving the goal of program improvement. Among principle-based integrated caching methods, our approach is not limited to using a fixed set of rules for program analysis and transformations. On the contrary, we can even use our approach to derive such rules when necessary. Compared to the general approaches advocated by KIDS or CIP, our approach is more algorithmic and automatable. A prototype system CACHET based on our approach is under development.

Although we present the approach in a first-order functional language, the underlying principle is general and can be applied to other languages as well, e.g., higher-oder functional languages, functional languages with lazy semantics, and especially imperative languages with complex data structures and side effects. We have given an example [26] where the principle is applied to improve imperative programs with arrays for the local neighborhood problems in image processing [49, 51]. Further application of our principles to language with these features is a subject for future study.

References

[1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1974.

[2] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers, Principles, Techniques, and Tools*. Addison-Wesley series in Computer Science. Addison-Wesley Publishing Company, Reading, Massachusetts, 1986.

[3] F. E. Allen, J. Cocke, and K. Kennedy. Reduction of operator strength. In S. Muchnick and N. Jones, editors, *Program Flow Analysis*, pages 79–101. Prentice-Hall, Englewood Cliffs, New Jersey, 1981.

[4] R. A. Ballance, S. L. Graham, and M. L. Van De Vanter. The *Pan* language-based editing system. *ACM Transactions on Software Engineering and Methodology*, 1(1):95–127, January 1992.

[5] R. S. Bird. Tabulation techniques for recursive programs. *ACM Computing Surveys*, 12(4):403–417, December 1980.

[6] R. Cartwright. Recursive programs as definitions in first order logic. *SIAM Journal on Computing*, 13(2):374–408, May 1984.

[7] W.-N. Chin. Towards an automated tupling strategy. In *Proceedings of the ACM Symposium on PEPM*, Copenhagen, Denmark, June 1993.

[8] W.-N. Chin and S.-C. Khoo. Tupling functions with multiple recursion parameters. In P. Cousot, M. Falaschi, G. Filè, and A. Rauzy, editors, *Proceedings of the 3rd International Workshop on Static Analysis*, pages 124–140. Springer-Verlag, September 1993. LNCS 724.

[9] J. Cocke and K. Kennedy. An algorithm for reduction of operator strength. *Communications of the ACM*, 20(11):850–856, November 1977.

[10] N. H. Cohen. Eliminating redundant recursive calls. *ACM Transactions on Programming Languages and Systems*, 5(3):265–299, July 1983.

[11] J. Earley. High level iterators and a method for automatically designing data structure representation. *Journal of Computer Languages*, 1:321–342, 1976.

[12] C. A. Gunter. *Semantics of Programming Languages*. The MIT Press, Cambridge, Massachusetts, 1992.

[13] R. J. Hall. Program improvement by automatic redistribution of intermediate results: An overview. In M. R. Lowry and R. D. McCartney, editors, *Automating Software Design*, chapter 14, pages 339–372. AAAI Press/The MIT Press, 1991. Proceedings of the Workshop on Automating Software Design, AAAI '88.

[14] R. Hoover. Alphonse: Incremental computation as a programming abstraction. In *Proceedings of the ACM SIGPLAN '92 Conference on PLDI*, pages 261–272, California, June 1992.

[15] J. Hughes. Lazy memo-functions. In *Proceedings of the 2nd Conference on FPCA*, pages 129–146, Nancy, France, September 1985. Springer-Verlag. LNCS 201.

[16] F. Jalili and J. H. Gallier. Building friendly parsers. In *Conference Record of the 9th Annual ACM Symposium on POPL*, pages 196–206, Albuquerque, New Mexico, January 1982.

[17] S. B. Jones and D. Le Métayer. Compile-time garbage collection by sharing analysis. In *Proceedings of the 4th International Conference on FPCA*, pages 54–74, London, September 1989.

[18] T. Katayama. Translation of attribute grammars into procedures. *ACM Transactions on Programming Languages and Systems*, 6(3):345–369, July 1984.

[19] R. M. Keller and M. R. Sleep. Applicative caching. *ACM Transactions on Programming Languages and Systems*, 8(1):88–108, January 1986.

[20] J. Knoop, O. Rüthing, and B. Steffen. Partial dead code elimination. In *Proceedings of the ACM SIGPLAN '94 Conference on PLDI*, pages 147–158, Orlando, Florida, June 1994.

[21] J. Launchbury. Projection factorisations in partial evaluation. Ph.d. thesis, Department of Computing, University of Glasgow, 1989.

[22] J. Launchbury. Strictness and binding-time analysis: Two for the price of one. In *Proceedings of the ACM SIGPLAN '91 Conference on PLDI*, pages 80–91, Toronto, Ontario, Canada, June 1991.

[23] J. L. Lawall and O. Danvy. Separating stages in the continuation-passing style transformation. In *Conference Record of the 20th Annual ACM Symposium on POPL*, pages 124–136, January 1993.

[24] P. Lipps, U. Möncke, M. Olk, and R. Wilhelm. Attribute (re)evaluation in OPTRAN. *Acta Informatica*, 26:213–239, 1988.

[25] Y. A. Liu and T. Teitelbaum. Caching intermediate results for program improvement. Technical Report TR 95-1498, Department of Computer Science, Cornell University, Ithaca, New York, March 1995.

[26] Y. A. Liu and T. Teitelbaum. Incremental computation for transformational software development. Technical Report TR 95-1499, Department of Computer Science, Cornell University, Ithaca, New York, March 1995.

[27] Y. A. Liu and T. Teitelbaum. Systematic derivation of incremental programs. *Science of Computer Programming*, 24(1):1–39, 1995.

[28] D. Michie. "memo" functions and machine learning. *Nature*, 218:19–22, April 1968.

[29] D. J. Mostow and D. Cohen. Automating program speedup by deciding what to cache. In *Proceedings of the Ninth IJCAI*, pages 165–172, Los Angeles, August 1985.

[30] R. Paige. Transformational programming – applications to algorithms and systems. In *Conference Record of the 10th Annual ACM Symposium on POPL*, pages 73–87, January 1983.

[31] R. Paige. Symbolic finite differencing - part I. In *Proceedings of the 3rd ESOP*, pages 36–56, Copenhagen, Denmark, May 1990. Springer-Verlag. LNCS 432.

[32] R. Paige and S. Koenig. Finite differencing of computable expressions. *ACM Transactions on Programming Languages and Systems*, 4(3):402–454, July 1982.

[33] H. A. Partsch. *Specification and Transformation of Programs - A Formal Approach to Software Development.* Texts and Monographs in Computer Science. Springer-Verlag, 1990.

[34] A. Pettorossi. A powerful strategy for deriving efficient programs by transformation. In *Proceedings of the ACM '84 Symposium on LFP*, Austin, Texas, August 1984.

[35] A. Pettorossi. Strategical derivation of on-line programs. In L. G. L. T.Meertens, editor, *Program Specification and Transformation*, pages 73–88, Amsterdam, 1987. The Netherlands: North-Holland. Proceedings of the IFIP TC2/WG 2.1 Working Conference on Program Specification and Transformation, April 1986.

[36] G. D. Plotkin. Call-by-name, call-by-value and the λ-calculus. *Theoretical Computer Science*, 1:125–159, 1975.

[37] W. Pugh. An improved cache replacement strategy for function caching. In *Proceedings of the ACM '88 Conference on LFP*, pages 269–276, 1988.

[38] W. Pugh and T. Teitelbaum. Incremental computation via function caching. In *Conference Record of the 16th Annual ACM Symposium on POPL*, pages 315–328, January 1989.

[39] T. Reps and T. Teitelbaum. *The Synthesizer Generator: A System for Constructing Language-Based Editors.* Springer-Verlag, New York, 1988.

[40] T. Reps, T. Teitelbaum, and A. Demers. Incremental context-dependent analysis for language-based editors. *ACM Transactions on Programming Languages and Systems*, 5(3):449–477, July 1983.

[41] M. Rosendahl. Automatic complexity analysis. In *Proceedings of the 4th International Conference on FPCA*, pages 144–156, London, September 1989.

[42] D. S. Scott. Lectures on a mathematical theory of computation. In M. Broy and G. Schmidt, editors, *Theoretical Foundations of Programming Methodology*, pages 145–292. D. Reidel Publishing Company, 1982. Lecture Notes of 1981 Marktoberdorf Summer School on Theoretical Foundations of Programming Methodology, directed by F.L. Bauer, E.W. Dijkstra, and C.A.R. Hoare.

[43] D. R. Smith. KIDS: A semiautomatic program development system. *IEEE Transactions on Software Engineering*, 16(9):1024–1043, September 1990.

[44] R. S. Sundaresh and P. Hudak. Incremental computation via partial evaluation. In *Conference Record of the 18th Annual ACM Symposium on POPL*, pages 1–13, January 1991.

[45] K. R. Traub. A compiler for the MIT tagged-token dataflow architecture. Master's Thesis LCS TR-370, Department of Electrical Engineering and Computer Science, MIT, August 1986.

[46] P. Wadler and R. J. M. Hughes. Projections for strictness analysis. In *Proceedings of the 3rd International Conference on FPCA*, pages 385–407, Portland, Oregon, September 1987. LNCS 274.

[47] A. B. Webber. A formal definition of unnecessary computation in functional programs. Technical Report TR 92-1260, Department of Computer Science, Cornell University, Ithaca, New York, January 1992.

[48] B. Wegbreit. Mechanical program analysis. *Communications of the ACM*, 18(9):528–538, September 1975.

[49] W. M. Wells, III. Efficient synthesis of Gaussian filters by cascaded uniform filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(2):234–239, March 1986.

[50] D. Yeh and U. Kastens. Improvements on an incremental evaluation algorithm for ordered attribute grammars. *SIGPLAN Notices*, 23(12):45–50, 1988.

[51] R. Zabih. Individuating unknown objects by combining motion and stereo. Ph.D. Thesis, Department of Computer Science, Stanford University, Stanford, California, 1994.