Computational Divided Differencing
and Divided-Difference Arithmetics

Thomas W. Reps* (reps@cs.wisc.edu)
Comp. Sci. Dept., Univ. of Wisconsin, 1210 W. Dayton St., Madison, WI 53706

Louis B. Rall (rall@math.wisc.edu)
Dept. of Mathematics, Univ. of Wisconsin, 480 Lincoln Dr., Madison, WI 53706

Abstract. Tools for *computational differentiation* transform a program that computes a numerical function $F(x)$ into a related program that computes $F'(x)$ (the derivative of F). This paper describes how techniques similar to those used in computational-differentiation tools can be used to implement other program transformations—in particular, a variety of transformations for *computational divided differencing*. The specific technical contributions of the paper are as follows:

 - It presents a program transformation that, given a numerical function $F(x)$ defined by a program, creates a program that computes $F[x_0, x_1]$, the *first divided difference* of $F(x)$, where

$$F[x_0, x_1] \stackrel{\text{def}}{=} \begin{cases} \frac{F(x_0) - F(x_1)}{x_0 - x_1} & \text{if } x_0 \neq x_1 \\ \frac{d}{dz}F(z), \text{evaluated at } z = x_0 & \text{if } x_0 = x_1 \end{cases}$$

 - It shows how computational first divided differencing *generalizes* computational differentiation.

 - It presents a second program transformation that permits the creation of *higher-order divided differences* of a numerical function defined by a program.

 - It shows how to extend these techniques to handle functions of several variables.

The paper also discusses how computational divided-differencing techniques could lead to faster and/or more robust programs in scientific and graphics applications.

   Finally, the paper describes how computational divided differencing relates to the numerical-finite-differencing techniques that motivated Robert Paige's work on finite differencing of set-valued expressions in SETL programs.

Keywords: divided differences, computational differentiation, interpolation, multivariate interpolation, program transformation, round-off error

Dedicated to the memory of Robert Paige, 1947-1999.

1. Introduction

A variety of studies in the field of programming languages have led to useful, high-level transformations that manipulate programs in semantically meaningful ways. In very general terms, these tools transform a program that performs a computation $F(x)$ into a program that performs a related computation $F^\sharp(x)$, for a variety of $F^\sharp$'s of interest.[1] (In some cases, an appropriate preprocessing operation h needs to be applied to the input; in such cases, the transformed program $F^\sharp$ is used to perform a computation of the form $F^\sharp(h(x))$.) Examples of such tools include *partial evaluators* and *program slicers*:

— A *partial evaluator* creates a specialized version of a program when only part of the program's input has been supplied (Futamura, 1971; Futamura, 1999; Bjørner et al., 1988; Jones et al., 1993). A partial evaluator transforms a program that performs a computation $F(\langle s, d \rangle)$, where F operates on a pair of inputs $\langle s, d \rangle$; when s is known, partial evaluation of the program with respect to s results in a program that computes the function $F^s(d)$ $(= F^s(\text{second}(\langle s, d \rangle)))$, such that

$$F^s(d) = F(\langle s, d \rangle). \tag{1}$$

(The mnemonic is that s and d stand for the "static" and "dynamic" parts of the input, respectively.)

Partial evaluation is useful for removing interpretive overhead, and can also speed up programs that have two arguments that change value at different rates (such as ray tracing (Mogensen, 1986)).

— The *slice* of a program with respect to a set of program elements S is a projection of the program that includes only program elements that might affect (either directly or transitively) the values of the variables used at members of S (Weiser, 1984; Ottenstein and Ottenstein, 1984; Horwitz et al., 1990). Given a program that computes a function F, one version of the slicing problem focuses on creating the slice by symbolically composing the original program with an appropriate *projection function* π, where π characterizes what part of F's output should be discarded and what part should

---

[1] In this paper, we do not generally make a distinction between *programs* and *procedures*. We use "program" both to refer to the program as a whole, as well as to refer to individual subroutines in a generic sense. We use "procedure" only in places where we wish to emphasize that the focus of interest is an individual subroutine *per se*.

be retained (Reps and Turnidge, 1996). Program slicing creates a program that computes F^π, where

$$F^\pi(x) = (\pi \circ F)(x). \tag{2}$$

Program-slicing tools allow one to find semantically meaningful decompositions of programs, where the decompositions consist of elements that are not textually contiguous. Slicing, and subsequent manipulation of slices, has applications in many software-engineering tools, including ones for program understanding, maintenance (Gallagher and Lyle, 1991), debugging (Lyle and Weiser, 1986), testing (Binkley, 1992; Bates and Horwitz, 1993), differencing (Horwitz et al., 1989; Horwitz, 1990), specialization (Reps and Turnidge, 1996), reuse (Ning et al., 1994), and merging (Horwitz et al., 1989).

Less well known in the programming-languages community is the work that has been done by numerical analysts on tools for *computational differentiation* (also known as *automatic differentiation* or *algorithmic differentiation*) (Wengert, 1964; Rall, 1981; Griewank and Corliss, 1992; Berz et al., 1996; Griewank, 2000):

— Given a program that computes a numerical function $F(x)$, a computational-differentiation tool creates a related program that computes $F'(x)$ (the derivative of F).

— Applications of computational differentiation include optimization, solving differential equations, curve fitting, and sensitivity analysis.

Although in each of the cases mentioned above, the appropriate restructuring of the program could be carried out by hand, hand transformation of a program is an error-prone process. The automated assistance that the aforementioned tools provide prevents errors from being introduced when these transformations are applied. The work described in this paper expands the set of tools that programmers have at their disposal for performing such high-level, semantically meaningful program manipulations.

Because so much scientific, engineering, and graphical software tries to predict and render modeled situations, such software often performs extrapolation and/or interpolation. These operations inevitably involve the computation of divided differences, which can suffer badly from round-off error. In some cases, which motivate this paper, numerically unstable algorithms can be stabilized by computing divided differences in a non-standard way. Because it can be a tedious task to restructure

a program to perform computations in this way, it is worthwhile developing an automated approach. The paper describes how techniques similar to those that have been developed for computational differentiation can be used to transform programs that compute numerical functions into ones that compute divided differences.

The specific technical contributions of the paper are as follows:

— We present a program transformation that, given a numerical function $F(x)$ defined by a program, creates a program that computes $F[x_0, x_1]$, the *first divided difference* of $F(x)$, where

$$F[x_0, x_1] \stackrel{\text{def}}{=} \begin{cases} \frac{F(x_0)-F(x_1)}{x_0-x_1} & \text{if } x_0 \neq x_1 \\ \frac{d}{dz}F(z), \text{evaluated at } z = x_0 & \text{if } x_0 = x_1 \end{cases}$$

— We show how computational first divided differencing *generalizes* computational differentiation.

— We present a second program transformation that permits the creation of *higher-order divided differences* of a numerical function defined by a program.

— We present a third program transformation that permits higher-order divided differences to be computed more efficiently. This transformation does not apply to all programs; however, we show that there is at least one important situation where this optimization is of use.

— We show how to extend these techniques to handle functions of several variables.

— Finally, we describe how our work on computational differencing relates to the numerical-finite-differencing techniques that motivated Robert Paige's work on finite differencing of set-valued expressions in SETL programs (Paige, 1981; Paige and Koenig, 1982).

These ideas are illustrated by means of numerous examples.

Such program transformations can be implemented either as a source-to-source translation, or by means of overloaded operators and reinterpreted operands (in which case the source code is changed very little). The examples in the paper primarily illustrate the latter approach; the paper presents sketches of implementations of the various transformations in the form of C++ class definitions ("divided-difference arithmetics").

The benefits gained from the techniques described in the paper include the following:

- Because divided differences are the basis for a wide variety of numerical techniques, including polynomial interpolation, numerical integration, and solving differential equations (Conte and de Boor, 1972), this work could lead to more robust programs in scientific and graphics applications, when the function of interest is one that is defined by a program.

- Finite differences on an evenly spaced grid can be used to quickly generate a function's values at any number of points that extend the grid (see (Goldstine, 1977) and (Paige and Koenig, 1982, pp. 403–404)). Because finite differences on an evenly spaced grid can be obtained from divided differences on an evenly spaced grid, our techniques may be useful in graphics applications for quickly plotting a function, while retaining reasonable accuracy.

- Because the divided-differencing problems that we address can be viewed as generalizations of problems such as differentiation, computation of Taylor coefficients, etc., some of our techniques—in particular, the divided-difference arithmetic presented in Sect. 7—represent new approaches that, with appropriate simplification, can also be applied in computational-differentiation tools.

Empirical results presented in Sects. 5 and 8 provide two concrete demonstrations of some of the benefits that can be gained via our methods.

The remainder of the paper is organized into eight sections: To make the paper self-contained, Sect. 2 provides a succinct review of the basic principles of computational differentiation that are relevant to our work, and of so-called "differentiation arithmetics" (Rall, 1986; Rall, 1990; Rall, 1992). Sect. 3 discusses the basic principle behind computational divided differencing. Sect. 4 shows how computational divided differencing generalizes computational differentiation. Sect. 5 extends the ideas introduced in Sect. 3 to higher-order computational divided differencing. Sect. 6 discusses techniques that apply to a useful special case. Sect. 7 extends the ideas from Sects. 3, 5, and 6 to functions of several variables. Sect. 8 describes how these ideas relate to the numerical-finite-differencing techniques that motivated Robert Paige's work on finite differencing of set-valued expressions in SETL programs. Sect. 9 discusses other related work.

2. Background on Computational Differentiation

One bright spot during the last thirty years with respect to the control of round-off error has been the emergence of tools for *computational differentiation*, which transform a program that computes a numerical function $F(x)$ into a related program that computes the derivative $F'(x)$.[2] These tools address the following issue: Suppose that you have a program $F(x)$ that computes a numerical function $F(x)$.[3] It is a very bad idea to try to compute $F'(x_0)$, the value of the derivative of F at x_0, by picking a small value `delta_x` and invoking the following program with the argument x_0:[4]

$$
\boxed{
\begin{array}{l}
\texttt{float delta\_x = ...}\langle\texttt{some small value}\rangle\texttt{ ...;} \\
\texttt{float F}'\texttt{\_naive(float x)\{} \\
\quad \texttt{return (F(x + delta\_x) - F(x))/delta\_x;} \\
\texttt{\}}
\end{array}
}
\tag{2}
$$

For a small enough value of `delta_x`, the values of $F(x_0\texttt{+delta\_x})$ and $F(x_0)$ will usually be very close. Round-off errors in the computation of $F(x_0\texttt{+delta\_x})$ and $F(x_0)$ are magnified by the subtraction of the two quantities, and further amplified by the division by the small quantity `delta_x`, which may cause the overall result to be useless. Computational differentiation sidesteps this problem by computing derivatives in another fashion.

Computational differentiation can be illustrated by means of the following example:

EXAMPLE 2.1. (Zippel, 1996). Suppose that we have been given a collection of programs f_i for the functions f_i, $1 \leq i \leq k$, together

---

[2] Another bright spot has been the application of interval arithmetic to the verification of the accuracy of computed results, for many basic numerical computations (Hammer et al., 1993; Hammer et al., 1995).

[3] Throughout the remainder of the paper, Courier Font is used to denote functions defined by programs, whereas *Italic Font* is used to denote mathematical functions. That is, $F(x)$ denotes a function (evaluated over real numbers), whereas $F(x)$ denotes a program (evaluated over floating-point numbers). We adhere to this convention both in concrete examples that involve C++ code, as well as in more abstract discussions in order to distinguish between a mathematical function and a program that implements the function.

[4] The example programs in the paper are all written in C++; however, the ideas described apply to other programming languages—including functional programming languages (cf. (Karczmarczuk, 1999))—as well as to other imperative languages.

To emphasize the links between mathematical concepts and their implementations in C++, we take the liberty of sometimes using $'$ and/or subscripts on C++ identifiers.

with the program `Prod` shown below, which computes the function $Prod(x) = \prod_{i=1}^{k} f_i(x)$. In addition, suppose that we have also been given programs $\mathtt{f'_i}$ for the functions f'_i, $1 \le i \le k$. Finally, suppose that we wish to obtain a program `Prod'` that computes the function $Prod'(x)$. Column two of the table given below shows mathematical expressions for $Prod(x)$ and $Prod'(x)$. Column three shows two C++ procedures: Procedure `Prod` computes $Prod(x)$; procedure `Prod'` is the procedure that a computational-differentiation system would create to compute $Prod'(x)$.

	Mathematical Notation	Programming Notation

Function	$$Prod(x) = \prod_{i=1}^{k} f_i(x)$$	<pre>float Prod(float x){
		   float ans = 1.0;
		   for (int i = 1; i <= k; i++){
		      ans = ans * f_i(x);
		   }
		   return ans;
		}</pre>

Derivative	$$Prod'(x) = \sum_{i=1}^{k} f'_i(x) * \prod_{j \ne i} f_j(x)$$	<pre>float Prod'(float x){
		   float ans' = 0.0;
		   float ans = 1.0;
		   for (int i = 1; i <= k; i++){
		      ans' = ans' * f_i(x) + ans * f'_i(x);
		      ans = ans * f_i(x);
		   }
		   return ans';
		}</pre>

Notice that program `Prod'` resembles program `Prod`, as opposed to `F'_naive` (see box (2)). `Prod'` preserves accuracy in its computation of the derivative because, as illustrated below in Example 2.2, it is based on the rules for the exact computation of derivatives, rather than on the kind of computation performed by `F'_naive`. $\square$

The transformation illustrated above is merely one instance of a general transformation that can be applied to any program: Given a program `G` as input, the transformation produces a derivative-computing program `G'`. The method for constructing `G'` is as follows:

— For each variable `v` of type `float` used in `G`, another `float` variable `v'` is introduced.

— Each statement in $\mathtt{G}$ of the form "$\mathtt{v}\ =\ exp;$", where exp is an arithmetic expression, is transformed into "$\mathtt{v}'\ =\ exp';\ \ \mathtt{v}\ =\ exp;$", where exp' is the expression for the derivative of exp. If exp involves calls to a function $\mathtt{g}$, then exp' may involve calls to both $\mathtt{g}$ and $\mathtt{g}'$.

— Each return statement in $\mathtt{G}$ of the form "`return v;`" is transformed into "`return v';`".

In general, this transformation can be justified by appealing to the chain rule of differential calculus (see below).

EXAMPLE 2.2.   For Example 2.1, we can demonstrate the correctness of the transformation by symbolically executing $\mathtt{Prod}'$ for a few iterations, comparing the values of $\mathtt{ans}'$ and $\mathtt{ans}$ (as functions of $\mathtt{x}$) at the start of each iteration of the for-loop:

Iteration	Value of $\mathtt{ans}'$ (as a function of x)	Value of $\mathtt{ans}$ (as a function of x)
0	0.0	1.0
1	$\mathtt{f}_1'(\mathtt{x})$	$\mathtt{f}_1(\mathtt{x})$
2	$\mathtt{f}_1'(\mathtt{x}) * \mathtt{f}_2(\mathtt{x}) + \mathtt{f}_1(\mathtt{x}) * \mathtt{f}_2'(\mathtt{x})$	$\mathtt{f}_1(\mathtt{x}) * \mathtt{f}_2(\mathtt{x})$
3	$\mathtt{f}_1'(\mathtt{x}) * \mathtt{f}_2(\mathtt{x}) * \mathtt{f}_3(\mathtt{x})$ $+\ \mathtt{f}_1(\mathtt{x}) * \mathtt{f}_2'(\mathtt{x}) * \mathtt{f}_3(\mathtt{x})$ $+\ \mathtt{f}_1(\mathtt{x}) * \mathtt{f}_2(\mathtt{x}) * \mathtt{f}_3'(\mathtt{x})$	$\mathtt{f}_1(\mathtt{x}) * \mathtt{f}_2(\mathtt{x}) * \mathtt{f}_3(\mathtt{x})$
$\dots$	$\dots$	$\dots$
k	$\displaystyle\sum_{i=1}^{k}\mathtt{f}_i'(\mathtt{x}) * \prod_{j\neq i}\mathtt{f}_j(\mathtt{x})$	$\displaystyle\prod_{i=1}^{k}\mathtt{f}_i(\mathtt{x})$

The loop maintains the invariant that, at the start of each iteration, $\mathtt{ans}'(\mathtt{x}) = \frac{d}{d\mathtt{x}}\mathtt{ans}(\mathtt{x})$.[5] $\square$

For the computational-differentiation approach, we did not really need to make the assumption that we were given programs $\mathtt{f}_i'$ for

---

[5] The value of $\mathtt{ans}'$ on the 3^{rd} iteration would actually be computed with the terms grouped as follows: $(\mathtt{f}_1'(\mathtt{x})*\mathtt{f}_2(\mathtt{x})+\mathtt{f}_1(\mathtt{x})*\mathtt{f}_2'(\mathtt{x}))*\mathtt{f}_3(\mathtt{x})+(\mathtt{f}_1(\mathtt{x})*\mathtt{f}_2(\mathtt{x}))*\mathtt{f}_3'(\mathtt{x})$. Terms have been expanded in the table given above to clarify how $\mathtt{ans}'$ builds up a value that is equivalent—from the standpoint of evaluation in real arithmetic—to
$$\mathtt{Prod}'(\mathtt{x}) = \sum_{i=1}^{k}\mathtt{f}_i'(\mathtt{x}) * \prod_{j\neq i}\mathtt{f}_j(\mathtt{x}).$$

the functions f'_i, $1 \leq i \leq k$; instead, the programs $\mathtt{f'_i}$ can be generated from the programs $\mathtt{f_i}$ by applying the same statement-doubling transformation that was applied to $\mathtt{Prod}$.

In languages that support operator overloading, such as C++, Ada, and Pascal-XSC, computational differentiation can be carried out by defining a new data type that has fields for both the value and the derivative, and overloading the arithmetic operators to carry out appropriate manipulations of both fields (Rall, 1983; Rall, 1984), along the lines of the definition of the C++ class $\mathtt{FloatD}$, shown in Fig. 1. A class such as $\mathtt{FloatD}$ is called a *differentiation arithmetic* (Rall, 1986; Rall, 1990; Rall, 1992).

The transformation then amounts to changing the types of each procedure's formal parameters, local variables, and return value (including those of the $\mathtt{f_i}$).[6]

EXAMPLE 2.3. Using class $\mathtt{FloatD}$, the $\mathtt{Prod}$ program of Example 2.1 can be handled as follows:

```
float f₁(float x){...}          ⇒   FloatD f₁(const FloatD &x){...}
 ⋮                                     ⋮
float fₖ(float x){...}          ⇒   FloatD fₖ(const FloatD &x){...}
float Prod(float x){                FloatD Prod(const FloatD &x){
  float ans = 1.0;                    FloatD ans(CONST,1.0); // ans = 1.0
  for (int i = 1; i <= k; i++){       for (int i = 1; i <= k; i++){
    ans = ans * fᵢ(x);        ⇒         ans = ans * fᵢ(x);
  }                                   }
  return ans;                         return ans;
}                                   }
                                    float Prod'(float x){
                                      FloatD xD(VAR,x);
                                ⇒     return Prod(xD).val';
                                    }
```

By changing the types of the formal parameters, local variables, and the return values of $\mathtt{Prod}$ and the $\mathtt{f_i}$ (and making a slight change

---

[6] We have referred to both computational differentiation and computational divided differencing as "program transformations", which may conjure up the image of tools that perform source-to-source rewriting fully automatically. Although this is one possible embodiment, in this paper the term "transformation" will also include the use of C++ classes in which the arithmetic operators have been overloaded. With the latter approach, rewriting might be carried out by a preprocessor, but might also be performed by hand, since usually only light rewriting of the program source text is required.

```
enum ArgDesc { CONST, VAR };
class FloatD {
 public:
  float val';
  float val;
  FloatD(ArgDesc,float);
};
// Constructor to convert a constant
// or a value for the independent
// variable to a FloatD
FloatD::FloatD(ArgDesc a, float v){
  switch (a) {
    case CONST:
      val' = 0.0;
      val = v;
    break;
    case VAR:
      val' = 1.0;
      val = v;
    break;
  }
}
FloatD operator+(FloatD a, FloatD b){
   FloatD ans;
   ans.val' = a.val' + b.val';
   ans.val = a.val + b.val;
   return ans;
}
FloatD operator*(FloatD a, FloatD b){
  FloatD ans;
  ans.val' = a.val * b.val' + a.val' * b.val;
  ans.val = a.val * b.val;
  return ans;
}
```

Figure 1. A differentiation-arithmetic class.

to the initialization of `ans` in `Prod`), the program now carries around
derivative values (in the `val'` field) in addition to performing all of the
work performed by the original program. Because of the C++ overload-
resolution mechanism, the f_i procedures invoked in the fourth line of
the transformed version of `Prod` are the *transformed* versions of the f_i
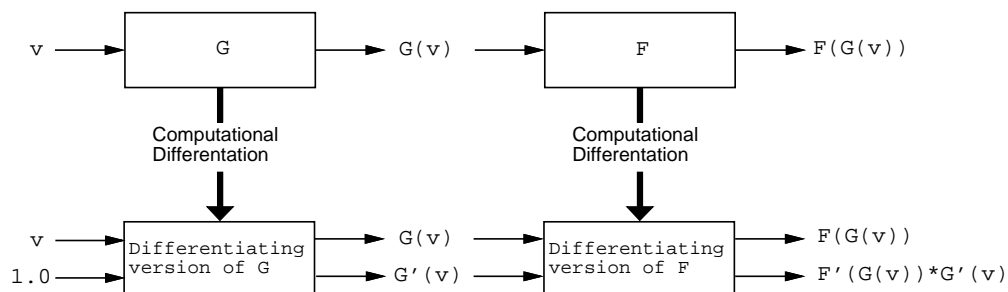(i.e., the f_i of type `FloatD` $\rightarrow$ `FloatD`).

The value of `Prod`'s derivative at `v` is obtained by calling `Prod'(v)`.
□

In a differentiation arithmetic, each procedure in the user's program,
such as `Prod` and the f_i in Example 2.3, can be viewed as a box that
maps two inputs to two outputs, as depicted below:

v ⟶ | F | ⟶ F(v)

Computational
Differentation

v ⟶ | Differentiating version of F | ⟶ F(v)
w ⟶ | | ⟶ F'(v)*w
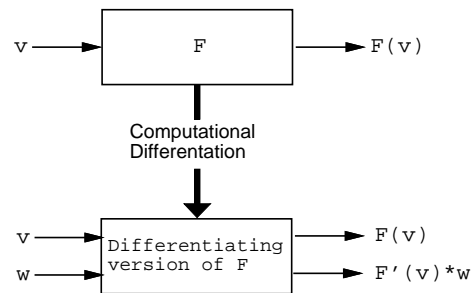
In particular, in each differentiating version of a user-defined or library
procedure `F`, the lower-right-hand output produces the value `F'(v)*w`.

An input value `v` for the formal parameter is treated as a pair
`(v,1.0)`. Boxes like the one shown above "snap together": when `F`
is composed with `G` (and the input is `v`), the output value on the
lower-right-hand side is `F'(G(v))*G'(v)`, which agrees with the usual
expression for the chain rule for the first-derivative operator:

v ⟶ | G | ⟶ G(v) ⟶ | F | ⟶ F(G(v))

Computational   Computational
Differentation   Differentation

v ⟶ | Differentiating version of G | ⟶ G(v) ⟶ | Differentiating version of F | ⟶ F(G(v))
1.0 ⟶ | | ⟶ G'(v) ⟶ | | ⟶ F'(G(v))*G'(v)

The computational-differentiation technique summarized above is
what is known as *forward-mode* differentiation. A different computational-
differentiation technique, *reverse mode* (Linnainmaa, 1976; Speelpen-
ning, 1980; Iri, 1984; Griewank, 1989; Griewank, 1991), is generally

preferable when the number of independent variables is much greater than the number of dependent variables. However, although it is possible to develop a reverse-mode version of computational divided differencing, it does not appear to offer the same potential savings in operations performed that reverse mode achieves for computational differentiation. Because the remainder of the paper concerns the generalization of forward-mode computational differentiation to forward-mode computational divided differencing, reverse mode is not summarized here.

The availability of overloading makes it possible to implement (forward-mode) computational differentiation conveniently, by packaging it as a differentiation-arithmetic class, as illustrated above. The alternative to the use of overloading is to build a special-purpose preprocessor to carry out the statement-doubling transformation that was illustrated in Examples 2.1 and 2.2. Examples of systems that use the latter approach include ADIFOR (Bischof et al., 1992; Bischof et al., 1996) and ADIC (Bischof et al., 1997).

2.1. LIMITATIONS OF COMPUTATIONAL DIFFERENTIATION

This section discusses certain limitations of the computational-differentiation transformation. First, it is worthwhile mentioning that the presence of aliasing (e.g., due to pointers or reference parameters) is *not* a limitation of computational differentiation (nor of computational divided differencing): The transformations presented above (as well as later in Sects. 3, 5, 6, and 7) work properly in the presence of aliasing (and are said to be *alias-safe* (Griewank, 2000).

One limitation of computational differentiation comes from the fact that a program $F'(x)$ that results from computational differentiation can perform additions and subtractions for which there are no analogues in the original program $F(x)$. For instance, in program Prod′, an addition is performed in the statement

$$\mathtt{ans'\ =\ ans'\ *\ f_i(x)\ +\ ans\ *\ f_i'(x);}$$

whereas no addition is performed in the statement

$$\mathtt{ans\ =\ ans\ *\ f_i(x);}$$

Consequently, the result of evaluating $F'(x)$ can be degraded by round-off error even when $F(x)$ is computed accurately. However, the accuracy of the result from evaluating $F'(x)$ can be verified by performing the same computation in interval arithmetic (Rall, 1983; Rall, 1992).

Another problem that arises is that the manner in which a function is programmed influences whether the results obtained from the

derivative program are correct. For instance, for programs that use a conditional expression or conditional statement in which the condition depends on the independent variable—i.e., where the function is defined in a piecewise manner—the derivative program may not produce the correct answer.

EXAMPLE 2.4.   (Fischer, 1992). Suppose that the function $F(x) = x^2$ is programmed using a conditional statement, as shown below on the left:

```
float F(float x){
  float ans;
  if(x == 1.0){
    ans = 1.0;
  }
  else{
    ans = x*x;
  }
  return ans;
}
```

$\Rightarrow$

```
float F'(float x){
  float ans';
  float ans;
  if(x == 1.0){
    ans' = 0.0;
    ans = 1.0;
  }
  else{
    ans' = x+x;
    ans = x*x;
  }
  return ans';
}
```

Computational differentiation would produce the program shown above on the right. With this program, F'(1.0) returns 0.0, rather than the correct value of 2.0 (i.e., correct with respect to the meaning of the program as the mathematical function $F(x) = x^2$). $\square$

The phenomenon illustrated in Example 2.4 has been called the *branch problem* or the *if problem* for computational differentiation. A more important example of the branch problem occurs in Gaussian elimination code, where pivoting introduces branches into the program (Fischer, 1992; Beck and Fischer, 1994; Griewank, 2000). Some additional problems that can arise with computational differentiation are identified in (Fischer, 1992). A number of different approaches to these problems have been discussed in the literature (Fischer, 1992; Beck and Fischer, 1994; Shamseddine and Berz, 1996; Kearfott, 1996; Griewank, 2000).

Computational divided differencing has some similar (or even worse) problems. All of these issues are outside the scope of the present paper;

the problem of finding appropriate ways to generalize the aforementioned techniques to handle the problems that arise with computational divided differencing is left for future work.

3.   Computational Divided Differencing

In this paper, we exploit the principle on which computational differentiation is based—namely, that it is possible to differentiate entire *programs*, not just *expressions*—to develop a variety of new *computational divided-differencing* transformations. We develop several transformations that can be applied to numerical programs. One of these corresponds to the *first-divided-difference operator*, denoted by $\cdot[x_0, x_1]$ and defined as follows:

$$F[x_0, x_1] \stackrel{\text{def}}{=} \begin{cases} \frac{F(x_0) - F(x_1)}{x_0 - x_1} & \text{if } x_0 \neq x_1 \\ \frac{d}{dz}F(z), \text{evaluated at } z = x_0 & \text{if } x_0 = x_1 \end{cases} \quad (3)$$

As with the differentiation operator, the problem that we face is that because division by a small value and subtraction are both operations that amplify accumulated round-off error, direct use of Eqn. (3) may lead to highly inaccurate results. In contrast, given a program that computes a numerical function $F(x)$, our technique for computational first divided differencing creates a related program that computes $F[x_0, x_1]$, *but without directly evaluating the right-hand side of Eqn. (3)*.

As we show below, the program transformation that achieves this goal is quite similar to the transformation used in computational-differentiation tools. The transformed program sidesteps the explicit subtraction and division operations that appear in Eqn. (3), while producing answers that are equivalent (from the standpoint of evaluation in real arithmetic). The program that results thereby avoids many operations that could potentially amplify round-off error, and hence retains accuracy when evaluated in floating-point arithmetic.

To understand the basis of the idea, consider the case in which $F(x) = x^2$ and $x_0 \neq x_1$:

$$F[x_0, x_1] = \frac{F(x_0) - F(x_1)}{x_0 - x_1} = \frac{x_0^2 - x_1^2}{x_0 - x_1} = x_0 + x_1. \quad (4)$$

That is, the first divided difference can be obtained by evaluating $x_0 + x_1$. In general, for monomials we have:

$F(x)$	c	x	x^2	x^3	...
$F[x_0, x_1]$	0	1	$x_0 + x_1$	$x_0^2 + x_0 x_1 + x_1^2$	...

Turning to programs, suppose that we are given the following program for squaring a number:

```
float Square(float x){
   return x * x;
}
```

The above discussion implies that to compute the first divided-difference of Square, we have our choice between the programs Square_1DD_naive and Square_1DD:

```
float Square_1DD_naive(float x0,float x1){
   return (Square(x0) - Square(x1))/(x0 - x1);
}
```

```
float Square_1DD(float x0,float x1){
   return x0 + x1;
}
```

However, the round-off-error characteristics of Square_1DD are much better than those of Square_1DD_naive.

The basis for creating expressions and programs that compute accurate divided differences is to be found in the basic properties of the first-divided-difference operator (Krawczyk and Neumaier, 1985), which closely resemble those of the first-derivative operator, as shown in Table I:

The program transformation for performing computational divided differencing can be explained by means of an example.

EXAMPLE 3.1. Suppose that we have a C++ class Poly that represents polynomials, and a member function Poly::Eval that evaluates a polynomial via Horner's rule; i.e., it accumulates the answer by repeatedly multiplying by x and adding in the current coefficient, iterating down from the high-order coefficient:[7]

---

[7] As a running example, the paper uses the evaluation of a polynomial in x via Horner's rule. It is well-known that Horner's rule can return inaccurate results when it is used to evaluate a polynomial in floating-point arithmetic (Hammer et al., 1995, pp. 65–67). Our examples are not meant to illustrate a way to circumvent this shortcoming. Horner's rule is used as the basis of our examples solely because of one virtue: it is a very simple procedure, which allows the various different computational-divided-differencing transformations to be illustrated in a succinct manner.

Table I. Basic properties of the first-derivative and first-divided-difference operators.

First Derivative	First Divided Difference (Krawczyk and Neumaier, 1985)
$c' = 0.0$	$c[x_0, x_1] = 0.0$
$x' = 1.0$	$x[x_0, x_1] = 1.0$
$(c + F)'(x) = F'(x)$	$(c + F)[x_0, x_1] = F[x_0, x_1]$
$(c * F)'(x) = c * F'(x)$	$(c * F)[x_0, x_1] = c * F[x_0, x_1]$
$(F + G)'(x) = F'(x) + G'(x)$	$(F + G)[x_0, x_1] = F[x_0, x_1] + G[x_0, x_1]$
$(F * G)'(x) = \begin{array}{l} F'(x) * G(x) \\ + \ F(x) * G'(x) \end{array}$	$(F * G)[x_0, x_1] = \begin{array}{l} F[x_0, x_1] * G(x_1) \\ + \ F(x_0) * G[x_0, x_1] \end{array}$
$\left(\frac{F}{G}\right)'(x) = \frac{F'(x)*G(x) - F(x)*G'(x)}{G(x)^2}$	$\left(\frac{F}{G}\right)[x_0, x_1] = \frac{F[x_0,x_1]*G(x_1) - F(x_1)*G[x_0,x_1]}{G(x_0)*G(x_1)}$

```
class Poly {
 public:
  float Eval(float);
 private:
  int degree;
  // Array coeff[0..degree]
  float *coeff;
};
```

```
// Evaluation via Horner's rule
float Poly::Eval(float x){
   float ans = 0.0;
   for (int i = degree; i >= 0; i--){
     ans = ans * x + coeff[i];
   } return ans;
}
```

A new member function, `Poly::Eval_1DD`, to compute the first divided difference can be created by transforming `Poly::Eval` as shown below:

```
class Poly {
 public:
  float Eval(float);
  float Eval_1DD(float,float);
 private:
  int degree;
  // Array coeff[0..degree]
  float *coeff;
};
```

```
float Poly::Eval_1DD(float x0,float x1){
   float ans_1DD = 0.0;
   float ans = 0.0;
   for (int i = degree; i >= 0; i--){
     ans_1DD = ans_1DD * x1 + ans;
     ans = ans * x0 + coeff[i];
   }
   return ans_1DD;
}
```

□

The transformation used to obtain `Eval_1DD` from (the text of) `Eval` is similar to the computational-differentiation transformation that would be used to create a derivative-computing program `Eval'` (`Eval'` appears in Fig. 2):

— `Eval_1DD` is supplied with an additional formal parameter (and the two parameters are renamed `x0` and `x1`).

— For each local variable `v` of type `float` used in `Eval`, an additional `float` variable `v_1DD` is introduced in `Eval_1DD`.

— Each statement of the form "`v = exp;`" in `Eval` is transformed into "`v_1DD = exp[x0,x1]; v = exp`$_0$`;`", where `exp[x0,x1]` is the expression for the divided difference of exp, and exp_0 is exp with `x0` substituted for all occurrences of `x`.

— Each statement of the form "`return v`" in `Eval` is transformed into "`return v_1DD`".

One caveat concerning the transformation presented above should be noted: the transformation applies only to procedures that have a certain special syntactic structure—namely, the only multiplication operations that depend on the independent variable `x` must be multiplications on the right by `x`. Procedure `Eval` is an example of a procedure that has this property.

A different, but similar, transformation can be used if all of the multiplication operations that depend on the independent variable `x` are multiplications on the left by `x`. (This point is discussed further in Example 6.1.) It is also possible to give a fully general first-divided-difference transformation; however, this transformation can be viewed as a special case of the material presented in Sect. 5. Consequently, we will not pause to present the general first-divided-difference transformation here.

Alternatively, as with computational differentiation, for languages that support operator overloading, computational divided differencing can be carried out with the aid of a new class, say `Float1DD`, for which the arithmetic operators are appropriately redefined. (We will call such a class a *divided-difference arithmetic*.) Computational divided differencing is then carried out by making appropriate changes to the types of each procedure's formal parameters, local variables, and return value.

Again, definitions of first-divided-difference arithmetic classes—both for the case of general first divided differences, as well as for the special case that covers programs like `Eval`—can be viewed as special cases of the divided-difference arithmetic classes `FloatDD` and `FloatDDR1`

discussed in Sects. 5 and 6, respectively. For this reason, we postpone giving a concrete example of a divided-difference arithmetic until Sects. 5 and 6.

4. Computational Divided Differencing as a Generalization of Computational Differentiation

In this section, we explain in what sense computational divided differencing can be said to generalize computational differentiation. First, observe that over real numbers, we have

$$\lim_{x_1 \to x_0} F[x_0, x_1] = \lim_{x_1 \to x_0} \frac{F(x_0) - F(x_1)}{x_0 - x_1} = F'(x_0). \tag{5}$$

However, although Eqn. (5) holds over reals, it does not hold over floating-point numbers: as x_1 approaches x_0, because of accumulated round-off error, the quantity $\frac{F(x_0) - F(x_1)}{x_0 - x_1}$ *does not*, in general, approach $F'(x_0)$. (Note the use of `Courier Font` here; this is a statement about quantities computed by *programs*.) This is why derivatives cannot be computed accurately by procedure `F'_naive` (see box (2)).

In contrast, for the *programs* $F[x_0, x_1]$ and $F'(x_0)$, we do have the property that

$$\lim_{x_1 \to x_0} F[x_0, x_1] = F'(x_0). \tag{6}$$

More precisely, the property that holds is that we have equality when x_1 equals x_0:

$$F[x_0, x_0] = F'(x_0). \tag{7}$$

EXAMPLE 4.1. To illustrate Eqn. (7), consider applying the two transformations to member function `Poly::Eval` of Example 3.1; the result is shown in Fig. 2. When formal parameters `x0`, `x1`, and `x` all have the same value—say `v`—then *exactly the same operations* are performed by `Eval_1DD(v,v)` and `Eval'(v)`. □

Because computations are carried out over floating-point numbers, the programs $F[x_0, x_1]$ and $F'(x_0)$ are only approximations to the functions that we actually desire. That is, $F[x_0, x_1]$ approximates the function $F[x_0, x_1]$, and $F'(x_0)$ approximates $F'(x_0)$. The relationships among these functions and programs are depicted below:
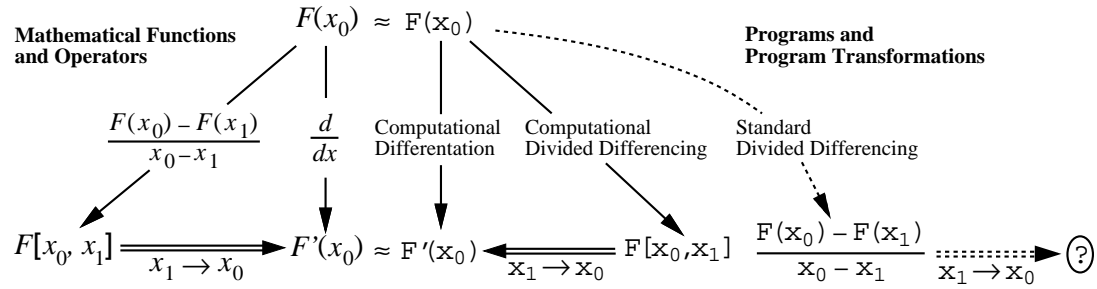
```
class Poly {
 public:
  float Eval(float);
  float Eval_1DD(float,float);
  float Eval'(float);
 private:
  int degree;
  // Array coeff[0..degree]
  float *coeff;
};

float Poly::Eval_1DD(float x0,float x1){
   float ans_1DD = 0.0;
   float ans = 0.0;
   for (int i = degree; i >= 0; i--){
     ans_1DD = ans_1DD * x1 + ans;
     ans = ans * x0 + coeff[i];
   }
   return ans_1DD;
}

float Poly::Eval'(float x){
   float ans' = 0.0;
   float ans = 0.0;
   for (int i = degree; i >= 0; i--){
     ans' = ans' * x + ans;
     ans = ans * x + coeff[i];
   }
   return ans';
}
```

Figure 2. The result of applying the computational-differentiation and first-divided-difference transformations to member function `Poly::Eval` of Example 3.1.

Mathematical Functions and Operators $F(x_0) \approx \mathtt{F(x_0)}$ **Programs and Program Transformations**

$$\frac{F(x_0) - F(x_1)}{x_0 - x_1} \qquad \frac{d}{dx} \qquad \text{Computational Differentation} \qquad \text{Computational Divided Differencing} \qquad \text{Standard Divided Differencing}$$

$$F[x_0, x_1] \xrightarrow[x_1 \to x_0]{} F'(x_0) \approx \mathtt{F''(x_0)} \xleftarrow[\mathtt{x_1 \to x_0}]{} \mathtt{F[x_0, x_1]} \qquad \frac{\mathtt{F(x_0) - F(x_1)}}{\mathtt{x_0 - x_1}} \xrightarrow[\mathtt{x_1 \to x_0}]{} \textcircled{?}$$

As the discussion above has noted, the interesting feature of this diagram is the relationship between $\mathtt{F[x_0, x_1]}$ and $\mathtt{F'(x_0)}$, on the side of the diagram labeled "Programs and Program Transformations". In particular, as $\mathtt{x_1}$ approaches $\mathtt{x_0}$, $\mathtt{F[x_0, x_1]}$ approaches $\mathtt{F'(x_0)}$. Consequently, a program produced by a system for computational divided differencing can be used to compute values of derivatives (in addition to divided differences) by feeding it duplicate actual parameters (e.g., `Eval_1DD(v,v)`). In contrast, a program created by means of computational differentiation can only produce derivatives (and not divided differences). In this sense, computational divided differencing can be said to generalize computational differentiation.

Computational divided differencing suffers from one of the same problems that arises with computational differentiation—namely that the program $\mathtt{F[x_0, x_1]}$ that results from the transformation can perform additions and subtractions that have no analogues in the original program $\mathtt{F(x)}$ (see the discussion in Sect. 2.1). Consequently, the result of evaluating $\mathtt{F[x_0, x_1]}$ can be degraded by round-off error even when $\mathtt{F(x)}$ is computed accurately. However, computational divided differencing is no worse in this regard than computational differentiation. Moreover, because of the fact that $\mathtt{F[x_0, x_1]}$ converges to $\mathtt{F'(x_0)}$ as $\mathtt{x_1}$ approaches $\mathtt{x_0}$, if $\mathtt{F'(x_0)}$ returns a result of sufficient accuracy, then $\mathtt{F[x_0, x_1]}$ will return a result of sufficient accuracy when $|\mathtt{x_0 - x_1}|$ is small. (In future work, we plan to investigate the use of interval arithmetic for providing interval bounds for computational divided differencing.)

5. Higher-Order Computational Divided Differencing

In this section, we show that the idea that was introduced in Sect. 3 can be generalized to define a transformation for *higher-order computational divided differencing*. To do so, we will define a divided-difference arithmetic that manipulates *divided-difference tables*.

Higher-order divided differences are divided differences of divided differences, defined recursively as follows:

$$F[x_i] \overset{\text{def}}{=} F(x_i) \tag{8}$$

$$F[x_0, x_1, \ldots, x_{n-1}, x_n] \overset{\text{def}}{=} \begin{cases} \frac{F[x_0, x_1, \ldots, x_{n-1}] - F[x_1, \ldots, x_{n-1}, x_n]}{x_0 - x_n} & \text{if } x_0 \neq x_n \\ \frac{\partial}{\partial z} F[z, x_1, \ldots, x_{n-1}]\big|_{z = x_0} & \text{if } x_0 = x_n \end{cases} \tag{9}$$

In our context, a divided-difference table for a function F is an upper-triangular matrix whose entries are divided differences of different orders, as indicated below:[8]

$$\begin{pmatrix} F(x_0) & F[x_0, x_1] & F[x_0, x_1, x_2] & F[x_0, x_1, x_2, x_3] \\ 0 & F(x_1) & F[x_1, x_2] & F[x_1, x_2, x_3] \\ 0 & 0 & F(x_2) & F[x_2, x_3] \\ 0 & 0 & 0 & F(x_3) \end{pmatrix}$$

Higher-order divided differences have numerous applications in interpolation and approximation of functions (Conte and de Boor, 1972).

We occasionally use $[x_{i,j}]$ as an abbreviation for $[x_i, \ldots, x_j]$. However, the reader should note that

$$\begin{aligned} F[x_{0,2}] &= F[x_0, x_1, x_2] \\ &= \frac{F[x_0, x_1] - F[x_1, x_2]}{x_0 - x_2}, \end{aligned}$$

which is not the same as $F[x_0, x_2]$:

$$F[x_0, x_2] = \frac{F(x_0) - F(x_2)}{x_0 - x_2}.$$

We use $\cdot^{\triangledown[x_0, \ldots, x_n]}$ to denote the operator that yields the divided-difference table for a function with respect to points $x_0, \ldots, x_n$. (We use $\cdot^{\triangledown}$ if the points $x_0, \ldots, x_n$ are clear from the context.)

A method for creating accurate divided-difference tables for rational expressions is found in Opitz (Opitz, 1964). This method is based on the properties of $\cdot^{\triangledown[x_0, \ldots, x_n]}$ given in the right-hand column of Table II: A few items in Table II require explanation:

— The symbol I denotes the identity matrix.

---

[8] Other arrangements of F's higher-order divided differences into matrix form are possible. For example, one could have a lower triangular matrix with the $F(x_i)$ running down the first column. However, the use of the arrangement shown above is key to being able to use simple notation—i.e., ordinary matrix operations—to describe our methods (Opitz, 1964).

Table II. Basic properties of two divided-difference operators.

First Divided Difference (Krawczyk and Neumaier, 1985)	Divided-Difference Table (Opitz, 1964)
$c[x_0, x_1] = 0.0$	$c^{\triangledown[x_0, \ldots, x_n]} = c * I$
$x[x_0, x_1] = 1.0$	$x^{\triangledown[x_0, \ldots, x_n]} = A_{[x_0, \ldots, x_n]}$
$(c + F)[x_0, x_1] = F[x_0, x_1]$	$(c + F)^{\triangledown[x_0, \ldots, x_n]} = c * I + F^{\triangledown[x_0, \ldots, x_n]}$
$(c * F)[x_0, x_1] = c * F[x_0, x_1]$	$(c * F)^{\triangledown[x_0, \ldots, x_n]} = c * F^{\triangledown[x_0, \ldots, x_n]}$
$(F + G)[x_0, x_1] = F[x_0, x_1] + G[x_0, x_1]$	$(F + G)^{\triangledown[x_0, \ldots, x_n]} = F^{\triangledown[x_0, \ldots, x_n]} + G^{\triangledown[x_0, \ldots, x_n]}$
$(F * G)[x_0, x_1] = \begin{array}{l} F[x_0, x_1] * G(x_1) \\ + \ \ F(x_0) * G[x_0, x_1] \end{array}$	$(F * G)^{\triangledown[x_0, \ldots, x_n]} = F^{\triangledown[x_0, \ldots, x_n]} * G^{\triangledown[x_0, \ldots, x_n]}$
$\left(\frac{F}{G}\right)[x_0, x_1] = \frac{F[x_0, x_1] * G(x_1) - F(x_1) * G[x_0, x_1]}{G(x_0) * G(x_1)}$	$\left(\frac{F}{G}\right)^{\triangledown[x_0, \ldots, x_n]} = \frac{F^{\triangledown[x_0, \ldots, x_n]}}{G^{\triangledown[x_0, \ldots, x_n]}}$

— The symbol $A_{[x_0, \ldots, x_n]}$ denotes the matrix

$$
\begin{pmatrix}
x_0 & 1 & 0 & \cdots & 0 \\
0 & x_1 & 1 & \cdots & 0 \\
\vdots & \ddots & \ddots & \ddots & \vdots \\
0 & \cdots & \ddots & x_{n-1} & 1 \\
0 & \cdots & \cdots & 0 & x_n
\end{pmatrix} .
\tag{10}
$$

— In the entry for $(F * G)^{\triangledown[x_0, \ldots, x_n]}$, the multiplication operation in $F^{\triangledown[x_0, \ldots, x_n]} * G^{\triangledown[x_0, \ldots, x_n]}$ is matrix multiplication.

— In the entry for $\left(\frac{F}{G}\right)^{\triangledown[x_0, \ldots, x_n]}$, the division operation in $\frac{F^{\triangledown[x_0, \ldots, x_n]}}{G^{\triangledown[x_0, \ldots, x_n]}}$ is matrix division (i.e., $\frac{P}{Q} = P * Q^{-1}$).

The two columns of Table II can be read as recursive definitions for the operations $\cdot[x_0, x_1]$ and $\cdot^{\triangledown[x_0, \ldots, x_n]}$, respectively. These have straightforward implementations as recursive programs that walk over an expression tree.

The reader should be aware that the $\cdot[x_0, x_1]$ operation defined in the first column of Table II creates an expression that computes *only* the first divided difference of the original expression $e(x)$, whereas the operation defined in the second column of Table II creates a (matrix) expression that computes *all* of the first, second, ..., n^{th} divided differences with respect to the points x_0, x_1, ..., x_n, *as well as $n + 1$* values of $e(x)$. In particular, in the matrix that results from evaluating the transformed expression, the values of $e(x)$ evaluated at the points

x_0, x_1, ..., x_n are found on the diagonal. For instance, in the case of an expression that multiplies two subexpressions $F(x)$ and $G(x)$, the elements on the diagonal are $F(x_0) * G(x_0)$, $F(x_1) * G(x_1)$, ..., $F(x_n) * G(x_n)$.

It is easy to verify (by means of induction) that the first and second columns of Table II are consistent with each other: in each case, the quantity $e[x_0, x_1]$ represents the (0,1) entry of the matrix $e^{\triangledown[x_0,...,x_n]}$.

The second column of Table II has an even more straightforward interpretation:

OBSERVATION 5.1. [Reinterpretation Principle]. *The divided-difference table of an arithmetic expression $e(x)$ with respect to the $n + 1$ points $x_0, \ldots, x_n$ can be obtained by reinterpreting $e(x)$ as a* matrix expression, *where the matrix $A_{[x_0,...,x_n]}$ is used at each occurrence of the variable x, and $c * I$ is used at each occurrence of a constant c.* □

That is, the expression tree for $e(x)$ is unchanged—except at its leaves, where $A_{[x_0,...,x_n]}$ is used in place of x, and $c*I$ is used in place of c—but the operators at all internal nodes are reinterpreted as denoting matrix operations. This observation is due to Opitz (Opitz, 1964).

With only a slight abuse of notation, we can express this as

$$e^{\triangledown[x_0,...,x_n]} = e(A_{[x_0,...,x_n]}).$$

Using this notation, we can show that the chain rule for the divided-difference operator $\cdot^{\triangledown[x_0,...,x_n]}$ has the following particularly simple form:

$$\begin{aligned}(F \circ G)^{\triangledown[x_0,...,x_n]} &= (F \circ G)(A_{[x_0,...,x_n]}) \qquad (11)\\ &= F(G(A_{[x_0,...,x_n]}))\\ &= F(G^{\triangledown[x_0,...,x_n]}).\end{aligned}$$

Opitz's idea can be extended to the creation of accurate divided-difference tables for functions defined by *programs* by overloading the arithmetic operators used in the program to be matrix operators—i.e., by defining a divided-difference arithmetic that manipulates divided-difference tables:

OBSERVATION 5.2. [Computational Divided-Differencing Principle]. *Rather than computing a divided-difference table with respect to the points x_0, x_1, ..., x_n by invoking the program $n + 1$ times and then applying Eqns. (8) and (9), we may instead evaluate the program (once) using a divided-difference arithmetic that overloads arithmetic operations as matrix operations, substituting $A_{[x_0,...,x_n]}$ for each occurrence of the formal parameter x, and $c * I$ for each occurrence of a constant c.* □

The reader should understand that the single invocation of the program using the divided-difference arithmetic will actually be more expensive than the $n + 1$ ordinary invocations of the program. The advantage of using divided-difference arithmetic is *not* that execution is speeded up because the program is only invoked once (in fact, execution is slower); the advantage is that the result computed using divided-difference arithmetic is much more *accurate*.

Because higher-order divided differences are defined recursively in terms of divided differences of lower order (cf. Eqns. (8) and (9)), it would be possible to define an algorithm for higher-order computational-divided-differencing using repeated applications of lower-order computational-divided-differencing transformations. However, with each application of the transformation for computational first divided differencing, the program that results performs (roughly) three times the number of operations that are performed by the program the transformation starts with. Consequently, this approach has a significant drawback: the final program that would be created for computing k^{th} divided differences could be $O(3^k)$ times slower than the original program. In contrast, the slow-down factor with the approach based on Observation 5.2 is $O(k^3)$.

We now sketch how a version of higher-order computational divided differencing based on Observation 5.2 can be implemented in C++. Below, we present highlights of a divided-difference arithmetic class, named `FloatDD`. We actually make use of two classes: (i) class `FloatDD`, the divided-difference arithmetic proper, and (ii) class `FloatV`, vectors of x_i values. These classes are defined as follows:

```
class FloatDD {
 public:
   int numPts; // Size is numPts-by-numPts
   float **divDiffTable; // Two-dimensional upper-triangular array
   FloatDD(ArgDesc ad, int N, float v); // N-by-N constant or variable: value v
   FloatDD(const FloatV &); // Construct a FloatDD from a FloatV
   FloatDD(int N); // Construct a zero-valued FloatDD of size N-by-N
   FloatDD& operator+ (const FloatDD &) const; // binary addition
   FloatDD& operator- (const FloatDD &) const; // binary subtraction
   FloatDD& operator* (const FloatDD &) const; // binary multiplication
   FloatDD& operator/ (const FloatDD &) const; // binary division
};
```

```
class FloatV {
 public:
   int numPts;
   float *val; // An array of values: val[0]..val[numPts-1]
   FloatV(int N, ...); // N points
   FloatV(float start, int N, float incr); // N equally spaced points
};
```

The constructor `FloatDD(const FloatV &);` plays the role of generating a matrix $A_{[x_0,...,x_n]}$ from a vector $[x_0,...,x_n]$ of values for the independent variable. It is defined as follows:

```
// Construct a FloatDD from a FloatV
FloatDD::FloatDD(const FloatV &fv) :
   numPts(fv.numPts),
   divDiffTable(calloc_ut(numPts)) // allocate upper-triangular matrix of zeros
{
   for (int i = 0; i < numPts; i++) {
     divDiffTable[i][i] = fv.val[i];
     if (i < numPts-1) divDiffTable[i][i+1] = 1.0;
   }
}
```

The constructor `FloatDD(ArgDesc ad, int N, float v);` generates either the matrix $A_{[v,...,v]}$ of size `N`-by-`N` or the matrix $v * I$ of size `N`-by-`N`, depending on the value of parameter `ad`. It is defined as follows:

```
FloatDD::FloatDD(ArgDesc ad, int N, float v) :
  numPts(N),
  divDiffTable(calloc_ut(numPts))
{
  int i;
  switch (ad) {
    case CONST:
      for (i = 0; i < numPts; i++) {
        divDiffTable[i][i] = v;
      }
      break;
    case VAR:
      for (i = 0; i < numPts; i++) {
        divDiffTable[i][i] = v;
        if (i < numPts - 1)
          divDiffTable[i][i+1] = 1.0;
      }
      break;
  }
}
```

The multiplication operator of class `FloatDD` simply performs matrix multiplication:

```
FloatDD& FloatDD::operator* (const FloatDD &fdd) const{
  assert(numPts == fdd.numPts);
  FloatDD *ans = new FloatDD(numPts);
  for (int r = 0; r < numPts; r++) {
    for (int c = r; c < numPts; c++) {
      float temp = 0.0;
      for (int k = r; k <= c; k++) {
        temp += divDiffTable[r][k] * fdd.divDiffTable[k][c];
      }
      ans->divDiffTable[r][c] = temp;
    }
  }
  return *ans;
}
```

The division operator of class `FloatDD` is implemented using back substitution. That is, suppose we wish to find the value of `A/B` (call this value `X`). `X` can be found by solving the system `X * B = A`. Because the divided-difference tables `A` and `B` are both upper-triangular matrices, this can be done using back substitution, as follows:

```
// Use back substitution
FloatDD& FloatDD::operator/ (const FloatDD &fdd) const{
  assert(numPts == fdd.numPts);
  assert(NonZeroDiagonal(fdd));
  FloatDD *ans = new FloatDD(numPts);
  for (int r = 0; r < numPts; r++) {
    for (int c = r; c < numPts; c++) {
      float temp = 0.0;
      for (int k = r; k < c; k++) {
        temp += ans->divDiffTable[r][k] * fdd.divDiffTable[k][c];
      }
      ans->divDiffTable[r][c] =
            (divDiffTable[r][c] - temp) / fdd.divDiffTable[c][c];
    }
  }
  return *ans;
}
```

(5)

EXAMPLE 5.3. To illustrate these definitions, consider again the function `Poly::Eval` that evaluates a polynomial via Horner's rule. Computational divided differencing is carried out by changing the types of `Eval`'s formal parameters, local variables, and return value from `float` to `FloatDD`:

```
// Evaluation via Horner's rule
float Poly::Eval(float x){
  float ans = 0.0;
  for (int i = degree; i >= 0; i--){
    ans = ans * x + coeff[i];
  }
  return ans;
}
// Evaluation via Horner's rule
FloatDD Poly::Eval(const FloatDD &x){
  FloatDD ans(x.numPts); // ans = 0.0
  for (int i = degree; i >= 0; i--){
    ans = ans * x + coeff[i];
  }
  return ans;
}
```

The transformed procedure can be used to generate the divided-difference table for the polynomial

$$P(x) = 2.1 * x^3 - 1.4 * x^2 - .6 * x + 1.1$$

with respect to the (unevenly spaced) points 3.0, 3.01, 3.02, 3.05 by performing the following operations:

```
Poly *P = new Poly(4,2.1,-1.4,-0.6,1.1);
FloatV x(4,3.0,3.01,3.02,3.05);
FloatDD A(x); // Corresponds to A[x0,...,x3]
FloatDD fdd = P->Eval(A);
```
(5.3)

We now present some empirical results that illustrate the advantages of the computational-divided-differencing method. In this experiment, we worked with the polynomial

$$P(x) = 2.1 * x^3 - 1.4 * x^2 - .6 * x + 1.1,$$

and performed computations using single-precision floating-point arithmetic on a Sun SPARCstation 20/61 running SunOS 5.6. Programs were compiled with the egcs-2.91.66 version of **g++** (egcs-1.1.2 release). The experiment compared the standard method for generating divided-difference tables—namely, the recursive definition given by Eqn. (8) and the first line of Eqn. (9)—against the overloaded version of function

`Poly::Eval` from Example 5.3 (which was invoked using code like the fragment shown in box (5.3)).

In each of the four examples shown below, the values used for the (unevenly spaced) points x_0, x_1, x_2, and x_3 are shown on the left. Note how the standard method for generating divided-difference tables degrades as the points move closer together. (Places where the results from the two methods differ are indicated in boldface.) In particular, because P is a cubic polynomial whose high-order coefficient is 2.1, the proper value of $P[x_0, x_1, x_2, x_3]$—the third divided difference of P—is 2.1, not 117,520! (Compare the entries that appear in the upper-right-hand corners of the fourth pair of divided-difference tables shown below.)

		Computational Divided Differencing					Standard Divided Differencing		
x_0 :	3.0	43.4	67.3	23.8	**2.1**	43.4	67.3	23.8	**2.09999**
x_1 :	4.0		110.7	114.9	32.2		110.7	114.9	32.2
x_2 :	5.0			225.6	211.5			225.6	211.5
x_3 :	7.0				648.6				648.6
x_0 :	3.0	43.4	47.8752	17.5**63**	**2.1**	43.4	47.8749	17.5858	**1.59073**
x_1 :	3.01		43.8787	48.2265	17.6**68**		43.8787	48.2266	17.66**53**
x_2 :	3.02			44.361	48.9332			44.361	48.9332
x_3 :	3.05				45.829				45.829
x_0 :	3.0	43.4	47.7175	17.5063	**2.1**	43.4	47.7177	**15.2886**	**754.685**
x_1 :	3.001		43.477	47.7**525**	17.5168		43.4477	47.7483	19.0621
x_2 :	3.002			43.4955	47.8**226**			43.4955	47.8245
x_3 :	3.005				43.6389				43.6389
x_0 :	3.0	43.4	47.**7017**	17.5006	**2.1**	43.4	47.**6945**	**3.62336**	**117520**
x_1 :	3.0001		43.4048	47.**7052**	17.5017		43.4048	47.**6952**	62.379
x_2 :	3.0002			43.4095	47.7**122**			43.4095	47.7**202**
x_3 :	3.0005				43.4238				43.4238

Finally, when we set all of the input values to 3.0, we obtain

		Computational Divided Differencing				Standard Divided Differencing			
x_0 :	3.0	43.4	**47.7**	**17.5**	**2.1**	43.4	**NaN**	**NaN**	**NaN**
x_1 :	3.0		43.4	**47.7**	**17.5**		43.4	**NaN**	**NaN**
x_2 :	3.0			43.4	**47.7**			43.4	**NaN**
x_3 :	3.0				43.4				43.4

With the standard divided-differencing method, division by 0 occurs and yields the exceptional value NaN. In contrast, computational divided differencing produces values for P's first, second, and third derivatives. More precisely, each k^{th} divided-difference entry in the computational-divided-differencing table equals

$$\frac{1}{k!} \left. \frac{d^k P(x)}{dx^k} \right|_{x=3.0} \tag{12}$$

The $k = 1$ case was already discussed in Sect. 4, where we observed that computational first divided differencing could be used to compute first derivatives. $\square$

EXAMPLE 5.4. (Kahan, 2000). Suppose that we wish to compute the future value of n monthly payments, each of 1 unit, paid at the end of each month into a savings account that compounds interest at the rate of α per month (where α is a small positive value and n is a positive integer). This answers the question "How many dollars are accumulated after n months, when you deposit \$1 per month for n months, into a savings account that pays annual interest at the rate of $(12 \times \alpha \times 100)\%$, compounded monthly?" Future value can be computed by the function

$$FutureValue(\alpha, n) \stackrel{\text{def}}{=} \frac{((1 + \alpha)^n - 1)}{\alpha}.$$ (13)

However, this can also be written as

$$FutureValue(\alpha, n) = \frac{((1 + \alpha)^n - 1^n)}{(1 + \alpha) - 1},$$

and thus is equal to the following first-divided difference of the power function:

$$FutureValue(\alpha, n) = (x^n)[1 + \alpha, 1].$$ (14)

The latter quantity can be computed to nearly full accuracy using computational divided differencing by computing $(x^n)^{\nabla[1+\alpha,1]}$, and then extracting the $(0, 1)$ entry. For instance, we can start with the following procedure power, which computes x^n via repeated squaring and multiplication by x, according to the bits of argument n:

```
const unsigned int num_bits = sizeof(unsigned int)*8;

float power(float x, unsigned int n) {
  unsigned int mask = 1 << (num_bits - 1);
  float ans = 1.0;
  for (unsigned int i = 0; i < num_bits; i++) {
    ans = ans * ans;
    if (mask & n)
      ans = ans * x;
    mask >>= 1;
  }
  return ans;
}
```

By changing the types of power's formal parameters, local variables, and return value, we create a version that computes a divided-difference table:

```
FloatDD power(FloatV &x, unsigned int n) {
   unsigned int mask = 1 << (num_bits - 1);
   FloatDD ans(CONST, 2, 1.0); // ans = 1.0
   for (unsigned int i = 0; i < num_bits; i++) {
     ans = ans * ans;
     if (mask & n)
       ans = ans * x;
     mask >>= 1;
   }
   return ans;
}
```

The transformed procedure can be used to compute the desired computation to nearly full accuracy by calling the procedure FutureValue that is defined below:

```
float FutureValue(float alpha, unsigned int n) {
   float w[2] = { 1+alpha, 1 };
   FloatV v(2, w);
   return power(v,n).divDiffTable[0][1];
}
```

We now present some empirical results that illustrate the advantages of this approach. In this experiment, we performed computations using single-precision floating-point arithmetic on a Sony VAIO PCG-Z505JSK (650 MHz Intel Pentium III processor) running Windows 2000. Programs were compiled with Microsoft Visual C++ 6.0. The table given below shows the future values of $1 deposits made for 360 months, as computed via Eqn. (13) versus procedure FutureValue, for a variety of interest rates. (Places where the results from the two methods differ are indicated in boldface.)

T.W. Reps and L.B. Rall

Interest Rate	Eqn. (13)	Procedure FutureValue
10%	2260.**5**	2260.**49**
8%	1490.36	1490.36
6%	1004.52	1004.52
4%	694.04**5**	694.048
2%	492.71**3**	492.7**22**
1%	419.6**56**	419.6**32**
0.8%	406.**692**	406.**718**
0.6%	394.**282**	394.**323**
0.4%	382.4	382.4**22**
0.2%	370.9**34**	370.9**86**
0.1%	365.**354**	365.**438**
0.08%	364.**142**	364.**34**
0.06%	36**2.873**	36**3.247**
0.04%	362.**595**	362.**165**
0.02%	361.**505**	361.**08**
0.01%	360.**961**	360.**54**
0.008%	360.**882**	360.**432**
0.006%	360.**751**	360.**324**
0.004%	360.**668**	360.**216**
0.002%	360.**56**	360.**108**
0.001%	360.**489**	360.**054**
0.0008%	**386.238**	**360.046**
0.0006%	**343.323**	**360.031**
0.0004%	**386.238**	**360.023**
0.0002%	**257.492**	**360.008**
0.0001%	**514.984**	**360.008**
0.00008%	**643.73**	**360.008**
0.00006%	**0**	**360**
0.00004%	**0**	**360**
0.00002%	**0**	**360**

□

6. A Special Case

A divided-difference table for a function F can be thought of as a (redundant) representation of an interpolating polynomial for F. For instance, if you have a divided-difference table T (and also know the appropriate vector of values x_0, x_1, ..., x_n), you can explicitly construct the Newton form of the interpolating polynomial for F according to the following definition (Conte and de Boor, 1972, pp. 197):

$$p_n(x) = \sum_{i=0}^{n} F[x_0, \ldots, x_i] * \prod_{j=0}^{i-1} (x - x_j) \qquad (15)$$

Note that to be able to create the Newton form of the interpolating polynomial for F via Eqn. (15), only the first row of divided-difference table T is required to be at hand—i.e., the values $F[x_0, \ldots, x_i]$, for $0 \leq i \leq n$—together with the values of x_0, x_1, ..., x_n. This observation suggests that we should develop an alternative divided-difference arithmetic that builds up and manipulates only first rows of divided-difference tables. We call this divided-difference arithmetic FloatDDR1 (for **D**ivided-**D**ifference **R**ow **1**). The motivation for this approach is that FloatDDR1 operations will be much faster than FloatDD ones, because FloatDD operations must manipulate upper-triangular matrices, whereas FloatDDR1 operations involve only simple vectors.

To achieve this, we define class FloatDDR1 as follows:

```
class FloatDDR1 {
  friend FloatDDR1& operator+ (const FloatDDR1 &, const float);
  friend FloatDDR1& operator+ (const FloatDDR1 &, const FloatV &);
  friend FloatDDR1& operator+ (const float, const FloatDDR1 &);
  friend FloatDDR1& operator+ (const FloatV &, const FloatDDR1 &);
  friend FloatDDR1& operator- (const FloatDDR1 &, const float);
  friend FloatDDR1& operator- (const FloatDDR1 &, const FloatV &);
  friend FloatDDR1& operator- (const float, const FloatDDR1 &);
  friend FloatDDR1& operator- (const FloatV &, const FloatDDR1 &);
  friend FloatDDR1& operator* (const FloatDDR1 &, const float);
  friend FloatDDR1& operator* (const FloatDDR1 &, const FloatV &);
  friend FloatDDR1& operator* (const float, const FloatDDR1 &);
  friend FloatDDR1& operator* (const FloatV &, const FloatDDR1 &);
  friend FloatDDR1& operator/ (const FloatDDR1 &, const float);
  friend FloatDDR1& operator/ (const FloatDDR1 &, const FloatV &);
  public:
  int numPts;
  float *divDiffTable; // One-dimensional array of divided differences
  FloatDDR1(int N) // Construct a zero-valued FloatDDR1 of length N
    : numPts(N), divDiffTable(new float[numPts])
    { }
  FloatDDR1& operator+ (const FloatDDR1 &) const; // binary addition
  FloatDDR1& operator- (const FloatDDR1 &) const; // binary subtraction
};
```

Compared with class `FloatDD`, class `FloatDDR1` is somewhat impoverished: we can add or subtract two arbitrary `FloatDDR1`'s; however, because we do not have full divided-difference tables available, we cannot multiply two arbitrary `FloatDDR1`'s; nor do we have the full $A_{[x_0,\dots,x_n]}$ matrices that are used at each occurrence of the independent variable. We finesse these difficulties by limiting the other operations of class `FloatDDR1` to those defined by the friend functions indicated in the class definition given above: (i) addition, subtraction, and multiplication on either side by a `float` or a `FloatV`; (ii) division on the right by a `float` or a `FloatV`.

The operations that involve a `float` argument `c` have their "obvious" meanings, if one bears in mind that a `float` value `c` serves as a stand-in for a full matrix `c*I`. For the addition (subtraction) operations, `c` is only added to (subtracted from) the `divDiffTable[0]` entry of the

FloatDDR1 argument. For the multiplication (division) operations, all
of the divDiffTable entries are multiplied by (divided by) c.

In the operations that involve a FloatV argument, the FloatV value
serves as a stand-in for a full $A_{[x_0,\ldots,x_n]}$ matrix. For instance, the op-
erator for multiplication on the right by a FloatV can be thought
of as performing a form of matrix multiplication—but specialized to
produce only the first row of the output divided-difference table (and
to use only values that are available in the given FloatDDR1 and FloatV
arguments):

```
FloatDDR1& operator* (const FloatDDR1 &fddr1, const FloatV &fv){
  FloatDDR1 *ans = new FloatDDR1(fddr1.numPts);
  ans->divDiffTable[0] = fddr1.divDiffTable[0] * fv.val[0];
  for (int c = 1; c < fddr1.numPts; c++) {
    ans->divDiffTable[c] =
          fddr1.divDiffTable[c-1] + fddr1.divDiffTable[c] * fv.val[c];
  }
  return *ans;
}
```

It might be thought that the operator for multiplication on the left
by a FloatV does not have the proper values available in the given
FloatV and FloatDDR1 arguments to produce the first row of the
product divided-difference table as output. (In particular, the second
argument, which is of type FloatDDR1, is a row vector, yet we want to
produce a row vector as the result.) However, it is easy to show that
divided-difference matrices are commutative:

$$F^{\triangledown} * G^{\triangledown} = (F * G)^{\triangledown} = (G * F)^{\triangledown} = G^{\triangledown} * F^{\triangledown}. \tag{16}$$

Consequently, the operator for multiplication on the left by a FloatV
can be treated as if the FloatV were on the right:

```
FloatDDR1& operator* (const FloatV &fv, const FloatDDR1 &fddr1){
  return fddr1 * fv;
}
```

As with class FloatDD, the division operator is implemented using a
form of back substitution—specialized here to compute just what is
needed for the first row of the divided-difference table:

```
FloatDDR1& operator/ (const FloatDDR1 &fddr1, const FloatV &fv){
   FloatDDR1 *ans = new FloatDDR1(fddr1.numPts);
   ans->divDiffTable[0] = fddr1.divDiffTable[0] / fv.val[0];
   for (int c = 1; c < fddr1.numPts; c++) {
     ans->divDiffTable[c] =
         (fddr1.divDiffTable[c] - ans->divDiffTable[c-1]) / fv.val[c];
   }
   return *ans;
}
```

Because only a limited set of arithmetic operations are available for objects of class `FloatDDR1`, this divided-difference arithmetic can only be applied to procedures that have a certain special syntactic structure, namely ones that are "accumulative" in the independent variable (with only "right-accumulative" quotients). In other words, the procedure must never perform arithmetic operations (other than addition or subtraction) on two local variables that both depend on the independent variable.

EXAMPLE 6.1. The procedure `Poly::Eval` for evaluating a polynomial via Horner's rule is an example of a procedure of the right form. Consequently, an overloaded version of the function `Poly::Eval` using `FloatDDR1` arithmetic can be written as shown below on the right:

```
// Evaluation via Horner's rule
float Poly::Eval(float x){
   float ans = 0.0;
   for (int i = degree; i >= 0; i--){
      ans = ans * x + coeff[i];
   }
   return ans;
}
// Evaluation via Horner's rule
FloatDDR1 Poly::Eval(const FloatV &x){
   FloatDDR1 ans(x.numPts); // ans = 0.0
   for (int i = degree; i >= 0; i--){
      ans = ans * x + coeff[i];
   }
   return ans;
}
```

In Sect. 3, Example 3.1 discussed the procedure `Poly::Eval_1DD`, a transformed version of `Poly::Eval` that computes the value of the first divided difference of a polynomial with respect to two values, x_0 and x_1. With the way that the overloaded operations are defined for class `FloatDDR1`, when the actual parameter supplied for `x` is a `FloatV` of length two consisting of x_0 and x_1, the procedure

```
FloatDDR1 Poly::Eval(const FloatV &x)
```

performs essentially the same steps as `Poly::Eval_1DD`.

One slight difference is that, in addition to returning the value of the first divided difference, the `FloatDDR1` version also returns the result of evaluating the polynomial on x_0.

Another difference is that with class `FloatDDR1`, because of our trick for handling multiplication by a `FloatV` on the left (cf. Eqn. 16 and the discussion that follows), `FloatDDR1` arithmetic can be used with programs that contain multiplications by the independent variable `x` on the left *as well as* on the right. (The transformation used in Example 3.1 could also be enhanced in this fashion.) □

Some empirical results that illustrate the advantages of `FloatDDR1` arithmetic in a useful application are presented at the end of Sect. 8.

As with the methods discussed in Sects. 3 and 5, `FloatDDR1` arithmetic can be used to produce values of interest for computational differentiation. For instance, suppose we have transformed procedure `F`:

```
float F(float x);   ⇒   FloatDDR1 F(const FloatV &x);
```

When all of the x_i values in the actual parameter supplied for `FloatV` `x` are the same value, say $\bar{x}$, then the `FloatDDR1` value returned as the output holds the Taylor coefficients for the expansion of `F` at $\bar{x}$ (cf. Formula (12)). Thus, the `FloatV` divided-difference arithmetic generalizes previously known techniques for producing accurate Taylor coefficients for functions defined by programs (Rall, 1983; Rall, 1984).

If we attempt to use `FloatDDR1` arithmetic in a procedure that is not "accumulative" in the independent variable, with only "right-accumulative" quotients, the overload-resolution mechanism of the C++ compiler will detect and report a problem. This is illustrated by the following example:

EXAMPLE 6.2.   Returning to the future-value calculation of Example 5.4, because the desired quantity in Eqn. (14) is merely a first divided difference, we might attempt to carry out the computation by

means of `FloatDDR1` arithmetic—hoping to save the cost of computing
and storing the $(1,1)$ entry of the divided-difference table—using the
following overloaded version of procedure `power`:

```
FloatDDR1 power(FloatV &x, unsigned int n) {
  unsigned int mask = 1 << (num_bits - 1);
  FloatDDR1 ans(CONST, 2, 1.0); // ans = 1.0
  for (unsigned int i = 0; i < num_bits; i++) {
    ans = ans * ans; // overload-resolution fails here
    if (mask & n)
      ans = ans * x;
    mask >>= 1;
  }
  return ans;
}
```

However, at the statement

$$\texttt{ans = ans * ans;}$$

procedure `power` multiplies two local variables that both depend on the
independent variable `x`. Consequently, the `FloatDDR1` version of `power`
is not accumulative in the independent variable.[9] In the case of the
Microsoft Visual C++ compiler, the following error message is issued:

```
binary '*' : no operator defined which takes a left-hand operand
                 of type 'class FloatDDR1'
```

□

7. Multi-Dimensional Computational Divided Differencing

In this section, we explain how to define a third divided-differencing
arithmetic that extends our techniques to handle multi-dimensional
computational divided differencing (i.e., computational divided differ-
encing of functions of several variables).

---

[9] Because the $(1,1)$ entry of `ans` is always `1.0` in the invocation of `power` from
`FutureValue`, a version of `power` specialized to this context would be able to save the
cost of computing and storing the $(1,1)$ entry. This example illustrates a limitation
of an approach based solely on overloading; however, in principle, this limitation
could be overcome using sufficiently powerful partial-evaluation methods.

7.1. Background Discussion

As background to the material that will be discussed in Sect. 7.2, let us reiterate a few points concerning the divided-difference tables that result from computational divided differencing of functions of a single variable. In the following discussion, we assume that the divided-difference table in question has been constructed with respect to some known collection of values x_0, x_1, ..., x_n.

As mentioned at the beginning of Sect. 6, a divided-difference table can be thought of as a (redundant) representation of an interpolating polynomial. For instance, if you have a divided-difference table T (and know the appropriate vector of values x_0, x_1, ..., x_n, as well), you can explicitly construct the interpolating polynomial in Newton form by using the values in the first row of T in accordance with Eqn. (15). One of the consequences of this point is so central to what follows in Sect. 7.2 that it is worthwhile to state it explicitly and to introduce some helpful notation:

OBSERVATION 7.1. [Representation Principle]. *A divided-difference table T is a finite representation of a function $Func[\![T]\!]$ defined by Eqn. (15). (Note that if $F = Func[\![T]\!]$, then $T = F^{\triangledown}$.)*

Given two divided-difference tables, T_1 and T_2, that are defined with respect to the same set of points x_0, x_1, ..., x_n, the operations of matrix addition, subtraction, multiplication, and division applied to T_1 and T_2 yield representations of the sum, difference, product, and quotient, respectively, of $Func[\![T_1]\!]$ and $Func[\![T_2]\!]$. □

In other words, the operations of class `FloatDD` provide ways to (i) instantiate representations of functions of one variable (by evaluating programs in which `float`s have been replaced by `FloatDD`s), and (ii) perform operations on function representations (i.e., by addition, multiplication, etc. of `FloatDD` values).

It is also worthwhile restating the Computational Divided-Differencing Principle (Observation 5.2), adding the additional remark given in the second paragraph:

OBSERVATION 7.2. [Computational Divided-Differencing Principle Redux]. *Rather than computing a divided-difference table with respect to the points x_0, x_1, ..., x_n by invoking the program $n + 1$ times and then applying Eqns. (8) and (9), we may instead evaluate the program (once) using a divided-difference arithmetic that overloads arithmetic operations as matrix operations, substituting $A_{[x_0,...,x_n]}$ for each occurrence of the formal parameter x, and $c * I$ for each occurrence of a constant c.*

Furthermore, this principle can be applied to divided-difference tables for functions on any field (because addition, subtraction, multiplication, and division operations are required, together with additive and multiplicative identity elements). □

7.2. Computational Divided Differencing of Functions of Several Variables

We now consider the problem of defining an appropriate notion of divided differencing for a function F of several variables. Observation 7.1 provides some guidance, as it suggests that the generalized divided-difference table for F that we are trying to create should also be thought of as a representation of a function of several variables that *interpolates* F. Such a generalized computational divided-differencing technique will be based on the combination of Observations 7.1 and 7.2.

Because we have already used the term *higher-order* to refer generically to second, third, ..., n^{th} divided differences, we use the term *higher-kind* to refer to the generalized divided-difference tables that arise with functions of several variables. In the remainder of this section, we make use of an alternative notation for the divided-difference operator $\cdot^{\triangledown[x_0,...,x_n]}$:

$$\mathbf{DD}^1_{[x_0,...,x_n]}[\![F]\!] \stackrel{\text{def}}{=} F^{\triangledown[x_0,...,x_n]}.$$

We use $\mathbf{DD}^1[\![F]\!]$ when the x_i are understood, and abbreviate ranges of variables in the usual way, e.g., $\mathbf{DD}^1_{[x_{0,3}]}[\![F]\!] = \mathbf{DD}^1_{[x_0,x_1,x_2,x_3]}[\![F]\!] = F^{\triangledown[x_0,x_1,x_2,x_3]}$. The notation $\mathbf{DD}^1[\![\cdot]\!]$ refers to divided-difference tables of kind 1 (the kind we are already familiar with from Sect. 5, namely `FloatDD` values). Below, we use $\mathbf{DD}^2[\![\cdot]\!]$ to refer to divided-difference tables of kind 2; in general, we use $\mathbf{DD}^k[\![\cdot]\!]$ to refer to divided-difference tables of kind k.

To understand the basic principle that underlies our approach, consider the problem of creating a surface that interpolates a two-variable function $F(x,y)$ with respect to a grid formed by three coordinate values x_0, x_1, x_2 in the x-dimension, and four coordinate values y_0, y_1, y_2, y_3 in the y-dimension. The clearest way to explain the technique in common programming-language terminology involves currying F. That is, instead of working with $F : float \times float \rightarrow float$, we work with $F : float \rightarrow float \rightarrow float$. We can create (a representation of) an interpolating surface for F (i.e., a divided-difference table of kind 2, denoted by $\mathbf{DD}^2_{[x_{0,2}],[y_{0,3}]}[\![F]\!]$) by building a divided-difference table of

kind 1 using the functions $F(x_0)$, $F(x_1)$, and $F(x_2)$, each of which is of type $float \rightarrow float$, as the "interpolation points".[10]

Note that this process requires that we be capable of performing addition, subtraction, multiplication, and division of *functions*. However, each of the functions $F(x_0)$, $F(x_1)$, and $F(x_2)$ is itself a one-argument function for which we can create a representation, namely by building the divided-difference tables $\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_0)]\!]$, $\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_1)]\!]$, and $\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_2)]\!]$ (with respect to the coordinate values y_0, y_1, y_2, and y_3). By Observation 7.2, the arithmetic operations on functions $F(x_0)$, $F(x_1)$, and $F(x_2)$ needed to create $\mathbf{DD}^2_{[x_{0,2}],[y_{0,3}]}[\![F]\!]$ can be carried out by performing matrix operations on the matrices $\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_0)]\!]$, $\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_1)]\!]$, and $\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_2)]\!]$. For instance,

$$
\begin{aligned}
\mathbf{DD}^1_{[y_{0,3}]}[\![F[x_0,x_1]]\!] &= \mathbf{DD}^1_{[y_{0,3}]}[\![\frac{F(x_0)-F(x_1)}{x_0-x_1}]\!] \\
&= \frac{\mathbf{DD}^1_{[y_{0,3}]}[\![F(x_0)]\!] - \mathbf{DD}^1_{[y_{0,3}]}[\![F(x_1)]\!]}{x_0-x_1}. \quad (17)
\end{aligned}
$$

In what follows, it is convenient to express functions using lambda notation (i.e., in $\lambda z.exp$, z is the name of the formal parameter, and exp is the function body). For instance, $\lambda x.\lambda y.x$ denotes the curried two-argument function (of type $float \rightarrow float \rightarrow float$) that merely returns its first argument. For our purposes, the advantage of lambda notation is that it provides a way to express the anonymous one-argument function that is returned when a curried two-argument function is supplied a value for its first argument (e.g., $(\lambda x.\lambda y.x)(x_0)$ returns $\lambda y.x_0$).

In short, the idea is that a divided-difference table of kind 2 for function F is a matrix of matrices:

$$
\begin{aligned}
\mathbf{DD}^2_{[x_{0,2}],[y_{0,3}]}[\![F]\!] &= \mathbf{DD}^1_{[x_{0,2}]}[\![\lambda x.\mathbf{DD}^1_{[y_{0,3}]}[\![F(x)]\!]]\!] \\
&= \begin{pmatrix} \mathbf{DD}^1_{[y_{0,3}]}[\![F(x_0)]\!] & \mathbf{DD}^1_{[y_{0,3}]}[\![F[x_0,x_1]]\!] & \mathbf{DD}^1_{[y_{0,3}]}[\![F[x_0,x_1,x_2]]\!] \\ & \mathbf{DD}^1_{[y_{0,3}]}[\![F(x_1)]\!] & \mathbf{DD}^1_{[y_{0,3}]}[\![F[x_1,x_2]]\!] \\ & & \mathbf{DD}^1_{[y_{0,3}]}[\![F(x_2)]\!] \end{pmatrix} \\
&= \text{ the matrix shown in Fig. 3} \quad (18)
\end{aligned}
$$

It is instructive to consider some concrete instances of $\mathbf{DD}^2_{[x_{0,2}],[y_{0,3}]}[\![F]\!]$ for various F's:

---

[10] This idea is a specific instance of the very general approach to surface approximation via the tensor-product construction given in (de Boor, 1978, Chap. XVII).

$$
\left(\begin{array}{cc}
\left(\begin{array}{cccc}
F(x_0)(y_0) & F(x_0)[y_{0,1}] & F(x_0)[y_{0,2}] & F(x_0)[y_{0,3}] \\
& F(x_0)(y_1) & F(x_0)[y_{1,2}] & F(x_0)[y_{1,3}] \\
& & F(x_0)(y_2) & F(x_0)[y_{2,3}] \\
& & & F(x_0)(y_3)
\end{array} \right) & \left. \begin{array}{c} \\ \\ \\ \\ \end{array} \right\} \\
\left. \begin{array}{c} \\ \\ \\ \\ \end{array} \right\}
\end{array} \right.
$$

The divided-difference table of kind 2 for function F.

EXAMPLE 7.3. Consider the function $\lambda x.\lambda y.x$. For $0 \le i \le 2$, we have

$$\mathbf{DD}^1_{[y_{0,3}]}[\![(\lambda x.\lambda y.x)(x_i)]\!] = \mathbf{DD}^1_{[y_{0,3}]}[\![\lambda y.x_i]\!] = \begin{pmatrix} x_i & 0 & 0 & 0 \\ & x_i & 0 & 0 \\ & & x_i & 0 \\ & & & x_i \end{pmatrix}$$

and, for $0 \le i \le 1$, we have

$$
\begin{aligned}
\mathbf{DD}^1_{[y_{0,3}]}[\![(\lambda x.\lambda y.x)[x_i, x_{i+1}]]\!] &= \mathbf{DD}^1_{[y_{0,3}]}[\![\frac{(\lambda x.\lambda y.x)(x_i) - (\lambda x.\lambda y.x)(x_{i+1})}{x_i - x_{i+1}}]\!] \\
&= \mathbf{DD}^1_{[y_{0,3}]}[\![\frac{\lambda y.x_i - \lambda y.x_{i+1}}{x_i - x_{i+1}}]\!] \\
&= \mathbf{DD}^1_{[y_{0,3}]}[\![\lambda y.1]\!] \\
&= \begin{pmatrix} 1 & 0 & 0 & 0 \\ & 1 & 0 & 0 \\ & & 1 & 0 \\ & & & 1 \end{pmatrix}
\end{aligned}
$$

Consequently,

$$\mathbf{DD}^2_{[x_{0,2}],[y_{0,3}]}[\![\lambda x.\lambda y.x]\!] = \begin{pmatrix} \begin{pmatrix} x_0 & 0 & 0 & 0 \\ & x_0 & 0 & 0 \\ & & x_0 & 0 \\ & & & x_0 \end{pmatrix} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ & 1 & 0 & 0 \\ & & 1 & 0 \\ & & & 1 \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix} \\ \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} x_1 & 0 & 0 & 0 \\ & x_1 & 0 & 0 \\ & & x_1 & 0 \\ & & & x_1 \end{pmatrix} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ & 1 & 0 & 0 \\ & & 1 & 0 \\ & & & 1 \end{pmatrix} \\ \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} x_2 & 0 & 0 & 0 \\ & x_2 & 0 & 0 \\ & & x_2 & 0 \\ & & & x_2 \end{pmatrix} \end{pmatrix}$$

$$(19)$$

$\square$

EXAMPLE 7.4. Consider the function $\lambda x.\lambda y.y$. For $0 \le i \le 2$, we have

$$\mathbf{DD}^1_{[y_{0,3}]}[\![(\lambda x.\lambda y.y)(x_i)]\!] = \mathbf{DD}^1_{[y_{0,3}]}[\![\lambda y.y]\!] = \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix}$$

and, for $0 \le i \le 1$, we have

$$\mathbf{DD}^1_{[y_{0,3}]}[\![(\lambda x.\lambda y.y)[x_i, x_{i+1}]]\!] = \mathbf{DD}^1_{[y_{0,3}]}[\![\frac{(\lambda x.\lambda y.y)(x_i) - (\lambda x.\lambda y.y)(x_{i+1})}{x_i - x_{i+1}}]\!]$$

$$= \mathbf{DD}^1_{[y_{0,3}]}[\![\frac{\lambda y.y - \lambda y.y}{x_i - x_{i+1}}]\!]$$

$$= \mathbf{DD}^1_{[y_{0,3}]}[\![\lambda y.0]\!]$$

$$= \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix}$$

Consequently, we have

$$\mathbf{DD}^2_{[x_{0,2}],[y_{0,3}]}[\![\lambda x.\lambda y.y]\!] = \begin{pmatrix} \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix} \\ \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix} \\ \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix} \end{pmatrix} \quad (20)$$

□

Moreover, Observation 7.2 tells us that divided-difference tables for functions of two variables can be built up by means of a divided-difference arithmetic that operates on matrices of matrices. That is, we can build up divided-difference tables of kind 2 for more complex functions of x and y by using operations on matrices of matrices, substituting $\mathbf{DD}^2[\![\lambda x.\lambda y.x]\!]$ for each occurrence of the formal parameter x in the function, and $\mathbf{DD}^2[\![\lambda x.\lambda y.y]\!]$ for each occurrence of the formal parameter y.

EXAMPLE 7.5. For the function $\lambda x.\lambda y.(x \times y)$, $\mathbf{DD}^2[\![\lambda x.\lambda y.(x \times y)]\!]$ can be created by *multiplying the matrices* $\mathbf{DD}^2[\![\lambda x.\lambda y.x]\!]$ and $\mathbf{DD}^2[\![\lambda x.\lambda y.y]\!]$ from Eqns. (19) and (20), respectively:

$$\mathbf{DD}^2[\![\lambda x.\lambda y.(x \times y)]\!] = \mathbf{DD}^2[\![(\lambda x.\lambda y.x) \times (\lambda x.\lambda y.y)]\!] \quad (21)$$
$$= \mathbf{DD}^2[\![\lambda x.\lambda y.x]\!] \times \mathbf{DD}^2[\![\lambda x.\lambda y.y]\!]$$

$$= \begin{pmatrix} \begin{pmatrix} x_0 y_0 & x_0 & 0 & 0 \\ & x_0 y_1 & x_0 & 0 \\ & & x_0 y_2 & x_0 \\ & & & x_0 y_3 \end{pmatrix} & \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix} \\ \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} x_1 y_0 & x_1 & 0 & 0 \\ & x_1 y_1 & x_1 & 0 \\ & & x_1 y_2 & x_1 \\ & & & x_1 y_3 \end{pmatrix} & \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix} \\ & \begin{pmatrix} & & & \\ & & & \\ & & & \\ & & & \end{pmatrix} & \begin{pmatrix} x_2 y_0 & x_2 & 0 & 0 \\ & x_2 y_1 & x_2 & 0 \\ & & x_2 y_2 & x_0 \\ & & & x_2 y_3 \end{pmatrix} \end{pmatrix}$$

Note, for example, that the (0,1) entry in the above matrix, namely

$$\begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix}$$

was obtained via the calculation

$$\begin{pmatrix} x_0 & 0 & 0 & 0 \\ & x_0 & 0 & 0 \\ & & x_0 & 0 \\ & & & x_0 \end{pmatrix} \times \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix}$$
$$+ \begin{pmatrix} 1 & 0 & 0 & 0 \\ & 1 & 0 & 0 \\ & & 1 & 0 \\ & & & 1 \end{pmatrix} \times \begin{pmatrix} y_0 & 1 & 0 & 0 \\ & y_1 & 1 & 0 \\ & & y_2 & 1 \\ & & & y_3 \end{pmatrix} \qquad (22)$$
$$+ \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix} \times \begin{pmatrix} 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix}$$

and *not* by the use of Eqn. (17), which involves a matrix subtraction, a scalar subtraction, and a scalar division:

$$\mathbf{DD}^1_{[y_0,3]}[\![(\lambda x.\lambda y.(x \times y))[x_0, x_1]]\!] = \frac{\mathbf{DD}^1_{[y_0,3]}[\![\lambda y.(x_0 \times y)]\!] - \mathbf{DD}^1_{[y_0,3]}[\![\lambda y.(x_1 \times y)]\!]}{x_0 - x_1}$$

$$= \frac{\begin{pmatrix} x_0 y_0 & x_0 & 0 & 0 \\ & x_0 y_1 & x_0 & 0 \\ & & x_0 y_2 & x_0 \\ & & & x_0 y_3 \end{pmatrix} - \begin{pmatrix} x_1 y_0 & x_1 & 0 & 0 \\ & x_1 y_1 & x_1 & 0 \\ & & x_1 y_2 & x_1 \\ & & & x_1 y_3 \end{pmatrix}}{x_0 - x_1} \qquad (23)$$

Expressions (22) and (23) are equivalent over real numbers, but not over floating-point numbers. By sidestepping the explicit subtraction and division operations in expression (23), expression (22) avoids the potentially disasterous magnification of round-off error that can occur with floating-point arithmetic. □

The principle illustrated in Example 7.5 gives us the machinery that we need to perform computational divided differencing for bivariate functions defined by programs. As usual, computational divided differencing is performed by changing the types of formal parameters, local variables, and return values to the type of an appropriate divided-difference arithmetic.

Furthermore, these ideas can be applied to a function F with an arbitrary number of variables: when F has k variables, $\mathbf{DD}^k[\![F]\!]$, F's divided-difference table of kind k, is a matrix of matrices of ... of matrices nested to depth k. Currying with respect to the first parameter of F "peels off" one dimension; $\mathbf{DD}^k[\![F]\!]$ is a matrix whose entries are divided-difference tables of kind $k-1$ (i.e., matrices of matrices of ... of matrices nested to depth $k-1$). For instance, the diagonal entries are the divided-difference tables of kind $k-1$ for the $(k-1)$-parameter functions $F(x_0)$, $F(x_1)$, ..., $F(x_n)$ (i.e., $\mathbf{DD}^{k-1}[\![F(x_0)]\!]$, $\mathbf{DD}^{k-1}[\![F(x_1)]\!]$, ..., $\mathbf{DD}^{k-1}[\![F(x_n)]\!]$).

To implement this approach in C++, we define two classes and one class template:

- Class template `template <int k> class DivDiffArith` can be instantiated with a value `k > 0` to represent divided-difference tables of kind `k`. Each object of class `DivDiffArith<k>` has links to sub-objects of class `DivDiffArith<k-1>`.

- Class `DivDiffArith<0>` represents the base case; `DivDiffArith<0>` objects simply hold a single `float`.

- Class `IntVector`, a vector of `int`'s, is used to describe the number of points in each dimension of the grid of coordinate points.

Excerpts from the definitions of these classes are shown below:

```
template <int k> class DivDiffArith {
 public:
  int numPts;
  DivDiffArith<k-1> **divDiffTable; // Two-dimensional upper-triangular array
  DivDiffArith(const FloatV &v, const IntVector &grid, int d);
  DivDiffArith(float, const IntVector &grid); // constant; shape conforms to grid
  DivDiffArith(float, const DivDiffArith<k-1> &dda); // constant; shape conforms to dda
  DivDiffArith<k>& operator+ (const DivDiffArith<k> &) const; // binary addition
  DivDiffArith<k>& operator- (const DivDiffArith<k> &) const; // binary subtraction
  DivDiffArith<k>& operator* (const DivDiffArith<k> &) const; // binary multiplication
  DivDiffArith<k>& operator/ (const DivDiffArith<k> &) const; // binary division
};
```

```
class DivDiffArith<0> {
 public:
  float value;
  DivDiffArith(float v = 0.0); // Default constructor
  DivDiffArith(const FloatV &v, const IntVector &grid, int d);
  DivDiffArith<0>& operator+ (const DivDiffArith<0> &) const; // binary addition
  DivDiffArith<0>& operator- (const DivDiffArith<0> &) const; // binary subtraction
  DivDiffArith<0>& operator* (const DivDiffArith<0> &) const; // binary multiplication
  DivDiffArith<0>& operator/ (const DivDiffArith<0> &) const; // binary division
};
```

```
class IntVector {
 public:
  int numPts;
  int *val; // An array of values: val[0]..val[numPts-1]
  IntVector();
  IntVector(int N, ...); // Construct IntVector given N values
  IntVector& operator<< (const int i); // left shift
};
```

The operations of class `DivDiffArith<k>` are overloaded in a fashion similar to those of class `FloatDD`. (Class `FloatDD` is essentially identical to `DivDiffArith<1>`.) For instance, the overloaded multiplication operator performs matrix multiplication:

```
template <int k>
DivDiffArith<k>& DivDiffArith<k>::operator* (const DivDiffArith<k> &dda) const{
  assert(numPts == dda.numPts);
  DivDiffArith<k> *ans = new DivDiffArith<k>(numPts);
  for (int r = 0; r < numPts; r++) {
    for (int c = r; c < numPts; c++) {
      DivDiffArith<k-1> temp((float)0.0, divDiffTable[r][c]); // temp = 0.0
      for (int j = r; j <= c; j++) {
        temp += divDiffTable[r][j] * dda.divDiffTable[j][c];
      }
      ans->divDiffTable[r][c] = temp;
    }
  }
  return *ans;
}
```

Class `DivDiffArith<k>` has two constructors for creating a `DivDiffArith<k>` object from a `float` constant. They differ only in their second arguments (an `IntVector` versus a `DivDiffArith<k>`), which are used to

determine the appropriate dimensions to use at each level in the nesting of matrices.

Suppose that in the function on which computational divided differencing is to be carried out, variable z is the independent variable associated with argument position $d+1$. To generate an appropriate DivDiffArith<k> object for z for a given set of grid values z_0, ..., z_m, a FloatV with the values z_0, ..., z_m is created, and then passed to the following DivDiffArith<k> constructor:

```
template <int k>
DivDiffArith<k>::DivDiffArith(const FloatV &v, const IntVector &grid, int d) :
  numPts(grid.val[0]),
  divDiffTable(calloc_ut< k >(numPts))
{
  assert(grid.val[d] == v.numPts);
  IntVector tail = grid << 1;
  for (int r = 0; r < numPts; r++) {
    for (int c = r; c < numPts; c++) {
      divDiffTable[r][c] = DivDiffArith<k-1>((float)0.0,tail);
    }
  }
  if(d == 0) {
    DivDiffArith<k-1> one((float)1.0, tail);
    for (int i = 0; i < numPts; i++) {
      divDiffTable[i][i] = DivDiffArith<k-1>(v.val[i],tail);
      if(i < numPts - 1) {
        divDiffTable[i][i+1] = one;
      }
    }
  }
  else {
    for (int i = 0; i < numPts; i++) {
      divDiffTable[i][i] = DivDiffArith<k-1>(v,tail,d-1);
    }
  }
}
```

EXAMPLE 7.6. The following code fragment generates two DivDiffArith<2> values, x and y, which correspond to the matrices shown in Eqns. (19) and (20), respectively:

```
IntVector grid(2,3,4);
FloatV fv_x(3,x_0,x_1,x_2);
DivDiffArith<2> x(fv_x,grid,0); // argument position 1
FloatV fv_y(4,y_0,y_1,y_2,y_3);
DivDiffArith<2> y(fv_y,grid,1); // argument position 2
```

□

EXAMPLE 7.7. Consider a C++ class `BivariatePoly` that represents bivariate polynomials, and a member function `BivariatePoly::Eval` that evaluates a polynomial via a bivariate version of Horner's rule:

```
class BivariatePoly {
 public:
  float Eval(float,float);
 private:
  int degree1,degree2;
  // Array coeff[0..degree1][0..degree2]
  float **coeff;
};


// Evaluation via bivariate Horner's rule
float BivariatePoly::Eval(float x,float y){
  float ans = 0.0;
  for (int i = degree1; i >= 0; i--){
    float temp = 0.0;
    for (int j = degree2; j >= 0; j--){
      temp = temp * y + coeff[i][j];
    }
    ans = ans * x + temp;
  }
  return ans;
}
```

Similar to what has been done in Examples 3.1, 4.1, 5.3, and 6.1, computational divided differencing is carried out on this version of `Eval` by changing the types of its formal parameters, local variables, and return value from `float` to `DivDiffArith<2>`.

```
// Evaluation via bivariate Horner's rule
DivDiffArith<2> BivariatePoly::Eval(const DivDiffArith<2> &x,
                                    const DivDiffArith<2> &y)
{
  DivDiffArith<2> ans(0.0,x); // ans = 0.0
  for (int i = degree1; i >= 0; i--){
    DivDiffArith<2> temp(0.0,y); // temp = 0.0
    for (int j = degree2; j >= 0; j--){
      temp = temp * y + coeff[i][j];
    }
    ans = ans * x + temp;
  }
  return ans;
}
```

To use this procedure to create a divided-difference table of kind 2 for a given variable P of type `BivariatePoly*`, with respect to the 3-by-4 grid $\{x_0,x_1,x_2\} \times \{y_0,y_1,y_2,y_3\}$, we would generate the `IntVector grid` and `DivDiffArith<2>` values `x` and `y` as shown in Example 7.6, and then invoke

$$P\text{->}Eval(x,y);$$

□

One final point concerning this approach: it is worthwhile noting that by generalizing the grid descriptors slightly, it is possible to devise an even more general divided-differencing arithmetic that is *heterogeneous* in shape with respect to different argument positions. By "heterogeneous", we mean that full two-dimensional (upper-triangular) divided-difference tables could be provided for some argument positions, while other argument positions could just provide a single row of divided differences (i.e., one-dimensional, `FloatDDR1`-like tables). By this means, when a procedure body is "accumulative" in certain of its formal parameters but not others, it would be possible to tailor the divided-differencing version of the procedure to improve its efficiency. (In the case of procedure

$$DivDiffArith<2> \ BivariatePoly::Eval,$$

it would be possible to specify that both argument positions provide `FloatDDR1`-like tables.)

8. Paige's Work on Finite Differencing of Computable Expressions

In earlier sections, and also in Sect. 9, we have attempted to place the ideas that are developed in this paper in their proper context by describing how they relate to previous work on computational-differentiation (Wengert, 1964; Rall, 1981; Griewank and Corliss, 1992; Berz et al., 1996; Griewank, 2000) and to previous work on the creation of accurate divided-difference tables for expressions (Opitz, 1964; Mc-Curdy, 1980). In the present section, we discuss some additional intellectual forbearers of our work, focusing particularly on how our ideas relate to Robert Paige's work on finite differencing of set-valued expressions in SETL programs.

Starting in the mid-70s, Paige studied how finite-differencing transformations of applicative set-former expressions could be exploited to optimize loops in very-high-level languages, such as SETL. References (Paige, 1981; Paige and Koenig, 1982) are just two of many works that Paige wrote about this subject, and these ideas were implemented in his RAPTS system (Paige and Koenig, 1982; Paige, 1983). Some of the techniques that Paige explored have their roots in earlier work by Earley (Earley, 1974; Earley, 1976). Independently of and contemporaneously with Paige, similar loop-optimization methods targeted toward very-high-level set-theoretic languages were investigated by Fong and Ullman (Fong and Ullman, 1976; Fong, 1977; Fong, 1979). More recently, Liu and Stoller have used some extensions of these ideas to optimize array computations (Liu and Stoller, 1998) and recursive programs (Liu and Stoller, 2000). Liu et al. have also shown how such transformations can be applied to derive algorithms for incremental-computation problems (i.e., problems in which the goal is to maintain the value of some function $F(x)$ as the input x undergoes small changes) (Liu, 1995; Liu and Teitelbaum, 1995a; Liu and Teitelbaum, 1995b; Liu et al., 1996; Liu and Stoller, 1999).

Both Paige (Paige, 1981) and Paige and Koenig (Paige and Koenig, 1982) provide lengthy discussions of the roots of the ideas that are developed in those papers. The basic idea for optimizing SETL loops is described as a generalization of strength reduction, an idea attributed to John Cocke from the 60s, whereby a loop is transformed so that a multiplication operation in the loop is eliminated in favor of an addition, as shown below:

T.W. Reps and L.B. Rall

```
                          i = ...;
                          T = i*c; // T depends on i
  i = ...;                deltaT = delta*c;
  while (...) {           while (...) {
     ... i*c ...;     ⇒      ... T ...; // T replaces i*c
    i = i + delta;           i = i + delta; // change to i
  }                          T = T + deltaT // update of T
                          }
```

This transformation improves the running time of the loop if the cost of the additions performed by the transformed loop are less than the cost of the multiplications performed in the original loop. In (Cocke and Schwartz, 1970), Cocke and Schwartz presented a variety of strength-reduction transformations for use in optimizing compilers.

Paige's work on loop optimization in SETL was based on the observation that a similar transformation could be applied to loops that involve set-former expressions. In this transformation, an expensive set-former expression in a loop is replaced by a set-initialization statement (placed before the loop) together with a set-update operation (placed inside the loop):

```
                                   A = ...;
                                   T = { x ∈ A | x%2 == 0 }; // T depends on A
  A = ...;                         while (...) {
  while (...) {                       ... T ...; // T replaces the set former
     ... { x ∈ A | x%2 == 0 } ...;    d = ...;
    d = ...;                 ⇒       A = A ∪ d; // change to A
    A = A ∪ d;                       if (d%2 == 0) T = T ∪ d; // update of T
  }                                }
```

In the transformed program shown above on the right, the expression { x ∈ A | x%2 == 0 } in the loop is replaced by a use of T. Because the statement

$$A = A \cup d;$$

may alter the value of variable A, just after this statement a new statement is introduced:

$$\text{if } (d\%2 == 0)\ T = T \cup d;$$

The latter statement updates the value of variable T to have the same value that the expression { x ∈ A | x%2 == 0 } has when evaluated with the new value of variable A.[11]

Of more direct relevance to the topic of the present paper is the discussion in Paige's papers in which he points out affinities between his work, on the one hand, and differentiation and finite differencing of numerical functions, on the other hand.[12] For instance, Paige and Koenig describe the relationship between their SETL finite-differencing methods and numerical finite-difference methods as follows (Paige and Koenig, 1982, pp. 403–404):

> It is interesting to note that the origins of our method may be traced back to the finite difference techniques introduced by the English mathematician Henry Briggs in the sixteenth century. His method, which can be used to generate a sequence of polynomial values $p(x_0)$, $p(x_0 + h)$, $p(x_0 + 2h)$, ..., hinges on the following idea. For a given polynomial $p(x)$ of degree n and an increment h, the first difference polynomial
>
> $$p_1(x) = p(x + h) - p(x)$$
>
> is of degree $n - 1$ or less, the second difference polynomial
>
> $$p_2(x) = p_1(x + h) - p_1(x)$$
>
> is of degree $n - 2$ or less, ..., and, finally, $p_n(x)$ must be a constant. Thus, to tabulate successive values of $p(x)$ starting with $x = x_0$, we can perform these two steps:
>
> **1.** Calculate initial values for $p(x_0)$, $p_1(x_0)$, ..., $p_n(x_0)$ and store them in $t(1)$, $t(2)$, ..., $t(n + 1)$.
>
> **2.** Generate the desired polynomial table by iterating over the following code block:

---

[11] The code fragment if (d%2 == 0) T = T ∪ d; is called a *post-derivative* of T = { x ∈ A | x%2 == 0 }; with respect to the change A = A ∪ d;. Similarly, a code fragment for updating variable T that is placed before the change A = A ∪ d; is called a *pre-derivative*.

[12] The (forward) *finite difference* of a function F with respect to h is defined as follows:
$$\Delta_h F(x) \stackrel{\text{def}}{=} F(x + h) - F(x).$$

> **print** x, $t(1)$;   \$ print x and $p(x)$
> $t(1) := t(1) + t(2)$;   \$ place new values for
> $t(2) := t(2) + t(3)$;   \$ $p(x)$, $p_1(x)$, ..., $p_{n-1}(x)$
> $\vdots$ \$ into
> $t(n) := t(n) + t(n+1)$;   \$ $t(1)$, $t(2)$, ..., $t(n)$.
> $x := x + h$;   \$

> Note that Briggs's method requires only n additions in step 2
> to compute each new polynomial value, while Horner's rule to
> compute a fresh polynomial value costs n additions and n mul-
> tiplications.

They relate Briggs's method to strength reduction in the following
passage (Paige and Koenig, 1982, pp. 404–405):

> Although Cocke's technique does not treat polynomials as special
> objects, strength reduction is sufficiently powerful to transform a
> program involving repeated calculations of a polynomial accord-
> ing to Horner's rule into an equivalent program that essentially
> uses the more efficient finite difference method of Briggs. Indeed,
> this is a surprising and important result that demonstrates that
> the success of polynomial evaluation by differencing results from
> properties of the elementary operations used to form polynomials
> rather than from properties exclusive to polynomials. In other
> words, Cocke's method works because the following distributive
> and associative laws hold for sums and products:
>
> $$(i \pm delta) * c \Rightarrow i * c \pm delta * c;$$
> $$(i \pm delta) + c \Rightarrow (i + c) \pm delta.$$
>
> In (Cocke and Schwartz, 1970) Cocke and Schwartz extend
> this idea to show how reduction in strength (which we call finite
> differencing) applies to a wide range of arithmetic operations that
> exhibit appropriate distributive properties.

Later in the paper, after Paige and Koenig have introduced their
rules for finite differencing of set-former expressions with respect to
changes in argument values (an operation that they sometimes call "dif-
ferentiation"), they return to the discussion of Briggs's method (Paige
and Koenig, 1982, pp. 421):

> Profitable differentiation of an expression f can sometimes be
> supported by differentiating f together with a chain of auxiliary
> expressions (as in Briggs's first, second, ..., difference polynomials
> ...). Thus, the prederivative $\nabla^- E\langle x +:= delta;\rangle$ of the n^{th} degree
> polynomial $E = P(x)$ is
>
> $$E +:= P_1(x)$$

where $P_1(x)$ is the first difference polynomial. However, for the prederivative code above to be inexpensive, we must also differentiate the second, third, ..., n^{th} difference polynomials, denoted $E_i = P_i(x)$, $i = 2..n$. To realize Briggs's efficient technique, we consider the extended prederivative (of expressions ordered carefully into a "differentiable chain") $\nabla^- E_{n-1}, \ldots, E_1, E\langle x +:= delta; \rangle$ that expands into

$$E +:= E_1;$$
$$E_1 +:= E_2;$$
$$\vdots$$
$$E_{n-1} +:= E_n;$$

Essentially all of the material that Paige and Koenig present in their paper to relate their work to Briggs's method has been quoted above. However, a few details about the derivation of Briggs's method were not spelled out in their treatment, which we now attempt to rectify. We will show below that computational divided differencing supplies a clean way to handle an important step in the derivation for which what Paige and Koenig say is ambiguous.

The initial program for tabulating a polynomial at a collection of equally spaced points can be written in C++ as follows:[13]

```
void Poly::Tabulate(float start, int numPoints, float h){
  float x = start;
  float y;
  for (int i = 1; i <= numPoints; i++) { // Tabulation loop
    y = Eval(x);
    cout << x << ": " << y << endl;
    x += h;
  }
}
```

In the program shown above, `Eval` is the member function of class `Poly` that evaluates a polynomial via Horner's rule (see Sect. 3), of type

```
float Poly::Eval(float x);
```

---

[13] It should be pointed out that in practice it is better to code the statement in the loop that changes the value of `x` as "`x = start + i * h;`", rather than as "`x += h;`", so that small errors in `h` do not accumulate in `x` due to repeated addition. We have chosen to use the latter form to emphasize the similarities between `Tabulate` and the two earlier strength-reduction examples.

The intention of Paige and Koenig is to transform procedure `Poly::Tabulate` into something like the following version:

```
void Poly::Tabulate(float start, int numPoints, float h){
  float x = start;
  float y;
  float E = Eval(start); // Call on Eval hoisted out of the loop
  float E₁ = ???; // Unspecified initialization
    ⋮
  float E_{n-1} = ???;
  float E_n = ???;
  for (int i = 1; i <= numPoints; i++) { // Tabulation loop
    y = E;
    cout << x << ": " << y << endl;
    E += E₁; // Extended pre-derivative w.r.t. x += h;
    E₁ += E₂;
      ⋮
    E_{n-1} += E_n;
    x += h;
  }
}
```

However, as indicated by the question marks in the above code, Paige and Koenig do not state explicitly how they plan to arrive at the proper initialization code that is to be placed just before the loop in the transformed program. It is unclear whether they intend to generate the values E_1, E_2, ..., E_{n-1}, E_n by evaluating polynomial P at `start`, `start+h`, ..., `start+(n-1)*h`, `start+n*h` and then create the E_i via subtraction operations, or whether they intend to generate the finite-difference polynomials P_1, P_2, ..., P_{n-1}, P_n symbolically and then apply each of them to `start`. The former method can lead to very inaccurate results (see below), whereas the latter method requires that a substantial amount of symbolic manipulation be performed to generate the P_i.

However, the divided-difference arithmetic `FloatDDR1` gives us an easy way to create suitable initialization code that produces accurate values for the E_i. This initialization code consists of three steps:

— Create a `FloatV` of equally spaced points, starting at `start` and separated by `h`, where the number of points is one more than the degree of the polynomial.

— Introduce a single call on the member function

```
FloatDDR1 Poly::Eval(const FloatV &x);
```

to create the first row of the divided-difference table for the polynomial with respect to the given `FloatV`.

— Convert the resulting `FloatDDR1` (which holds *divided-difference* values) into the first row of a *finite-difference* table for the polynomial by multiplying each entry by an appropriate adjustment factor (see (Conte and de Boor, 1972, Lemma 4.1)).

This initialization method is used in the version of `Tabulate` shown below. (In this version of `Tabulate`, the E_i are renamed `diffTable[i]`.)

```
void Poly::Tabulate(float start, int numPoints, float h){
   float x = start;
   float y;
   // Create accurate divided-difference table
   FloatV fv(x, degree+1, h);
   FloatDDR1 fddr1 = Eval(fv); // Calls FloatDDR1 Poly::Eval(const FloatV &);
   // Convert divided-difference entries to finite-difference entries
   float *diffTable(new float[degree+1]);
   float adjustment = 1.0;
   for (int i = 0; i <= degree; i++) {
      diffTable[i] = fddr1.divDiffTable[i] * adjustment;
      adjustment *= (h * (i+1));
   }
   for (int i = 1; i <= numPoints; i++) { // Tabulation loop
      y = diffTable[0];
      cout << x << ": " << y << endl;
      for (int j = 0; j < degree; j++) { // Pre-derivative w.r.t. x += h;
         diffTable[j] += diffTable[j+1];
      }
      x += h;
   }
}
```

EMPIRICAL RESULTS: TABULATION OF A POLYNOMIAL VIA
BRIGGS'S METHOD

We now present some empirical results that illustrate the advantages of
the final version of `Poly::Tabulate` presented above. Again, we work
with the polynomial $P(x) = 2.1 * x^3 - 1.4 * x^2 - .6 * x + 1.1$, and perform
computations using single-precision floating-point arithmetic on a Sun
SPARCstation 20/61 running SunOS 5.6. Programs were compiled with
the egcs-2.91.66 version of `g++` (egcs-1.1.2 release) with optimization at
the `-O1` level.

The final version of `Poly::Tabulate` uses the divided-difference
arithmetic `FloatDDR1` in the initialization step that creates the ini-
tial finite-difference vector `diffTable`. An alternative way to gener-
ate `diffTable` is to evaluate polynomial P at `start`, `start+h`, ...,
`start+(n-1)*h`, `start+n*h` and then create `diffTable` via subtraction
operations, according to the standard definition (Conte and de Boor,
1972, pp. 214). However, the latter way of generating `diffTable` in-
volves subtraction operations, and hence may magnify any round-off
errors in the $n + 1$ values computed for P. In contrast, the method
using computational divided differencing yields a way to create a more
accurate initial finite-difference table.

To give a concrete illustration of the benefits, the following table
shows what happens when $P(x)$ is evaluated at the 10,001 points in
the interval $[0.0, 1.0]$ with a grid spacing of .0001. (Places where the
results from the three methods differ are indicated in boldface.)

	Evaluate via Horner	Comp. Div. Diff. + Briggs	Standard Finite Diff. + Briggs
x	P(x)	P(x)	P(x)
0.0000	1.10000	1.10000	1.10000
0.0001	1.09994	1.09994	1.09994
0.0002	1.09988	1.09988	1.09988
...	...	...	...
0.9998	1.19942	1.19941	**19844.3**
0.9999	1.19971	1.19970	**19850.3**
1.0000	1.**20000**	1.19999	**19856.2**
Time (milliseconds)	7.64	5.62	5.49

The numbers that appear in the rightmost column for P(.9998), P(.9999),
and P(1.0000) are *not* typographical errors. What happens is that

round-off errors in the computation of the initial values for P are magnified by the standard method for computing the initial finite-difference table, and after 10,000 iterations of the Briggs calculation, the accumulated effect has caused the values produced to diverge widely from the correct answers.

Overall, the method based on computational divided differencing is far more accurate than the one in which the vector needed for Briggs's method is obtained by subtraction operations (and only 2% slower). Furthermore, the results from the method based on computational divided differencing are nearly as accurate as those obtained by reevaluating the polynomial at each point, but the reevaluation method is 36% slower.

9. Other Related Work

Sect. 8 described how our ideas relate to Robert Paige's work on finite differencing of set-valued expressions in SETL programs. This section concerns other related work, which falls into four categories:

COMPUTATIONAL DIFFERENTIATION

Computational differentiation is a well-established area of numerical analysis, with its own substantial literature (Wengert, 1964; Rall, 1981; Griewank and Corliss, 1992; Berz et al., 1996; Griewank, 2000). The importance of the subject is underscored by the fact that the 1995 Wilkinson Prize for outstanding contributions to the field of numerical software was awarded to Chris Bischof (then at Argonne) and Alan Carle (Rice) for the development of the FORTRAN computational-differentiation system ADIFOR 2 (Bischof et al., 1996). As discussed in Sect. 4, computational divided differencing is a *generalization* of computational differentiation: a program resulting from computational divided differencing can be used to obtain derivatives (as well as divided differences), whereas a program resulting from computational differentiation can only produce derivatives (and not divided differences).

OTHER WORK ON ACCURATE DIVIDED DIFFERENCING

The program-transformation techniques for creating accurate divided differences described in this paper are based on a 1964 result of Opitz's (Opitz, 1964), which was later rediscovered in 1980 by McCurdy (McCurdy, 1980) and again in 1998 by one of us (Reps). However, Opitz and McCurdy both discuss how to create accurate divided differences only

for *expressions*. In this paper, the idea has been applied to the creation of accurate divided differences for functions defined by *programs*.

As already mentioned in Sect. 2, the computational-differentiation technique that was summarized there is what is known as *forward-mode* differentiation. A different computational-differentiation technique, *reverse mode* (Linnainmaa, 1976; Speelpenning, 1980; Iri, 1984; Griewank, 1989; Griewank, 1991), is generally preferable when the number of independent variables is much greater than the number of dependent variables. However, although it is possible to develop a reverse-mode version of computational divided differencing, it does not appear to offer the same potential savings in operations performed that reverse mode achieves for computational differentiation.

McCurdy, and later Kahan and Fateman (Kahan and Fateman, 1985) and Rall and Reps (Rall and Reps, 2000), have looked at ways to compute accurate divided differences for library functions (i.e., sin, cos, exp, etc.).

Kahan and Fateman have also investigated how similar techniques can be used to avoid unsatisfactory numerical answers when evaluating formulas returned by symbolic-algebra systems. In particular, their work was motivated by the observation that naive evaluation of a definite integral $\int_a^b f(x)dx$ can sometimes produce meaningless answers: When a symbolic-algebra system produces a closed-form solution for the indefinite integral $\int f(x)dx$, say $G(x)$, the result of the computation $G(b) - G(a)$ may have no significant digits. Kahan and Fateman show that divided differences can be used to develop accurate numerical formulas that sidestep this problem.

One of the techniques developed by McCurdy for computing accurate divided-difference tables involved first computing just the first row of the table and then generating the rest of the entries by a backfilling algorithm. He studied the conditions under which this technique maintained sufficient accuracy. However, his algorithm for accurately computing the first row of the divided-difference table was based on a series expansion of the function, rather than a divided-difference arithmetic, such as the `FloatDDR1` arithmetic developed in Sect. 6.

Divided-difference arithmetic for first divided differences has also been called *slope arithmetic*, and an interval version of it has been investigated previously as a way to obtain an interval enclosure for the range of a function evaluated over an interval (Krawczyk and Neumaier, 1985; Zuhe and Wolfe, 1990; Ratz, 1996). It has been shown that interval slope arithmetic can yield tighter enclosures than methods based on derivatives (Krawczyk and Neumaier, 1985; Zuhe and Wolfe, 1990; Ratz, 1996). Zuhe and Wolfe have also shown that a form of interval divided-difference arithmetic involving second divided differ-

ences can provide even tighter interval enclosures (Zuhe and Wolfe, 1990). In the present paper, we have confined ourselves to point (non-interval) divided-difference arithmetics, but have explored the use of a divided-difference arithmetic for divided differences of arbitrary order (originally due to Opitz (Opitz, 1964)). We have shown how to extend the latter divided-difference arithmetic to the case of multi-variate functions involving a fixed, but arbitrary, number of variables, and have also proposed various specializations of it, which would improve run-time efficiency in certain situations.

Other Work on Controlling Round-Off Error in Numerical Computations

Computational differentiation and computational divided differencing are methods for controlling round-off error that can arise in two types of numerical computations. An extensive collection of methods for controlling round-off error for a wide variety of numerical computations has been developed by Kulisch's group at the University of Karlsruhe (Hammer et al., 1993; Hammer et al., 1995).

In future work, we plan to investigate the use of such techniques to achieve greater accuracy (and to validate results) in steps of the computational-divided-differencing transformations where operations are implemented by solving systems of linear equations, such as division operations (cf. box (5)).

Other Work that Exploits Operator Overloading

The operator-overloading feature of C++ provides a convenient mechanism for implementing the divided-difference arithmetics that are described in this paper. Compared to other methods of providing polymorphism in programming languages, such as *parametric polymorphism* (Milner, 1978) and *inheritance polymorphism* (Cardelli, 1984), operator overloading has always been something of a poor cousin. In the programming-languages community, operator overloading is sometimes referred to as "*ad hoc* polymorphism" (Strachey, 2000), a term that has pejorative overtones. Nevertheless, in addition to its application in computational differentiation and computational divided differencing, operator overloading provides a convenient implementation mechanism for a wide variety of other interesting applications, including partial evaluation (Andersen, 1994), safe pointers (Austin et al., 1994), and expression templates (Veldhuizen, 1995).

Acknowledgements

References

Andersen, L. O.: 1994, 'Program Analysis and Specialization for the C Programming Language'. Ph.D. thesis, DIKU, Univ. of Copenhagen. (DIKU report 94/19).

Austin, T., S. Breach, and G. Sohi: 1994, 'Efficient Detection of all Pointer and Array Access Errors'. In: *SIGPLAN Conference on Programming Languages Design and Implementation*. New York, NY, pp. 290–301, ACM Press.

Bates, S. and S. Horwitz: 1993, 'Incremental Program Testing Using Program Dependence Graphs'. In: *Symp. on Princ. of Prog. Lang.* pp. 384–396.

Beck, T. and H. Fischer: 1994, 'The If-Problem in Automatic Differentiation'. *J. Comp. and Appl. Math.* **50**, 119–131.

Berz, M., C. Bischof, G. Corliss, and A. Griewank (eds.): 1996, *Computational Differentiation: Techniques, Applications, and Tools*. Philadelphia, PA: Soc. for Indust. and Appl. Math.

Binkley, D.: 1992, 'Using Semantic Differencing to Reduce the Cost of Regression Testing'. In: *Proceedings of the 1992 Conference on Software Maintenance* (Orlando, FL, November 9-12, 1992). pp. 41–50.

Bischof, C., A. Carle, G. Corliss, and A. Griewank: 1992, 'ADIFOR: Automatic Differentiation in a Source Translator Environment'. In: *Proc. of ISSAC 1992: Int. Symp. on Symb. and Alg. Comp.* New York, NY, pp. 294–302, ACM Press.

Bischof, C., A. Carle, P. Khademi, and A. Mauer: 1996, 'ADIFOR 2.0: Automatic Differentiation of Fortran 77 Programs'. *IEEE Comp. Sci. and Eng.* **3**, 18–32.

Bischof, C., L. Roh, and A. Mauer: 1997, 'ADIC: An Extensible Automatic Differentiation Tool for ANSI-C'. Tech. Rep. ANL/MCS-P626-1196, Math. and Comp. Sci. Div., Argonne Nat. Lab., Argonne, IL.

Bjørner, D., A. Ershov, and N. Jones (eds.): 1988, 'Partial Evaluation and Mixed Computation: Proc. of the IFIP TC2 Workshop on Partial Evaluation and Mixed Computation'. New York, NY: North-Holland.

Cardelli, L.: 1984, 'A Semantics of Multiple Inheritance'. In: G. Kahn, D. MacQueen, and G. Plotkin (eds.): *Semantics of Data Types: Proc. of the Int. Symp.*, Vol. 173 of *Lecture Notes in Computer Science*. New York, NY, pp. 51–67, Springer-Verlag.

Cocke, J. and J. Schwartz: 1970, *Programming Languages and Their Compilers: Preliminary Notes, 2nd Rev. Version*. New York, NY: Courant Inst. of Math. Sci., New York Univ.

Conte, S. and C. de Boor: 1972, *Elementary Numerical Analysis: An Algorithmic Approach, 2nd. Ed.* New York, NY: McGraw-Hill.

de Boor, C.: 1978, *A Practical Guide to Splines*, Vol. 27 of *Appl. Math. Sciences*. New York, NY: Springer-Verlag.

Earley, J.: 1974, 'High-Level Operations in Automatic Programming'. In: *Proc. of the ACM SIGPLAN Symp. on Very High Level Languages.* New York, NY, ACM Press.

Earley, J.: 1976, 'High-Level Iterators and a Method for Automatically Designing Data Structure Representation'. *J. Comp. Lang.* **1**(4), 321–342.

Fischer, H.: 1992, 'Special Problems in Automatic Differentiation'. In (Griewank and Corliss, 1992), pp. 43–50.

Fong, A.: 1977, 'Elimination of Common Subexpressions in Very High Level Languages'. In: *ACM Symposium on Principles of Programming Languages.* New York, NY, pp. 48–57, ACM Press.

Fong, A.: 1979, 'Inductively Computable Constructs in Very High Level Languages'. In: *ACM Symposium on Principles of Programming Languages.* New York, NY, pp. 21–28, ACM Press.

Fong, A. and J. Ullman: 1976, 'Induction Variables in Very High Level Languages'. In: *ACM Symposium on Principles of Programming Languages*. New York, NY, pp. 104–112, ACM Press.

Futamura, Y.: 1971, 'Partial Evaluation of Computation Process – An Approach to a Compiler-Compiler'. *Systems · Computers · Controls* **2**(5).

Futamura, Y.: 1999, 'Partial Evaluation of Computation Process – An Approach to a Compiler-Compiler'. *Higher-Order and Symbolic Computation* **12**(4). Reprinted from Systems · Computers · Controls 2(5), 1971.

Gallagher, K. and J. Lyle: 1991, 'Using Program Slicing in Software Maintenance'. *IEEE Transactions on Software Engineering* **SE-17**(8), 751–761.

Goldstine, H.: 1977, *A History of Numerical Analysis*. New York, NY: Springer-Verlag.

Griewank, A.: 1989, 'On Automatic Differentiation'. In: M. Iri and K. Tanabe (eds.): *Mathematical Programming: Recent Developments and Applications*. Boston, MA: Kluwer Academic Publishers, pp. 83–108.

Griewank, A.: 1991, 'The Chain Rule Revisited in Scientific Computing'. *SIAM News* **24**.

Griewank, A.: 2000, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, Vol. 19 of *Frontiers in Applied Mathematics*. Philadelphia, PA: Soc. for Indust. and Appl. Math.

Griewank, A. and G. Corliss (eds.): 1992, *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*. Philadelphia, PA: Soc. for Indust. and Appl. Math.

Hammer, R., M. Hocks, U. Kulisch, and D. Ratz: 1993, *Numerical Toolbox for Verified Computing I: Basic Numerical Problems*, Vol. 21 of *Springer Ser. in Comp. Math.* New York, NY: Springer-Verlag.

Hammer, R., M. Hocks, U. Kulisch, and D. Ratz: 1995, *C++ Toolbox for Verified Computing I: Basic Numerical Problems*. New York, NY: Springer-Verlag.

Horwitz, S.: 1990, 'Identifying the Semantic and Textual Differences Between Two Versions of a Program'. In: *SIGPLAN Conf. on Prog. Lang. Design and Impl.* pp. 234–245.

Horwitz, S., J. Prins, and T. Reps: 1989, 'Integrating Non-Interfering Versions of Programs'. *Trans. on Prog. Lang. and Syst.* **11**(3), 345–387.

Horwitz, S., T. Reps, and D. Binkley: 1990, 'Interprocedural Slicing Using Dependence Graphs'. *Trans. on Prog. Lang. and Syst.* **12**(1), 26–60.

Iri, M.: 1984, 'Simultaneous Computation of Functions, Partial Derivatives and Estimates of Rounding Errors: Complexity and Practicality'. *Japan J. Appl. Math.* **1**(2), 223–252.

Jones, N., C. Gomard, and P. Sestoft: 1993, *Partial Evaluation and Automatic Program Generation*. Prentice-Hall International.

Kahan, W.: 2000, 'Personal Communication to Thomas Reps and Louis Rall'.

Kahan, W. and R. Fateman: 1985, 'Symbolic Computation of Divided Differences'. Technical report, Dept. of Elec. Eng. and Comp. Sci., Univ. of Calif.–Berkeley, Berkeley, CA. Unpublished report; available at http://www.cs.berkeley.edu/ fateman/papers/divdiff.pdf.

Karczmarczuk, J.: 1999, 'Functional Differentiation of Computer Programs'. In: *Proc of the ACM SIGPLAN Int. Conf. on Func. Prog. (ICFP '98)*. New York, NY, pp. 195–203, ACM Press.

Kearfott, R.: 1996, 'Automatic Differentiation of Conditional Branches in an Operator Overloading Context'. In (Berz et al., 1996), pp. 75–81.

Krawczyk, R. and A. Neumaier: 1985, 'Interval Slopes for Rational Functions and Associated Centered Forms'. *SIAM J. Numer. Anal.* **22**(5), 604–616.

Linnainmaa, S.: 1976, 'Taylor Expansion of the Accumulated Rounding Error'. *BIT* **16**(1), 146–160.

Liu, Y.: 1995, 'Incremental Computation: A Semantics-Based Systematic Transformation Approach'. Ph.D. diss. and Tech. Rep. TR95-1551, Dept. of Comp. Sci., Cornell Univ., Ithaca, NY.

Liu, Y. and S. Stoller: 1998, 'Loop Optimization for Aggregate Array Computations'. In: *Proc. of the IEEE 1998 Int. Conf. on Comp. Lang.* New York, NY, ACM Press.

Liu, Y. and S. Stoller: 1999, 'Dynamic Programming via Static Incrementalization'. In: S. Swierstra (ed.): *Proc. of ESOP '99: European Symp. on Programming,*

Vol. 1576 of *Lecture Notes in Computer Science*. New York, NY, pp. 288–305, Springer-Verlag.

Liu, Y. and S. Stoller: 2000, 'From Recursion to Iteration: What are the Optimizations?'. In: *Proc. of the 2000 ACM SIGPLAN Workshop on Part. Eval. and Sem.-Based Prog. Manip.* New York, NY, pp. 73–82, ACM Press.

Liu, Y., S. Stoller, and T. Teitelbaum: 1996, 'Discovering Auxiliary Information for Incremental Computation'. In: *ACM Symposium on Principles of Programming Languages.* New York, NY, pp. 157–170, ACM Press.

Liu, Y. and T. Teitelbaum: 1995a, 'Caching Intermediate Results for Program Improvement'. In: *Proc. of the ACM SIGPLAN Symp. on Part. Eval. and Sem.-Based Prog. Manip.* New York, NY, ACM Press.

Liu, Y. and T. Teitelbaum: 1995b, 'Systematic Derivation of Incremental Programs'. *Science of Computer Programming* **24**, 1–39.

Lyle, J. and M. Weiser: 1986, 'Experiments on Slicing-Based Debugging Tools'. In: *Proceedings of the First Conference on Empirical Studies of Programming.*

McCurdy, A.: 1980, 'Accurate Computation of Divided Differences'. Ph.D. diss. and Tech. Rep. UCB/ERL M80/28, Univ. of Calif.–Berkeley, Berkeley, CA.

Milner, R.: 1978, 'A Theory of Type Polymorphism in Programming'. *J. Comp. Syst. Sci.* **17**, 348–375.

Mogensen, T.: 1986, 'The Application of Partial Evaluation to Ray-Tracing'. Master's thesis, Datalogisk Institut, Univ. of Copenhagen, Copenhagen, Denmark.

Ning, J., A. Engberts, and W. Kozaczynski: 1994, 'Automated Support for Legacy Code Understanding'. *Commun. ACM* **37**(5), 50–57.

Opitz, G.: 1964, 'Steigungsmatrizen'. *Zeitschrift für Angewandte Mathematik und Mechanik* **44**, T52–T54. In German. English translation available at http://www.cs.wisc.edu/wpis/papers/opitz.zamm64.ps.

Ottenstein, K. and L. Ottenstein: 1984, 'The Program Dependence Graph in a Software Development Environment'. In: *Proceedings of the ACM SIGSOFT/SIGPLAN Software Engineering Symposium on Practical Software Development Environments.* New York, NY, pp. 177–184, ACM Press.

Paige, R.: 1981, *Formal Differentiation – A Program Synthesis Technique.* Ann Arbor, MI: UMI Research Press.

Paige, R.: 1983, 'Transformational Programming – Applications to Algorithms and Systems'. In: *ACM Symposium on Principles of Programming Languages.* New York, NY, pp. 73–87, ACM Press.

Paige, R. and S. Koenig: 1982, 'Finite Differencing of Computable Expressions'. *ACM Transactions on Programming Languages and Systems* **4**(3), 402–454.

Rall, L.: 1981, *Automatic Differentiation: Techniques and Applications, Lecture Notes in Computer Science, Vol. 120.* Springer-Verlag.

Rall, L.: 1983, 'Differentiation and Generation of Taylor Coefficients in Pascal-SC'. In: U. Kulisch and W. Miranker (eds.): *A New Approach to Scientific Computation.* New York, NY: Academic Press, pp. 291–309.

Rall, L.: 1984, 'Differentiation in Pascal-SC: Type GRADIENT'. *ACM Trans. Math. Softw.* **10**, 161–184.

Rall, L.: 1986, 'The Arithmetic of Differentiation'. *Mathematics Magazine* **59**, 275–282.

Rall, L.: 1990, 'Differentiation Arithmetics'. In: C. Ullrich (ed.): *Computer Arithmetic and Self-Validating Numerical Methods.* New York, NY: Academic Press, pp. 73–90.

Rall, L.: 1992, 'Point and Interval Differentiation Arithmetics'. In (Griewank and Corliss, 1992), pp. 17–24.

Rall, L. and T. Reps: 2000, 'Algorithmic Differencing'. In: *SCAN 2000: 9th GAMM-IMACS Int. Symp. on Sci. Comput., Comp. Arith., and Validated Numerics.*

Ratz, D.: 1996, 'An Optimized Interval Slope Arithmetic and its Application'. Bericht 4/1996, Institut für Angewandte Mathematik, Universität Karlsruhe, Karlsruhe, Germany.

Reps, T. and T. Turnidge: 1996, 'Program Specialization Via Program Slicing'. In: O. Danvy, R. Glück, and P. Thiemann (eds.): *Proc. of the Dagstuhl Seminar on Partial Evaluation,* Vol. 1110 of *Lec. Notes in Comp. Sci.* Schloss Dagstuhl, Wadern, Germany, pp. 409–429, Springer-Verlag.

Shamseddine, K. and M. Berz: 1996, 'Exception Handling in Derivative Computation with Nonarchimedean Calculus'. In (Berz et al., 1996), pp. 37–51.

Speelpenning, B.: 1980, 'Compiling Fast Partial Derivatives of Functions Given by Algorithms'.

Strachey, C.: 2000, 'Fundamental Concepts in Programming Languages'. *Higher-Order and Symbolic Computation* **13**(1/2), 1–49.

Veldhuizen, T.: 1995, 'Expression Templates'. *C++ Report* **7**(5), 26–31.

Weiser, M.: 1984, 'Program Slicing'. *IEEE Transactions on Software Engineering* **SE-10**(4), 352–357.

Wengert, R.: 1964, 'A Simple Automatic Derivative Evaluation Program'. *Commun. ACM* **7**(8), 463–464.

Zippel, R.: 1996, 'Personal Communication to Thomas Reps'.

Zuhe, S. and M. Wolfe: 1990, 'On Interval Enclosures Using Slope Arithmetic'. *Applied Mathematics and Computation* **39**(1), 89–105.