

MODULAR PROGRAMMING OF RECURRENCES: A COMPARISON OF TWO APPROACHES*

Vytautas Čyras

Department of Informatics, Vilnius University
2600 Vilnius, Naugarduko St. 24, Lithuania
and
Institute of Mathematics and Informatics
2600 Vilnius, Akademijos St. 4, Lithuania

Magne HAVERAAEN

Department of Informatics, University of Bergen
Høyteknologisenteret, N-5020 Bergen, Norway

Abstract. We present two methods for expressing computations based on recurrence relations and discuss their relative merits. One method, the *structural blanks* approach, is built on top of traditional programming languages like Fortran or Pascal. It aims at program reuse and bases a certain architecture of software packages. The other method, the *constructive recursive* approach, is based on recursive relations over graphs.

Key words: recurrence relation, loop program synthesis, data dependency graph, architecture of software packages, program module specification.

1. Introduction. Recurrence relations play an important role in the formulation of many problems, such as partial differential equations in mathematical physics or dynamic programming in operations research. A recurrence may be viewed as composed from three parts: a structural part, an expression part, and an initialization part. In this paper we discuss two programming methods for solving generalized recurrences: the *structural blanks* approach and the *con-*

* This paper extends the comparison of two approaches first contributed to the 6th NWPT in Aarhus, Denmark (Čyras and Haveraaen, 1994). The research was supported in part by the Research Council of Norway under the Nordic–Baltic scholarship programme, and in part by the University of Bergen.

structive recursive approach. Both approaches rely on the separation between the structural part describing the *data dependency pattern* of the recurrence, the definition of the expressions to be computed as the *computational aspect*, and the initialization and definition of outputs.

The structural blanks approach was first presented by Čyras (1983), then by Greshnev, Lyubimskii and Čyras (1985), and later by Čyras (1986, 1988). The approach was inspired by the computation of finite difference solutions of partial differential equations (PDE), where driver routines for sets of mutually dependent recurrences were needed. One of the aims was to develop a framework where the correctness of the driver routine need only be proved once, while the scheduling it defines may be reused for different problems with the same basic dependency structures. The solution to this was to define driver routines (S-modules) based on the structure of the recurrence, and requiring that the routines (F-modules) for solving each recurrence included a declaration of its dependency structure. The driver routine could then be applied to all recurrences with a compatible structure. Compatibility was shown by exhibiting an injective function from the S-module to the global arrays underlying the F-modules.

The constructive recursive approach was developed by Haverlaen (1990, 1993) from a programming perspective. Traditionally a recursive program implicitly defines an exponentially growing tree-structured graph. In the constructive recursive case the graph is explicitly defined by the user, allowing linear solution time for recurrences, even for higher order recurrences. The information in the graph also allows the translation of the recursive program to recursion free loop programs.

This paper is structured as follows. First we discuss some basic properties of recurrences. Then we present the structural blanks approach followed by a presentation of the constructive recursion case in Section 4. In Section 5 we show similarities and differences of the two approaches by comparing them on some examples. Finally we discuss the relative merits of the approaches.

2. Generalized recurrences. An *order k* linearly dependent recurrence r with the natural numbers as *index domain* is a relation defined by a set of equations

$$r_n = \phi(r_{n-1}, r_{n-2}, \dots, r_{n-k}), \quad (1)$$

$$r_{k-1} = \varepsilon_{k-1},$$

$$\dots$$

$$r_0 = \varepsilon_0,$$

where the indices are natural numbers, ϕ is a k -ary expression not referring to r , and the ε_i , representing initial values, are expressions not referring to r . The choice of r_0, \dots, r_{k-1} as initial elements is arbitrary, and one may even envision cases where the recurrence is infinite or where different recurrences only differ in the choice of initial elements. The archetypical second order recurrence relation is the Fibonacci function

$$F_n = F_{n-1} + F_{n-2},$$

$$F_1 = 1,$$

$$F_0 = 0,$$
(2)

defining the sequence 0, 1, 1, 2, 3, 5, 8, 13, 21, ... The dependency pattern of this function is illustrated in Fig. 1.




Fig. 1. Data dependency graph of a *second order one-dimensional* recurrence, such as the Fibonacci function. The numbers in circles label the two arcs from a node. The nodes are enumerated by the plain numbers underneath them.

A straight forward method to compute the n 'th value of the recurrence (1) is to start with an array $R[0:k-1]$ with the k initial values, *i.e.* $R[j] = \varepsilon_j, j = 0, 1, \dots, k - 1$. Then for each $j = k, k + 1, \dots, n$, compute

$$R[j \bmod k] := \phi(R[(j-1) \bmod k], R[(j-2) \bmod k], \dots, R[(j-k) \bmod k]),$$
(3)

where the value $R[n \bmod k]$ represents r_n . If all values r_0, r_1, \dots, r_n are needed, the array should be declared $R[0:n]$, and the computations be

$$R[j] := \phi(R[j-1], R[j-2], \dots, R[j-k]),$$
(4)

where $R[j]$ will then contain r_j for $0 \leq j \leq n$. Other result sets may also be defined, and have to be mirrored in the declaration and use of the array R . More efficient computation techniques exist for special cases of recurrences, e.g., if the expression $\phi(\cdot_1, \dots, \cdot_k)$ is linear, the recurrence may be reformulated as a matrix exponentiation problem. Such techniques will not be discussed further in this paper.

Recurrences may be generalized to arbitrary index domains. Given a sufficient set of initial values $\varepsilon_{i_1, \dots, i_m}$, the m -dimensional order k general recurrence has the form

$$r_{n_1, \dots, n_m} = \phi(r_{\delta_1(n_1, \dots, n_m)}, \dots, r_{\delta_k(n_1, \dots, n_m)}), \quad (5)$$

where the m -ary functions δ_i , each returning an m -tuple of indices, have to be well founded with respect to the set of initial values. This scheme is more powerful than that of conventional recurrences, and algorithms such as the Fast Fourier Transform (Cooley and Tukey, 1965) belong to the class of general recurrences. Since the δ_i have a more complex relationship than the linear dependency in (1), it is impossible to give a general algorithm for computing r_{n_1, \dots, n_m} . Moreover, finding such an algorithm for a given set of δ_i , even if they are affine, may be difficult. But the structure of the algorithm to compute the recurrence is dependent only on the δ_i , the *data dependency pattern* of the recurrence, and is independent of the actual ϕ , known as the *computational aspect* of the recurrence.

Sometimes we will be working with a set of recurrences, all mutually dependent on each other. A *set of mutually dependent recurrences* is a set of ℓ recurrences r^1, \dots, r^ℓ , the recurrence r^i being of dimensionality m_i and order k_i , of the form

$$\begin{aligned} r_{n_1, \dots, n_{m_1}}^1 &= \phi_1(r_{\delta_{1,1}^{i_1,1}(n_1, \dots, n_{m_1})}, \dots, r_{\delta_{1,k_1}^{i_1,k_1}(n_1, \dots, n_{m_1})}), \\ r_{n_1, \dots, n_{m_2}}^2 &= \phi_2(r_{\delta_{2,1}^{i_2,1}(n_1, \dots, n_{m_2})}, \dots, r_{\delta_{2,k_2}^{i_2,k_2}(n_1, \dots, n_{m_2})}), \\ &\vdots \\ r_{n_1, \dots, n_{m_\ell}}^\ell &= \phi_\ell(r_{\delta_{\ell,1}^{i_\ell,1}(n_1, \dots, n_{m_\ell})}, \dots, r_{\delta_{\ell,k_\ell}^{i_\ell,k_\ell}(n_1, \dots, n_{m_\ell})}) \end{aligned} \quad (6)$$

together with a suitable set of initial values. Here $i_{j,q} \in \{1, \dots, \ell\}$, and $\delta_{j,q}$ is an m_j -ary function returning an $m_{i_{j,q}}$ -tuple of indices. Without loss of

generality we can assume that all the dimensionalities are equal: $m_1 = \dots m_\ell = m$. The mutually dependent recurrences correspond to *course of value recursion* in the terminology of Tucker and Zucker (1988).

3. Structural blanks. The *structural blanks* (SB) approach (Čyras, 1983; Greshnev *et al.*, 1985; Čyras, 1986; Čyras, 1988) was developed to express solutions to mutually dependent recurrences in the form of reusable program components defining loops over arrays. The problem of synthesizing a right sequence of array element updates in order to compute a set of mutually dependent recurrences was formulated by Lyubimskii as early as in 1958 (published in (Lyubimskii, 1960)), and later on investigated by Zadykhailo (1963). The organization of computations for linear recurrences over multidimensional arrays was studied by Karp, Miller and Winograd (1967) independently of the earlier research. The presentation here represents a further development of the SB approach, so the notation and definitions differ from the older papers. We will use a mixed Fortran/Pascal notation in the examples.

The SB approach distinguishes between *structural components* (*S-modules*) and *functional components* (*F-modules*). It is well suited to define mutually dependent recurrences (6), and F- and S-modules derived directly from such recurrences are called *elementary* F- and S-modules. Each module contains a data dependency part and a procedure part. The S-module describes the data dependencies, the set of initial elements and the set of output elements, and in the *S-procedure* it defines a driver algorithm for recurrences with this dependency structure.

An *F-procedure* defines the algorithm to compute one step of one recurrence expression r^j of (6), and the containing F-module describes the data dependencies of this step. An S-module is *applied* to a collection of F-modules by matching the dependencies of the F-modules with those of the S-module as defined by a substitution Ξ on the S-module. The application produces a new F-module containing an algorithm to compute the full recurrence. Normally this F-module will not be elementary.

In the case of an order k linear recurrence (1) an elementary structural module would capture the computational idea of (4) by

```
S-module LDEP (Fmod  $\Phi$ (integer); k, N: integer) ==
formal x: array [*]
```

```

internal-template
  (var q: integer;  $\Phi(q) == (x[t], t=q-k..q-1) \longrightarrow x[q]$ )
external-template
  (x[t], t=0..k-1)  $\longrightarrow$  (x[t], t=k..N) (7)
procedure
  var q: integer;
  for q := k to N do
    call  $\Phi(q)$ 
end

```

This is to be interpreted as: given a one-dimensional (one argument in the declaration of formal F-module Φ) order k recurrence over the array x (as declared in the internal template), the S-module defines a procedure that will invoke Φ to compute all elements $x[k], \dots, x[N]$ given that $x[0], \dots, x[k-1]$ are defined (external template). The set of array elements to the left of the “ \longrightarrow ” (gives) in the external template is the *set of initial elements*, and the set to the right is the *set of output elements*. The parameters to the formal F-module Φ range over the index domain of the recurrence. The formal array x will be part of the environment for the argument F-module “ Φ ”. The S-module only needs size information for the formal array x since it is only used in the templates to declare the dependencies. The parameters – formal arrays – of the S-module are not parameters in the traditional sense, but they will be matched by the substitution rules. The data dependency graph of the computation organized by the S-module LDEP when $k = 2$ is shown in Fig. 1, where square nodes mean that the nodes here have initial values, while the disc nodes represent nodes that will be computed.

The elementary functional module giving the computational aspect of each step of the Fibonacci function is

```

F-module FIBSTEP (q: integer) ==
  global X: array[*] of integer
  template X[q-1], X[q-2]  $\longrightarrow$  X[q] (8)
  procedure X[q] := X[q-1] + X[q-2]
end

```

This is to be interpreted as: FIBSTEP contains a one-dimensional (index domain parameter q) second order recurrence expression over the array X (as can be

seen from the template). The size of the array X is not declared in the F-module, but it will be declared in the program unit that uses the modules. The base type of X is declared since the operations on the elements require this knowledge. We view X as being declared in a global external environment with respect to FIBSTEP (8). This environment can be, for example, COMMON area in Fortran, STATIC in PL/1, etc.

To be able to use FIBSTEP to compute the Fibonacci function, we need a driver procedure that will schedule the computations of its F-procedure. Driver procedures are part of the S-modules, and are applicable if the internal template of the S-module matches the template of the F-module. This occurs when the dependency pattern $\mathcal{I}_{in} \longrightarrow \mathcal{I}_{out}$ of the S-module's corresponding internal template is equal to the pattern $\mathcal{F}_{in} \longrightarrow \mathcal{F}_{out}$ of the F-module's template. In our example we obtain an equality by substituting

$$\mathbf{k} \mapsto 2; \quad \mathbf{x}[\cdot]; \mapsto X[\cdot]; \quad \Phi(\cdot) \mapsto \text{FIBSTEP}(\cdot). \quad (9)$$

This shows the three kinds substitution rules:

- The *binding substitution* β that replaces an argument of the S-module by a constant, i.e., it removes \mathbf{k} from the parameter list, and replaces all occurrences of \mathbf{k} in the body of the S-module by 2.
- The *array domain substitution* $\xi(\cdot)$ that embeds the formal array \mathbf{x} and its index domain into the global array X and its index domain.
- The *formal F-module domain substitution* $\tau(\cdot)$ that changes the formal F-module Φ and its parameters in all calls in the S-module, allowing the embedding of the formal index domain in a higher dimensional domain as well as other manipulations. In this case changing the name to FIBSTEP.

Calling the substitution (9) for Ξ , we denote the application

$$\text{FIB} = \text{LDEP} \Big|_{\Xi} (\text{FIBSTEP}).$$

The actual parameter FIBSTEP indicates that the internal template's pattern Φ in LDEP should match that of FIBSTEP. The actual application is defined by the use of the substitution Ξ , and this substitution must be compatible with the argument of the application.

Unfolding the application above we get a new F-module

```

F-module FIB (N: integer) ==
  global X: array[*] of integer
  template X[0], X[1] → X[2..N]
  procedure
    var q:integer;
    for q := 2 to N do
      X[q] := X[q-1] + X[q-2]
  end

```

The resulting F-module FIB is not an elementary one. The template of FIB specifies that X contains Fibonacci numbers numbered from 0 to N, where X[2..N] are regarded as output, based on the initial values of X[0] and X[1]. We may now extract a normal procedure that does the computation by extracting the F-procedure from the F-module FIB and adding the parameters. This yields

```

procedure FIB (N: integer);
  global X: array[*] of integer
  var q:integer;
  for q := 2 to N do
    X[q] := X[q-1] + X[q-2]
  end

```

3.1. Development methodology. The development methodology of the structural blanks approach can be formulated as three steps. In the **first step** a domain expert, e.g., a physicist, formulates the problem as a set of mutually dependent recurrence equations, which is encoded as a collection of F-modules and global array declarations, comprising the *computational model* for the problem. The sizes of the arrays may be dependent on the size of input data, the number of time-steps to be used, or may be forced by numerical properties of the discretization technique involved in formulating the recurrence equations.

As an example take the problem that can be formulated as the real valued general recurrence equation on the exponential scale

$$\begin{aligned}
 g(2^{i+2}) &= \gamma (g(2^{i+2}/2), g(2^{i+2}/4)), \\
 g(2^1) &= \varepsilon_1, \\
 g(2^0) &= \varepsilon_0,
 \end{aligned}
 \tag{10}$$

where we want to find $g(2^i)$ for $i = 0, 1, 2, \dots, N$. This may be formulated as the declaration of "Y: array[1..2**N] of real" together with the F-module

```

F-module GSTEP (i: integer) ==
    global Y : array[*] of real
    template Y[2**i], Y[2**(i+1)]  $\rightarrow$  Y[2**(i+2)]      (11)
    procedure Y[2**(i+2)] :=  $\gamma$ (Y[2**(i+1)], Y[2**i])
    end
    
```

The data dependency graph of this recurrence is shown in Fig. 2.




Fig. 2. Data dependency graph of the recurrence g defined in (10).

The second step is to devise a driver routine for the computational model, i.e., to find an appropriate S-module. For this purpose there may be a library of S-modules, and one of them may be adapted to the problem at hand by using a substitution.

In the case of the recurrence (10) we may reuse the S-module LDEP with the substitution

$$\Xi = [k \mapsto 2; \quad x[\cdot] \mapsto Y[2**\cdot]; \quad \Phi(\cdot) \mapsto \text{GSTEP}(\cdot - 2)], \quad (12)$$

involving all three substitution rules. Here the array domain substitution does the exponential expansion, while the formal F-module domain substitution, shifts the formal F-module parameters two positions in order to adjust the starting point of the loop in the S-procedure to the indices used by the F-module. This yields the application $G = \text{LDEP}|_{\Xi}(\text{GSTEP})$:

```

F-module G (N: integer) ==
  global Y : array[*] of real
  template Y[1], Y[2]]  $\longrightarrow$  (Y[2**t], t=2..N)
  procedure (13)
    var q:integer;
    for q := 2 to N do
      call GSTEP(q-2)
    end

```

The **third step** is to show that an application is correct by proving that the substitutions are safe. In this case it is obvious since the function on the array index domain, $j \mapsto 2^j$ as embodied in “ $x[\cdot] \mapsto Y[2^{**}\cdot]$ ”, is injective.

Note that only N elements of the array Y are involved in the computation. The array Y is treated as a part of the environment and has to have at least 2^N elements.

Suppose that we shift the index of the initial values in (10) to that of $g(2^{N_0}) = \varepsilon'_0$, $g(2^{N_0+1}) = \varepsilon'_1$, for some N_0 . In this case, the same F-module GSTEP (11) can be used to compute $g(2^{N_0+t})$ for $t = 2, 3, \dots, N$. The S-module LDEP (7) can be reused, but with the substitution Ξ'

$$\Xi' = [k \mapsto 2; \quad x[\cdot] \mapsto Y[2^{**}(N_0+\cdot)]; \quad \Phi(\cdot) \mapsto \text{GSTEP}(N_0 + \cdot - 2)].$$

This yields the application $\text{GNO} = \text{LDEP}|_{\Xi'}(\text{GSTEP})$:

```

F-module GNO (N0, N: integer) ==
  global Y: array[*] of real
  template Y[2**N0], Y[2**(N0+1)]  $\longrightarrow$  (Y[2**(N0+t)], t=2..N)
  procedure
    var q:integer;
    for q := 2 to N do
      call GSTEP(N0+q-2)
    end

```

Both the input/output notation and the notion of computational model are influenced by Tyugu, his method of *structural synthesis of programs*, and the NUT system (Tyugu, 1987). NUT supports the automatic synthesis of “**for**

$i := m$ to n ” loops (Harf, 1994). The structural blanks approach, when compared with the structural synthesis of programs, aims at (i) the reuse of complicated loop combinations, and (ii) the transformation of the loop parameter i allowing embeddings in larger dimensions.

3.2. Dependency patterns. Dependency patterns are defined in the template of an F-module and the internal templates and the external template of an S-module. They describe the data dependencies being assumed by the modules. The dependency pattern describes a pair of non-intersecting sets (of array elements) called *input* and *output*, and is of the form

$$\mathcal{P}_{in} \longrightarrow \mathcal{P}_{out}, \quad \text{where} \quad \mathcal{P}_{in} \cap \mathcal{P}_{out} = \emptyset.$$

The interpretation at the dependency pattern is that the array elements identified on the left hand side, \mathcal{P}_{in} , are needed in the computation of the array elements identified on the right hand side, \mathcal{P}_{out} . The elements on the right hand side will be defined (assigned to) by some expression of the array elements on the left hand side in the body of the procedure.

For the case of the equations on form (6) the dependencies would have the form

$$\begin{aligned} & \mathbf{X}^{i_j,1}[\delta_{j,1}(n_1, \dots, n_m)], \dots, \mathbf{X}^{i_j,k_j}[\delta_{j,k_j}(n_1, \dots, n_m)] \\ & \longrightarrow \mathbf{X}^j[n_1, \dots, n_m]. \end{aligned} \quad (14)$$

This is interpreted as: if the array elements on the left hand side of the “ \longrightarrow ” arrow,

$$\mathbf{X}^{i_j,1}[\delta_{j,1}(n_1, \dots, n_m)], \dots, \mathbf{X}^{i_j,k_j}[\delta_{j,k_j}(n_1, \dots, n_m)],$$

are appropriately initialized, then the F-procedure F_j will compute the array element $\mathbf{X}^j[n_1, \dots, n_m]$. The dependency of the form (14), where exactly one array element is in \mathcal{P}_{out} , is called an *elementary dependency*. In the case of elementary module, the number of elements k on the left hand side of the arrow represent the *order* of the dependency.

The general form of dependency pattern that is used in the F- and S-modules allow more than one element to be computed: both sides of the arrow “ \longrightarrow ” contain a list of one or more array elements. These may be listed explicitly, or a group of array elements may be enclosed by an implicit DO-loop of the

form used in Fortran. The implicit DO will allow easy identification of array segments or more scattered array element patterns. Another possible notation for dependency pattern is the shape declaration of Fortran-90, which cuts out a segment of a (multidimensional) array.

3.3. The F-module. The elementary F-module defines the dependency pattern and the computational aspect of a step of the recurrence equation. When programming recurrences using the structural blanks approach, the set of mutually dependent recurrence relations (6) is taken as starting point. Global arrays $X^1, \dots, X^{\ell \times}$ with the corresponding dimensions are declared for each of the recurrences r^1, \dots, r^{ℓ} , and an F-module F_j has to be declared for each of the recurrence equations ϕ_j of the set.

The parameters of the F-module belong to one of two different groups.

- Parameters declared in the parameter list of the F-module. This group is divided into two subgroups. The first subgroup, the index domain parameters: reflect the number of dimensions m of the index domain of the recurrence equation. These parameters are free in the index expressions of the array elements at the template declaration. The second subgroup comprises (i) constants to the procedure representing the computational aspect, and (ii) loop boundaries (in the case of non-elementary F-modules). In general case an array domain parameter can play the role of a loop boundary.
- Global array names corresponding to each recurrence of the set. The bounds of these arrays are in the declaration of the *computational model*.

The computational model is treated as an environment. It is viewed as the set of all the global arrays a collection of F-modules operates with. The computational model can be viewed as a graph: each node corresponds to one element of a certain global array.

The basic form of the elementary F-module referring to one recurrence in the set of mutually dependent recurrences (6) is

```

F-module FNAME ( $n_1, n_2, \dots, n_m$ : integer, <other parameters>) ==
    global  $X^1$ : array[*,...,*] of <type1>;
            $X^2$ : array[*,...,*] of <type2>;
           ⋮

```

(15)

```

XℓX : array[*, ..., *] of <typeℓX>;
template
   $\mathcal{F}_{in} \longrightarrow \mathcal{F}_{out}$ 
procedure
  <statements>
end

```

where n_1, \dots, n_m are index domain parameters, X^1, \dots, X^{ℓ_X} are global array names, and the number of stars of an array X^i corresponds to its number of dimensions. Normally the $\langle type_i \rangle$ will all be the same, namely the type of the recurrence (typically real or complex numbers). In the case of an elementary F-module FNAME, the $\langle statements \rangle$ are the program statements defining the actual expression

$$\phi_j \left(r_{\delta_1^j(n_1, \dots, n_m)}^{i,j,1}, \dots, r_{\delta_k^j(n_1, \dots, n_m)}^{i,j,k} \right),$$

with the appropriate array elements replacing the r^i expressions, and assigning this value to the array element corresponding to r_{n_1, \dots, n_m}^j . The statements may be in any suitable programming language, e.g., Fortran or Pascal. We have chosen a Pascal-like language with some Fortran-90 extensions and notation for the examples given here. We will be following the Pascal conventions of interpreting a multidimensional array declaration

$$X : \mathbf{array}[l_1 : u_1, l_2 : u_2, \dots, l_m : u_m],$$

as equal to the declaration

$$X : \mathbf{array}[l_1 : u_1] \mathbf{of} \mathbf{array}[l_2 : u_2] \mathbf{of} \mathbf{array}[\dots] \mathbf{of} \mathbf{array}[l_m : u_m].$$

This also applies to indexing operations, where a multidimensional index is considered equivalent to a sequence of indices:

$$X[i_1, i_2, \dots, i_m] \Leftrightarrow X[i_1][i_2][\dots][i_m].$$

This convention gives a greater flexibility when combining S-modules and F-modules.

To illustrate the notation, we will develop the F-module HEATSTEP. It defines the recurrence of the classic explicit finite difference approximation to solve a

parabolic partial difference equation (see, e.g., Lapidus and Pinder, 1982). The F-module HEATSTEP corresponds to the two-dimensional, order 3 recurrence equation given in Fig. 3.

$$\begin{aligned} x_{q_1, q_2} &= \varphi(x_{q_1-1, q_2-1}, x_{q_1, q_2-1}, x_{q_1+1, q_2-1}), & (16) \\ x_{t, 0} &= \varepsilon_t, \quad t = 1, 2, \dots, N_1, \\ x_{0, t} &= \varepsilon'_t, \quad t = 0, 1, \dots, N_2 - 1, \\ x_{N_1+1, t} &= \varepsilon''_t, \quad t = 0, 1, \dots, N_2 - 1. \end{aligned}$$




Fig. 3 The classic explicit finite difference approximation of heat flow in the dimension q_1 , where the q_2 axis represents time. The points 0 and $N_1 + 1$ on the q_1 axis represent the end points. On the top the actual two-dimensional, order 3 recurrence relation with initial values and values at the end points is shown. The dependency pattern of the equation is shown to the left. On the right a global picture of how the initial values (squares) relate to the interior domain (discs).

We translate the recurrence x into the two-dimensional array X , and will also need two arguments for the F-procedure. The dependency pattern of the equation (16) tells us that $X[q_1-1, q_2-1]$, $X[q_1, q_2-1]$, and $X[q_1+1, q_2-1]$ are all needed to compute $X[q_1, q_2]$ and we state this

```
F-module HEATSTEP (q1, q2: integer) ==
  global X: array[*,*] of real
  template X[q1-1, q2-1], X[q1, q2-1], X[q1+1, q2-1] → X[q1, q2]
  procedure X[q1, q2] := φ(X[q1-1, q2-1], X[q1, q2-1], X[q1+1, q2-1])
end
```

(17)

The F-procedure body contains the assignment to $X[q1, q2]$ based on computing the expression φ with the appropriate array element arguments.

The F-procedure may be extracted from the F-module as a normal procedure in the chosen programming language. The F-procedure extracted from the general form (15) of the F-module, has the following form

```

procedure FNAME ( $n_1, n_2, \dots, n_m$ : integer, <other parameters>);
  global  $X^1$ : array[*, ..., *] of <type1>;
            $X^2$ : array[*, ..., *] of <type2>;
           ⋮
            $X^{\ell_X}$ : array[*, ..., *] of <type $\ell_X$ >;
  <statements>
end

```

(18)

Note that the statements and the declarations are the same as those in the F-module. Different programming constructs can be used in different programming languages to implement global arrays.

3.4. The S-module. The purpose of the *elementary* S-module is to organize the computations needed to solve a recurrence equation. The S-module declares a set of arrays x^1, \dots, x^{ℓ_S} , but for the S-module, however, the types of the array elements are immaterial, while the number of dimensions still is important. Thus the S-module array declarations need only emphasize this. The internal templates of the S-module serve the same purpose as the template of the F-module: to identify the data dependencies of the computation steps. The external template of the S-module states which elements of the arrays must be initialized in order to compute the recurrences for a specific set of index domain points. It is defined using a dependency pattern $\mathcal{E}_{in} \longrightarrow \mathcal{E}_{out}$, where \mathcal{E}_{in} to the left of the arrow “ \longrightarrow ” describes the initial values, while the elements \mathcal{I}_{out} to the right of the arrow identify the values being computed.

The S-module itself does not depend on a specific recurrence (i.e., functions ϕ_j), but rather on the dependency pattern of a recurrence (i.e., functions $\delta_{j,i}$, $i = 1, \dots, k_j$). Thus the specific F-modules Φ_j , $j = 1, \dots, \ell$ associated with each recurrence are parameters to the S-module. The F-module parameters are declared with only the index domain parameters. This convention applies to all uses of the F-modules within the S-module.

The dependency pattern embedded in each F-module parameter is described in the internal template. For every procedure Φ_j the pattern is declared using

$$(\text{var } q_{j,1}, \dots, q_{j,m_j} : \text{integer}; \Phi_j(q_{j,1}, \dots, q_{j,m_j}) == \mathcal{I}_{j,in} \longrightarrow \mathcal{I}_{j,out})$$

where the $q_{j,i}$ denote index domain variables. In this presentation they will be reaching over the full Cartesian product domain integer^{m_j} , but in the general setting they may be constrained to some subdomain. The interpretation of the pattern is similar to the F-module case: the call to $\Phi_j(q_{j,1}, \dots, q_{j,m_j})$ will use the array elements in $\mathcal{I}_{j,in}$ to compute the array elements in $\mathcal{I}_{j,out}$.

The S-procedure is a driver routine that will call the F-procedures in a pre-determined order, so that the computation successively will define new elements of the arrays until the entire output has been computed.

Although the actual parameter declarations and their ordering may vary, the *recurrence form* of the S-module is based on the pattern of (6) and has the form

$$\begin{aligned} \text{S-module SNAME (Fmod } \Phi_1 (\text{integer}, \dots, \text{integer}); \\ \text{Fmod } \Phi_2 (\text{integer}, \dots, \text{integer}); \\ \vdots \\ \text{Fmod } \Phi_\ell (\text{integer}, \dots, \text{integer}); \\ \text{< other parameters >)} == \\ \text{formal } x^1 : \text{array}[* , \dots , *]; \\ \text{ } x^2 : \text{array}[* , \dots , *]; \\ \vdots \\ \text{ } x^{\ell_S} : \text{array}[* , \dots , *]; \end{aligned} \quad (19)$$

internal-template

$$(\text{var } q_{1,1}, \dots, q_{1,m_1} : \text{integer}; \Phi_1(q_{1,1}, \dots, q_{1,m_1}) == \mathcal{I}_{1,in} \longrightarrow \mathcal{I}_{1,out})$$

$$(\text{var } q_{2,1}, \dots, q_{2,m_2} : \text{integer}; \Phi_2(q_{2,1}, \dots, q_{2,m_2}) == \mathcal{I}_{2,in} \longrightarrow \mathcal{I}_{2,out})$$

\vdots

$$(\text{var } q_{\ell,1}, \dots, q_{\ell,m_\ell} : \text{integer}; \Phi_\ell(q_{\ell,1}, \dots, q_{\ell,m_\ell}) == \mathcal{I}_{\ell,in} \longrightarrow \mathcal{I}_{\ell,out})$$

external-template

$$\mathcal{E}_{in} \longrightarrow \mathcal{E}_{out}$$

procedure

< statements >

end

where the *<statements>* are the program statements defining the driver algorithm, and *<other parameters>* are other parameters the S-module may need. In our examples *<other parameters>* are named N_1, \dots, N_{ℓ_N} and play the role of loop boundaries.

In some cases the templates of the S-module may leave unspecified which F-module Φ_j is being used to generate a specific output element value. This can be remedied by splitting the output dependency pattern of the Φ_j templates into different subdomains, e.g., by introducing different formal arrays or by constraining the index domain of the variables. We only explore the former option in this presentation.

To illustrate the notation, we develop the S-module FAN3 corresponding to the two-dimensional recurrence equation given in Fig. 3. We translate the recurrence x into the two-dimensional array x . We also see the F-module argument Φ of the S-module needs two arguments, and the equation (16) has a fan-like pattern so the internal template reflects this. The external template defines the border elements are needed as inputs, and that the interior will be computed. The lower bounds of the border and the array dimensions are fixed to 0, but the upper bounds depend on the constants N_1 and N_2 , which will be declared as parameters to the S-module.

```

S-module FAN3 (Fmod  $\Phi$ (integer, integer); N1, N2: integer) ==
  formal x: array[*,*]
  internal-template
    (var q1, q2: integer;  $\Phi$ (q1,q2) ==
      x[q1-1,q2-1], x[q1,q2-1], x[q1+1,q2-1]  $\longrightarrow$  x[q1,q2])
  external-template
    (x[t1,0], t1=1..N1),
    (x[0,t2], t2=0..N2-1),
    (x[N1+1,t2], t2=0..N2-1)  $\longrightarrow$ 
    (x[t1,t2], t1=1..N1, t2=1..N2)
  procedure
    var q1, q2: integer;
    for q2:=1 to N2 do
      for q1:=1 to N1 do
        call  $\Phi$ (q1,q2)
end

```

(20)

3.5. Substitution rules. The F-module and the S-module capture different aspects of how to compute a recurrence. In order to compute the values of an actual recurrence, the expressions encoded in the F-procedures must be combined with the driver routine of a compatible S-module. An S-module is compatible with a list of F-modules, if the individual internal templates of the S-module match the templates of the corresponding F-modules. The application yields a new F-module.

Since F- and S-modules may be programmed independently of each other, different programmers may choose different names for the same entities, or be working on more or less specific instances of the equations for a problem. In order to combine such modules, they must be made to agree with each other, hence certain substitution rules are needed for the S-modules. In order to avoid unintentional variable capture, none of the free variables must be equal to variables declared in a local context in the S-module.

DEFINITION 3.1. A substitution Ξ is a string of atomic substitutions

$$[\langle \text{atomic substitution} \rangle; \dots; \langle \text{atomic substitution} \rangle],$$

each atomic substitution having the general form

$$\langle \text{pattern} \rangle \mapsto \langle \text{pattern} \rangle .$$

Variables introduced in the pattern to the left of “ \mapsto ” are bound in the substitution, those introduced on the right are free in the substitution.

DEFINITION 3.2. The *binding substitution* β is of the form

$$\mathbb{N} \mapsto e,$$

where \mathbb{N} is a parameter to the S-module, and the e is an expression of the same type. The effect is to replace all occurrences of \mathbb{N} in the body of the S-module with the expression e , and to remove the declaration of \mathbb{N} from the parameter list, and adding declarations for the free variables of e to the parameter list of the S-module.

DEFINITION 3.3. The *array domain substitution* is of the form

$$\mathbf{x}[_1, \dots, _n] \mapsto \mathbf{X}[\xi(\cdot_1, \dots, \cdot_n)],$$

where \mathbf{x} is a formal array of at least n dimensions in the S-module, and \mathbf{X} must be a global array, of at least d dimension, and $\xi = \langle \xi_1, \dots, \xi_d \rangle$ is a d -tuple of n -ary functions such that ξ is injective.

The effect is to take all occurrences of $\mathbf{x}[p_1, \dots, p_n]$ and replace them with $\mathbf{X}[\xi_1(p_1, \dots, p_n), \dots, \xi_d(p_1, \dots, p_n)]$, doing the required manipulations of all index expressions in all occurrences of \mathbf{x} . Finally, the \mathbf{x} is removed from the formal array declarations of the S-module, and declarations for any free variables of $\xi(\cdot_1, \dots, \cdot_n)$ being added to the parameter list of the S-module.

Since this substitution is on the formal arrays, only the array expressions in the internal template of the S-module is affected. This is a very general substitution rule that allows the renaming of arrays, array index domain embeddings, and also array embeddings. The substitution may be applied even if the dimensionality of the index domain is lower than that of the array \mathbf{x} itself, as the rest of the dimensions may be treated as part of the type declaration of the array (see the Pascal array declaration convention earlier). The requirement that the function ξ is injective means that all distinct old elements must be mapped to distinct new elements. This could be relaxed so that several read-only locations of the array \mathbf{x} could be mapped to the same read-only location of the array \mathbf{X} , but the benefits of this are not quite clear, and the constraints to be checked have not been worked out.

DEFINITION 3.4. The formal F-module index domain substitution is of the form

$$\Phi(\cdot_1, \dots, \cdot_m) \mapsto \mathbf{F}(\tau(\cdot_1, \dots, \cdot_m)),$$

where Φ has m arguments and is a formal F-module parameter to the S-module, and \mathbf{F} is an actual F-module. τ must be injective.

The effect is to take all occurrences of $\Phi(\cdot_1, \dots, \cdot_m)$, throughout the templates and S-procedure, and replace them with $\mathbf{F}(\tau(\cdot_1, \dots, \cdot_m))$, doing the required manipulations of all index expressions in all occurrences. Finally, the Φ declaration is removed from the parameter list of the S-module, and declarations for any free variables of $\tau(\cdot_1, \dots, \cdot_m)$ being added to the parameter list of the S-module.

This substitution allows the change of the number of arguments to an F-module parameter, as well as changing the expressions used in calls of the F-module. The purpose of this rule is to allow greater flexibility in the use

of S-modules. With this substitution it is possible to let a two-dimensional S-module drive the computations of a three-dimensional F-module along a hyperplane, or shift the indexing conventions, e.g., rotation, of a formal F-module, as well as add other parameters being used by the actual F-module. The function τ is required to be injective to avoid the danger that a call to the actual F-module will overwrite previous results.

The individual substitution rules have their requirements on the substitution functions, but there are also some global criteria that affect the whole substitution Ξ .

DEFINITION 3.5. A substitution Ξ is *safe* if it does not merge any distinct array elements of the S-module's output set.

FACT 3.6. A substitution Ξ is safe if all array domain substitutions have different global arrays on the right hand side of the \mapsto arrow.

3.6. Application of an S-module to F-modules. Given a declaration of an S-module of the form (19), it may be applied to an argument list of ℓ F-modules F_1, \dots, F_ℓ .

An application is denoted by

$$\tilde{F} = S|_{\Xi}(F_1, \dots, F_\ell), \quad (21)$$

where S is an S-module and F_1, \dots, F_ℓ are F-modules, and it yields a new F-module \tilde{F} .

DEFINITION 3.7. The application (21) is *legal* if

- The substitution Ξ removes all formal arrays from the S-modules;
- All global arrays introduced by Ξ are declared by at least one actual F-module;
- The number of actual F-module arguments in the application are the same as the number of formal F-module parameters of the S-module;
- For all $j = 1, \dots, \ell$, the formal F-module index domain substitutions in Ξ bind the formal F-module Φ_j to the actual F-module F_j given in the argument list of the application;
- For all $j = 1, \dots, \ell$, the declaration of F_j matches its use in S as given after the formal F-module index domain substitutions have been performed;
- For all $j = 1, \dots, \ell$, the templates

$$\mathcal{F}_{j,in}(\vec{i}) \longrightarrow \mathcal{F}_{j,out}(\vec{i})$$

as defined by the call $F_j(\tau_j(q_{j,1}, \dots, q_{j,m_j}))$ matches the internal template of the S-module as defined after the array domain substitutions have been performed. By match we mean that the corresponding sets of array elements, for \bar{q} being the variables declared in the internal template of the S-module, must satisfy

$$\xi(\mathcal{I}_{j,in}(\bar{q})) = \mathcal{F}_{j,in}(\tau(\bar{q})) \quad \text{and} \quad \xi(\mathcal{I}_{j,out}(\bar{q})) = \mathcal{F}_{j,out}(\tau(\bar{q})); \quad (22)$$

- The substitution is safe.

As can be seen, τ changes the templates of the argument Φ_j F-modules, while ξ changes the formal arrays in the S-module. This effect can be summarized by

$$\xi(\text{internal-template } \Phi_j(\bar{q})) = \text{template } F_j(\tau_j(\bar{q})). \quad (23)$$

The effect of the τ_j will show up in the code of the resulting F-module, while the ξ play a role in the template definition.

DEFINITION 3.8. For a legal application (21) the resulting F-module is given by

- The parameters of the F-module are the parameters of the S-module that remain when all substitutions in Ξ have been performed;
- The global arrays of the resulting F-module are the union of the global arrays of the actual F-module arguments;
- The template of the F-module is the external template of the S-module after substitutions in Ξ have been performed. The resulting template of \tilde{F} can be summarized by

$$\tilde{\mathcal{F}}_{in} = \xi(\mathcal{E}_{in}(\beta(\vec{N}))) \quad \text{and} \quad \tilde{\mathcal{F}}_{out} = \xi(\mathcal{E}_{out}(\beta(\vec{N}))), \quad (24)$$

where $\beta(\vec{N})$ is the total effect of all binding substitutions;

- The statements of the F-procedure are the statements of the S-procedure that result when the substitution Ξ has been performed. The calls to the F_j now refer to the actual F-procedures, and not to an F-module as such.

We are now ready to formulate the central consistency theorem for the reuse of the computational structures as embodied in the F- and S-modules.

Theorem 3.9. *The generated template of \tilde{F} in the application (21) defines the dependencies of the F-module's F-procedure, provided the ap-*

plication is legal, all F-modules have correct templates, and the template specification of the S-module is correct.

It should be possible to check the matching of templates syntactically given a rewrite system that includes the arithmetic of the index domain types. Furthermore, this rewrite-system should be able to generate the resulting F-module of the application. However, further investigations into this have to be performed.

The definition of matching and application is illustrated in the computation of the recurrence (10). The S-module LDEP, (7), with linear internal template is here applied to the F-module GSTEP, (11). The result of the application $G = \text{LDEP}|_{\Xi}(\text{GSTEP})$, as given in (13), provides the Fibonacci-like computation, but on the exponential scale. The substitution Ξ , (12), defines an exponential expansion by ξ and a shift adjustment by τ .

3.7. Example: inverting the computation ordering. Let us illustrate an application by looking at the recurrence defined in Fig. 3. The F-module HEATSTEP (17) and the S-module FAN3 (20) are both defined based on the simple, two-dimensional recurrence in Fig. 3. Thus FAN3 may be applied to HEATSTEP using the identity substitution Ξ_{id} yielding the F-module

$$\text{SIMPLEHEATFLOW} = \text{FAN3}|_{\Xi_{id}}(\text{HEATSTEP}),$$

which may be expanded to

```

F-module SIMPLEHEATFLOW (N1, N2: integer) ==
  global X: array[*,*] of real
  template
    X[1..N1,0], X[0,0..N2-1], X[N1+1,0..N2-1]  $\longrightarrow$ 
    X[1..N1,1..N2]
  procedure (25)
    var q1, q2: integer;
    for q2:=1 to N2 do
      for q1:=1 to N1 do
        X[q1,q2] :=  $\varphi$ (X[q1-1,q2-1], X[q1,q2-1],
          X[q1+1,q2-1])
  end

```

SIMPLEHEATFLOW computes the interior of the domain using the equation (16), and leaves the appropriate values in the array X provided the borders are properly initialized.

When applying an S-module to an F-module a programmer should care not to exceed the bounds of global arrays.

When we applied the S-module FAN3 (20) to the F-module HEATSTEP (17) (see Fig. 3) no substitutions were necessary. Fig. 4 shows another two-dimensional, order 3 recurrence. The F-module will be very different from the previous case, but we may use the same S-module, if we choose the right substitutions to make the templates match. From the equations (26) in the figure, we write the F-module

$$\begin{aligned}
 y_{i_1, i_2} &= \psi(y_{i_1-1, i_2+1}, y_{i_1, i_2+1}, y_{i_1+1, i_2+1}), \\
 y_{t, 4} &= \epsilon_t, \quad t = 1, 2, \dots, 7, \\
 y_{0, t} &= \epsilon'_t, \quad t = 4, 3, 2, 1, \\
 y_{8, t} &= \epsilon''_t, \quad t = 4, 3, 2, 1,
 \end{aligned}
 \tag{26}$$




Fig. 4. A two-dimensional order 3 recurrence. The computation can be organized by the S-module FAN3 (20) which traverses the dependency graph shown in Fig. 3. The transformation to be considered is $i_1 = q_1$, $i_2 = 4 - q_2$ or $i_1 = 8 - q_1$, $i_2 = 4 - q_2$.

```

F-module DOWNSTEP (i1, i2: integer) ==
  global Y: array[*,*] of real
  template
    Y[i1-1, i2+1], Y[i1, i2+1], Y[i1+1, i2+1] → Y[i1, i2]
  procedure
    Y[i1, i2] := ψ(Y[i1-1, i2+1], Y[i1, i2+1], Y[i1+1, i2+1])
  end
    
```

To match this template with that of the S-module FAN3, we need the following substitution Ξ

$$\begin{aligned}x[\cdot_1, \cdot_2] &\mapsto Y[\cdot_1, 4 - \cdot_2], \\ \Phi(\cdot_1, \cdot_2) &\mapsto \text{DOWNSTEP}(\cdot_1, 4 - \cdot_2), \\ N1 &\mapsto 7; \quad N2 \mapsto 4.\end{aligned}$$

Since the dimensionality of the index domain is the same after the substitution, no redeclaration of F-module parameters or global arrays are needed, but the size of the global arrays get fixed since N1 and N2 are replaced by constants. The big changes occur in the templates and the body of the S-procedure, where the second index expression q2 in both procedure "calls" and array indexing operations are replaced by the index expression 4-q2. Thus the internal template $\Phi(q1, q2)$ gets substituted to

$$\begin{aligned}Y[q1-1, 4-(q2-1)], Y[q1, 4-(q2-1)], Y[q1+1, 4-(q2-1)] &\longrightarrow \\ Y[q1, 4-q2],\end{aligned}$$

which is equal to $\text{DOWNSTEP}(q1, 4-q2)$. The external template changes to

$$\begin{aligned}(Y[t1, 4-0], t1=1..7), (Y[0, 4-t2], t2=0..3), \\ (Y[8, 4-t2], t2=0..3) \longrightarrow (Y[t1, 4-t2], t1=1..7, t2=1..4),\end{aligned}$$

and this may be simplified to

$$\begin{aligned}(Y[t1, 4], t1=1..7), (Y[0, t2'], t2'=4..1), \\ (Y[8, t2'], t2'=4..1) \longrightarrow (Y[t1, t2'], t1=1..7, t2'=3..0),\end{aligned}$$

by replacing t2 with 4-t2' and normalizing the form of the implicit DO-loops. With this formulation it is easy to verify that the internal templates match, so the application

$$\text{DOWNHEATFLOW} = \text{FAN3} \Big|_{\Xi} (\text{DOWNSTEP})$$

may be expanded to

```
F-module DOWNHEATFLOW() ==
  global Y: array[0..8, 0..4] of real
  template Y[1..7, 4], Y[0, 1..4], Y[8, 1..4]  $\longrightarrow$  Y[1..7, 0..3]
```



```

procedure
  var q1, q2: integer;
  for q2:=1 to 4 do
    for q1:=1 to 7 do
      call DOWNSTEP(q1,4-q2)
    end
  end
  
```

(28)

showing that we may reuse the same S-module even for apparently quite different recurrences. We just have to develop a case specific transformation.

3.8. Example: affine transformation in the application of an S-module in 2D to an F-module in 3D. In this example, the S-module FAN3 (20) is applied to an F-module to organize the computation on a two-dimensional plane in three-dimensional space (see Fig. 5). The F-module PLANESTEP represents the recurrence (29) on the three-dimensional array Z

$$Z_{i_1, i_2, i_3} = \phi(Z_{i_1-2, i_2, i_3-1}, Z_{i_1-1, i_2-1, i_3-1}, Z_{i_1, i_2-2, i_3-1}). \quad (29)$$




Fig. 5. The computation on a two-dimensional plane in three-dimensional space can be organized by the S-module FAN3 (20). The transformation $i_1 = -q_1 - q_2 + h_1$, $i_2 = q_1 - q_2 + h_2$, $i_3 = q_2 + h_3$ embeds the S-module's FAN3 (20) data dependency graph which is shown in Fig. 3. Here \vec{v}_1 and \vec{v}_2 are the generating vectors for the plane, and \vec{h} is a shift vector.

```

F-module PLANESTEP (i1, i2, i3 : integer) ==
  global Z: array[*,*,*] of <type>
  template
    Z[i1-2,i2,i3-1], Z[i1-1,i2-1,i3-1],
    Z[i1,i2-2,i3-1] → Z[i1,i2,i3]
  procedure
    Z[i1,i2,i3] := φ(Z[i1-2,i2,i3-1], Z[i1-1,i2-1,i3-1],
    Z[i1,i2-2,i3-1])
end

```

(30)

The equation of a two-dimensional plane is a linear combination of two generating vectors $\vec{v}_1 = (-1, 1, 0)$ and $\vec{v}_2 = (-1, -1, 1)$

$$\{ \vec{i} \mid \vec{i} = \vec{v}_1 * t_1 + \vec{v}_2 * t_2 + \vec{h}, \quad t_1, t_2 = 0, 1, 2, \dots \},$$

where the vector $\vec{h} = (h_1, h_2, h_3)$ plays the role of a shift. The application of the S-module FAN3 (20) to the above F-module PLANESTEP (30) yields a new F-module

$$\text{PLANEFLOW} = \text{FAN3}|_{\Xi}(\text{PLANESTEP}) \quad (31)$$

where we need the substitution $\Xi = [\xi, \tau]$

$$\begin{aligned} \xi : \mathbf{x}[\cdot_1, \cdot_2] &\mapsto Z[-\cdot_1 - \cdot_2 + h_1, \cdot_1 - \cdot_2 + h_2, \cdot_2 + h_3], \\ \tau : \Phi(\cdot_1, \cdot_2) &\mapsto \text{PLANESTEP}(-\cdot_1 - \cdot_2 + h_1, \cdot_1 - \cdot_2 + h_2, \cdot_2 + h_3). \end{aligned}$$

The yielded F-module PLANEFLOW (31) has parameters $h_1, h_2, h_3, N_1,$ and N_2

```

F-module PLANEFLOW (h1, h2, h3, N1, N2: integer) ==
  global Z: array [*,*,*] of <type>
  template (Z[-t1+h1,t1+h2,h3], t1=1..N1),
    (Z[-t2+h1,-t2+h2,t2+h3], t2=0..N2-1),
    (Z[-N1-1-t2+h1,N1+1-t2+h2,t2+h3], t2=0..N2-1) →
    (Z[-t1-t2+h1,t1-t2+h2,t2+h3], t1=1..N1, t2=1..N2)
  procedure
    var q1,q2: integer;
    for q2:=1 to N2 do
      for q1:=1 to N1 do
        call PLANESTEP(-q1-q2+h1, q1-q2+h2, q2+h3)
  end

```

4. Constructive recursion. *Constructive recursion* (CR) is an extension of the primitive recursive and μ -recursive schemes over the integers to recursion schemes over any graph structure. It was developed by Haveraaen and early ideas are presented in (Haveraaen 1990; 1993). Although constructive recursion was developed independently from the work of Tucker and Zucker (1988), CR can be seen as a generalization of the theory of computations on arbitrary algebraic structures developed there. Tucker and Zucker restrict their recursive structures to discrete recursion schemes, while CR can be defined for denser graphs. The development of constructive recursion was inspired by ideas in the programming language Crystal (Chen *et al.*, 1991), which in turn was inspired by systolic algorithms.

The CR approach defines general recurrences (5) by distinguishing between the definition of the *data dependency graph*, the *recursive relation* defining a value on each node of the data dependency graph, and the specification of initial values at *input nodes* and the interesting set of *output nodes*. A data dependency graph is a directed multigraph with two edge-labeling functions that satisfy certain injectivity properties. The graph is defined as an algebraic structure, which allows for the concise expression of repetitive structures. We will refer to the graph as a *data dependency algebra (dda)*. The recursive functions define a relation between the value of a node and its neighbors in the direction of the arcs of the dda. Thus the value at any node P may be computed if the values at the nodes that P depends on are known. This generalizes to the whole graph, so that given a certain set of nodes with initial values, it will be possible to compute the values of (some subset of) the dda nodes. The set of interesting outputs is a subset of the set of nodes whose values may be computed.

DEFINITION 4.1. A *data dependency algebra (dda)* is given by the data of an 8-tuple

$$\langle P, D, \text{canR}, \text{canS}, r, s, \text{dirS}, \text{dirR} \rangle, \quad (32)$$

where

- | | | |
|------------------------------------|---|---------------------------------------|
| P | – | is a set of points, |
| D | – | is a set of directions, |
| $\text{canR} \subseteq P \times D$ | – | is the relation <i>can receive</i> , |
| $\text{canS} \subseteq P \times D$ | – | is the relation <i>can send</i> , |
| $r : \text{canR} \rightarrow P$ | – | is the function <i>receive from</i> , |

$\text{dirS} : \text{canR} \rightarrow D$ - is the function *direction sent to*,
 $\text{s} : \text{canS} \rightarrow P$ - is the function *send to*, and
 $\text{dirR} : \text{canS} \rightarrow D$ - is the function *direction to receive from*,

so that the following equations are satisfied (writing $\text{canR}(p,d)$ for $p,d \in \text{canR}$, etc.)

$$\begin{aligned}
 \text{canR}(p,d) &\Rightarrow \text{canS}(r(p,d), \text{dirS}(p,d)); \\
 \text{canR}(p,d) &\Rightarrow \text{s}(r(p,d), \text{dirS}(p,d)) = p; \\
 \text{canR}(p,d) &\Rightarrow \text{dirR}(r(p,d), \text{dirS}(p,d)) = d; \\
 \text{canS}(p,d) &\Rightarrow \text{canR}(s(p,d), \text{dirR}(p,d)); \\
 \text{canS}(p,d) &\Rightarrow \text{r}(s(p,d), \text{dirR}(p,d)) = p; \\
 \text{canS}(p,d) &\Rightarrow \text{dirS}(s(p,d), \text{dirR}(p,d)) = d;
 \end{aligned}$$

Let $D_p = \{d \in D \mid \langle p, d \rangle \in \text{canR}\}$.

There is a symmetry between the r - and the s -notions of a dda, and this is captured by

COROLLARY 4.2

$$\langle P, D, \text{canR}, S, r, s, \text{dirS}, \text{dirR} \rangle$$

is a dda if and only if

$$\langle P, D, \text{canS}, \text{canR}, s, r, \text{dirR}, \text{dirS} \rangle$$

is a dda.

There is a close correspondence between dda's and a special class of edge-labeled multi-graphs, just think of P as the set of nodes and canR as the arcs with D as edge-labels. This allows us to use graph terminology when discussing dda's.

To concretize the definition, let's take a look at the order k linear recurrence (1). We can define the recurrence data dependency algebra RDDA with P being the set of natural numbers, $D = \{1, 2, \dots, k\}$ and the labels of an arc being the number in D representing the distance between two nodes. As is common in programming, we define relations (canR and canS) as boolean functions.

$$\begin{aligned}
 \text{canR}(p,d) &= (1 \leq d \leq k) \text{ and } (k \leq p); \\
 r(p,d) &= p-d;
 \end{aligned}$$

$$\begin{aligned}
\text{dirS}(p,d) &= d; \\
\text{canS}(p,d) &= (1 \leq d \leq k) \text{ and } (k \leq p+d); \\
s(p,d) &= p+d; \\
\text{dirR}(p,d) &= d.
\end{aligned}
\tag{33}$$

In addition we can define two injective node-labeling functions

$$\begin{aligned}
\text{llab}(p) &= p, \\
\text{elab}(p) &= 2**p.
\end{aligned}
\tag{34}$$

The graph defined by the RDDA when $k = 2$ is depicted in Fig. 1 using the `llab` node labeling scheme and in Fig. 2 using the `elab` node labeling scheme.

Given a set of points in a dda, there is a notion of the points which can be reached using the `s`-operations of the dda.

DEFINITION 4.3. A *maximal domain* A over $I \subseteq P$, given a dda

$$H = \langle P, D, \text{canR}, \text{canS}, r, s, \text{dirS}, \text{dirR} \rangle$$

is defined by

- $p \in I$ then $p \in A$,
- $p \in A \setminus I$ then for all $d \in D_p$, $r(p, d) \in A$, and
- $p \in P \setminus A$ then for some $d \in D_p$, $r(p, d) \notin A$.

The *least maximal domain*, denoted P_I , over $I \subseteq P$ is such that $P_I \subseteq A$ for all A that are maximal domains over I .

There may be several maximal domains over a set I , e.g., P itself is always a maximal domain over any $I \subseteq P$.

PROPOSITION 4.4. *The least maximal domain P_I over $I \subseteq P$ always exists, and is the set of points solely reachable with paths in the `s`-direction from the points I .*

We say that P_I is *generated* by `s`-paths from I .

In order to define recurrence relations over a dda, we need to have a notion of consistently assigning values to the points of the dda. This will allow us to define a computation generating those values.

DEFINITION 4.5. A *P -indexed relation* is given by a 7-tuple

$$C = \langle P, D, R, T, x, x^D, E \rangle,$$

where P , D and T are sets, $R \subseteq P \times D$ such that for any $p \in P$ there is a subset $D_p = \{d_p \mid \langle p, d_p \rangle \in R\} \subseteq D$, x is a variable over the set P , $x^D = \{x_d \mid d \in D\}$ is a set of variables over the set T , and $E = \langle E_1, \dots, E_m \rangle$ is a sequence of conditional expressions

$$E_j = c_{j,1} \ \& \ \dots \ \& \ c_{j,n_j} \ \rightarrow \ e_j,$$

where the $c_{j,i}$ are truth expressions and e_j is an expression with value in T . The expressions may be written in some programming language extended with expressions with values in P and D , but without using loops or recursion, and such that all free variables in the expressions are from $\{x\} \cup x^D$.

C defines a family $C_P = \langle C_p \mid p \in P \rangle$ of partial functions

$$C_p = T^{D_p} \rightarrow T$$

where $T^{D_p} = \{V \mid V = \langle v_d \in T \mid d \in D_p \rangle\}$ is a set of families of values from T . $C_p(V)$, for $V \in T^{D_p}$, is defined when at least one of the conditional expressions E_j is such that all expressions $c_{j,i}$ and e_j are defined, all the conditions $c_{j,i}$ are true, and there are no free variables remaining, when given the assignment $x = p$ and $x_d = v_d$ for every $d \in D_p$. The value of $C_p(V)$ is then the value of the expression in the first of those E_j .

The conditional expression E may easily be thought of as an if-expression, so C_P may be conceived of as a procedure with the variables x and x^D as parameters. This captures the ϕ of (5), and corresponds roughly to the F-modules. Now we should combine this definition with that of the dda, in order to assign values to the points of a dda and thus describing the whole recurrence.

DEFINITION 4.6. Given a dda $H = \langle P, D, \text{canR}, \text{canS}, \text{r}, \text{s}, \text{dirS}, \text{dirR} \rangle$ and a P -indexed relation $C = \langle P, D, \text{canR}, T, x, x^D, E \rangle$, a consistent assignment of values T to a set $A \subseteq P$ is a mapping $f : A \rightarrow T$ such that for every $p \in A$ either

- for some $d \in D_p$, $\text{r}(p, d) \notin A$, or
- for all $d \in D_p$, $\text{r}(p, d) \in A$ and $C_p(V)$ is defined and $f(p) = C_p(V)$, where $V = \langle v_d \mid v_d = f(\text{r}(p, d)), d \in D_p \rangle$.

DEFINITION 4.7. A maximally extended consistent assignment of a consistent assignment $f : A \rightarrow T$ is a mapping $g : B \rightarrow T$, such that $A \subseteq B \subseteq P$,

for every $p \in A$, $g(p) = f(p)$, $g : B \rightarrow T$ is a consistent assignment, and for every point $p \in P \setminus B$,

- for some $d \in D_p$, $\mathbf{r}(p, d) \notin B$, or
- for all $d \in D_p$, $\mathbf{r}(p, d) \in B$ and $C_p(V)$, where $V = \langle v_d \mid v_d = f(\mathbf{r}(p, d), d \in D_p) \rangle$, is undefined.

There may exist many maximally extended consistent assignments given a consistent assignment, but since C_P is a partial function, they are identical on a subset of the points generated by the set of initial points.

PROPOSITION 4.8. *For every consistent assignment $f : A \rightarrow T$ for $A \subseteq P$, there exists a least maximally extended consistent assignment $\hat{f} : A_C \rightarrow T$ such that for any maximally extended consistent assignment $g : B \rightarrow T$ of $f : A \rightarrow T$, then $A_C \subseteq B$ and for all $p \in A_C$, $g(p) = \hat{f}(p)$.*

Moreover, $A_C \subseteq P_A$, where P_A is the least maximal domain over $A \subseteq P$.

The interesting case is that we can force the existence of a unique such set given initial values for a collection of points.

Theorem 4.9. *Given a dda $H = \langle P, D, \text{canR}, \text{canS}, \mathbf{r}, \mathbf{s}, \text{dirS}, \text{dirR} \rangle$ and a P -indexed relation $C = \langle P, D, \text{canR}, T, x, x^D, E \rangle$, for every non-empty $I \subseteq P$ and $\alpha : I \rightarrow T$, there is a least maximally extended consistent assignment $\bar{\alpha} : P_{C,\alpha} \rightarrow T$, where $P_{C,\alpha} \subseteq P_I \setminus I$, such that when $p \in P_I \setminus I$ and all $d \in D_p$, $\mathbf{r}(p, d) \in I$ then either*

- $C_p(V)$, where $V = \langle v_d \mid v_d = \alpha(\mathbf{r}(p, d), d \in D_p) \rangle$, is undefined and $p \notin P_{C,\alpha}$, or
- $C_p(V)$, where $V = \langle v_d \mid v_d = \alpha(\mathbf{r}(p, d), d \in D_p) \rangle$, is defined and $p \in P_{C,\alpha}$.

We are now able to define a constructive recursive programming style with a semantics as defined above.

DEFINITION 4.10. A *constructive recursive relation* $f : P \rightarrow T$ over a dda $H = \langle P, D, \text{canR}, \text{canS}, \mathbf{r}, \mathbf{s}, \text{dirS}, \text{dirR} \rangle$, is a function definition

$$E = [\mathbf{f}(x) = \text{if } E_1 \mid E_2 \mid \dots \mid E_m \text{ fi}],$$

where

$$E_j = c_{j,1} \ \& \ \dots \ \& \ c_{j,n_j} \ \rightarrow \ e_j,$$

where the $c_{j,i}$ are truth expressions and e_j is an expression with value in T . The expressions may be written in some programming language extended with variable p , expressions with values in P and D , but without using loops and restricting recursion to the form $\mathbf{f}(x(d))$, where $d \in D$.

A *constructive recursive program* $\mathbf{F} = \langle E, \alpha : I \rightarrow T, Q \rangle$ is a constructive recursive relation $f : P \rightarrow T$ given by E , with the initial value assignments $\alpha : I \rightarrow T$, and a set of demanded output points $Q \subseteq P$.

The set Q is the set of interesting output values, and does not have to reflect the points that need to be computed in order to obtain these values, and it need not reflect a set of values that can be computed using the given α . The actual result will be a computable subset of Q . We can define the initial value assignments α by writing $f(p) = \alpha(p)$ for all $p \in I$. The output set Q , and the dependency on input values, could be defined using a notation with \longrightarrow and lists of function elements, similar to that of dependency patterns in SB.

DEFINITION 4.11. The semantics of constructive recursive program $\langle E, \alpha : I \rightarrow T, Q \rangle$ over a dda $H = \langle P, D, \text{canR}, \text{canS}, \mathbf{x}, \mathbf{s}, \text{dirS}, \text{dirR} \rangle$ is the map $\Omega : P_{C,\alpha} \cap Q \rightarrow T$, $\Omega(p) = \bar{\alpha}(p)$ for every $p \in P_{C,\alpha} \cap Q$, as given by the least maximal consistent assignment of the P -indexed relation

$$C = \langle P, D, \text{canR}, T, \mathbf{x}, x^D, \hat{E} \rangle,$$

with initial values $\alpha : I \rightarrow T$ and

$$\hat{E}_j = \hat{c}_{j,1} \ \& \ \dots \ \& \ \hat{c}_{j,n_j} \ \rightarrow \ \hat{e}_j,$$

for every $E_j = c_{j,1} \ \& \ \dots \ \& \ c_{j,n_j} \ \rightarrow \ e_j$, such that $\hat{c}_{j,i}$ is $c_{j,i}$ and \hat{e}_j is e_j , where, for every $d \in D$, $\mathbf{f}(x(d))$ has been replaced by $x_d \in x^D$.

This provides a compilation strategy, where we translate the equations E into programming language functions with headers

```
function CPb(x:P,V:array[D] of T): boolean;
function CP(x:P,V:array[D] of T): T;
```

such that, when $x = p$, CPb tells when $C_p(V)$ is defined, and, CP returns the value in T given by $C_p(V)$ when this value is defined. Both these procedures can easily be created by a compiler based on the if-statement embedded in \hat{E} by replacing the free variables $x^D = \langle x_d \rangle$ of \hat{E} by $V[d]$.

A recurrence of the form (1) can easily be expressed as a constructive recursive relation. In the Fibonacci (2) case, we choose the graph nodes of the RDDA guided by the function `lpro`, and arrive at the following CR relation

$$\begin{aligned} \text{Fib}(x) = & \text{if canR}(x,1) \ \& \ \text{canR}(x,2) \ \rightarrow \\ & \text{Fib}(r(x,1)) + \text{Fib}(r(x,2)) \ \text{fi}, \quad (35) \\ \text{Fib}(0)=0, \ \text{Fib}(1)=1 \ \rightarrow & \ (\text{Fib}(t), \ t=0..N). \end{aligned}$$

This means that whenever there is an arc in direction 1 and an arc in direction 2 from a point p , the value at p is the sum of the values at $r(p,1)$ and $r(p,2)$.

So far we have only showed that constructive recursive programs are well defined. In order to execute them on a computer, the following is needed.

Theorem 4.12. *A constructive recursive program F is computable by a wave-front algorithm, if the dda is computable, the set I of initial values are finite, all expressions in E are computable, and $Q \cap P_{C,\alpha}$ is generated by finite s -paths for the P -indexed relation C derived from F .*

This provides us with a very general execution strategy. It does not claim anything about the order of evaluation, and there is no means of controlling the size of the temporary storage needed by the algorithm.

We may now use this result and compute the Fibonacci numbers by traversing the graph in the s -direction, from the inputs at 0 and 1, till we have computed all values in the output set.

4.1. Space-time structured computations. The previous development provided us with a wave-front algorithm to compute any constructive recursive program. This algorithm had the drawback that the storage space requirements were not under control. In order to achieve this we need to provide the dda with extra structure.

DEFINITION 4.13. *A space-time algebra (sta) is given by the data of a 7-tuple*

$$\langle H, S, R, <, \text{space}, \text{time}, \text{point} \rangle,$$

where H is a dda as in (32), $<$ is an ordering relation on R , and $\text{space} : P \rightarrow S$, $\text{time} : P \rightarrow R$ and $\text{point} : S \times R \rightarrow P$ are functions, such that

$$\begin{aligned} \text{point}(\text{space}(p), \text{time}(p)) &= p, \\ \text{time}(r(p, d)) &< \text{time}(p). \end{aligned}$$

An immediate consequence of this is that $\text{time}(p) < \text{time}(s(p, d))$.

DEFINITION 4.14. A computable space-time dependency algebra of height k is a sta where the dda and the sta functions are computable, S is finite, R is a subset of the integers, and $\text{time}(p) - \text{time}(r(p, d)) \leq k$.

We can amend the dda RDDA, (33), with $S = \{1\}$, $R = \{0, 1, \dots\}$,

$$\begin{aligned} \text{space}(x) &= 1 \\ \text{time}(x) &= 2 \\ \text{point}(i, j) &= j \end{aligned} \tag{36}$$

in order to get an sta. The sta has height k when $D = \{1, 2, \dots, k\}$.

Theorem 4.15. A constructive recursive program $F = \langle E, \alpha : I \rightarrow T, Q \rangle$ over the dda part of a computable sta with height at most k is, given some constant c , computable with fixed execution time requirement proportional to $(c + k) * s$ and storage space requirements proportional to $k * s$, where s is the cardinality of S , if $\text{time}(p) \in \{0, 1, \dots, k - 1\}$ for all $p \in I$, all expressions in E are computable, and $\text{time}(p) - c \in \{0, 1, \dots, k\}$ for all $p \in Q$, the set of output values.

Proof. The CR program F may be compiled into the recursion free loop procedure \mathcal{F} shown in Fig. 6. This procedure computes F efficiently by traversing $k * s$ points of the array U .

The algorithm in Fig. 6 will visit every element of the temporary storage as many times as required to compute the output set. If we have the situation that $c = 0$, $I \cap Q = \emptyset$, $\text{space}(I \cup Q) = S$, and $\text{time}(I \cup Q) = \{0, 1, \dots, k\}$, we have a situation very similar to that used in the structural blanks approach.

For many CR programs, such as `Fib`, (35), the cardinality of the space component is 1. In this case we could remove the first dimension of the array U in Fig. 6. In the computation of (35) we should set $c = 0$ and $k = N$ in order to return the whole output set.

If we want to embed the array $U[S, R]$ into a larger array $A[B_1, B_2, \dots, B_n]$, this can easily be done by providing injective functions $\text{ind}_1 : S \times R \rightarrow B_1$, $\text{ind}_2 : S \times R \rightarrow B_2$, ..., $\text{ind}_n : S \times R \rightarrow B_n$, and replacing the references to $U[j, i]$ with references $A[\text{ind}_1(j, i), \text{ind}_2(j, i), \dots, \text{ind}_n(j, i)]$.

The space-time structure can also be used to map computations on parallel

```

procedure  $\mathcal{F}$  (c, k: R; var  $\Omega$ : array[Q] of T);
var U: array[S,0:k] of T;

forall p $\in$  I
do U[space(p),time(p)] =  $\alpha$ (p) od;

for i := 0 to c+k
do forall j $\in$  S
  do var b: boolean;
  if point(j,i) $\notin$  I
  then b := true;
  forall d  $\in$   $D_p$ 
  do q := r(point(j,i),d);
  b := b and (U[space(q),time(q) mod (k+1)]
     $\neq$  undefined)
  od
  if b
  then var V: array[D] of T;
  forall d  $\in$   $D_p$ 
  do q := r(point(j,i),d);
  V[d] := U[space(q),time(q)]
  od;
  if CPb(point(j,i),V)
  then U[j,i mod (k+1)] = CP(point(j,i),V)
  else U[j,i mod (k+1)] = undefined
  fi;
  fi
  od
od;

forall p $\in$  Q
do  $\Omega$ [p] = U[space(p),time(p)] od
end;

```

Fig. 6. A recursion free loop program that computes the CR program $F = \langle E, \alpha : I \rightarrow T, Q \rangle$ over an sta, where Ω gives the outputs. Computations in the S coordinate are unordered, while those over $R = \{0, 1, \dots, c+k\}$ are ordered. The functions CPb and CP are compiled from E , so this is not an interpreter. The choice of values for c, k depends on the sets I and Q and the height of the sta.

computers by defining the sta based on the communication structure of a parallel computer. See Haveraaen (1990; 1993) for more information.

4.2. Development methodology The development methodology for defining constructive recursive programs is quite similar to that of the structural blanks approach, see Section 3.1. Going from a recurrence of the form (5) to a constructive recursive program may be described as a 3-step approach.

In **step 1** we define the constructive recursive function based on the recurrence equation. Simply use the set $D = \{1, \dots, k\}$ as given by the occurrences of the recurrence variable on the right-hand side of the equation. Then define the CR program by

$$\begin{aligned} f(x) = & \text{if canR}(x,1) \ \& \ \text{canR}(r(x,2)) \ \& \ \dots \ \& \ \text{canR}(x,k) \\ & \rightarrow \phi(f(r(x,1)), f(r(x,2)), \dots, f(r(x,k)) \text{ fi,} \end{aligned}$$

and the appropriate set of input values and desired outputs. This corresponds to the task of writing an F-module in the SB development methodology.

Step 2 is to define the dda based on the structure of the recurrence. In the normal case an appropriate dda will be available in a library of dda's, just as one expect to find a suitable S-module in a library in the SB case. A scheduling algorithm, given in the form of an sta connected to the dda, will also be in the CR library.

If the dda is not in the library, it must be developed. The sets $P = X$, where X is the index domain of the recurrence, and $D = \{1, \dots, k\}$ are straight forward. The r and canR functions follow directly from the recurrence equations, as they are defined from the δ -functions, although the form chosen may vary (see Lundevik (1994) for details). dda-functions are more difficult and may require ingenuity if it is to be done efficiently. For affine dependencies and certain other special cases this can be automated. Finally the scheduling task embedded in an sta definition should be provided. This corresponds closely to the development of a S-module in the SB case.

Step 3 is to verify that the recursive relation, given the dda, defines the recurrence. This involves defining an injective function $\text{nlab} : P \rightarrow X$, where X is the index domain of the recurrence, such that $\delta_i(\text{nlab}(p)) = \text{nlab}(r(p,i))$, for $i = 1, \dots, k$.

Let us apply this programming technique on the recurrence defined in (10),

using RDDA defined in (33). The CR relation and input/output sets are

$$g(x) = \text{if canR}(x,1) \ \& \ \text{canR}(x,2) \rightarrow \gamma(g(r(x,1)),g(r(x,2))) \text{ fi,}$$

$$g(0) = \varepsilon_0, \ g(1) = \varepsilon_1 \longrightarrow g(t), \ t=2..N.$$

for $g: P \rightarrow \text{real}$. The sta is given in (36). Also in this case $c=0$ and $k=N$ for the algorithm in Fig. 6. We use the `epr` function to label the points consistently with the domain used in (10), and show that $\text{epr}(p)/2^d = \text{epr}(r(p,d))$ for all $p \in P$ and $d \in D_p$.

5. Comparing the approaches. The structural blanks and constructive recursive approaches are similar in that they break a recurrence equation in a structural part (templates of the S- and F-modules or a graph defined by a dda), a computational part (F-module and recursive relation), and a defined initial value set and output set (external template of S-module or separate input/output description in the CR case).

Looking at how the input/output description is handled, we see that the CR approach is slightly more flexible. There we define the output set independently from the graph, so we may specify that only the N 'th value, or some other subset of the computed values, is to be output. In the SB approach we will have to modify the driver routine to reflect and take advantage of any change in the external template of an S-module. But, as in the CR case, any subset of the computed values may of course be specified as the set of outputs.

In the SB case the data array is declared by the user, separately from the modules, while in the CR approach the points to be computed are given implicitly by the dda graph. In the latter case we may provide the points with a space-time mapping, assigning the computations to explicit time-steps on the processors of a parallel computer. Thus data parallelism is inherent in the CR approach. The data distribution information needed in the SB case is dependent on the driver procedure of the S-module, but has to be declared by the user of the module independently of what S-module that is to be chosen, hence data parallelism is difficult to achieve in the SB approach.

When it comes to programmability the SB approach probably has the edge on the CR approach. SB is based on conventional languages such as Fortran, a notation the practitioner is familiar with. The practitioner is also familiar with the task of adding *pragmas* and additional information to make

a program run faster or exploit properties of specific architectures. The CR approach, although clearly related to the mathematical structure of the recurrence, introduces concepts not used in traditional programming, and has a less familiar notation. In the next subsections we show the two approaches on some more complex examples to illustrate the practical similarities and differences.

5.1. Example: mutually dependent equations. Given a set of mutually dependent recurrences

$$\begin{aligned}
 x_q &= \varphi_1(x_{q-1}, y_q), \\
 y_q &= \varphi_2(y_{q-1}, z_q), \\
 z_q &= \varphi_3(z_{q-1}, y_{q-1}), \\
 x_0 &= \xi_0, \quad y_0 = \nu_0, \quad z_0 = \zeta_0.
 \end{aligned}
 \tag{37}$$

The dependency pattern is shown in Fig. 7. Expressing this in the structural blanks approach we will define three arrays $X[0..N]$, $Y[0..N]$ and $Z[0..N]$, where N is a constant defining how much we want computed. For each of the equations define an F-module

```

F-module FX (q: integer) ==
    global X, Y: array[*] of <type>
    template X[q-1], Y[q] → X[q]
    procedure X[q] :=  $\varphi_1(X[q-1], Y[q])$ 
end

F-module FY (q: integer) ==
    global Y, Z: array[*] of <type>
    template Y[q-1], Z[q] → Y[q]
    procedure Y[q] :=  $\varphi_2(Y[q-1], Z[q])$ 
end

F-module FZ (q: integer) ==
    global Y, Z: array[*] of <type>
    template Y[q-1], Z[q-1] → Z[q]
    procedure Z[q] :=  $\varphi_3(Y[q-1], Z[q-1])$ 
end

```




Fig. 7. The dependencies of the mutually dependent recurrence in (37) shown separately (on the top) and combined into one large graph (bottom).

We assume that elements $x[0], y[0], z[0]$ are initialized. Elements $x[1], \dots, x[N], y[1], \dots, y[N], z[1], \dots, z[N]$ are to be computed.

The S-module that will structure the computation of the array elements must describe a computation that starting from the initial elements in $x[0], y[0]$ and $z[0]$ will compute the next element based on already computed values. Denoting the patterns Φ_1, Φ_2 and Φ_3 , we see that this can be achieved by repeating the sequence $\Phi_3(q); \Phi_2(q); \Phi_1(q)$

```

S-module S3 ( $\Phi_1, \Phi_2, \Phi_3$ : Fmod(integer);  $N$ : integer) ==
    formal x, y, z: array[*]
internal-template    --- Three internal templates:
    (var q: integer;  $\Phi_1(q) == x[q-1], y[q] \longrightarrow x[q]$ )
    (var q: integer;  $\Phi_2(q) == y[q-1], z[q] \longrightarrow y[q]$ )
    (var q: integer;  $\Phi_3(q) == z[q-1], y[q-1] \longrightarrow z[q]$ )
    
```

```

external-template (38)
  x[0], y[0], z[0]  $\longrightarrow$  x[1..N], y[1..N], z[1..N]
procedure
  var q: integer;
  for q:=1 to N do
    begin call  $\Phi_3(q)$ ; call  $\Phi_2(q)$ ; call  $\Phi_1(q)$  end
  end

```

Using the substitution

$$\Xi = [x[\cdot] \mapsto X[\cdot]; \quad y[\cdot] \mapsto Y[\cdot]; \quad z[\cdot] \mapsto Z[\cdot]; \\ \Phi_1(\cdot) \mapsto FX(\cdot); \quad \Phi_2(\cdot) \mapsto FY(\cdot); \quad \Phi_3(\cdot) \mapsto FZ(\cdot)],$$

we see that S3 may readily be applied to FX, FY and FZ yielding FXYZ = S3| Ξ (FX, FY, FZ). The data dependency graph of FXYZ is illustrated in Fig. 7, with array names X, Y and Z being substituted for x, y and z.

In the constructive recursion approach the graph depicted in Fig. 7 may be defined by taking the set of points $P = \{1, 2, 3\} \times \mathcal{N}$, where $\mathcal{N} = \{0, 1, 2, \dots\}$, and the set of directions $D = \{1, 2, 3, 4, 5, 6\}$. The functions required to define a dda are

$$\begin{aligned}
 \text{canR}(\langle t, q \rangle, d) &= q > 0 \text{ and (case } t \text{ of 1: } d \text{ in } [1, 2]; \\
 &\quad 2: d \text{ in } [3, 4]; \\
 &\quad 3: d \text{ in } [5, 6]); \\
 \text{canS}(\langle t, q \rangle, d) &= (q = 0 \text{ and (case } t \text{ of 1: } d \text{ in } [1]; \\
 &\quad 2: d \text{ in } [3, 6]; \\
 &\quad 3: d \text{ in } [5])) \text{ or} \\
 &\quad (q > 0 \text{ and (case } t \text{ of 1: } d \text{ in } [1]; \\
 &\quad 2: d \text{ in } [2, 3, 6]; \\
 &\quad 3: d \text{ in } [4, 5])); \\
 r(\langle t, q \rangle, d) &= (\text{case } d \text{ of 1: } \langle 1, q-1 \rangle; 2: \langle 2, q \rangle; \\
 &\quad 3: \langle 2, q-1 \rangle \quad 4: \langle 3, q \rangle; \\
 &\quad 5: \langle 3, q-1 \rangle; 6: \langle 2, q-1 \rangle); \\
 s(\langle t, q \rangle, d) &= (\text{case } d \text{ of 1: } \langle 1, q+1 \rangle; 2: \langle 1, q \rangle; \\
 &\quad 3: \langle 2, q+1 \rangle \quad 4: \langle 2, q \rangle; \\
 &\quad 5: \langle 3, q+1 \rangle; 6: \langle 3, q+1 \rangle); \\
 \text{dirS}(\langle t, q \rangle, d) &= d; \\
 \text{dirR}(\langle t, q \rangle, d) &= d.
 \end{aligned} \tag{39}$$

Since this dda has a strict linear dependency, the sta has $S = \{1\}$, $R = \mathcal{N}$,

$$\begin{aligned} \text{space}(\langle t, q \rangle) &= 1, \\ \text{time}(\langle t, q \rangle) &= 3 * q + 3 - t, \\ \text{point}(i, j) &= \langle 3 - (j \bmod 3), j \text{ div } 3 \rangle, \end{aligned}$$

which gives the sta a height of 3. The functions needed to show the correspondence with the graph in Fig. 7 are

$$\begin{aligned} \text{labxyz}(\langle t, q \rangle) &= t, \\ \text{labN}(\langle t, q \rangle) &= q, \end{aligned}$$

where the return values of labxyz are 1 for array x, 2 for array y and 3 for array z. Taken as a pair labxyz and labN define an injective mapping that is preserved over the r-function. The function xyz that defines the recurrence in (37) is

$$\begin{aligned} \text{xyz}(p) &= \text{if } \text{labxyz}(p) = 1 \rightarrow \varphi_1(\text{xyz}(r(p,1)), \text{xyz}(r(p,2))) \\ & \quad | \text{labxyz}(p) = 2 \rightarrow \varphi_2(\text{xyz}(r(p,3)), \text{xyz}(r(p,4))) \\ & \quad | \text{labxyz}(p) = 3 \rightarrow \varphi_3(\text{xyz}(r(p,5)), \text{xyz}(r(p,6))) \text{ fi} \end{aligned}$$

The definition of xyz is split into three cases, each corresponding to one equation of (37). Then

$$\begin{aligned} \text{xyz}(\langle 1, 0 \rangle) &= \xi_0, \text{xyz}(\langle 2, 0 \rangle) = v_0, \text{xyz}(\langle 3, 0 \rangle) = \zeta_0 \longrightarrow \\ \text{xyz}(\langle t, q \rangle), \quad t &= 1..3, \quad q = 1..N \end{aligned} \quad (40)$$

is the input/output specification. In order to retain the full set of computed values, we let $c=0$ and $k=N$ in the algorithm in Fig. 6.

5.2. Example: embedding in higher-dimensional domain. To illustrate the flexibility of these approaches, we will define a diagonal c2omputation in a two-dimensional domain W (see Fig. 8). The computational aspects are defined by the F-modules F1 and F2

```
F-module F1 (i1, i2: integer) ==
  global W: array[*,*] of <type>
  template W[i1-1,i2-1], W[i1+1,i2-1]  $\longrightarrow$  W[i1,i2]
  procedure W[i1,i2] :=  $\gamma_1$ (W[i1-1,i2-1], W[i1+1,i2-1])
end
```

(41)




Fig. 8. Three mutually dependent two-dimensional order 2 recurrences as defined by F-modules F1 (41) (used twice) and F2 (42). The structural similarity with the recurrence (37) as depicted in Fig. 7 is obvious: think of the upper diagonal as *x*, the intermediate one as *y* and the lower one as *z*, and check the labeling of the arcs.

```

F-module F2 (i1, i2: integer) ==
    global W: array[*,*] of <type>
    template W[i1-1,i2-1], W[i1-2,i2] → W[i1,i2]           (42)
    procedure W[i1,i2] := γ2(W[i1-1,i2-1], W[i1-2,i2])
end
    
```

Where F1 is to be used twice in order to define the diagonals marked *x* and *y*. Due to the structural similarity with (37), we may use the S-module S3 (38) as the driver, giving the application $FDIAG = S3|_{\Xi}(F1, F1, F2)$. The substitution needed is

$$\begin{aligned}
 \Xi = \{ & \xi_x(\cdot) : x[\cdot] \mapsto W[-2, \cdot + 2]; \\
 & \xi_y(\cdot) : y[\cdot] \mapsto W[-1, \cdot + 1]; \\
 & \xi_z(\cdot) : z[\cdot] \mapsto W[\cdot, \cdot];
 \end{aligned}$$

$$\begin{aligned}\tau_{\Phi_1}(\cdot) &: \Phi_1(\cdot) \mapsto F_1(\cdot-2, \cdot+2); \\ \tau_{\Phi_2}(\cdot) &: \Phi_2(\cdot) \mapsto F_1(\cdot-1, \cdot+1); \\ \tau_{\Phi_3}(\cdot) &: \Phi_3(\cdot) \mapsto F_2(\cdot, \cdot)\end{aligned}$$

Writing the result of the application out in full we get the F-module

```
F-module FDIAG (N: integer) ==
  global W: array[*,*] of <type>
  template W[-2,2], W[-1,1], W[0,0]  $\longrightarrow$ 
    (W[t-2,t+2], t=1..N), (W[t-1,t+1], t=1..N),
    (W[t,t], t=1..N)
  procedure (43)
    var q: integer;
    for q := 1 to N do
      begin call F2(q,q); call F1(q-1,q+1); call F1(q-2,q+2) end
  end
```

We now need to show that the substitution Ξ is safe, and this follows since each formal array is mapped to a different diagonal of W .

Let us provide more detail for the requirement that S3 was applied to F1 and F2 correctly, specifically that the templates of F1 and F2 match the internal templates S3, (38). We have to prove that the mapping $\xi = \langle \xi_x, \xi_y, \xi_z \rangle$ maps x , y and z nodes to W nodes preserving the isomorphism between the internal templates Φ_1 , Φ_2 , and Φ_3 and templates of F1 and F2. Formally, we have to prove that for $q = 1, \dots, N$, the mapping ξ maps

- the internal template $\Phi_1(q)$ to the template of $F_1(\tau_{\Phi_1}(q))$, and
- the internal template $\Phi_2(q)$ to the template of $F_1(\tau_{\Phi_2}(q))$, and
- the internal template $\Phi_3(q)$ to the template of $F_2(\tau_{\Phi_3}(q))$.

For short we demonstrate only the last one, i.e., the matching to Φ_3 . Let us take the internal template Φ_3 from (38) and map it. The mapping is treated as the substitution

$$\begin{array}{ccc} \text{internal-template } \Phi(q) == & z[q-1], & y[q-1] & \longrightarrow & z[q] \\ & \xi_z \downarrow & \xi_y \downarrow & & \xi_z \downarrow \\ & \underbrace{W[q-1, q-1], \quad W[q-2, q]} & \longrightarrow & & W[q, q] \\ & & & & \parallel \\ & & & & \text{template F2 with } i_1=q \text{ and } i_2=q \end{array}$$

And the last term is exactly $F2(\tau_{\Phi_3}(q))$. Q. e. d.

In the constructive recursive solution we need to define a reindexing of (39) into the indices of the array W , the array where the computation should be embedded, as shown in Fig. 8. This is given by $\text{ind1} : P \rightarrow \{-2, -1, 0, \dots\}$ and $\text{ind2} : P \rightarrow \{0, 1, \dots\}$.

$$\text{ind1}(\langle t, q \rangle) = q - 3 + t,$$

$$\text{ind2}(\langle t, q \rangle) = q + 3 - t,$$

and the function f1f2 on the graph corresponding to $F1$ and $F2$ (definitions (41) and (42)) is simply

$$\begin{aligned} \text{f1f2}(p) = & \text{if } \text{canR}(p,1) \ \& \ \text{canR}(p,2) \rightarrow \gamma_1(\text{f1f2}(r(p,1)), \\ & \text{f1f2}(r(p,2)) \) \\ & | \text{canR}(p,3) \ \& \ \text{canR}(p,4) \rightarrow \gamma_1(\text{f1f2}(r(p,3)), \\ & \text{f1f2}(r(p,4)) \) \\ & | \text{canR}(p,5) \ \& \ \text{canR}(p,6) \rightarrow \gamma_2(\text{f1f2}(r(p,5)), \\ & \text{f1f2}(r(p,6)) \) \text{fi} \end{aligned}$$

with input–output specification

$$\begin{aligned} \text{f1f2}(1,0) = \xi_0, \quad \text{f1f2}(2,0) = v_0, \quad \text{f1f2}(3,0) = \zeta_0 \longrightarrow \\ \text{f1f2}(t,q), \quad t=1..3, \quad q=1..N \end{aligned}$$

Note how we do the cases on the canR functions since xpoint etc. are not defined in this view of the graph. The input/output sets relate to the points of the dda , not of the embedding into W , thus the same sets of input/output points is needed as in (40).

6. Summary. We have presented two approaches, structural blanks and constructive recursion, and shown how they may be applied for the transcription of generalized recurrence relations to computer programs.

The structural blanks approach extends a traditional imperative programming language with constructs for defining explicitly the dependency pattern of a recurrence. The program to compute the recurrence is defined as a collection of global arrays and several program components: one for each equation of the recurrence (6), and a scheduler for the entire computation. These components

may be reused, and especially the scheduler may be applied on many different recurrence relations. Since the notation used is based on well known programming languages, it should be fairly easy to start using it for a practitioner in a field where recurrences are used. In SB the time axis is explicit. This is because a data dependency graph is explicitly represented in computer memory. This explicit representation allows the usage of matrix mathematics in affine graph transformations. The whole array representing the nodes of explicit data dependency graph is viewed as the output. The SB approach provides an architecture of software packages in the numerically oriented domain.

The constructive recursive approach is a functional programming language where the structure of the directed graph implicitly defined by a recursive expression is made explicit. In most functional languages, such as Haskell and Standard ML, the graph is a tree. Using memoization, nodes of this tree may be merged, but the graph can only be traversed in the direction that the functional expression defines it. In the CR approach the arcs can be traversed both in this and in the opposite direction. The latter traversal scheme translates into an efficient loop program to compute the recurrence. The graph may be defined with an assignment of the nodes to the space-time of a parallel computer. The result is then a data parallel program distributed on the processors of the parallel computer. In para-functional programming (Hudak, 1991) it is possible to schedule computations explicitly on a parallel computer, but the underlying graph will be a tree as given by the standard semantics of the programming language, preventing the efficiency obtained here.

Even though this presentation has focused on recurrence relations, the programming techniques presented are not restricted to recurrences in the classical form (1). Most problems with a repetitive or recursive structure can be expressed with the notation presented. Future work includes demonstrating these techniques on a broader set of examples, defining the underlying mathematics of the approaches, and building tools to facilitate the practical use of these approaches.

REFERENCES

- Chen, M., Y.Choo and J.Li (1991). Crystal: theory and pragmatics of generating efficient parallel code. In B.K. Szymanski (Ed.), *Parallel Functional Languages and*

- Compilers*. pp. 255–308.
- Cooley, J.W, and J.W.Tukey (1965). An algorithm for the machine computation of complex Fourier series. In *Mathematics of Computation*, Vol. 19. pp. 297–301.
- Čyras (1983). Loop synthesis over data structures in program packages. In *Computer Programming*, Vol. 7. Institute of Mathematics and Cybernetics, Vilnius. pp. 27–50 (in Russian).
- Greshnev, S.N., E.Z.Lyubimskii and V.A.Chiras (1985). Synthesis of programs on data structures. *Programming and Computer Software*, 11(5), 282–291, 1986. Translated from *Programmirovaniye* (in Russian), No. 5, 44–54, 1985.
- Čyras, V. (1986). Loop program synthesis in the system that separates functional modules from data structure traversing modules. *Lietuvos Matematikos Rinkiny*s, 26(4), 636–655 (in Russian).
- Čyras, V. (1988). Loop program synthesis using array traversing modules. Forschungsberichte Künstliche Intelligenz, Report FKI-93-88, Technische Universität München, Institut für Informatik, 25 p. Also In S. Meldal and M. Haveraaen (Eds.), *Proceedings of the 4th Nordic Workshop on Program Correctness*. Bergen, Norway. University of Bergen, Reports in Informatics, No. 78, 1993, pp. 97–109.
- Čyras, V., and M.Haveraaen (1994). Programming with data dependencies: a comparison of two approaches. In U.H. Engberg, K.G. Larsen and P.D. Mosses (Eds.), *Proceedings of the 6th Nordic Workshop on Programming Theory*. Aarhus, Denmark. BRICS Notes Series, NS-94-6, University of Aarhus, pp. 112–126.
- Haveraaen, M. (1990). Distributing programs on different parallel architectures. In *Proc. of the 1990 International Conference on Parallel Processing*, ICPP, Vol. II. Software pp. 288–289.
- Haveraaen, M. (1993). How to create parallel programs without knowing it. In S. Meldal and M. Haveraaen (Eds.), *Proceedings of the 4th Nordic Workshop on Program Correctness*. Bergen, Norway. University of Bergen, Reports in Informatics No. 78, pp. 165–176.
- Harf, M. (1994). Structural synthesis of programs using regular data structures. In U.H. Engberg, K.G. Larsen and P.D. Mosses (Eds.), *Proceedings of the 6th Nordic Workshop on Programming Theory*. Aarhus, Denmark. BRICS Notes Series, NS-94-6, University of Aarhus, pp. 112–126.
- Hudak, P. (1991). Para-functional Programming in Haskell. In B.K. Szymanski (Ed.), *Parallel Functional Languages and Compilers*. Addison Wesley, New York. pp. 159–196.
- Karp, R.M., R.E.Miller, S.Winograd (1967). The organization of computations for uniform recurrence equations. *Journal of the ACM*, 14(3), 563–590.
- Lapidus, L., and G.F.Pinder (1982). *Numerical Solution of Partial Differential Equations in Science and Engineering*. John Wiley, New York. 677 pp.
- Lundevik, R. (1994). *Translation from Standard ML to Sapphire*. Thesis, Department of Informatics, University of Bergen, Norway. 104 pp. (in Norwegian).

- Lyubimskii, E.Z. (1960). Issues of automatic programming. *Vestnik Akademii Nauk SSSR*, **8**, 47–55 (in Russian).
- Tucker, J.V., and J.I. Zucker (1988). *Program Correctness over Abstract Data Types, with Error-State Semantics*. North-Holland, Amsterdam, 212 pp.
- Tyugu, E. (1987). *Knowledge-Based Programming*. Turing Institute with Addison-Wesley, Wokingham.
- Zadykhailo, I.B. (1963). The organization of a cyclical computing process using a parametric representation of special form. *U.S.S.R. Computational Mathematics and Mathematical Physics*, **3**(2), 442–468. Translated from *Zhurnal vychislitel'noi matematiki i matematicheskoi fiziki*, **3**(2), 337–357 (in Russian).

Received November 1995

V. Čyras is a senior lecturer in computer science at the Vilnius University and a researcher at the Institute of Mathematics and Informatics, Vilnius. In 1979 he graduated from the Vilnius University. In 1985 he received the Degree of Doctor of sciences of physics and mathematics from the M.V. Lomonosov Moscow State University. Current research interests include theoretical computer science and loop program synthesis.

M. Haveraaen is a lecturer and researcher at the University of Bergen, Norway. In 1983 he obtained the Degree of Doctor at the University of Bergen. Current research interests include theoretical and practical programming, specification and modularization of programs.

**MODULINIS REKURENTINIŲ SANTYKIŲ PROGRAMAVIMAS:
DVIEJŲ POŽIŪRIŲ PALYGINIMAS****Vytautas ČYRAS, Magne HAVERAAEN**

Pristatomi ir palyginami du požiūriai į programas, operuojančias rekurentiniais santykiais. Tyrimo objektas yra programų modulių neprocedūrinis aprašymas. Siūlomi formalūs aparatai programų specifikacijų pavaizdavimui. *Struktūrizuotų paruošų* metodas akcentuoja ciklinių programų daugkartinio panaudojimo galimybę. Siūloma programų paketų architektūra skaičiavimo matematikos probleminėms sritims. *Konstruktivosios rekursijos* metodas grindžiamas rekursyvių santykių grafe formalizavimu.