


• 14000/1453



**An approach to correctness
of data parallel algorithms**

Joaquim Gabarró
Ricard Gavaldá

Report LSI-91-19



An approach to correctness of data parallel algorithms

Joaquim Gabarró * Ricard Gavaldà *

Dep. de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya
Pau Gargallo 5, 08028 Barcelona
Spain
gabarro@lsi.upc.es gavalda@lsi.upc.es

Abstract: The design of data parallel algorithms for fine-grained SIMD machines is a fundamental domain in today computer science. High standards in the specification and resolution of problems have been achieved in the sequential case. It seems reasonable to apply the same level of quality to data parallel programs.

It appears that most of the data parallel problems can be specified in terms of post and preconditions. These conditions characterize the overall state of the fine-grained processors in the initial and final states. In this paper:

- We present an axiomatic system to prove correctness of data parallel algorithms on fine-grained SIMD machines.
- We specify some data parallel problems like tree sum, radix sorting, and dynamic memory allocation.
- With this set of axioms we prove the correctness of programs solving the above problems.

It seems that the framework to deal with data parallel problems is quite different from the other one dealing with problems of parallelism with multiple threads of control, like those solvable in CSP.

Keywords: SIMD machines, programming languages, axiomatic semantics, stepwise refinement.

* Research supported by the ESPRIT II Basic Research Actions Program of the EC under contract No. 3075 (project ALCOM).

1. Introduction

The study of PRAM algorithms implemented in fine-grained machines is a fundamental domain in today computer science [GR88, KR90]. A lot of effort has been done to find optimal and efficient algorithms to solve different problems. However much less effort has been done to have a good programming methodology. This can be dangerous if we look back at what happened in the sequential case [Yo82].

Building the Connection Machine, Hillis [Hi85] seems to have fixed a new paradigm in programming fine-grained machines. On the other hand, active data structures [HS86] seem to be also an important and unexplored tool. In order to have a clear, usable, and complete program methodology for PRAM algorithms it seems unavoidable to:

- 1) Consider what kind of high-level instructions are needed to construct high-level programs on fine-grained machines.
- 2) Give an axiomatic semantics for these instructions.
- 3) Develop the notion of data structure adequate to this kind of machines.

This paper intends to be a first step towards points 1) and 2). With respect to point 1), let us be more precise about our goals: We want to find a set of instructions that can express most of the SIMD algorithms described, for instance, in [GR88, KR90]. Furthermore, we want to define them with no reference to the PRAM model or any other low-level model. More important, let us make clear which *are not* our goals:

We do not claim that our instructions are sufficient for general-purpose programming. We do not discuss input/output, data typing and other features that any true programming language should have.

We do not claim either that the instructions can be efficiently implemented with current technology. For instance, our high-level view ignores crucial issues such as:

- Mapping between data structures and processors: Similarly, when describing formally high-level sequential languages, it is not reasonable to explain how global and local variables, procedure parameters, and dynamic memory have to be mapped into a machine's memory, that is usually a linear array.
- Communication between processors, including network topology, broadcasting of values, routing... In our high-level view it is not even possible to say if processors communicate through shared memory or by exchanging messages. In principle, every data item is visible to every processor.
- Scheduling of tasks to individual processors: A typical consideration that we want to exclude from our semantics is whether the parallel assignment

$$x[k + 1] := x[k] + x[k - 1]$$

is executed by the processor "numbered" k , $k + 1$, or $k - 1$, be this number physical or virtual.

Of course, ignoring these items is unrealistic with present machines and compiler techniques. However, by comparison with the history of sequential programming, we hope

that these problems will be less and less important in the far future. Otherwise, there is little hope of ever having high-level, machine-independent parallel software.

2. Preliminaries

In order to present the semantics of a SIMD language we start with an informal overview of it. It is very close to that used in [HS86]. We describe an abstract parallel machine that supports it and how each language construct is implemented. We do this for clarity; in the following sections, we will give a semantics not tied to any machine model.

We want to describe programs written for a machine with a large number of processor+memory cells. Every processor has an internal state bit, called the *context flag*. All processors receive a common instruction stream from a host and examine it synchronously. For each of them, the context flag determines whether the incoming instruction is actually executed or ignored. There is also a set of instructions that are executed unconditionally, i.e., even if the context flag is off. Such instructions are needed, at least, for setting and clearing context flags. We call a processor *selected* or *active* at a given moment if its context flag is on. The *context* at that moment is the set of selected processors. Note again that the next instruction in the stream will be executed identically by all processors in the context, while the rest of processors will remain idle.

Two classes of variables are considered: *Global* variables are not attached to any particular processor. One can think of them as located in the host's memory. *Local* variables are replicated in each processor. We use array notation to define and use local variables. A declaration such as

$$x : \text{array } [1..N] \text{ of proc of integer}$$

creates an integer variable belonging to each processor k , that can be referenced as $x[k]$. Processors can exchange values of local variables (in a way not specified).

Language constructs

- The main parallel construct is the for-all, denoted as

$$\text{for all } k \text{ do in parallel } S(k) \text{ end}$$

It indicates that statement $S(k)$ will be executed by all active processors, each with a different value of variable k . We will also use the construction

$$\text{for all } k : E(k) \text{ do in parallel } S(k) \text{ end}$$

to indicate that only active processors k such that condition $E(k)$ is true will actually execute $S(k)$. This statement is used for notational convenience. It can be viewed as an unrestricted for-all enclosing an if statement (described below).

- Inside a for-all we can have parallel assignments of the form $x[F_1(k)] := F_2(k)$ giving statements like

$$\text{for all } k : E(k) \text{ do in parallel } x[F_1(k)] := F_2(k) \text{ end}$$

This statement instructs each selected processor k to compute $F_2(k)$ and update variable x of processor $F_1(k)$. This requires some communication if $F_1(k) \neq k$. Note that expression $F_2(k)$ can be in fact independent of k (for example, it can be a global variable).

- We also have a conditional statement of the form

if $E(k)$ then $S(k)$ end

It indicates that all currently selected processors evaluate $E(k)$. Those that evaluate to true will execute $S(k)$ and all others are deselected during that time.

- But remember that the instruction stream is still sent to all processors in the machine. This is necessary because S might contain some unconditional instructions that require activating all of them. Instructions that are executed unconditionally are indicated as

unconditionally S end

Here all processors are temporarily selected for execution of S and after termination the original context is restored.

- The while statement has some important differences in relation to the sequential one. A statement like

while $E(k)$ do $S(k)$ end

should be interpreted as following: at each iteration, a processor for which $E(k)$ is false becomes deselected. The loop terminates when all processors have become deselected, and the original context is then restored.

- We use two functions called count and enumerate as primitives of our language. If c is a global variable and X is a local variable, then the statement

$c := \text{count}$

will place in c the number of elements in the current context, and

$x[k] := \text{enumerate}$

will place in variable x of each selected processor its relative position in the context. That is, the selected processor with lowest value of k will receive 0, the immediate next will receive 1, etc. [HS86] and [Ble86] argue that these operations can be regarded as primitive because they can be implemented very efficiently on real machines and allow to write many algorithms concisely. We will support this opinion by showing that they also have a very simple and elegant definition in our semantics.

- Finally, we inherit all constructs common in sequential programming, such as sequential composition (denoted $;$) and sequential for loop.

With all these constructs we can code algorithms easily. This is specially true when have right to consider procedures and functions. For example, Figure 1 presents an algorithm

```

procedure Fast_Binary_Sum
  (a, b: array [N - 1..0] of proc of bit): array [N..0] of proc of bit;

import
  Tree_And (a: array [N - 1..0] of proc of bit): bit;
  Tree_Or (b: array [N - 1..0] of proc of bit): bit;

var
  p: array [N - 1..0] of proc of bit; { carry propagate }
  g: array [N - 1..0] of proc of bit; { carry generate }
  c: array [N - 1..0] of proc of bit; { carry }
  x: array [N - 1..0, N - 1..0] of proc of bit;

begin
  { compute the carry }
  for all i : N > i ≥ 0 do in parallel p[i] := a[i] + b[i]; g[i] := a[i] * b[i] end;
  for all i, j : N > i > j ≥ 0 do in parallel x[i, j] := Tree_And (p[i..j + 1]) * g[j] end;
  c[0] := g[0];
  for all i : N > i ≥ 0 do in parallel c[i] := g[i] + Tree_Or (x[i..0]) end;
  { compute the sum }
  Fast_Binary_Sum[0] := a[0] + b[0];
  for all i : 0 < i < N do in parallel Fast_Binary_Sum[i] := a[i] ⊕ b[i] ⊕ c[i - 1] end;
  Fast_Binary_Sum[N] := c[N - 1]
end Fast_Binary_Sum

```

Figure 1. Fast binary sum.

implementing a fast adder. In this paper we deal with procedures informally. However, their axioms can be translated directly from those of sequential languages.

Notation

We introduce some notations to deal with finite sets of integers. Let A be a subset of \mathbb{N} such that $A = \{a_0, a_1, a_2, \dots\}$. We assume A ordered, this means $a_0 < a_1 < a_2 < \dots$. Let $\#A$ be the cardinal of A . Given $k \geq 0$, we define the following subsets of A :

$$A_{<k} = \{a \in A \mid a < k\}$$

$$first_k A = \begin{cases} \{a_0, \dots, a_{k-1}\} & \text{if } k < \#A \\ A & \text{otherwise.} \end{cases}$$

Given a mapping $y : A \rightarrow A'$ and two sets $B \subseteq A$ and $B' \subseteq A'$ we denote as $B \xrightarrow{y} B'$ the restriction of this mapping to the sets B and B' . Note that given two sets A and B the following conditions are equivalent

- $first_{\min(\#A, \#B)} A \xrightarrow{y} first_{\min(\#A, \#B)} B$ is a bijection.
- $first_{\#B} A \xrightarrow{y} first_{\#A} B$ is a bijection.

3. Semantics

In this section we present an axiomatic semantics, based on the notion of state and pre- and post- conditions for the language described. This extends the work done by A. Hoare [Ho69] and E. Dijkstra [Di75] for sequential programming.

To see why this type of semantics is natural in this framework, recall that we assume a single instruction stream that is executed in perfect synchrony by all processors. This gives us two immediate advantages at the time of designing and proving programs correct:

- First, there is no need to synchronize processors explicitly. In particular, there is no danger of deadlock: the state of each processor (selected or not) is internal to it and can only be changed by the instruction stream, not by any other processor's signal. Thus, we can prove termination of these programs using the same techniques as in sequential programming.
- Second, programs are completely deterministic. There are no invisible state changes or nondeterministic choices. As a consequence, there is a precise definition of the "state" of a program at a given moment.

For these reasons, it is possible to specify the behaviour of one of these programs *completely* in terms of its initial and final state. Note that such specification is not possible for concurrent MIMD programs, because one must also specify how and when each program is ready to communicate.

States, processors, and assertions

In the sequential case the state of a computation is determined by the value of the variables. The relationship between these values is described by a predicate P . To describe the state of a SIMD algorithm we need to take care of two different facts.

- As in the sequential case we need to describe the relationship between the values of the local and global variables. This can be done by a predicate P .
- Recall that each processor can be active or inactive. Any description of the state of a program must take care of this, because the same program, started with different active processors, will have different results. We assume that processors are denoted by natural numbers; our axioms can be easily generalized if any other countable type is used. The set of active processors will be described by a set A .

Summarizing these two aspects, our assertions about states will be a pair of the form $\{A; P\}$ where A is a set of natural numbers and P is a predicate describing the relationship between the values of all program variables.

One can argue that the set of active processors A can be regarded as just another variable and handled by the predicate P . We will not adopt this view for two reasons:

- First, our axioms are much simpler if the set A is displayed clearly in each assertion and not hidden inside P .
- Second, processors should be selected or deselected on entry or exit of the block-structured statements only. In this sense, changing context flags explicitly is the

parallel analog of the sequential `goto` statement. In particular, no program expression or assignment should use the context flags as variables.

Therefore, our specification of a program S will have the form

$$\{ A; P \} S \{ A; Q \}$$

meaning that S transforms predicate P into predicate Q if it is started with processors in A active. Note that with our program constructs the active processors are always the same before and after S . However, they can change during the execution of S . For example, S can contain `if` statements that deselect some processors temporarily.

Axioms

Global Assignment. The axiom for assignment to a global variable is identical to that of the sequential case. If c is a global variable and E a global expression, then:

$$\{ A; P \} c := E \{ A; Q \}$$

iff P implies $\text{domain}(E) \wedge Q_E^c$. Here, Q_E^c denotes the predicate that results of the *textual substitution* of c by E in Q ; see [Gr81] for details on textual substitution. In the following axioms, conditions such as $\text{domain}(E)$ will not appear explicitly in order to lighten the notation.

Local Assignment. We present two axioms for assignment to local variables that correspond to assignment in Owner Write and Exclusive Write modes. The first axiom is a particular case of the second. We include both because parallel programs usually contain many assignments in Owner Write and we can then use a considerably simpler axiom. Let x be a local variable. We consider that x contains a function of finite domain, so each local instance of x , $x[k]$, is in fact the application of the function in x to argument k . Thus, we can view assignments to components of x as a change of the function x .

- *Owner Write.* Let $E(k)$ be an expression that may depend on k . Let us give an axiom to formalize the idea “every active processor k does some local computation (represented by $E(k)$) and memorizes the result on its local variable $x[k]$ ”. We call this kind of assignment Owner Write assignment. If we consider the most inner for-all statement enclosing an Owner Write assignment $x[k] := E(k)$ we have the following axiom.

$$\text{for all } k \text{ do in parallel } \dots \{ A; P \} x[k] := E(k) \{ A; Q \} \dots \text{end}$$

iff P implies $Q_{(x,A,E)}^x$ and (x, A, E) is the array (or function) defined as

$$(x, A, E)[k] = \begin{cases} E(k) & \text{if } k \in A \\ x[k] & \text{otherwise} \end{cases}$$

Remark that (x, A, E) describes the global modification of x and $Q_{(x,A,E)}^x$ describes the global modification of the state described by Q .

• **Exclusive Write.** The axiom for Exclusive Write assignment is slightly more complicated because it must deal with indirection between processors as in $x[F(k)] := E(k)$. If $F(k)$ and $E(k)$ are expressions that may depend on k and, as we want exclusive writes, $F(k)$ must be injective on the domain of k . After executing the assignment, we have for all l and k that $x[l] = E(k)$, if $F(k) = l$ and k is active. The processor k can be identified as $k = F^{-1}(l)$ where $F^{-1}(j) = \{k : F(k) = j\}$, so the expression $E(F^{-1}(l))$ makes sense here. Then, the assignment $x[F(k)] := E(k)$ is equivalent to $x[l] = E(F^{-1}(l))$. As k and l are dummy variables we get the following axiom

for all k do in parallel $\dots \{ A; P \} x[F(k)] := E(k) \{ A; Q \} \dots$ **end**

iff P implies

$$(F(k) \text{ as a function of } k \text{ is injective}) \wedge (Q_{(x,A,E,F)}^z).$$

Here, (x, A, E, F) is the function that is equal to x in each component that is not altered, and that has value $E(k)$ at the component altered by processor k :

$$(x, A, E, F)[k] = \begin{cases} E(F^{-1}(k)) & \text{if } F^{-1}(k) \in A \\ x[k] & \text{otherwise} \end{cases}$$

We can give a slightly stronger axiom noticing the following: to have Exclusive Write, it is enough to prove that F is injective on the domain A . Then, $F^{-1}(k) \cap A$ has at most one element even if $F^{-1}(k)$ is larger, and the component k of X receives the value $E(F^{-1}(k) \cap A)$ if $F^{-1}(k) \cap A \neq \emptyset$.

Counting active processors. This instruction assigns to a global variable c the number of currently active processors. Thus, we have an axiom similar to that of global assignment. If $\#A$ denotes the cardinality of set A , then

$$\{ A; P \} c := \text{count} \{ A; Q \}$$

iff P implies $Q_{\#A}^c$.

Enumerating active processors. This instruction enumerates from 0 onwards the currently active processors. Each active processor receives its relative position into its own copy of a local variable. Thus, the axiom for **enumerate** is similar to local assignment in Owner Write. For every index i and a set A we define $\#A_{<i} = \#\{j \in A \mid j < i\}$. Then $\#A_{<i}$ is 0 if i is the minimum of set A , 1 if it is the second element in A , and so on. Then the axiom is:

$$\{ A; P \} x[k] := \text{enumerate} \{ A; Q \}$$

iff P implies $Q_{(x,A,\#A_{<})}^z$. In this predicate, $\#A_{<}$ is interpreted as an expression that depends on one argument.

Sequential composition. The axiom for sequential composition is identical to that of the sequential case.

$$\{ A; P \} S_1; S_2 \{ A; Q \}$$

iff there exists a predicate R such that

$$\{ A; P \} S_1 \{ A; R \} \wedge \{ A; R \} S_2 \{ A; Q \}$$

Conditional statement. The following holds

$$\{ A; P \} \text{ if } E(k) \text{ then } S(k) \text{ end } \{ A; Q \}$$

iff

$$\{ \{ k \mid E(k) \wedge k \in A \}; P \} S(k) \{ \{ k \mid E(k) \wedge k \in A \}; Q \}$$

When the expression $E(k)$ does not depend on k the axiom for sequential conditional statement can be derived.

Unconditional statement. The following holds

$$\{ A; P \} \text{ unconditionally } S \text{ end } \{ A; Q \}$$

iff

$$\{ \mathcal{D}; P \} S \{ \mathcal{D}; Q \}$$

where \mathcal{D} denotes the domain of processors of the program.

For-all statement. Let us consider how a for-all statement affects the state of a computation. If the following holds

$$\text{for all } k \text{ do in parallel } \{ A; P \} S(k) \{ A; Q \} \text{ end}$$

it also holds

$$\{ A; P \} \text{ for all } k \text{ do in parallel } S(k) \text{ end } \{ A; Q \}$$

It seems that, in general, the for-all instruction can be considered redundant. It only declares the variable that will be used to index parallel constructs.

From another point of view the statement **for all k do in parallel $S(k)$ end** is often intuitively interpreted as “every active processor k does some easy local computation. This computation does not interfere with the computation of its neighbours”. In these cases, the postcondition usually has the form $\forall k : k \in A \Rightarrow Q(k)$, and we can give a useful rule.

Given u and v in A such that $u \neq v$, the statements $S(u)$ and $S(v)$ *do not interact* if they act over different sets of variables, this means:

- $S(u)$ cannot read from any variable written by $S(v)$, and reciprocally,
- $S(u)$ and $S(v)$ cannot write into the same variable, but
- $S(u)$ and $S(v)$ can read from the same variable.

Then, the following holds

$$\{ A; \forall k : k \in A \Rightarrow P(k) \} \text{ for all } k \text{ do in parallel } S(k) \text{ end } \{ A; \forall k : k \in A \Rightarrow Q(k) \}$$

if

$$\begin{aligned} & \forall k : k \in A \Rightarrow \{A; P(k)\} S(k) \{A; Q(k)\} \\ & \text{and} \\ & \forall u, v : u \neq v \wedge u, v \in A \Rightarrow S(u) \text{ and } S(v) \text{ do not interact} \end{aligned}$$

We would like to remark that, in some cases, programming with for-all statements having only non-interacting $S(k)$ can be a serious handicap. In these cases the previous rule is useless.

To see what kind of computation is possible with no interaction, suppose that a program has the following form, where the different $S(k)$ do not interact:

for all $k \in \{1, \dots, l\}$ do in parallel $S(k)$ end

Then, any arbitrary interleaving of the statements $S(1), \dots, S(l)$ will give us a correct computation. In CSP notation we could characterize the meaning of the for-all as:

$$S(1) \parallel S(2) \parallel \dots \parallel S(l)$$

In particular the sequential computation $S(1); S(2); \dots; S(l)$ gives us a correct result.

4. Examples

As examples of use of our axioms, we present now three programs of increasing complexity:

- 1) Computing the sum of the elements of an array (Figure 2).
- 2) Radix sorting an array of integers (Figure 3).
- 3) Dynamic allocation of processors and memory (Figure 4).

The three programs are taken directly from [HS86]. We have only added variable declarations and split long programs into procedures for clarity.

Example 1: Tree Sum of N numbers

Program `Tree.Sum` in Figure 2 computes the sum of all components of array X in time $O(\log N)$ by the well known binary-tree technique. This program illustrates the use of Owner Write assignment within for-all's, and is an example of small program where considering mainly "local states" is still feasible. The set of active processors is $A = [0 \dots N]$. It can be specified as:

$$\begin{aligned} & \{A; (x[0], \dots, x[N-1]) = (X_0, \dots, X_{N-1}) \} \\ & \text{Tree.Sum} \\ & \{A; x[N-1] = \sum_{i=0}^{N-1} X_i \} \end{aligned}$$

To specify the procedure `Partial.Sums` it is interesting to consider the set

- $active(j) = \{k \in A : (k+1) \bmod 2^j = 0\}$

such that

Recalling that we view restricted for-all's as a composition of one for-all and one if statement, we must show that:

for all $k : (k + 1) \bmod 2^j = 0$ do in parallel
 $\{ \mathit{active}(j); \mathit{SUMS}(j - 1) \}$
 $x[k] := x[k - 2^{j-1}] + x[k]$
 $\{ \mathit{active}(j); \mathit{SUMS}(j) \}$
end

We call $e(k)$ the expression $x[k - 2^{j-1}] + x[k]$ that depends on k . Then

$$(x, \mathit{active}(j), e)[k] = \begin{cases} x[k - 2^{j-1}] + x[k] & \text{if } k \in \mathit{active}(j) \\ x[k] & \text{otherwise.} \end{cases}$$

By the axiom of owner write assignment we must show that

$$\mathit{SUMS}(j - 1) \implies \mathit{SUMS}(j)_{(x, \mathit{active}(j), e)}^x$$

As the predicate $\mathit{SUMS}(j)_{(x, \mathit{active}(j), e)}^x$ is equivalent to

$$k \in \mathit{active}(j) \implies x[k - 2^{j-1}] + x[k] = \sum_{i=k-2^j+1}^k X_i$$

we must prove

$$\mathit{SUMS}(j - 1) \implies (k \in \mathit{active}(j) \implies x[k - 2^{j-1}] + x[k] = \sum_{i=k-2^j+1}^k X_i).$$

Let $k \in \mathit{active}(j)$. Then $k \in \mathit{active}(j - 1)$ and $k - 2^{j-1} \in \mathit{active}(j - 1)$. Assuming $\mathit{SUM}(j - 1)$ we have

$$k \in \mathit{active}(j - 1) \implies x[k] = \sum_{i=k-2^{j-1}+1}^k X_i$$

$$k - 2^{j-1} \in \mathit{active}(j - 1) \implies x[k - 2^{j-1}] = \sum_{i=k-2^j+1}^{k-2^{j-1}} X_i$$

and this implies $\mathit{SUMS}(j)_{(x, \mathit{active}(j), e)}^x$. The proof is done. \square

```

procedure Tree_Sum ( $x$  : array [0.. $N - 1$ ] of integer): integer;
var  $j$ : integer;
procedure Partial_Sums ( $i$ : integer),
  for all  $k$  :  $(k + 1) \bmod 2^i = 0$  do in parallel
     $x[k] := x[k - 2^{i-1}] + x[k]$ 
  end
end Partial_Sums;
begin
  for  $j := 1$  to  $\log_2 N$  do
    Partial_Sums( $j$ )
  end;
  Tree_Sum :=  $x[N - 1]$ 
end Tree_Sum

```

Figure 2. Sum of N numbers.

- $active(j) \subseteq active(j - 1)$.
- $k \in active(j) \implies k \in active(j - 1) \wedge k + 2^{j-1} \in active(j - 1)$.

With this set let us define the invariant of the sequential for:

$$SUMS(j) \equiv (\forall k \in active(j) \implies x[k] = \sum_{i=k-2^j+1}^k X_i)$$

Let us prove the following:

Lemma 1: The procedure Partial_Sums verifies for $j > 0$:

$$\begin{aligned} & \{A; SUMS(j - 1)\} \\ & \text{Partial_Sums } (j) \\ & \{A; SUMS(j)\} \end{aligned}$$

Proof. Substituting the call Partial_Sums by its body we have:

$$\begin{aligned} & \{A; SUMS(j - 1)\} \\ & \text{for all } k : (k + 1) \bmod 2^j = 0 \text{ do in parallel} \\ & \quad x[k] := x[k - 2^{j-1}] + x[k] \\ & \text{end} \\ & \{A; SUMS(j)\} \end{aligned}$$

Lemma 2: The program `Tree_Sum` satisfies its specification.

Proof. As

$$\begin{aligned} SUMS(0) &\equiv ((x[0], \dots, x[N-1]) = (X_0, \dots, X_{N-1})) \\ SUMS(\log_2 N) &\equiv x[N-1] = \sum_{i=0}^{N-1} X_i \end{aligned}$$

we must prove

$$\begin{aligned} &\{A; SUMS(0)\} \\ &\text{Tree_Sum} \\ &\{A; SUMS(\log_2 N)\}. \end{aligned}$$

Substituting `Tree_Sum` by its body we must prove

$$\begin{aligned} &\{A; SUMS(0)\} \\ &\mathbf{for} \ j := 1 \ \mathbf{to} \ \log_2 N \ \mathbf{do} \\ &\quad \text{Partial_Sums}(j) \\ &\mathbf{end} \\ &\{A; SUMS(\log_2 N)\} \end{aligned}$$

But this is done by the previous lemma and the well known axiom for sequential for. \square

Example 2: Radix Sort

Figure 3 presents a parallel implementation of Radix Sort. This algorithm provides us with the occasion to deal formally with the axioms for **count**, **enumerate**, and **exclusive write**. Let us introduce the following notations:

If V is a vector, $PERM(V)$ denotes the set of vectors that are permutations of V .

Let A be the domain of array x , that is, the set $\{0 \dots N\}$.

A specification of `Radix_Sort` (in fact, of any sorting program) is then:

$$\begin{aligned} &\{A; (x = X)\} \\ &\quad \text{Radix_Sort} \\ &\{A; (x \text{ is ordered}) \wedge x \in PERM(X)\} \end{aligned}$$

To show that `Radix_Sort` satisfies this specification, we must express that, after the j th iteration, vector x is already sorted if only the j less significant bits of its elements are considered. So, let us denote with “ $x \bmod i$ ” the following vector

$$x \bmod i = [x[0] \bmod i, x[1] \bmod i, \dots, x[N-1] \bmod i].$$

And finally, define predicate $SORTED(j)$ to be

$$SORTED(j) = ((x \bmod 2^j) \text{ is ordered}) \wedge (x \in PERM(X)).$$

```

procedure Radix_Sort
  (var  $x$ : array [0 ..  $N - 1$ ] of proc of {0, ...,  $N - 1$ });
var  $y$ : array [0 ..  $N - 1$ ] of proc of {0, ...,  $N - 1$ };
   $c, j$ : integer;

procedure Enum_Zeros ( $i$  : integer);
  for all  $k$  : ( $x[k] \bmod 2^i$ ) <  $2^{i-1}$  do in parallel
     $c :=$  count;
     $y[k] :=$  enumerate
  end
end Enum_Zeros;

procedure Enum_Ones ( $i$  : integer);
  for all  $k$  : ( $x[k] \bmod 2^i$ )  $\geq 2^{i-1}$  do in parallel
     $y[k] :=$  enumerate +  $c$ 
  end
end Enum_Ones;

procedure Exchange;
  for all  $k$  do in parallel
     $x[y[k]] := x[k]$ 
  end
end Exchange;

begin
  for  $j := 1$  to  $\log(\text{maxint})$  do
    Enum_Zeros( $j$ );
    Enum_Ones( $j$ );
    Exchange
  end
end Radix_Sort

```

Figure 3. Radix Sort.

It is easy to see that the precondition to Radix_Sort implies $SORTED(0)$, and that

$$\begin{aligned}
 SORTED(\log(\text{maxint})) &= ((x \bmod \text{maxint}) \text{ is ordered}) \wedge (x \in PERM(X)) \\
 &= (x \text{ is sorted}) \wedge (x \in PERM(X)).
 \end{aligned}$$

so it is exactly the postcondition. Therefore, using the axiom of the sequential for, we must show that for $j > 0$ we have

$$\{ A; SORTED(j - 1) \}$$

Enum_Zeros(j);
 Enum_Ones(j);
 Exchange
 { A ; *SORTED*(j) }

We do this in the following lemmas. Informally, Enum_Zeros(j) and Enum_Ones(j) divide the elements of x into two subsets and enumerate in y each of them. Let us give a name to the elements that each procedure enumerates at iteration j .

- $small(j) = \{ k \in A : x[k] \bmod 2^j < 2^{j-1} \}$.
- $big(j) = A \setminus small(j)$.

We have the following easy properties:

- $[0 \dots N] = small(j) \cup big(j)$.
- $\emptyset = small(j) \cap big(j)$.
- $k \in small(j), k' \in big(j) \implies x[k] < 2^{j-1} \leq x[k']$.
- $k \in small(j) \implies x[k] \bmod 2^j = x[k] \bmod 2^{j-1}$.
- $k \in big(j) \implies x[k] \bmod 2^j = 2^{j-1} + x[k] \bmod 2^{j-1}$.

Procedure Enum_Zeros(j) deals with the elements of $small(j)$. It counts the elements in $small(j)$ and enumerates them. To deal with the number of elements we introduce the predicate

$$COUNTED_{small}(j, c) \equiv (c = \#small(j)).$$

To deal with the enumeration we consider a mapping $small(j) \xrightarrow{y} [0 \dots \#small(j)]$ which is the (unique) increasing bijection defined as $y[k] = \#small(j)_{<k}$. We describe y by the predicate

$$\begin{aligned}
 ENUMERATED_{small}(j, y) \equiv & \\
 & (small(j) \xrightarrow{y} [0 \dots \#small(j)] \text{ is the increasing bijection}) \equiv \\
 & (\forall k \in small(j) : y[k] = \#small(j)_{<k}).
 \end{aligned}$$

We can now specify procedure Enum_Zeros(j).

Lemma 3: Procedure Enum_Zeros(j) satisfies the following specification:

$$\begin{aligned}
 & \{ A; \text{true} \} \\
 & \text{Enum_Zeros}(j) \\
 & \{ A; COUNTED_{small}(j, c) \wedge ENUMERATED_{small}(j, y) \}
 \end{aligned}$$

Proof. Substituting $\text{Enum_Zeros}(j)$ by its definition we obtain:

```

{ A; true }
for all k : (x[k] mod 2j) < 2j-1 do in parallel
  c := count;
  y[k] := enumerate
end
{ A; COUNTEDsmall(j, c) ∧ ENUMERATEDsmall(j, y) }

```

When we go inside the **for all** statement the set of active processors is $\text{small}(j)$. Then we have

```

for all k : (x[k] mod 2j) < 2j-1 do in parallel
  { small(j); true }
  c := count;
  y[k] := enumerate
  { small(j); COUNTEDsmall(j, c) ∧ ENUMERATEDsmall(j, y) }
end

```

To deal with the statement $y[k] := \text{enumerate}$ we need $(y, \text{small}(j), \# \text{small}(j) <)$ defining the mapping

$$(y, \text{small}(j), \# \text{small}(j) <)[k] = \begin{cases} \# \text{small}(j) <_k & \text{if } k \in \text{small}(j) \\ y[k] & \text{otherwise.} \end{cases}$$

As $\text{ENUMERATED}_{\text{small}}(j, y)_{(y, \text{small}(j), \# \text{small}(j) <)}^y \equiv \text{true}$, by the axiom of enumeration we get

```

{ small(j); true }
y[k] := enumerate
{ small(j); ENUMERATEDsmall(j, y) }

```

As $\text{COUNTED}_{\text{small}}(j, c)_{\# \text{small}(j)}^c \equiv \text{true}$, by the axiom of counting we get

$$\{ \text{small}(j); \text{true} \} c := \text{count} \{ \text{small}(j); \text{COUNTED}_{\text{small}}(j, c) \}.$$

and the lemma is done. \square

Procedure $\text{Enum_Ones}(j)$ enumerates the elements of $\text{big}(j)$ starting from the number $\# \text{small}(j)$. Rather than introducing a predicate giving us the enumeration of $\text{big}(j)$ we would like to enumerate all the elements of $[0 \dots N] = \text{small}(j) \cup \text{big}(j)$.

- The enumeration of $\text{small}(j)$ can be done as above. We consider the (unique) increasing bijection $\text{small}(j) \xrightarrow{y} [0 \dots \# \text{small}(j)]$ defined as $y[k] = \# \text{small}(j) <_k$.
- To enumerate $\text{big}(j)$ we start from $\# \text{small}(j)$. We take the (unique) increasing bijection $\text{big}(j) \xrightarrow{y} [\# \text{small}(j) \dots N]$ defined as $y[k] = \# \text{small}(j) + \# \text{big}(j) <_k$.

To describe this enumeration we introduce the predicate

$$\begin{aligned} \text{ENUMERATED}(j, y) \equiv & \\ & (\text{small}(j) \xrightarrow{y} [0 \dots \#\text{small}(j)]) \text{ is the increasing bijection} \\ & \wedge \text{big}(j) \xrightarrow{y} [\#\text{small}(j) \dots N] \text{ is the increasing bijection} \equiv \\ & (\forall k \in \text{small}(j) : y[k] = \#\text{small}(j)_{<k}) \wedge (\forall k \in \text{big}(j) : y[k] = \#\text{small}(j) + \#\text{big}(j)_{<k}). \end{aligned}$$

As

$$\begin{aligned} \text{ENUMERATED}(j, y) \stackrel{y}{\equiv} & \text{COUNTED}_{\text{small}(j), c + \#\text{big}(j)_{<}} \equiv \\ & \text{COUNTED}_{\text{small}(j), c} \wedge \text{ENUMERATED}_{\text{small}(j), y} \end{aligned}$$

by the same techniques as in the above lemma we obtain

Lemma 4: Procedure `Enum_Ones(j)` satisfies the following specification:

$$\begin{aligned} & \{ A; \text{COUNTED}_{\text{small}(j), c} \wedge \text{ENUMERATED}_{\text{small}(j), y} \} \\ & \quad \text{Enum_Ones}(j) \\ & \{ A; \text{ENUMERATED}(j, y) \} \end{aligned}$$

As `Enum_Zeros(j)` and `Enum_Ones(j)` do not modify the array x , the above lemmas give

Lemma 5: The sequential composition of `Enum_Zeros` and `Enum_Ones` satisfies:

$$\begin{aligned} & \{ A; \text{SORTED}(j - 1) \} \\ & \quad \text{Enum_Zeros}(j); \\ & \quad \text{Enum_Ones}(j) \\ & \{ A; \text{SORTED}(j - 1) \wedge \text{ENUMERATED}(j, y) \} \end{aligned}$$

To deal with procedure `Exchange` we consider the mapping y^{-1} where y is defined by $\text{ENUMERATED}(j, y)$. This mapping $[0 \dots N] \xrightarrow{y^{-1}} [0 \dots N]$ verifies:

- $[0 \dots \#\text{small}(j)] \xrightarrow{y^{-1}} \text{small}(j)$ is the increasing bijection.
- $[\#\text{small}(j) \dots N] \xrightarrow{y^{-1}} \text{big}(j)$ is the increasing bijection.

Let us specify and prove procedure `Exchange`.

Lemma 6: Procedure `Exchange` satisfies:

$$\begin{aligned} & \{ A; \text{SORTED}(j - 1) \wedge \text{ENUMERATED}(j, y) \} \\ & \quad \text{Exchange} \\ & \{ A; \text{SORTED}(j) \} \end{aligned}$$

Proof. Substituting the procedure call by its body we come to

$$\{ A; \text{SORTED}(j - 1) \wedge \text{ENUMERATED}(j, y) \}$$

$$\begin{aligned} x[y[k]] &:= x[k] \\ \{ A; \text{SORTED}(j) \} \end{aligned}$$

In order to deal with the statement $x[y[k]] := x[k]$ we need to consider the array (x, A, x, y) . We have

$$(x, A, x, y)[k] = x[y^{-1}[k]] = x \circ y^{-1}[k].$$

By the axiom of exclusive write we have

$$\{ A; \text{SORTED}_{(x, A, x, y)}^x(j) \} x[y[k]] := x[k] \{ A; \text{SORTED}(j) \}$$

As

$$\begin{aligned} \text{SORTED}_{(x, A, x, y)}^x(j) &\equiv \text{SORTED}_{x \circ y^{-1}}^x(j) \equiv \\ &((x \circ y^{-1} \bmod 2^j) \text{ is ordered}) \wedge (x \circ y^{-1} \in \text{PERM}(X)) \end{aligned}$$

we need to prove

$$\begin{aligned} \text{SORTED}(j-1) \wedge \text{ENUMERATED}(j, y) &\implies \\ ((x \circ y^{-1} \bmod 2^j) \text{ is ordered}) &\wedge (x \circ y^{-1} \in \text{PERM}(X)) \end{aligned}$$

First we prove

$$\text{SORTED}(j-1) \wedge \text{ENUMERATED}(j, y) \implies (x \circ y^{-1} \in \text{PERM}(X)).$$

As $\text{SORTED}(j-1)$ is assumed x is a permutation of X . $\text{ENUMERATED}(j, y)$ implies that y^{-1} is a bijection. Therefore $x \circ y^{-1}$ is a permutation of X .

Secondly we prove

$$\text{SORTED}(j-1) \wedge \text{ENUMERATED}(j, y) \implies ((x \circ y^{-1} \bmod 2^j) \text{ is ordered}).$$

We can rewrite $((x \circ y^{-1} \bmod 2^j) \text{ is ordered})$ as

$$(\forall k, k' \in A \wedge k < k' \implies x \circ y^{-1}[k] \bmod 2^j \leq x \circ y^{-1}[k'] \bmod 2^j).$$

There are three cases:

Case 1: Both k and k' are in $[0 \dots \#small(j))$. As we assume $\text{ENUMERATED}(j, y)$ the mapping $[0 \dots \#small(j)) \xrightarrow{y^{-1}} small(j)$ is the increasing bijection and

- $k < k' \implies y^{-1}[k] < y^{-1}[k']$.
- $k, k' \in [0 \dots \#small(j)) \implies y^{-1}[k], y^{-1}[k'] \in small(j)$.
- $y^{-1}[k] \in small(j) \implies (x[y^{-1}[k]] \bmod 2^{j-1} = x[y^{-1}[k]] \bmod 2^j)$ and similarly for k' .

As we assume $\text{SORTED}(j-1)$, then

$$y^{-1}[k] < y^{-1}[k'] \implies x \circ y^{-1}[k] \bmod 2^{j-1} \leq x \circ y^{-1}[k'] \bmod 2^{j-1}.$$

Putting together all these results we have

$$\begin{aligned} k < k' &\implies y^{-1}[k] < y^{-1}[k'] \implies x \circ y^{-1}[k] \bmod 2^{j-1} \leq x \circ y^{-1}[k'] \bmod 2^{j-1} \\ &\implies x \circ y^{-1}[k] \bmod 2^j \leq x \circ y^{-1}[k'] \bmod 2^j. \end{aligned}$$

Case 2: Both k and k' are in $[\#small \dots N)$. In this case we have the increasing bijection $[\#small(j) \dots N) \xrightarrow{y^{-1}} big(j)$. As $y^{-1}[k] \in big(j)$ we have

$$x \circ y^{-1}[k] \bmod 2^j = 2^{j-1} + x \circ y^{-1}[k] \bmod 2^{j-1}.$$

Other steps are as in case 1.

Case 3: Consider $k \in [0 \dots \#small(j))$ but $k' \in [\#small(j) \dots N)$. As y^{-1} is bijective we have $y^{-1}[k] \in small(j)$ but $y^{-1}[k'] \in big(j)$. Therefore

$$x \circ y^{-1}[k] < 2^{j-1} \leq x \circ y^{-1}[k'].$$

And the proof is done. \square

By the above lemmas we have that

Lemma 7: The sequential composition of `Enum_Zeros`, `Enum_Ones`, and `Exchange` satisfies

$$\begin{aligned} &\{ A; SORTED(j-1) \} \\ &\quad Enum_Zeros(j); \\ &\quad Enum_Ones(j); \\ &\quad Exchange \\ &\{ A; SORTED(j) \} \end{aligned}$$

The above lemma and the axiom of (sequential) for statement give us the correctness of the `Radix_Sort`.

Example 3: Dynamic processor/memory allocation

In machines with a large number of processors, it may be useful to create the high-level illusion that processors and memory are created dynamically on program demand. Figure 4 presents a program that allows an arbitrary set of processors to set up pointers to new processors, each with its own share of free memory. This program is based on a rendezvous technique: Demanding processors and free processors exchange their addresses through variables that they all know.

We present this example as an exercise in specification and refinement. First, we give an initial, very general specification of the problem. Then, we refine it adopting various

```

procedure Parallel_Cons
  (var free : array [0..MaxProc] of proc of boolean;
   var new : array [0..MaxProc] of proc of integer  $\cup$  {null});

var available : integer;
    rv : array [0..MaxProc] of proc of integer;

procedure Give_Rv;
var required : integer;
    free_proc : array [0..MaxProc] of proc of integer;
begin
  for all k do in parallel
    required := count;
    unconditionally
      if free [k] then
        available := count;
        free_proc [k] := enumerate;
        if free_proc [k] < required then
          free [k] := false;
          rv [ free_proc [k] ] := k
        end
      end
    end
  end
end Give_Rv;

procedure Accept_Rv;
var requestor : array [0..MaxProc] of proc of integer;
begin
  for all k do in parallel
    requestor [k] := enumerate;
    if requestor [k] < available then
      new [k] := rv [ requestor [k] ]
    end;
    if requestor [k]  $\geq$  available then
      new [k] := null
    end
  end
end Accept_Rv;

begin
  Give_Rv;
  Accept_Rv
end Parallel_Cons

```

Figure 4. Dynamic processor/memory allocation.

decisions on implementation, until we can prove that the particular program in Figure 4 is correct.

Assume there is a set F of processors having available free memory. This information is initially stored in the array

free: array [0..MaxProc] of proc of boolean

Therefore the set of initially free processors is defined as $F = \text{free}^{-1}(\text{true})$. Let A be the set of processors that are active when program `Parallel_Cons` is called. We will assume that A is the set of processors that require extra memory.

`Parallel_Cons` assigns free memory to processors in A : It sets up a total mapping $\text{new} : A \rightarrow F \cup \{\text{null}\}$ ($\text{new}(k)$ is defined for all $k \in A$) giving to each processor $k \in A$ the address $\text{new}(k)$ of a free processor if possible, otherwise $\text{new}(k) = \text{null}$. More precisely:

- When not so many processors ask for extra memory, all of them get it.

$$\#A \leq \#F \implies A \xrightarrow{\text{new}} \text{new}(A) \text{ is a bijection.}$$

- When there are too many processors, a maximum number is satisfied.

$$\#A > \#F \implies \text{new}^{-1}(F) \xrightarrow{\text{new}} F \text{ is a bijection.}$$

To deal with the mapping new we introduce the following predicate

$$\begin{aligned} \text{ALLOCATED}(\text{new}) \equiv & (\text{new} : A \rightarrow F \cup \{\text{null}\} \text{ is a total mapping} \wedge \\ & \#A \leq \#F \implies A \xrightarrow{\text{new}} \text{new}(A) \text{ is a bijection} \wedge \\ & \#A > \#F \implies \text{new}^{-1}(F) \xrightarrow{\text{new}} F \text{ is a bijection}). \end{aligned}$$

After the execution of `Parallel_Cons`, the set of free processors has diminished. Array *free* verifies:

$$\text{free}[k] = \text{true} \iff \text{new}^{-1}[k] = \emptyset.$$

To deal with this set we consider the following predicate, where the whole set of processors is $\text{IP} = [0..\text{MaxProc}]$.

$$\begin{aligned} \text{FREE}(\text{new}, \text{free}) \equiv & \\ & (\text{free} : \text{IP} \rightarrow \{\text{true}, \text{false}\} \text{ is total} \wedge \text{free}[k] = (k \in F \wedge \text{new}^{-1}(k) = \emptyset)). \end{aligned}$$

The dynamic processor/memory allocation would meet:

Specification 8: `Parallel_Cons` satisfies the following specification:

$$\begin{aligned} & \{A; free : \mathbb{P} \rightarrow \{\mathbf{true}, \mathbf{false}\} \text{ is total} \wedge F = free^{-1}(\mathbf{true})\} \\ & \quad \text{Parallel_Cons} \\ & \{A; ALLOCATED(new) \wedge FREE(new, free)\} \end{aligned}$$

To develop any implementation we need to assume some extra hypothesis about *new*. We can implement $new : A \rightarrow F \cup \{\mathbf{null}\}$ in a variety of ways. For example, take the processors of *A* and *F* in ascending order

$$A = \{a_0, a_1, a_2, \dots\} \quad , \quad F = \{f_0, f_1, f_2, \dots\}$$

and define *new* as

$$new(a_0) = f_0, new(a_1) = f_1, \dots$$

Of course there exists many other possible definitions. However, to develop `Parallel_Cons` by refinements it is better to hide as many details as possible of *new*. Let us consider implementations of *new* giving priority to the processors with small identifiers (as in the previous example). Then we consider functions *new* such that

$$new : first_{\#F}A \rightarrow first_{\#A}F \text{ is a bijection} \quad \wedge \quad new(A \setminus first_{\#F}A) = \{\mathbf{null}\}.$$

To describe this approach we introduce the predicate

$$\begin{aligned} LOW_ALLOCATED(new) \equiv \\ (new : A \rightarrow F \cup \{\mathbf{null}\} \text{ is total} \wedge first_{\#F}A \xrightarrow{new} first_{\#A}F \text{ is a bijection}). \end{aligned}$$

By construction $LOW_ALLOCATED(new)$ implies $ALLOCATED(new)$.

Under these hypotheses the free processors are those having large identifiers.

$$\begin{aligned} HIGH_FREE(free) \equiv \\ (free : \mathbb{P} \rightarrow \{\mathbf{true}, \mathbf{false}\} \text{ is total} \wedge (\forall k \in A : free[k] = (k \in F \wedge \#F_{<k} \geq \#A))) \end{aligned}$$

Then $LOW_ALLOCATED(new) \wedge HIGH_FREE(free)$ implies

$$ALLOCATED(new) \wedge FREE(new, free)$$

We can refine the above specification to the following one:

Specification 9: With the idea that the first elements of A will pair with the first elements of F we get

$$\begin{aligned} & \{A; \text{free} : \mathbb{IP} \rightarrow \{\mathbf{true}, \mathbf{false}\} \text{ is total} \wedge F = \text{free}^{-1}(\mathbf{true}) \} \\ & \quad \text{Parallel_Cons} \\ & \{A; \text{LOW_ALLOCATED}(\text{new}) \wedge \text{HIGH_FREE}(\text{free})\} \end{aligned}$$

As $\text{Parallel_Cons} = \text{Give_Rv} ; \text{Accept_Rv}$, we should find the right specifications for these two procedures. Give_Rv will give addresses of free processors in an array rv (rendezvous).

rv : array $[0..\text{MaxProc}]$ of proc of integer.

Formally a rendezvous is a bijection

$$rv : [0..\min(\#A, \#F)] \rightarrow \text{first}_{\#A} F.$$

There are many possibilities to construct such a bijection. For example $rv[\#F_{<k}] = k$, such that

$$rv[0] = f_0, rv[1] = f_1, rv[2] = f_2, \dots$$

But in fact the specification of procedure Give_Rv asks just for a bijection, any bijection will work. We introduce the predicate

$$\text{RENDEZ_VOUS}(rv) \equiv (rv : [0..\min(\#A, \#F)] \rightarrow \text{first}_{\#A} F \text{ is a bijection}).$$

Then we can prove that Give_Rv given in Figure 4 works.

Lemma 10: The procedure Give_Rv satisfies the following specification:

$$\begin{aligned} & \{A; \text{free} : \mathbb{IP} \rightarrow \{\mathbf{true}, \mathbf{false}\} \text{ is total} \wedge F = \text{free}^{-1}(\mathbf{true}) \} \\ & \quad \text{Give_Rv} \\ & \{A; (\text{available} = \#F) \wedge \text{RENDEZ_VOUS}(rv) \wedge \text{HIGH_FREE}(\text{free})\}. \end{aligned}$$

Proof. After procedure Give_Rv we have the following functions

$$\text{free_proc} : F \rightarrow \mathbb{IN} \quad \text{such that} \quad \text{free_proc}[k] = \#F_{<k}.$$

We have also

$$\text{free} : F \rightarrow \{\mathbf{true}, \mathbf{false}\} \quad \text{such that} \quad \text{free}[k] = \mathbf{false} \text{ iff } k < \#A.$$

When we consider the overall set F the function free can be rewritten as

$$\text{free}[k] = (k \in F \wedge \#F_{<k} \geq \#A)$$

The function $rv : [0..\min(\#A, \#F)] \rightarrow F$ satisfies $rv[\#F_{<k}] = k$. In fact the range of this function is $\text{first}_{\#A} F$. Then

$$rv : [0..\min(\#A, \#F)] \rightarrow \text{first}_{\#A} F \text{ such that } rv[\#F_{<k}] = k \text{ is a bijection.}$$

This finishes the lemma. \square

Finally the following lemma gives us an specification of Accept_Rv .

Lemma 11: The procedure `Accept_Rv` satisfies the following specification

$$\{A; (\text{available} = \#F) \wedge \text{RENDEZ_VOUS}(rv)\} \\ \text{Accept_Rv} \\ \{A; \text{LOW_ALLOCATED}(new)\}.$$

Proof. In the procedure `Accept_Rv` we have the following functions

$$\text{requestor} : A \rightarrow \mathbb{N} \quad \text{such that} \quad \text{requestor}[k] = \#A_{<k}$$

In particular there is a bijection

$$\text{requestor} : \text{first}_{\#F}A \rightarrow [0..\min(\#A, \#F))$$

The statement “`new[k] := rv[requestor[k]]`” gives a bijection

$$\text{first}_{\#F}A \xrightarrow{\text{requestor}} [0..\min(\#A, \#F)) \xrightarrow{rv} \text{first}_{\#A}F$$

Then `new` is a bijection on

$$\text{new} : \text{first}_{\#F}A \rightarrow \text{first}_{\#A}F$$

such that `new[k] = rv[#A<k]` and

$$\text{new}(A \setminus \text{first}_{\#F}A) = \{\text{null}\}$$

This is all. \square

5. Conclusions

Parallel machines are becoming cheaper and more flexible everyday. It is almost certain that the design of large software systems and the development of portable languages for them will be an important issue in a few years. All this is difficult if the only way to describe parallel programs is describing a machine that runs them.

As a small step, we have shown that it is possible to give a formal meaning to the notations used to describe PRAM algorithms. It seems that the theory and tools required are not so different from those used in the sequential case.

There is an important difference, though. The borderline between high level and low level is nowadays quite clear in sequential programming. We believe that there is a long way to go before we have such a clear distinction in parallel programming. Take as examples the points • mentioned in the Introduction. Or, concerning our notation, consider the following question: does it make sense to nest for-all statements in a program?

- If we take a low-level view, a statement such as

for all k do in parallel $S(k)$ end

indicates that an index k is given to each processor, and the instruction S is broadcast to all. Under this approach it seems nonsense a piece of code like

```
forall  $k : E_1(k)$  do in parallel
  for all  $j : E_2(j)$  do in parallel  $S(j)$  end
end
```

because it is rather unclear which instructions should be sent to each individual processor. Then it seems that for-all statements cannot be nested. But this fact is far from being obvious.

- Let us take another view. Suppose we have an array

```
 $x$  : array [0.. $N - 1$ , 0.. $N - 1$ ] of proc of boolean
```

and consider the following program:

```
forall  $i : 0 \leq i < N$  do in parallel
  forall  $j : 0 \leq j < N$  do in parallel
     $x[i, j] := \text{false}$ 
  end;
   $x[i, i] := \text{true}$ 
end
```

In a high level this text would suffice to indicate that all the assignments $x[i, j] := \text{false}$ can be done safely in parallel and the same happens for assignments $x[i, i] := \text{true}$. Therefore, this program computes the identity matrix and makes sense.

However, it does not seem wise to nest for-all's except in simple cases like this.

- There is one case where the nesting of for-all's is clearly a good program methodology. It appears when we hide the for-all statement inside a procedure. Let us consider as example the matrix product, Figure 5.

In this case the unfolding of the structure

```
forall  $i, j : 1 \leq i, j < N$  do in parallel
  Matrix_Product[ $i, j$ ] := Tree_Sum( $x[i, 0..N - 1, j]$ )
end
```

implies the nesting of two for-all's. Figure 6 gives us a version of the reflexive transitive closure with a high nesting of for-all instructions.

Let us give a tentative list of proverbs that may be useful in the design and proof of parallel algorithms.

- Try to avoid unconditional instructions. We believe that unconditionals are very low level statements because they break the block structure of the program. They should

```

procedure Matrix_Product;
  (a, b: array [0..N - 1, 0..N - 1] of proc of integer):
  array [0..N - 1, 0..N - 1] of proc of integer;

import
  Tree_Sum (a: array [0..N - 1] of proc of integer): integer;
  { returns  $\sum_{0 \leq i < N} a[i]$  in time  $O(\log N)$  }

var
  x: array [0..N - 1, 0..N - 1, 0..N - 1] of proc of integer;

begin
  for all i, j, k :  $0 \leq i, j, k < N$  do in parallel
    x[i, j, k] := a[i, j] * b[j, k]
  end;
  for all i, j :  $0 \leq i, j < N$  do in parallel
    Matrix_Product[i, j] := Tree_Sum(x[i, 0..N - 1, j])
  end;
end Matrix_Product

```

Figure 5. Matrix product.

be used in exceptional cases that occur asynchronously, such as dynamic memory management.

- Try to avoid the nesting of for-all's. When this nesting is necessary deal with it clearly. Hide it into procedures. A good idea is that different procedures act over different active data structures.
- Try to avoid $S(k)$ with complicated structure in a statement like

for all $k : E(k)$ **do in parallel** $S(k)$ **end**

- Try to program with small for-all pieces with clearly defined global intermediate states, like

```

{ clearly defined global state }
for all  $k : E_1(k)$  do in parallel  $S_1(k)$  end;
{ clearly defined global state }
for all  $k : E_2(k)$  do in parallel  $S_2(k)$  end
⋮

```

- When you have a statement

for all $k : E(k)$ **do in parallel** $S(k)$ **end**

```

procedure Transitive_Reflexive_Closure;
  (a: array [0..N - 1, 0..N - 1] of proc of boolean):
  array [0..N - 1, 0..N - 1] of proc of boolean;

import
  Bool_Mat_Product (a, b: array [0..N - 1, 0..N - 1] of proc of boolean):
    array [0..N - 1, 0..N - 1] of proc of boolean;
    { returns the matrix product  $a \cdot b$  in time  $O(\log N)$  }

var
  x: array [0..N - 1, 0..N - 1] of proc of boolean;

begin
  for all i, j :  $0 \leq i, j < N$  do in parallel
    if i = j then x[i, j] := true else x[i, j] := a[i, j] end
  end;
  for k := 1 to  $\log N$  do
    for all i, j :  $0 \leq i, j < N$  do in parallel
      x[i, j] := Bool_Mat_Product (x, x)[i, j]
    end
  end;
  for all i, j :  $0 \leq i, j < N$  do in parallel
    Transitive_Reflexive_Closure[i, j] := x[i, j]
  end;
end Transitive_Reflexive_Closure

```

Figure 6. Transitive reflexive closure.

with a local structure of $S(k)$ try to express the global state after the execution of the for-all as a union of local states. As an example we can take some well-known algorithms for numerical simulation.

- When you have a statement

for all $k : E(k)$ **do in parallel** $S(k)$ **end**

and different $S(k)$ interact heavily, try to give directly the global state.

6. References

- [Ble86] Bletloch, G.: Parallel prefix vs. concurrent memory access. Rep. Thinking Machines Corp., Cambridge, Mass., 1986.

- [Di75] Dijkstra, E.W.: Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM*, **18**, 453–457 (1975). Reprinted in *Programming Methodology, a Collection of Articles by Members of WG2.3*. editor Gries, D., Springer-Verlag 1978.
- [Gr81] Gries, D.: *The Science of Programming*. Springer-Verlag 1981.
- [Hi85] Hillis, D.W.: *The Connection Machine*. MIT Press 1985.
- [Ho69] Hoare, C.A.R.: An axiomatic basis for a computer programming. *Communications of the ACM*, **12**, 576–580 (1969). Reprinted in *Essays in Computer Science*. 45–58 edited by Hoare, C.A.R. and Jones, C.B., Prentice-Hall 1989.
- [HS86] Hillis, D.W.; Steele, G.L.: Data parallel algorithms. *Communications of the ACM*, **29**, 1170–1183 (1986).
- [KR90] Karp, R.M.; Ramachandran, V.: Parallel algorithms for shared-memory machines. In *Handbook of Theoretical Computer Science* (vol A), 869–941, editor Jan Van Leeuwen, Elsevier and MIT Press 1990.
- [Yo82] Yourdon, E.: *Writings of the Revolution*. Yourdon Press 1982.