

Fair Termination Revisited - With Delay

(Abstract)

K.R. Apt
LITP, Université Paris 7
2, Place Jussieu 75221
Paris, France

A. Pnueli
The Weizmann Institute of Science

J. Stavi
Bar-Ilan University

August 1982

Summary.

A proof method for establishing the fair termination and total correctness of both nondeterministic and concurrent programs is presented. The method calls for the extension of states by auxiliary delay variables which count down to the instant in which a certain action will be scheduled. It then uses well founded ranking to prove fair termination allowing nested fair selection and loops.

The work reported here was partly done while the first author was visiting the Weizmann Institute.

It is supported in part by the Israeli Academy of Sciences, the Basic Research Foundation.

INTRODUCTION.

The problem of termination of nondeterministic and concurrent programs under the assumption of fairness has recently been receiving considerable attention (see e.g. [Aol], [CFMR], [LPS], [P]).

The basic method for proving invariant properties, such as partial correctness, was developed by Floyd ([F]) and Hoare ([H]) for sequential programs. It is based on the idea of finding an inductive property which is preserved by every basic action of the program. When we consider nondeterministic and concurrent programs, the method of invariance is still applicable with very minor modifications.

By comparison, the suggested method for proving termination properties (total correctness for example, [M]), is not directly extendable to concurrent and nondeterministic programs when we stipulate fair executions. The method, as developed in [F], [MP] is based on establishing a mapping from the program states to some well-founded domain (a rank) such that any program action causes a decrease in the rank. That this method does not apply to fair termination is obvious from the following trivial example:

```
while b do if[b + skip  $\square$  b + b: = false]fi
```

A fair execution of this program must eventually choose the second branch of the conditional, causing b to be set to false and terminating the program. However any choice of the first branch preserves the program state. Correspondingly no mapping which always decreases can exist.

The study of fair executions is mainly motivated by concurrent programs. For concurrent computations fairness or its weaker version-justice ([LPS]) - is a most general modelling of the fact that the ratio of speeds between cooperating processors may be arbitrarily large and varying but is always finite. The study of fairness in the context of nondeterministic but sequential programs is motivated in part by the use of nondeterminism to model concurrency and also as a more restricted interpretation of nondeterminism.

Answering to the challenge of extending the method of well founded orderings to fair termination, several suggestions were made.

One approach, represented by [LPS] and [GFMR] is to relax the requirement that every action causes a decrease in the rank of the state. By this methodology, for each state there always exist some helpful actions which decrease the rank of the state, and some other actions, termed indifferent (steady in [GFMR]), which at least do not increase this rank. By fairness (and some additional requirements of the method), a helpful action must eventually be chosen which causes the rank to decrease and thus excludes infinite computations. This method was applied in [LPS] to concurrent programs represented in an abstract framework, and in [GFMR] to nondeterministic programs in a more syntax directed style. An interesting point is that the method of [GFMR] can only be applied to programs which terminate due to fairness on the top level, i.e. fair choice between the branches of an encompassing loop and not between branches of an enclosed conditional statement.

Thus the following example:

```

while b do
  if [ b + skip
    □ b + if [ b + skip
      □ b + b: = false ] fi
  ] fi od

```

cannot be proven fairly terminating by the method of [GFMR].

Another approach to fair termination developed in [AO] suggests modifying the program by the construction of an explicit fair scheduler for the program. This reduces the problem of fair termination to that of the termination of a deterministic program in which random assignments $x := ?$ of unbounded natural numbers are allowed. Such assignments are used by the scheduler to implement fair scheduling. By [AP] the termination of such programs can always be proved by well founded ranking, provided we allow ordinals higher than ω - the first countable ordinal. Once the proof rules are obtained for the program augmented by the scheduler statements, these statements can be eliminated. Thus, we do not have to actually construct the scheduler in order to apply the derived proof rules. They are directly applicable to the program as originally presented. In [AO] this method was developed again only for top level fairness in nondeterministic programs.

In this paper we present another approach to the termination of fair programs, covering both concurrent and nondeterministic programs. We believe it to be much simpler and more natural than any of the approaches discussed above, and as we will illustrate below, directly applicable. While the method can also be justified by program transformations, as in [AO], the presented justification does not call for program modification but extends instead the states by adding auxiliary variables. In a certain sense this extension parallels the introduction of auxiliary variables in [OG] providing a natural method for invariance properties of concurrent programs. As will be shown below our method provides proofs for termination under the assumption of overall fairness and not only top level fairness. Thus, in comparison with previous proof methods the approach suggested here is more general, is simpler to apply and justify and forms a natural generalization of the method of well founded ranking successfully used for sequential programs.

Similarly to [AO] we will show that the problem of fair total correctness of a nondeterministic program is reducible to that of the ordinary total correctness of a program which allows unbounded random assignments $x := ?$. Such programs were studied in [AP]. In our paper we will show the following additional result concerning such reductions in the other direction.

Given a program Π which allows random assignments, it is possible to construct a nondeterministic program Π_1 with no random assignments such that the fair total correctness of Π_1 is equivalent to the ordinary total correctness of Π . Furthermore, it is sufficient to require top level fairness in the computations of Π' . This result allows us to resolve the issue raised in [AO] by showing that all recursive ordinals (order types of recursive well ordering of sets of natural numbers) are required to establish fair termination of programs with top level fairness only. This of course is a significant increase in complexity over the sequential deterministic case where ω is the highest ordinal ever needed.

2. CONCURRENT PROGRAMS.

The method is illustrated first for concurrent programs represented in an unstructured framework. The framework is taken from [LPS] and we repeat its basic definitions here:

A concurrent system is a triple:

$$P = \langle S, F, I \rangle$$

where

S is a set of execution states.

$I \subseteq S$ is the set of initial states.

$F = \langle f_1, \dots, f_m \rangle$ is a set of transition functions associates with m processes. Each $f_i: S \rightarrow 2^S$ maps a state s into a set $f_i(s) \subseteq S$ which is the set of possible outcomes when the process P_i executes an atomic instruction on the state s .

If $f_i(s) \neq \emptyset$ we say that f_i is enabled on s , otherwise we say that it is disabled on s . A state s which is disabled for all $i = 1, \dots, m$ is called terminal. Let T denote the set of terminal states.

An execution sequence of P is a maximal sequence:

$$s_0 \xrightarrow{f_{i_1}} s_1 \xrightarrow{f_{i_2}} s_2 \longrightarrow .$$

such that $s_0 \in I$ and for each j , $s_{j+1} \in f_{i_{j+1}}(s_j)$. A state is accessible if it occurs in an execution sequence. The set of accessible states is denoted by $\text{Acc}(I)$.

An execution sequence is fair if it is either finite or if every transition f_k which is enabled infinitely many times in the sequence is also scheduled infinitely many times, i.e. $i_j = k$ for infinitely many j 's.

We say that a program P is fairly convergent if every fair execution sequence of P is finite.

We propose the following proof method for proving the fair convergence of con-

current systems.

The Delay Variables Method.

1. Choose a state predicate $Q \subseteq S$ such that

$$A. s \in I \Rightarrow (s \in T) \vee Q(s)$$

$$B. Q(s) \wedge s' \in f_i(s) \Rightarrow (s' \in T) \vee Q(s') \text{ for } i=1, \dots, m.$$

(T being the set of terminal states)

This ensures that the predicate Q holds for all accessible non terminal states.

2. Choose a well founded set $(W, >)$, i.e. a set W with an ordering relation $>$ such that every W-sequence $w_0 > w_1 > \dots w_i \in W$ is finite.

3. Find a ranking function

$$\rho : S \times N^m \rightarrow W$$

mapping extended states into the well founded domain W. An extended state consists of a state $s \in S$ augmented by m scheduling (or delay) variables z_1, \dots, z_m (also referred to as counters). The role of the delay variable z_i is to count how many steps will pass in which f_i is enabled but not yet scheduled. By fairness there can be only a finite number of them.

The ranking function must satisfy:

$$Q(s) \wedge s' \in f_i(s) \wedge \bigwedge_{j \neq i} [(f_j(s) \neq \emptyset \Rightarrow z_j = z'_j + 1) \wedge (f_j(s) = \emptyset \Rightarrow z_j = z'_j)] \Rightarrow$$

$$\Rightarrow \rho(s, z_1, \dots, z_{i-1}, 0, z_{i+1}, \dots, z_m) > \rho(s', z'_1, \dots, z'_m).$$

All the free variables in the above, i.e. s, s', \bar{z}, \bar{z}' are considered to be universally quantified. To justify the method, consider an infinite fair execution sequence

$$\sigma : s_0 \xrightarrow{f_{i_1}} s_1 \xrightarrow{f_{i_2}} s_2 \longrightarrow \dots$$

For each $j = 0, 1, \dots$ we define the vector of delay values $\bar{u}^j = (u_1^j, \dots, u_m^j) \geq 0$ as follows:

u_k^j = number of distinct $j' \geq j$ such that f_k is enabled on s_j , but not scheduled for any $j \leq l \leq j'$, i.e. number of contiguous steps from j on where f_k is enabled but not yet scheduled.

By fairness the \bar{u}^j 's are well defined, nonnegative and finite.

In addition they have the following properties:

Consider a transition $s_j \xrightarrow{f_k} s_{j+1}$, i.e. $i_{j+1} = k$.

a. $u_k^j = 0$.

b. For every $l \neq k$ such that $f_l(s_j) \neq \emptyset$ $u_l^j = u_l^{j+1} + 1$.

c. For every $l \neq k$ such that $f_l(s_j) = \emptyset$ $u_l^j = u_l^{j+1}$.

Let σ now represent an infinite fair sequence. Assume that Q , W and ρ have been found satisfying the method's requirements. Consider the sequence of augmented states $(s_0, \bar{u}^0), (s_1, \bar{u}^1), \dots$

By comparing properties a, b and c to the requirements on Q and ρ we obtain

$$\rho(s_0, \bar{u}^0) > \rho(s_1, \bar{u}^1) > \dots$$

Contradicting the well foundedness of W . This shows that a successful choice of Q , W and ρ guarantees fair termination.

Conclusion: The delay variable method is sound.

Completeness is even more trivial. Assume that P is fairly convergent.

Take $Q = \text{Acc}(I)$ - i.e. true for all accessible states.

Take W to be $S \times N^m$, ρ the identity mapping

$$\rho(s, z_1, \dots, z_m) = \langle s, z_1, \dots, z_m \rangle$$

The relation $>$ over W is defined by

$$\langle s, z_1, \dots, z_m \rangle > \langle s', z'_1, \dots, z'_m \rangle$$

↔

$$\exists i \ Q(s) \wedge s' \in f_i(s) \wedge \bigwedge_{j \neq i} [(f_j(s) \neq \emptyset \Rightarrow z_j = z'_j + 1) \wedge (f_j(s) = \emptyset \Rightarrow z_j = z'_j)] \wedge z_i = 0$$

i.e. it holds exactly between two accessible augmented states that can appear contiguously in a computation. That this relation is well founded follows immediately from the fact that P is fairly convergent.

We conclude this section by an example of proving fair termination of the following distributed GCD program:

```

while  $y_1 \neq y_2$  do if  $[y_1 > y_2 \rightarrow y_1 := y_1 - y_2 \square y_1 < y_2 \rightarrow \text{skip}] f_1$  od
                    ||
while  $y_1 \neq y_2$  do if  $[y_1 > y_2 \rightarrow \text{skip} \square y_1 < y_2 \rightarrow y_2 := y_2 - y_1] f_2$  od

```

In our framework

$$I = S = \{(y_1, y_2) \mid y_1 > 0, y_2 > 0\} .$$

The transition functions f_1, f_2 are given by

$$f_1(\{y_1, y_2\}) = \text{if } y_1 > y_2 \text{ then } [y_1 - y_2, y_2] \text{ else if } y_1 < y_2 \text{ then } [y_1, y_2]$$

$$f_2(\{y_1, y_2\}) = \text{if } y_1 > y_2 \text{ then } [y_1, y_2] \text{ else if } y_1 < y_2 \text{ then } [y_1, y_2 - y_1]$$

Note that both functions are disabled on states of the form $[y, y]$ which are therefore terminal.

To apply the delay variables method we choose Q as

$$Q(\{y_1, y_2\}) = (y_1 \neq y_2) \wedge (y_1, y_2 > 0) .$$

$W = \mathbb{N} \times \mathbb{N}$ - the set of pairs of nonnegative integers with the lexicographic ordering on pairs given by:

$$(m_1, n_1) > (m_2, n_2) \iff (m_1 > m_2) \text{ or } (m_1 = m_2 \text{ and } n_1 > n_2)$$

$$\rho(\{y_1, y_2\}, z_1, z_2) = (y_1 + y_2, \text{if } y_1 > y_2 \text{ then } z_1 \text{ else } z_2)$$

we have to show that the value of ρ decreases on each transition. Take for example the case of $y_1 > y_2$. We consider separately $i=1$ and $i=2$. We have to show

$$[y'_1, y'_2] = f_1([y_1, y_2]) \wedge z_2 = z'_2 + 1 \Rightarrow \rho([y_1, y_2], 0, z_2) > \rho([y'_1, y'_2], z'_1, z'_2) .$$

But certainly in this case $y_1 + y_2 > y'_1 + y'_2$ so that $\rho > \rho'$. For $i=2$ and $y_1 > y_2$ we have to show

$$[y'_1, y'_2] = f_2([y_1, y_2]) \wedge z_1 = z'_1 + 1 \Rightarrow \rho([y_1, y_2], z_1, 0) > \rho([y'_1, y'_2], z'_1, z'_2) .$$

But in this case $[y'_1, y'_2] = [y_1, y_2]$ so that

$$\rho([y_1, y_2], z_1, 0) = (y_1 + y_2, z_1 + 1) > (y_1 + y_2, z'_1) = \rho([y'_1, y'_2], z'_1, z'_2) .$$

This proves that the distributed GCD program is indeed fairly terminating.

3. NONDETERMINISTIC PROGRAMS.

In this section we develop the variant of the method appropriate to nondeterministic programs. The programs considered here will be presented in a structured language, and the method will lead to the establishment of overall fair total correctness.

The syntax of our programs is given by the following grammar:

$$\Pi ::= \text{skip} \mid x := t \mid \Pi; \Pi \mid \text{if } \bigoplus_{i=1}^m B_i \rightarrow \Pi_i \text{ fi} \mid \text{while } B \text{ do } \Pi \text{ od}$$

Here x is a program variable, t a term, B and B_i are boolean expressions. The boolean B_i in the context of the if-fi construct are called guards. Our language differs from that of Dijkstra [D] in that the loop is always guarded by a single condition.

3.1. Semantics

Let Var denote the set of program variables and \mathcal{D} a domain of an interpretation. By a state we mean a mapping $s : \text{Var} \rightarrow \mathcal{D}$.

Following [HP] we define a simple operational semantics for programs based on a transition relation " \rightarrow " between configurations, that is pairs $\langle \Pi, s \rangle$ consis-

ting of a program and a state. In addition we consider the two special states \perp standing for divergence and fail standing for abortion.

In general $\langle \Pi, s \rangle \rightarrow \langle \Pi', s' \rangle$ means that one step of execution of Π applied to s can lead to a state s' with Π' being the remainder of Π yet to be executed.

It is convenient to assume the empty program E . Then Π' is E if Π terminates in s' . We assume that for any program Π $E; \Pi = \Pi; E = \Pi$.

We define the above relation by the following clauses:

- i) $\langle \text{skip}, s \rangle \rightarrow \langle E, s \rangle$
- ii) $\langle x := t, s \rangle \rightarrow \langle E, s' \rangle$
 where $s'(x) = s(t)$ and $s'(y) = s(y)$ for $y \neq x$
- iii) $\langle \text{if } \square_{i=1}^m B_i \rightarrow \Pi_i \text{ fi}, s \rangle \rightarrow \langle \Pi_i, s \rangle$ if $\models B_i(s)$
- iv) $\langle \text{if } \square_{i=1}^m B_i \rightarrow \Pi_i \text{ fi}, s \rangle \rightarrow \langle E, \text{fail} \rangle$ if $\models \bigwedge_{i=1}^m \neg B_i(s)$
- v) $\langle \text{while } B \text{ do } \Pi \text{ od } s \rangle \rightarrow \langle \Pi; \text{while } B \text{ do } \Pi \text{ od}, s \rangle$ if $\models B(s)$
- vi) $\langle \text{while } B \text{ do } \Pi \text{ od}, s \rangle \rightarrow \langle E, s \rangle$ if $\models \neg B(s)$
- vii) if $\langle \Pi, s \rangle \rightarrow \langle \Pi', s' \rangle$ then
 $\langle \Pi; \Pi_1, s \rangle \rightarrow \langle \Pi'; \Pi_1, s' \rangle$

Let " \rightarrow^* " stand for the reflexive transitive closure of " \rightarrow ". We say that Π_0 can diverge from s_0 if there exists an infinite sequence

$$\langle \Pi_0, s_0 \rangle \rightarrow \langle \Pi_1, s_1 \rangle \rightarrow \dots$$

We say that Π can fail from s if for some Π_1

$$\langle \Pi, s \rangle \rightarrow^* \langle \Pi_1, \text{fail} \rangle.$$

We may now define various semantics of programs by putting

$$M_P[\Pi](s) = \{s' \mid \langle \Pi, s \rangle \rightarrow^* \langle E, s' \rangle\}$$

$$M_{wt}[\Pi](s) = M[\Pi](s)$$

$$\cup \{\perp \mid \Pi \text{ can diverge from } s\}$$

$$M_{\text{wt}}^{\circ}[\Pi](s) = M_{\text{wt}}[\Pi](s) \cup \{\text{fail} \mid \Pi \text{ can fail from } s\}$$

We now proceed to define yet another semantics of programs - the one taking under consideration the assumption of fairness.

A computation sequence

$$\sigma : \langle \Pi_0, s_0 \rangle \rightarrow \langle \Pi_1, s_1 \rangle \rightarrow \dots$$

is said to be fair if it is either finite or for every program $\Pi : \text{if}_{i=1}^m B_i \rightarrow \Pi_i \text{ fi}$; Π' and each $i = 1, \dots, m$, if there are infinitely many j 's for which $\langle \Pi, s_j \rangle$ appears in σ and $\models B_i(s_j)$, then there are infinitely many j 's among them such that the transition

$$\langle \Pi, s_j \rangle \rightarrow \langle \Pi_i; \Pi', s_j \rangle$$

appears in σ .

This again captures the idea that every guard associated with a fixed location in the program, which is tested and found enabled an infinite number of times will be selected an infinite number of times.

To avoid confusion resulting from the fact that various occurrences of Π in σ do not need to correspond with the same program, we should actually label each program with a unique label. It is clear how to perform this process and we leave it to the reader.

We may now define $M_{\text{fair}}[\Pi](s)$ analogously as $M_{\text{wt}}[\Pi](s)$ by allowing \perp in it only if Π can diverge from s by a fair computation sequence.

Let P, Q, R stand for formulae (assertions) in an assertion language which contains all program variables, terms and boolean expressions. We put

$[P] = \{s \mid \models P(s)\}$. Note that for any assertion P , $\perp \in [P]$ and fail $\in [P]$.

For any $f \in \{p, \text{wt}, t, \text{fair}\}$, assertions P, Q and a program we define

$$M_f[\Pi]([P]_j) = \bigcup_{s \in [P]_j} M_f[\Pi](s)$$

The statement of program correctness is defined by:

$$\models_f \{P\} \Pi \{Q\} \text{ iff } M_f \Pi \Pi \{P\}_j \subseteq \{Q\}_j .$$

We thus have few types of program correctness:

\models_p - partial correctness

\models_{wt} - weak total correctness

\models_t - total correctness

\models_{fair} - total correctness under the assumption of fairness (fair total correctness)

The weak total correctness and the corresponding M_{wt} semantics are less often considered in the literature. We need these notions in the next section. We call the constructs $\{P\} \Pi \{Q\}$ the correctness formulae.

3.2 A Transformation Realizing Fairness

In the subsequent considerations we need an atomic program $x:=?$ called a random assignment which sets x to an arbitrary nonnegative integer. The semantics of random assignment is defined by adopting the clause

$$\langle x:=?, s \rangle \rightarrow \langle E, s' \rangle$$

for any state s' such that $s'(y) = s(y)$ for $y \neq x$. We assume that x ranges over natural numbers which form a subset of the domain D of the interpretation.

Programs allowing random assignment have been extensively studied in [AP]. In particular a system for proving total correctness of these programs has been presented there and we shall value use of it in order to develop proof rules for fair total correctness.

To this purpose we provide first a transformation of an arbitrary program Π into a program Π_{fair} allowing random assignments which realizes exactly all fair computations of Π . We proceed by the following successive steps:

1. replace each subprogram $\text{if } \bigcap_{i=1}^m B_i \rightarrow \Pi_i \text{ fi}$ of Π by the following subprogram

$\text{for } j:=1 \text{ to } m \text{ if } B_j \text{ then } z_j := z_j - 1;$
 $\text{if } \bigcap_{i=1}^m B_i \wedge z_i = 0 \wedge \bar{z} \geq 0 \rightarrow z_i := ?; \Pi_i \text{ fi}$

where \bar{z} stands for z_1, \dots, z_m .

2. rename all variables z_1, \dots, z_m appropriately so that each if-fi construct has its "own" set of these variables.

The variables z_1, \dots, z_m play here exactly the same role as in Section 2 - they count down how many times the corresponding guard is enabled but not yet selected. The corresponding actions on these variables are incorporated in the program text.

The following lemma relates Π to Π_{fair} .

Lemma 1. For any state s

$$M_{\text{fair}}[\Pi](s) = M_{\text{wt}}[\Pi_{\text{fair}}](s)$$

Here and later we disregard the problem that Π_{fair} can change the initial values of the (auxiliary) delay variables z_1, \dots whereas Π cannot. It is easy to remedy this difficulty by retaining the initial values of these variables before the execution of Π_{fair} and restore them after the execution of Π . We ignore this issue here since it is not relevant in the further discussion.

Proof. a) We prove the \subseteq -inclusion. Let $\sigma = \langle \Pi_0, s_0 \rangle + \langle \Pi_1, s_1 \rangle + \dots$ be a fair computation of Π . We extend it to a computation of Π_{fair} by assigning in each state of σ the values to the delay variables z_i -s. Given a state s_j there are two cases.

Case 1. For no state s_k ($k > j$) the guard corresponding with z_j is selected.

Then by the assumption of Fairness this guard is enabled only finitely many

times in this computation. We put $s_j(z_i)$ to be equal 1 + the number of times this guard will be enabled beyond s_j .

Case II. For some state s_k ($k > j$) the guard corresponding with z_j is selected.

Then we put $s_j(z_i)$ to be equal 1 + the number of times this guard will be enabled before being next time selected.

b) We prove the \supseteq -inclusion.

Let σ be a computation of Π_{fair} . Then its restriction to the computation steps dealing with Π is a computation sequence of Π . We show that it is a fair computation sequence. Suppose otherwise. Then behind some point in this computation a guard would be infinitely many times enabled and yet never chosen. By the construction of Π_{fair} the corresponding variable z_i would become arbitrarily small. This is however impossible because as soon as z_i becomes negative a failure will arise. \square

Corollary. Suppose that none of the delay variables occurs free in assertions P and Q . Then

$$\begin{aligned} \models_{\text{fair}} \{P\} \Pi \{Q\} \text{ iff } \forall s [\models P(s) \rightarrow \Pi \text{ cannot fail from } s] \\ \text{and } \models_{\text{wt}} \{P\} \Pi_{\text{fair}} \{Q\} \end{aligned}$$

3.3 A Proof System for Fair Total Correctness

The above corollary indicates that in order to prove fair total correctness of Π it is sufficient to prove weak total correctness of Π_{fair} provided the absence of failure in Π can be established.

To prove weak total correctness of Π_{fair} we can use the proof system introduced in [AP] slightly modified for our purposes. The following axioms and proof rules are adopted

1. Random assignment axiom

$$\{P\} x := ? \{P'\}$$

provided x is not free in P

2. Skip axiom

$\{P\} \text{ skip } \{P\}$

3. Assignment axiom

$\{P[t/x]\} x:=t \{P\}$

where $P[t/x]$ stands for a substitution of t for all free occurrences of x in P .

4. Composition rule

$$\frac{\{P\} \Pi_1 \{Q\}, \{Q\} \Pi_2 \{R\}}{\{P\} \Pi_1; \Pi_2 \{R\}}$$

5. Selection rule

$$\frac{\{P \wedge B_i\} \Pi_i \{Q\} \quad i=1, \dots, m}{\{P\} \text{ if } \square_{i=1}^m B_i \rightarrow \Pi_i \text{ fi } \{Q\}}$$

6. While rule

$$\frac{\{P(\alpha) \wedge B\} \Pi \{ \exists \beta < \alpha \Pi(\beta) \}}{\{P(\alpha)\} \text{ while } B \text{ do } \Pi \text{ od } \{ \exists \beta \leq \alpha \Pi(\beta) \wedge \neg B \}}$$

where α, β are variables ranging over ordinals (or more generally well founded sets).

7. Consequence rule

$$\frac{P \rightarrow P_1, \{P_1\} \Pi \{Q_1\}, Q_1 \rightarrow Q}{\{P\} \Pi \{Q\}}$$

The above system is appropriate for proving weak total correctness of Π_{fair} . We call it WTC.

Consider now a proof of a correctness formula $\{P_1\} \Pi_{\text{fair}} \{Q_1\}$ in the above system. Due to the form of Π_{fair} this proof can be transformed into a proof of the correctness formula $\{P_1\} \Pi \{Q_1\}$ provided we use the following transformed version of the selection rule

$$\frac{\{P\} \text{ if}_{\text{fair}} \text{ fi } \{Q\}}{\{P\} \text{ if } \square_{i=1}^m B_i \rightarrow \Pi_i \text{ fi } \{Q\}}$$

where if_{fair} fi stands for the subprogram introduced in step 1 of the transformation from section 3.2.

The hypothesis of this rule can be simplified if we "absorb" all assignments to delay variables into the assertion P and apply "backwards" the original selection rule. In such a way we obtain a proof rule which deals exclusively with the if-construct and its components. It has the following form

$$\frac{\{P[B_j \rightarrow z_j + 1, z_j/z_j]_{j \neq i} [1/z_i] \wedge B_i \wedge \bar{z} \geq 0\} \Pi_i \{Q\}_{i=1, \dots, m}}{\{P\} \underline{\text{if}} \square_{i=1}^m B_i \rightarrow \Pi_i \underline{\text{fi}} \{Q\}}$$

where $B_j \rightarrow t_1, t_2$ stands for the conditional expression if B_j then t_1 else t_2 fi

According to Corollary we still have to deal with the issue of freedom of failure. This problem can be taken care of in the usual way, i.e. by simply adding to the premises of the above rule the assertion

$$P \rightarrow \bigvee_{i=1}^m B_i.$$

Summarizing, the final version of the rule has the following form

8. Fair selection rule

$$P \rightarrow \bigvee_{i=1}^m B_i,$$

$$\frac{\{P[B_j \rightarrow z_j + 1, z_j/z_j]_{j \neq i} [1/z_i] \wedge B_i \wedge \bar{z} \geq 0\} \Pi_i \{Q\}_{i=1, \dots, m}}{\{P\} \underline{\text{if}} \square_{i=1}^m B_i \rightarrow \Pi_i \underline{\text{fi}} \{Q\}}$$

We have thus obtained a proof system for proving fair total correctness of programs. It consists of the axioms 2,3 and proof rules 4, 6-8. Note that the random assignment axiom is not needed - it was used only to derive the final form of the fair selection rule. Call this proof system FTC (for fair total correctness).

3.4. Soundness and Completeness of FTC

Before we dwell on the issue of soundness and completeness of FTC we have to specify for which assertion languages and their interpretations FTC is an appropriate proof system.

We assume that the assertion language L contains two sorts: data and ord. We have a constant 0 of type ord and a binary predicate symbol $<$ over ord. Additionally we assume that L includes second order variables of arbitrary arity and sort. The second order variables can be bound only by the *least fixed point operator* μ provided the bound variable occurs positively in the considered formula. (Here a variable occurs *positively* in a formula if none of its occurrences in a disjunctive normal form of the formula is in the scope of a negation sign). Thus if the set variable a occurs positively in $p(a)$ then $\mu a.p$ is a well formed formula. The free variables of $\mu a.p$ are those of p other than a .

An interpretation J for this type of assertion language is an ordinary two-sorted second order structure subject to the following four conditions

1. The domain J_{data} of sort data is countable and contains all natural numbers.
2. The domain J_{ord} of sort ord is an initial segment of ordinals (to ensure a proper interpretation of the while rule).
3. The constant 0 denotes the least ordinal and the predicate symbol $<$ denotes the strict ordering of the ordinals, restricted to J_{ord} .
4. The domains of each of the set sorts contain all sets of the appropriate kind (to ensure the existence of the fixed points considered below).

The truth with respect to an interpretation J is denoted by \models_J . The truth of the formulae of L (assertions) with respect to J is defined in a standard way. The only nonstandard case is when a formula is of the form $\mu a.P$. We put then $\models_J \mu a.P$ iff $\models_J P[A/a]$ where A is the least fixed point of an operator naturally induced by P .

The truth of the correctness formulae with respect to J is defined as before. We only need to indicate the dependence of the appropriate program semantics

on the interpretation J .

By Tr_J denote the set of all formulae of L which are true with respect to J . Given a set of assertions AS and a proof system G for proving correctness formulae we denote by $AS \vdash_G \varphi$ the fact that the correctness formula φ can be proved in G from the set of assumptions AS which can be used in the consequence rule.

After having introduced all these notions we can now state a lemma which is a proof theoretic counterpart of Corollary from section 3.2.

Lemma 2. Suppose that none of the delay variables introduced in Π_{fair} occurs free in the assertions P and Q . Then for any interpretation J of the above kind

$$\begin{aligned} Tr_J \vdash_{FTC} (P) \Pi (Q) \text{ iff } Tr_J \vdash_{WTC} (P) \Pi_{fair} (Q) \\ \text{and } \forall s \{ \models_J P(s) \rightarrow \Pi \text{ cannot fail from } s \}. \end{aligned}$$

Proof. The proof is based on the analysis of the proofs in the corresponding proof systems and makes use of the Corollary from section 3.2. We leave the details to the reader. \square

Corollary and Lemma 2 reduce the question of soundness and completeness of the proof system FTC to that of WTC. But the results of [AP] show that the proof system WTC is sound and complete for all interpretations J of the above kind. This shows that the proof system FTC is also sound and complete in the sense of the following theorem

THEOREM 1. For all interpretations J of the above kind and all correctness formulae φ

$$Tr_J \vdash_{FTC} \varphi \text{ iff } \models_{J, fair} \varphi. \quad \square$$

3.5 An Example of a Proof in FTC

We conclude with an example, which can be dealt with using our system but not by any previous method. The program was suggested by Shmuel Katz.

Let

```

Π : while x > 0 do
  if (z1) true → if (z3) B → x := x-1
    □ (z4) B → b := false
    □ (z5) ¬B → skip      fi
  □ (z2) true → b := true
  fi
od

```

we want to prove

$$\models_{\text{fair}} \{\text{true}\} \Pi \{\text{true}\} \quad (1)$$

i.e. that Π always terminates under the assumption of fairness.

The well founded set we will consider is N^4 under lexicographic ordering. We have annotated each guard with an appropriate delay variable. There is a ranking function which underlies our formal proof which is given by

$$\rho(x, B, \bar{z}) = (x, z_3, 1-B, (B \rightarrow z_1, z_2))$$

In the expression $1-B$, true is interpreted as 1, false as 0.

The crucial fact upon which the proof depends is that in a fair execution the value of ρ decreases on each iteration of the loop. We first demonstrate this fact informally providing the formal proof later. An iteration of the loop can be characterized by the guards which are selected.

Consider first the z_1, z_3 path. Here x is decremented so that ρ certainly decreases.

Along the z_1, z_4 path, the z_3 guard was enabled since b must have been true for z_1 to be selected. Consequently z_3 is decremented, being an enabled

but unselected guard. Since x remains the same ρ again decreases. Along the z_1, z_5 path, B must have been false so that the fourth component of ρ is z_2 which is decremented when its guard is not selected.

In the z_2 path we have to distinguish between the case that B is initially false in which case $1-B$ drops from 1 to 0, and the case that B was initially true in which case the last component of ρ is z_1 which is decremented since z_2 is selected.

We now present a formal proof of (1). Let Π' be the body of the loop. We have to find an assertion $P(\alpha)$ such that

$$\{P(\alpha) \wedge \alpha > 0 \wedge x > 0\} \Pi' \{ \exists \beta < \alpha P(\beta) \} \quad (2)$$

and

$$\exists \alpha P(\alpha) . \quad (3)$$

We define

$$P(\alpha) \equiv x, \bar{z} \geq 0 \wedge \alpha = \rho(x, B, \bar{z})$$

It is clear that (3) holds. To prove (2) we have to apply the fair selection rule so we have first to prove the premises

$$\{(P(\alpha) \wedge \alpha > 0 \wedge x > 0) [z_2+1/z_2] [1/z_1] \wedge z_1, z_2 \geq 0\} \Pi_1 \{ \exists \beta < \alpha P(\beta) \} \quad (4)$$

and

$$\{(P(\alpha) \wedge \alpha > 0 \wedge x > 0) [z_1+1/z_1] [1/z_2] \wedge z_1, z_2 \geq 0\} B := \text{true} \{ \exists \beta < \alpha P(\beta) \} \quad (5)$$

as the first premise of the fair selection rule is obviously satisfied. Here

$$\begin{aligned} \Pi_1 &\equiv \text{if } B \rightarrow x := x-1 \\ &\quad \square B \rightarrow B := \text{false} \\ &\quad \square \neg B \rightarrow \text{skip } fi . \end{aligned}$$

To prove (4) we once again wish to apply the fair selection rule. The premises to prove are

$$\{P_1 [B \rightarrow z_4+1, z_4/z_4] [\neg B \rightarrow z_5+1, z_5/z_5] [1/z_3] \wedge B \wedge z_3, z_4, z_5 \geq 0\} x := x-1 \{ \exists \beta < \alpha P(\beta) \} \quad (6)$$

$\{P_1[B \rightarrow z_3+1, z_3/z_3] \{ \neg B \rightarrow z_5+1, z_5/z_5 \} [1/z_4] \wedge B \wedge z_3, z_4, z_5 \geq 0\} B := \underline{\text{false}} \{ \exists \beta < \alpha P(\beta) \}$ (7)

and

$\{P_1[B \rightarrow z_1+1, z_1/z_1]_{i=3,4} [1/z_5] \wedge \neg B \wedge z_3, z_4, z_5 \geq 0\} \underline{\text{skip}} \{ \exists \beta < \alpha P(\beta) \}$ (8)

where $P_1 \equiv (P(\alpha) \wedge \alpha > 0 \wedge x > 0) [z_2 + 1/z_2] [1/z_1] \wedge z_1, z_2 \geq 0$.

We have by the assignment axiom

$\{\rho(x, 1, 0, 1) = \alpha \wedge B \wedge x > 0 \wedge \bar{z} \geq 0\}$

$x := x - 1$

$\{\rho(x+1, 1, 0, 1) = \alpha \wedge B \wedge x \geq 0 \wedge \bar{z} \geq 0\}$

which implies by the consequence rule (6) as the necessary implications clearly hold.

To prove (7) we note that by the assignment axiom and the consequence rule

$\{\rho(x, z_3, 0, 1) = \alpha \wedge B \wedge x > 0 \wedge \bar{z} \geq 0\}$

$B := \underline{\text{false}}$

$\{\rho(x, z_3, 1, z_2) = \alpha \wedge \neg B \wedge x \geq 0 \wedge \bar{z} \geq 0\}$

so (7) holds by the consequence rule.

Finally, to prove (8) we note that

$P_1 [B \rightarrow z_1+1, z_1/z_1]_{i=3,4} [1/z_5] \wedge \neg B \wedge z_3, z_4, z_5 \geq 0$ implies

$\rho(x, z_1, 1, z_2+1) = \alpha \wedge \neg B \wedge \bar{z} \geq 0 \wedge x > 0$ which in turn implies $\exists \beta < \alpha P(\beta)$.

Hence (8) holds by the skip axiom.

Now, from (6) - (8) we get (4) by the fair selection rule.

To prove (5) we note that by the assignment axiom and the consequence rule

$\{\rho(x, z_3, 1-b, (b+z_1+1, 1)) = \alpha \wedge x, \bar{z} \geq 0\}$

$B := \underline{\text{true}}$

$\{\rho(x, z_3, 0, z_1+1) = \alpha \wedge B \wedge x, z \geq 0\}$

so (5) by the consequence rule.

We may now have proved both (4) and (5) and we get (2) by the fair selection rule. Now (2) and (3) imply by the while rule $\{\underline{\text{true}}\} \bar{\Gamma} \{\underline{\text{true}}\}$ so by virtue of the soundness of the system FIC we get (1). This concludes the proof.

4. ON THE SIZE OF NEEDED ORDINALS.

In the preceding sections we have presented methods for proving fair termination of (concurrent or structured non-deterministic) programs, using ranking functions into well-founded sets or predicates of ordinals. It is well-known that any well-founded set $\langle W, \succ \rangle$ has an order preserving mapping into $\langle W_\alpha, \succ \rangle$ for some ordinal α , where $W_\alpha = \{\beta \mid \beta < \alpha\}$ (see [LPS] for details). Thus, one measure of the "complexity" of fair termination of a concurrent program P is the least ordinal α for which there exist Q, W and ρ as in the delay variables method with $W = W_\alpha$. Let us call this ordinal α the "fair ordinal" of P and denote it by α_P ($\alpha_P = 0$ in case P is not fairly convergent). A similar measure of complexity can be associated with structured nondeterministic programs by studying ordinals needed for applying the while rule.

Consider bounds on α_P for natural classes of programs P . For definiteness we consider programs operating on natural numbers, i.e. the state space S is N^k for some k . In the case of concurrent programs each transition function corresponds to some recursive subset of N^k . (In fact, it suffices to look at transitions corresponding to assignments of the form $x:=0, x:=y+1, x:=y-1$ and guarded by tests of the form $x=0?$, without affecting the following theorem.) Call such programs "concurrent numerical programs".

In the case of structured nondeterministic programs assume that all functions and relations used in the expressions are recursive (i.e. effectively calculable) and the usual functions and relations of Peano arithmetic are available in the language. Call such programs "nondeterministic numerical programs".

In the subsequent discussion we restrict our attention to nondeterministic numerical programs. Similar results can be proved for concurrent numerical programs. The complexity of fair termination of nondeterministic numerical programs is closely related to the complexity of numerical (nondeterministic) programs with random assignments.

The translation presented in section 3.2 and the converse one replacing $x:=?$ by

```

x:=0; while B do if B → x:=x+1
      □ B → B:=false
      fi
od

```

show that both classes of programs are reducible to each other.

Since the proof rules for fair termination were obtained through the first translation, the ordinals α_p for both classes of programs are in fact the same. In [AP] it was proved that exactly all recursive ordinals are needed to prove total correctness of numerical programs with random assignments. Hence the same result holds for the ordinals α_p associated with nondeterministic numerical programs.

We now prove the following stronger theorem concerning top level fairness only.

THEOREM 2. For any recursive ordinal α there exists a nondeterministic numerical program P with nondeterminism on a top level only, with α_p satisfying $\alpha_p \geq \alpha$.

This theorem should be compared with [AO], where the authors prove an analogous statement for $\alpha < \omega^w$ only.

Proof. We prove that each numerical program with random assignments which is otherwise deterministic is equivalent to a nondeterministic numerical program with top level fairness only. More precisely we show that for each program Π of the first type there exists a nondeterministic numerical program Π_1 with nondeterminism on a top level only such that $M_t[\Pi] = M_{\text{fair}}[\Pi_1]$. The result then follows, since by [AP] exactly all recursive ordinals are needed for proving total correctness of the programs of the first type.

Let Π be a program of the first type. Insert before each random assignment of the form $x:=?$ the assignment $x:=0$. By a well known theorem Π is equivalent to a program Π' which contains one while loop only and makes use of the auxiliary

variable C ranging over labels attached to atomic programs and tests.

Assume that the labels form the set $\{1, \dots, \text{halt}-1\}$ and that \bar{x} is a vector of all variables of Π . Then we can assume that Π' is of the form

```
c:=1 ;  $\bar{x}:=\bar{t}$  ; while  $c \neq \text{halt}$  do
    halt-1
    if  $\square_{i=1} c = i \rightarrow$  execute statement with label  $i$  ; update  $c$ 
    fi .
```

If the statement is a test then its execution is void but updating the counter c is performed accordingly to the value of the test. Replace now each part of the if-fi construct of the form $\square_{i=1} c = i \rightarrow i : x := ? ; \text{update } c$ by $\square_{i=1} c = i \rightarrow x := x + 1 \square_{i=1} c = i \rightarrow \text{update } c$. Call the resulting program Π_1 .

By the construction the value of x just before updating the value of c to i is 0. It is now clear that Π_1 is the required program. \square

Conclusions.

We hope to have shown here that the issue of fair termination admits a simple approach being a straightforward extension of the usual method based on the use of well founded ranking.

The main novelty of the delay variables approach is the augmentation of the states by the delay variables which are not directly manipulated by the program but can be computed for each fair computation sequence. In the case of nondeterministic program the method was justified using program translations but it is certainly a simpler version than any of the previous methods based on such transformations. Its ability to deal naturally with fairness on all levels is again a proof of its strength and appropriateness. As we have shown it provides a unified principle upon which proofs of fairness properties of both nondeterministic and concurrent programs can be based.

The method was applied here to deal with the issue of fairness only. It is however clear that it can be also used to deal with the issue of justice (see [LPS])

by simply refining the way the values of delay variables are changed.

Another advantage of the method is that it can be applied to study the issue of justice and fairness in the context of structured concurrent programs considered e.g. in [OG]. The appropriate proof rules can be obtained by applying appropriate transformations similar to that studied in section 3.2. This subject will be dealt with in another paper which will form a sequel to the present one.

References.

- [AO] Apt, K.R., Olderog, E.R. Proof Rules Dealing with Fairness, in: Workshop on Logic of Programs, Springer Verlag, Lecture Notes in Computer Science, 131 pp. 1-8, 1982. (to appear in Science of Computer Programming).
- [AP] Apt, K.R., Plotkin, G.D. A Cook's Tour of Countable Nondeterminism, in: Proc. 8th Colloquium on Automata Languages and Programming, Acre 1981 Springer Verlag Lecture Notes in Computer Science, 115, pp. 477-493, 1981. (Full version appeared as: Technical Report of Dept. of Computer Science, Edinburgh University, 1980).
- [C] Chandra, A. Computable Non-Deterministic Functions, in: Proc. 19th Annual Symposium on Foundations of Computer Science, pp. 127-131, 1978.
- [D] Dijkstra, E.W. A Discipline of Programming, Prentice Hall, 1976.
- [F] Floyd, R.W. Assigning Meanings to Programs, in: Proc. AMS Symposium in Applied Mathematics 19, pp. 19-31, 1967.
- [FMR] Crumberg, C., Francez, N., Makowsky, J.A., Röver, W.P. A Proof Rule for Fair Termination of Guarded Commands, in: Algorithmic Languages (eds. J.W. de Bakker, J.C. van Vliet), pp. 399-416, IFIP, North Holland, 1981.
- [NP] Hennessy, M.C.B., Plotkin, G.D. Full Abstraction for a Simple Programming Language, in: Proc. 8th Symposium on Mathematical Foundations of Computer Science, Springer Verlag Lecture Notes in Computer Science, 74, pp. 108-120, 1979.
- [H] Hoare, C.A.R. An Axiomatic Basis of Computer Programming, CACM 12(10), 1969.
- [LPS] Lehmann, D., Pnueli, A., Stavi, J. Impartiality, Justice and Fairness: The Ethics of Concurrent Termination, in: Proc. 8th Coll. on Automata Languages and Programming, Springer Verlag Lecture Notes in Computer Science, 115, pp. 264-277, 1981.
- [M] Manna, Z. Mathematical Theory of Computation, McGraw Hill, 1974.
- [MP] Manna, Z., Pnueli, A. Axiomatic Approach to Total Correctness, Acta Informatica 3, pp. 243-263, 1974.
- [OG] Owicki, S., Gries, D. An Axiomatic Proof Technique for Parallel Programs, Acta Informatica 6, pp. 319-339, 1976.