

PFPL Supplement: Modernized Algol, Concurrently*

Robert Harper

Spring, 2023

1 Introduction

The scoped dynamics of Modernized Algol in **PFPL** Chapter 34 is given by transitions of the form

$$\mu \parallel m \xrightarrow[\Sigma]{} \mu' \parallel m'$$

in which Σ determines the assignables in scope, μ determines their contents, and m is any command that is well-formed in that scope. The scope-free dynamics given in Chapter 35 is given by transitions of the form

$$\nu \Sigma \{ \mu \parallel m \} \longmapsto \nu \Sigma' \{ \mu' \parallel m' \}$$

which gives global scope to all assignables, and permits assignables to persist beyond the body of their declaration.

These disparate formulations may be unified by using process calculus to manage the state of a computation. The main idea is to think of the assignable cells as “servers” that interact with the main program during its execution, and to introduce “continuations” to manage the nested control structure of commands.

2 States as Processes

The execution states of **MA** are replaced by *processes* given by the following grammar:

$$p ::= \mathbf{1} \mid p_1 \otimes p_2 \mid \nu a \sim \tau . p \mid a \hookrightarrow v \mid \mathbf{run}[a](m) \mid \mathbf{recv}[a](x . p)$$

The first two process forms represent concurrent composition. The third process form represents the allocation of a channel for use within a process. A memory is a concurrent composition of *cells*, $a \hookrightarrow v$, each of which are processes that respond to requests for their current contents and to signals to change their contents. The program under execution is a *command process*, $\mathbf{run}[a](m)$, where a is a channel on which the returned value is sent upon completion. Sequencing of commands is managed using *continuation processes*, $\mathbf{recv}[a](x . p)$, which accept a value v on channel a from the command process and continue the process $\{v/x\}p$.

The statics of processes is defined in Figure 1, and their dynamics is given in Figure 2.

*© 2023 Robert Harper. All Rights Reserved.

$$\begin{array}{c}
\text{S-STOP} \\
\hline
\Gamma \vdash_{\Sigma} \mathbf{1} \text{ proc}
\end{array}
\quad
\begin{array}{c}
\text{S-CONC} \\
\hline
\frac{\Gamma \vdash_{\Sigma} p_1 \text{ proc} \quad \Gamma \vdash_{\Sigma} p_2 \text{ proc}}{\Gamma \vdash_{\Sigma} p_1 \otimes p_2 \text{ proc}}
\end{array}
\quad
\begin{array}{c}
\text{S-NEW} \\
\hline
\frac{\Gamma \vdash_{\Sigma, a \sim \tau} p \text{ proc}}{\Gamma \vdash_{\Sigma} \nu a \sim \tau . p \text{ proc}}
\end{array}$$

$$\begin{array}{c}
\text{S-MEM} \\
\hline
\frac{\Gamma \vdash_{\Sigma} e : \tau \quad \Sigma \vdash a \sim \tau}{\Gamma \vdash_{\Sigma} a \hookrightarrow e \text{ proc}}
\end{array}
\quad
\begin{array}{c}
\text{S-CMD} \\
\hline
\frac{\Gamma \vdash_{\Sigma} m \dot{\sim} \rho \quad \Sigma \vdash a \sim \rho}{\Gamma \vdash_{\Sigma} \text{run}[a](m) \text{ proc}}
\end{array}
\quad
\begin{array}{c}
\text{S-STK} \\
\hline
\frac{\Gamma, x : \tau \vdash_{\Sigma} p \text{ proc} \quad \Sigma \vdash a \sim \tau}{\Gamma \vdash_{\Sigma} \text{recv}[a](x.p) \text{ proc}}
\end{array}$$

Figure 1: Process Forms: Statics

$$\begin{array}{c}
\text{D-CONC} \\
\hline
\frac{p_1 \xrightarrow[\Sigma]{\alpha} p'_1}{p_1 \otimes p_2 \xrightarrow[\Sigma]{\alpha} p'_1 \otimes p_2}
\end{array}
\quad
\begin{array}{c}
\text{D-SYNC} \\
\hline
\frac{p_1 \xrightarrow[\Sigma]{\alpha} p'_1 \quad p_2 \xrightarrow[\Sigma]{\bar{\alpha}} p'_2}{p_1 \otimes p_2 \xrightarrow[\Sigma]{} p'_1 \otimes p'_2}
\end{array}
\quad
\begin{array}{c}
\text{D-NEW} \\
\hline
\frac{p \xrightarrow[\Sigma, a \sim \tau]{\alpha} p' \quad \vdash_{\Sigma} \alpha \text{ action}}{\nu a \sim \tau . p \xrightarrow[\Sigma]{\alpha} \nu a \sim \tau . p'}
\end{array}$$

$$\begin{array}{c}
\text{D-GET} \\
\hline
\frac{v \text{ val}_{\Sigma}}{a \hookrightarrow v \xrightarrow[\Sigma]{a!v} a \hookrightarrow v}
\end{array}
\quad
\begin{array}{c}
\text{D-SET} \\
\hline
\frac{v' \text{ val}_{\Sigma}}{a \hookrightarrow v \xrightarrow[\Sigma]{a?v'} a \hookrightarrow v'}
\end{array}$$

Figure 2: Process Forms: Dynamics

3 Concurrent Dynamics of **MA**

The concurrent dynamics of **MA** is given in Figure 3. The main feature of the dynamics is that each rule is given in “local form,” in the sense that it need not specify the entire state, but only those parts that are relevant to each command. This is a prime advantage afforded by the use of process calculus for execution states.

An initial state is a command to be executed, $\text{run}[a](m)$ for some channel, a , whose associated type is that of m . A final state is one in which the running command is returning a value; it may be surrounded by any number of cells, but not by another running command or a continuation.

The return and bind commands are managed using the continuation process and a chosen channel on which to communicate results. Execution of a bind executes the encapsulated command, and spawns a continuation process expecting its result. Execution of a return transmits the return value along the channel naming the command process, and the continuation process is prepared to receive that value and continue by passing that value to the body of the bind that created it.

The declaration command allocates a new assignable and spawns a cell process governing it. Cell processes synchronize with the running process along the assignable. A cell process is always prepared to send its current value to the running command, and to receive a new value from the running command. Correspondingly, the get and set commands are executed by posting the appropriate message to obtain the contents of a cell, or to update its contents, respectively. It is impossible for a cell to interact with itself, because processes cannot be duplicated.

The choice of structural congruence for processes is influenced by whether assignables are scoped

or free. The allocation of assignables is formulated in Rule D-DCL, which allocates a fresh assignable and associates a cell process with it. Structural congruence rules are used to manage their scope and their deallocation.

When assignables are scoped, the mobility restriction on the type of their contents ensures that they may only occur within the running command, and not within its returned value or within any other cell. Bearing this in mind, the stack allocation of scoped assignables is expressed by the following rule of structural congruence for processes:

$$\text{POP} \quad \frac{v \text{ val}_{\Sigma} \quad e \text{ val}_{\emptyset}}{\nu a \sim \tau . a \hookrightarrow v \otimes \text{run}[b](\text{ret}[\tau](e)) \equiv_{\Sigma, b \sim \tau} \text{run}[b](\text{ret}[\tau](e))}$$

No other principles of structural congruence are necessary. Rule D-DCL may then be read as allocating a frame on the data stack that is then deallocated by Rule POP.

When assignables are scope-free, they may occur within the contents of any cell, including its own, and may escape the scope of the declaration that introduces it. To account for the global scope of assignables, it is necessary to postulate the following principle of *scope extrusion* for channels:

$$\text{SCOPE} \quad \frac{(a \notin p_2)}{\nu a \sim \tau . \{p_1 \otimes p_2\} \equiv_{\Sigma} \{\nu a \sim \tau . p_1\} \otimes p_2}$$

When read from right-to-left, Rule SCOPE states that the scope of an assignable may be *widened* to include any other active cell, thereby permitting cycles as would arise by backpatching, or in any other form of self-referential data structure such as a circularly linked list.

When read from left-to-right, the same rule states that the scope can be *narrowed* to exclude any process (command, continuation, or cell) that does not mention it. This is important for deallocation of free assignables. For example, the following rule states that any isolatable cell may be reclaimed:

$$\text{DEALLOC-ONE} \quad \frac{v \text{ val}_{\emptyset}}{\nu a \sim \tau . a \hookrightarrow v \equiv_{\Sigma} \mathbf{1}}.$$

A cell process may be isolated using Rule SCOPE to narrow its scope to the extent that it is not used in the running command or continuation, nor in any other cell.

Exercise 3.1. Rule (3) permits the deallocation of a single isolatable cell. Such a cell may be self-referential (construct an example!), but larger cycles are not isolatable in this sense, and hence not collectible. Generalize this rule to permit deallocation of cycles in memory using a similar principle of structural congruence.

The dynamics enjoys a local form of safety analogous to that for Concurrent Algol:

Theorem 3.1 (Weak Safety). *If $\vdash_{\Sigma} p \text{ proc}$, then either $p \text{ final}_{\Sigma}$, or there exists unique $\vdash_{\Sigma} \alpha \text{ action}$ and $\vdash_{\Sigma} p' \text{ proc}$ such that $p \xrightarrow[\Sigma]{\alpha} p'$.*

Exercise 3.2. Prove Theorem 3.1.

Theorem 3.1 permits states (well-formed processes) that cannot make an unlabelled transition (true progress). It may be strengthened to exclude this possibility by observing that the dynamics given in Figure 3 ensures the following additional properties of the state that are not captured by the statics alone:

1. There is exactly one running command, whose name is distinct from any assignable or any continuation.
2. There is exactly one cell process associated to each assignable whose name is distinct from that of any command or continuation process.
3. There are any number of continuation processes connected in a “chain” of the form

$$\mathbf{recv}[a_{k-1}](x_{k-1} . \mathbf{run}[a_{k-1}](m_{k-1})) \otimes \cdots \otimes \mathbf{recv}[a_0](x_0 . \mathbf{run}[a_0](m_0))$$

arising from nested bind commands, each in the first argument of the next.

Intuitively, the chain of stack processes corresponds to a control stack with pending commands arranged in the appropriate order of execution.

Exercise 3.3. *State and prove a strengthening of Theorem 3.1 that ensures progress, as would be expected for **MA**.*

References

- Robert Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, Cambridge, England, Second edition, 2016.
- Robert Harper. Concurrent Algol, Modernized. Supplement to Harper (2016), Fall 2019. URL <https://www.cs.cmu.edu/~rwh/pfpl/supplements/mca.pdf>.

$$\begin{array}{c}
\text{D-INIT} \\
\frac{\vdash_{\emptyset} m \dot{\sim} \tau}{\text{run}[a](m) \text{ initial}_{a \sim \tau}} \\
\\
\text{D-FINAL} \\
\frac{e \text{ val}_{\Sigma}}{\text{run}[a](\text{ret}[\tau](e)) \text{ final}_{\Sigma, a \sim \tau}} \\
\\
\text{D-FINAL-CELL} \\
\frac{p \text{ final}_{\Sigma}}{p \otimes a \hookrightarrow v \text{ final}_{\Sigma, a \sim \tau}} \\
\\
\text{D-FINAL-ALLOC} \\
\frac{p \text{ final}_{\Sigma, a \sim \tau}}{\nu a \sim \tau . p \text{ final}_{\Sigma}} \\
\\
\text{D-RET} \\
\frac{e \text{ val}_{\Sigma}}{\text{run}[a](\text{ret}[\tau](e)) \xrightarrow[\Sigma]{a!e} \mathbf{1}} \\
\\
\text{D-CONT} \\
\frac{v \text{ val}_{\Sigma}}{\text{recv}[a](x.p) \xrightarrow[\Sigma]{a?v} \{v/x\}p} \\
\\
\text{D-BND} \\
\frac{}{\text{run}[a](\text{bnd}[\tau_1](\text{cmd}[\tau_1](m_1); x.m_2)) \xrightarrow[\Sigma]{} \nu a_1 \sim \tau_1 . \{\text{run}[a_1](m_1) \otimes \text{recv}[a_1](x.\text{run}[a](m_2))\}} \\
\\
\text{D-DCL} \\
\frac{v \text{ val}_{\Sigma}}{\text{run}[a](\text{dcl}[\tau](v; b.m)) \xrightarrow[\Sigma]{} \nu b \sim \tau . \{b \hookrightarrow v \otimes \text{run}[a](m)\}} \\
\\
\text{D-GET} \\
\frac{v \text{ val}_{\Sigma}}{\text{run}[a](\text{get}[b]) \xrightarrow[\Sigma]{b?v} \text{run}[a](\text{ret}[\tau](v))} \\
\\
\text{D-SET} \\
\frac{v \text{ val}_{\Sigma}}{\text{run}[a](\text{set}[b](e)) \xrightarrow[\Sigma]{b!v} \text{run}[a](\text{ret}[\tau](v))}
\end{array}$$

Figure 3: Concurrent Dynamics of **MA** (Selected Rules)