

PFPL Supplement: Concurrent Algol, Modernized*

Robert Harper

Spring, 2024

1 Introduction

The technical formulation of Concurrent Algol (**CA**) in Chapter 40 of **PFPL** is an uncomfortable compromise of competing considerations. There is a tension between the desire to maintain the dynamics of **MA** largely intact and the desire to integrate mechanisms for allocating channels, spawning new processes, and emitting and accepting messages.

The dynamics given in **PFPL** makes use of a judgment defining the “local steps” of execution that are integrated into the “global steps” governing the program as a whole, in particular the definition of synchronization among processes. A local step transforming one command into another can have one of three ancillary effects:

1. It can allocate a fresh channel.
2. It can spawn a new process.
3. It can engender a synchronization action.

Each local step determines a global step that integrates the effects into the global state. The method works, but cannot be commended for its elegance.

Another approach is to consider that the global state is a *process* in the sense of the π -calculus. The virtue of this formulation is that all global steps are specified locally, relying on the machinery of process calculus to integrate effects as appropriate. As a bonus the same methods may be used to define the dynamics of **MA**, regarding assignable cells as “servers” that respond to get and put requests; see Harper (2020) for the particulars.

2 Processes for CA

The language of processes for **CA** is essentially that given in **PFPL**, but with the change that all messages are of type `clsfd`, with receive events matching against the class of the message to determine whether they are eligible to process that message. For clarity, the statics of the `emit` primitive and for a running command are given in Figure 1, along with the labelled transition for emitting a value of type `clsfd`, and selectively receiving such a value according to an event with matching class.

*Copyright © Robert Harper. All Rights Reserved.

$$\begin{array}{c}
\text{S-EMIT} \\
\frac{\vdash_{\Sigma} e : \text{clsfd}}{\vdash_{\Sigma} \text{emit}(e) \text{ proc}} \\
\\
\text{S-RUN} \\
\frac{\vdash_{\Sigma} m \dot{\sim} \tau \quad \Sigma \vdash a \sim \tau}{\vdash_{\Sigma} \text{run}[a](m) \text{ proc}} \\
\\
\text{D-EMIT} \\
\frac{e \text{ val}_{\Sigma}}{\text{emit}(e) \xrightarrow{\Sigma} \mathbf{1}} \\
\\
\text{D-AWAIT} \\
\frac{}{\text{await}(\text{rcv}[a](x.P) + \dots) \xrightarrow{\Sigma} \{v/x\}P}
\end{array}$$

Figure 1: Special Process Forms for **CA**

$$\begin{array}{c}
\text{D-RET} \\
\frac{e \Downarrow_{\Sigma} v}{\text{run}[a](\text{ret}(e)) \xrightarrow{\Sigma} \text{emit}(a.v)} \\
\\
\text{D-BND} \\
\frac{e_1 \Downarrow_{\Sigma} \text{cmd}[\tau_1](m_1)}{\text{run}[a](\text{bnd}(e_1; x.m_2)) \xrightarrow{\Sigma} \nu a_1 \sim \tau_1 . \text{run}[a_1](m_1) \otimes \text{await}(\text{rcv}[a_1](x.\text{run}[a](m_2)))} \\
\\
\text{D-NEWCH} \\
\frac{}{\text{run}[a](\text{newch}[\tau]) \xrightarrow{\Sigma} \nu b \sim \tau . \text{run}[a](\text{ret}(\&b))} \\
\\
\text{D-SPAWN} \\
\frac{e \Downarrow_{\Sigma} \text{cmd}[\tau](m)}{\text{run}[a](\text{spawn}(e)) \xrightarrow{\Sigma} \nu b \sim \tau . \text{run}[b](m) \otimes \text{run}[a](\text{ret}(\langle \rangle))} \\
\\
\text{D-SYNC} \\
\frac{e \Downarrow_{\Sigma} v \quad v \equiv \text{or}(\text{rcv}[a_1](x.e_1); \text{or}(\dots; \text{rcv}[a_n](x.e_n)))}{\text{run}[a](\text{sync}(e)) \xrightarrow{\Sigma} \text{await}(\text{rcv}[a_1](x.\text{run}[a](\text{ret}(e_1))) + \dots + \text{rcv}[a_n](x.\text{run}[a](\text{ret}(e_n))))}
\end{array}$$

Figure 2: **CA** Dynamics (Selected Rules)

3 Revised Dynamics of CA

As a technical convenience, the receive event of **CA** is herein replaced by a more general form, written $\text{rcv}[a](x.e)$, which indicates willingness to receive a message with class a , binding its associated value to x within e to determine the resulting value of that event. By choosing e to be x , the original form of receive event is obtained. It is then a simple exercise to define $\text{wrap}(e_1; x.e_2)$ in terms of the other event constructs (nullary and binary choice, the new form of receive). Importantly, **CA** event values are taken modulo the commutativity, associativity, and unit laws for the nullary and binary disjunctions, so that any value of event type is, essentially, a disjunction of receive events of the form just described.

With this in mind the dynamics of **CA** is given in Figure 2. There are several notable aspects of this dynamics. First, in contrast to **PFPL**, the dynamics of **bnd** and **ret** is given in terms of process synchronization. Rule **D-BND** allocates a private channel and creates a **run** process executing the first command in the sequence, then arranges for the second to await the completion of the first by querying the private channel. Correspondingly, a execution of a **ret** command sends the returned value along the channel associated with it. Once the first command has completed, the send synchronizes with the receive, and continues with the appropriate instance of the second command in the sequence.

The concurrency primitives are managed using the corresponding process calculus mechanisms. Spawning an encapsulated command creates a new process, with no destination, returning the unit value to the spawning process.¹ Allocation of a new channel allocates a new channel and returns a reference to it to the allocating process. Emitting a message emits the message concurrently with the return of the unit value to the emitter. Synchronizations turns into process synchronizations, albeit with the continuation being a process that returns the provided value.

References

Robert Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, Cambridge, England, Second edition, 2016.

Robert Harper. Modernized Algol, Concurrently. Supplement to Harper (2016), December 2020. URL <https://www.cs.cmu.edu/~rwh/pfpl/supplements/cma.pdf>.

¹Alternatively, **spawn** could return a reference to the channel that names it, so that another process can await its termination for receiving on that channel.