

Recitation 13: Concurrency

15-312: Foundations of Programming Languages

Serena Wang

April 18th, 2018

1 Modeling Concurrency with Programming Languages

We have previously modelling parallelism in programming languages, and we will now discuss concurrency using another language model. Keep in mind that parallelism is not concurrency!

| Parallelism | Concurrency |
|------------------------|--------------------------|
| Deterministic behavior | Allows non-determinism |
| Focus on dependencies | Focus on synchronization |
| Abstraction | Implementation |

Process calculus is a freestanding language model for concurrency and is the central machinery behind Concurrent Algol, or **CA**. Process calculus is defined by processes, which compute things, and events, which let processes communicate with each other. Chapter 39 in the textbook explains process calculus in detail, but the key idea is that processes in process calculus are identified up to structural congruence. This means that we can re-order processes at will, which is very helpful for implementing non-deterministic dynamic rules. By extending Modernized Algol with process calculus, we get **CA!**.

2 Concurrent Algol

Concurrent Algol, or **CA**, incorporates the typing system and modal distinction from **MA** with the mechanisms of process calculus for dealing with concurrent computation. For simplicity, we don't have assignables anymore in **CA**, but we can actually still define free assignables in **CA** using processes as cells (see the end of Chapter 40 for details).

The syntax for **CA** is derived from **MA**, but without assignables and with a new syntactic level of processes for representing the global state of a program:

| | | | | |
|-----------|----------|-----|---------------------------------|-------------------------------|
| Typ | τ | ::= | cmd | command |
| Exp | e | ::= | cmd (m) | encapsulation |
| Cmd | m | ::= | ret e | return |
| | | | bnd $x \leftarrow e ; m$ | sequence |
| Proc | p | ::= | stop | idle |
| | | | run (m) | atomic |
| | | | $p_1 \otimes p_2$ | concurrent |
| | | | $\nu a \sim \tau . p$ | creation of new channels |
| Signature | Σ | ::= | \cdot | empty |
| | | | $\Sigma, a \sim \tau$ | type of values in channel a |
| Action | α | ::= | ϵ | none |
| | | | $a ! e$ | send |
| | | | $a ? e$ | receive |

All of the **MA** constructs remain, but we now also have the ability to define and run commands that declare variables and assign to them. The modal separation ensures that expressions do not directly cause the effects of variable state update, but commands when they are executed may do so. **CA** thus introduces new concepts such as processes, messages, actions, events, and dynamic classification in combination with the previously seen separation between pure and impure computation.

In **CA**, processes are the outermost level of computation. You can execute commands in processes and have concurrent subprocesses. You can also create new channels within processes with $\nu a \sim \tau . p$ such that the new channel a can be used only within process p .

3 Shared Skeleton

There are different approaches that we can take to construct Concurrent Algol depending on the type of communication we want in our concurrent system. However, they all extend this shared skeleton to provide different types of communication. In the shared skeleton, actions then refer to sending and receiving messages at runtime.

3.1 Statics

For the statics of **CA**, we have four judgments:

| | |
|--|-------------------|
| $\Gamma \vdash_{\Sigma} e : \tau$ | expression typing |
| $\Gamma \vdash_{\Sigma} m \dot{\sim} \tau$ | command typing |
| $\Gamma \vdash_{\Sigma} p \text{ proc}$ | process formation |
| $\Gamma \vdash_{\Sigma} \alpha \text{ action}$ | action formation |

We have these rules for process formation. Again, processes are identified up to structural congruence.

$$\frac{}{\Gamma \vdash_{\Sigma} \text{stop proc}} \quad \frac{\Gamma \vdash_{\Sigma} m \dot{\sim} \tau}{\Gamma \vdash_{\Sigma} \text{run}(m) \text{ proc}} \quad \frac{\Gamma \vdash_{\Sigma} p_1 \text{ proc} \quad \Gamma \vdash_{\Sigma} p_2 \text{ proc}}{\Gamma \vdash_{\Sigma} p_1 \otimes p_2 \text{ proc}} \quad \frac{\Gamma \vdash_{\Sigma, a \sim \tau} p \text{ proc}}{\Gamma \vdash_{\Sigma} \nu a \sim \tau . p \text{ proc}}$$

We have these rules for action formation. Notice that messages don't necessarily need to contain values.

$$\frac{}{\Gamma \vdash_{\Sigma} \epsilon \text{ action}} \quad \frac{\Gamma \vdash_{\Sigma, a \sim \tau} e : \tau}{\Gamma \vdash_{\Sigma, a \sim \tau} a ! e \text{ action}} \quad \frac{\Gamma \vdash_{\Sigma, a \sim \tau} e : \tau}{\Gamma \vdash_{\Sigma, a \sim \tau} a ? e \text{ action}}$$

The statics for the expressions and commands stay the same as in **MA**.

3.2 Dynamics

For the dynamics of **CA**, since we can now step expressions, step processes, and spawn processes in commands, we have three judgments. Notice that processes and commands can produce actions but expressions cannot.

$$\begin{array}{ll} e \xrightarrow[\Sigma]{} e' & \text{expression stepping} \\ p \xrightarrow[\Sigma]{} p' & \text{process stepping} \\ m \xrightarrow[\Sigma]{} m' & \text{command stepping} \end{array}$$

We have these rules for stepping processes.

$$\frac{m \xrightarrow[\Sigma]{} m'}{\text{run}(m) \xrightarrow[\Sigma]{} \text{run}(m')} \quad \frac{m \xrightarrow[\Sigma]{} \nu \Sigma' \{m' || p\}}{\text{run}(m) \xrightarrow[\Sigma]{} \nu \Sigma' \{\text{run}(m') \otimes p\}} \quad \frac{e \text{ val}_{\Sigma}}{\text{run}(\text{ret}(e)) \xrightarrow[\Sigma]{} \text{stop}} \quad \frac{p_1 \xrightarrow[\Sigma]{} p'_1}{p_1 \otimes p_2 \xrightarrow[\Sigma]{} p'_1 \otimes p_2}$$

$$\frac{p_1 \xrightarrow[\Sigma]{} p'_1 \quad p_2 \xrightarrow[\Sigma]{} p'_2}{p_1 \otimes p_2 \xrightarrow[\Sigma]{} p'_1 \otimes p'_2} \quad \frac{p \xrightarrow[\Sigma, a \sim \tau]{} p' \quad \vdash_{\Sigma} \alpha \text{ action}}{\nu a \sim \tau. p \xrightarrow[\Sigma]{} \nu a \sim \tau. p'}$$

Notice in the second rule that the command m created a new channel contained in the signature Σ' and also spawned a new process p . We then still run the process for m' and run the new process p in parallel.

In the third rule, we see p_1 send a message that p_2 receives. Because of this rule, we have synchronous communication in this system. We need p_1 and p_2 to be on the same page before we can send or receive a message.

In the last rule, we step the process p within the signature where we can use the channel a . However, we check that the action α does not use channel a to prevent stealing of messages. Any message that gets leaked through α can only use channels that already exist in the external signature Σ . Since a is only bound within p , the only process that can use the channel a is p . Thus, if p has some external action (sending or receiving a message), that action should not use channel a .

4 Dynamic Classification for Messages

Messages are values that are sent or received by processes. A message consists of a channel, which is its class and a payload, which is a value of the type associated with the channel. Messages thus have the type `clsfd`, and channels are classes of messages. An event is the type of combined channels.

| | | |
|-----|-------------------------------------|----------------------------|
| Typ | $\tau ::= \text{clsfd}$ | type of values in messages |
| | $\text{cls}(\tau)$ | type of channels |
| Exp | $e ::= \&a$ | channel |
| | $e_1 \cdot e_2$ | message |
| | $\text{isof}(e_0; e_2; x.e_2; e_3)$ | check message |

We now have expressions for dynamically creating classes and classified values of type `clsfd`. Using a fresh channel a created by $\nu a \sim \tau.p$, we can use the $\&a$ expression to refer to that channel, and we can use $\&a \cdot e$ to attach the tag a with the expression e . $\text{isof}(e_0; e_2; x.e_2; e_3)$ checks whether e_0 is tagged with the class given by e_1 , evaluating e_2 if so and e_3 otherwise. We can thus use $e_1 \cdot e_2$ to create payloads for messages and use $\text{isof}(e_0; e_2; x.e_2; e_3)$ to check message payloads. Refer to Chapter 33 for full statics rules for dynamic classification in general.

5 Synchronous Broadcast Communication

In addition to all of the constructs defined above, we also add some new commands for sending and receiving messages to get synchronous broadcast communication.

| | | |
|-----|-------------------------|----------------|
| Cmd | $m ::= \text{spawn}(e)$ | spawn |
| | $\text{emit}(e)$ | emit message |
| | acc | accept message |
| | $\text{newch}\{\tau\}$ | new channel |

To have synchronous communication, we use commands to send messages. In asynchronous communication, we would use processes to send messages. The key point is that a command can send or receive messages on a channel if and only if they have access to that class of messages. By dynamically generating tags for message payloads, we create a secret-sharing system by wrapping and unwrapping expressions inside message payloads.

5.1 Statics

Refer to Chapter 33 in the textbook for full statics rules for dynamic classification in general. The typing rules for expressions in the synchronous broadcast version of **CA** follows the statics rules for dynamic classification exactly.

The statics rules for commands in this programming language are as follows:

$$\frac{\Gamma \vdash_{\Sigma} e : \text{cmd}(\text{unit})}{\Gamma \vdash_{\Sigma} \text{spawn}(e) \dot{\sim} \text{unit}} \quad \frac{\Gamma \vdash_{\Sigma} e : \text{clsfd}}{\Gamma \vdash_{\Sigma} \text{emit}(e) \dot{\sim} \text{unit}} \quad \frac{}{\Gamma \vdash_{\Sigma} \text{acc} \dot{\sim} \text{clsfd}} \quad \frac{}{\Gamma \vdash_{\Sigma} \text{newch}\{\tau\} \dot{\sim} \text{cls}(\tau)}$$

Note that since all messages have type `clsfd`, you can only send and receive messages of type `clsfd`.

5.2 Dynamics

Since commands can now spawn processes, we have a new judgment for when the command m transitions to the command m' while creating new channels Σ' and new processes p' .

$$m \xrightarrow[\Sigma]{\alpha} \nu \Sigma' \{m' \otimes p'\}$$

We have these rules for stepping commands in the synchronous broadcast version of **CA**.

$$\frac{m \xrightarrow[\Sigma]{\alpha} m'}{\text{spawn}(\text{cmd}(m)) \xrightarrow[\Sigma]{\xi} \nu \cdot \{\text{ret } \langle \rangle \otimes \text{run}(m)\}} \quad \frac{e \mapsto e'}{\text{spawn}(e) \xrightarrow[\Sigma]{\xi} \text{spawn}(e')} \quad \frac{}{\text{emit}(a \cdot e) \xrightarrow[\Sigma, a \sim \tau]{a!e} \text{ret } \langle \rangle}$$

$$\frac{a \cdot e \text{ val}_{\Sigma, a \sim \tau}}{\text{acc} \xrightarrow[\Sigma, a \sim \tau]{a?e} \text{ret } a \cdot e} \quad \frac{}{\text{newch}\{\tau\} \xrightarrow[\Sigma]{\xi} \nu a \sim \tau \{\text{ret } \& a \otimes \text{stop}\}}$$

Note that once we get a command inside **spawn**, we create a new process for running that encapsulated lazy computation. **spawn**(cmd(m)) doesn't create new any channels, however, so there are no new channels to put in a signature.

ret $\langle \rangle$ means that we stop executing the command and return unit. Thus, we step to this once we spawn a new process or finish sending a message.

We also see in the rules for **emit**($a \cdot e$) and **acc** that you must have a in your signature (i.e. know the channel a) to be able to send or receive a message on channel a . Since only commands within processes that have access to a can send messages on a , we are able to maintain an integrity invariant for that channel. Since only commands within processes with access to a can listen to that channel, we also have a confidentiality property for the channel a .

5.3 Safety

The preservation theorem for **CA** ensures that well-typed programs remain well-typed during execution. The actual proof for this theorem also requires a lemma about command execution, which you can see in Chapter 40 of the textbook, and some strengthening to deal with actions.

Theorem 40.2 (Preservation). If $\vdash_{\Sigma} p \text{ proc}$ and $p \xrightarrow[\Sigma]{\alpha} p'$, then $\vdash_{\Sigma} p' \text{ proc}$.

However, our preservation theorem does not guarantee progress if there's a process waiting to receive a message on an empty channel or if there's a process trying to send a message on a channel no one is listening to. Thus, our progress theorem states that the only way to "get stuck" is to wait for another process to have a complementary action on some external channel.

Theorem 40.2 (Progress). If $\vdash_{\Sigma} p \text{ proc}$, then either $p \equiv \text{stop}$ or $p \equiv \nu \Sigma' \{p'\}$ where $p' \xrightarrow[\Sigma, \Sigma']{\alpha} p''$ for some $\vdash_{\Sigma, \Sigma'} p'' \text{ proc}$ and some $\vdash_{\Sigma, \Sigma'} \alpha \text{ action}$.

6 Asynchronous Broadcast Communication

To allow asynchronous communication in **CA**, we simply add a way to send messages in processes.

$$\text{Proc } p ::= \text{send}[a](e) \text{ send message}$$

We still have everything previously mentioned in the shared skeleton and synchronous broadcast version of **CA**. The main change is that because we allow asynchronous communication, we can now step a process that sends a message without waiting for a complementary action for receiving that message.

6.1 Statics

$$\frac{\vdash_{\Sigma, a \sim \tau} e : \tau}{\vdash_{\Sigma, a \sim \tau} \text{send}[a](e) \text{ proc}}$$

Again, you can only send messages on channel a if you have access to channel a (i.e. having a in Σ).

6.2 Dynamics

$$\frac{}{\text{send}[a](e) \xrightarrow[\Sigma]{a!e} \text{stop}} \quad \frac{a \cdot e \text{ val}_{\Sigma}}{\text{emit}(a \cdot e) \xrightarrow[\Sigma]{\epsilon} \nu \cdot \{\text{ret } \langle \rangle \otimes \text{send}[a](e)\}}$$

The first rule above tells us that sending messages is now asynchronous. $\text{send}[a](e)$ can step to a stopped process without waiting for another process to receive a message on the same channel a .

What does the second dynamics rule tell us? Instead of the `emit` command resulting in an action explicitly, the `emit` command now spawns a concurrent process that sends the message.

7 Synchronous Selective Communication

Broadcast communication does not give us a way to receive messages of just a particular class. If we want an ability to wait for specific events, we need to add some new constructs to **CA**. Specifically, we'll take out our receive commands and replace them with receive expressions. This version of **CA** matches the programming language that you'll work with and implement in Assignment 6 almost exactly. The version of **CA** in Assignment 6 also includes channel references, whose rules you can see in the writeup.

| | | | | |
|-----|--------|-----|--------------------------------|---------------------------|
| Typ | τ | ::= | <code>event</code> (τ) | type of combined channels |
| Exp | e | ::= | <code>rcv</code> [a] | receive event |
| | | | <code>never</code> { τ } | null event |
| | | | <code>or</code> ($e_1; e_2$) | binary disjunction event |
| Cmd | m | ::= | <code>sync</code> (e) | wait for event |

`rcv`[a] allows us to receive messages on channel a , while `or`($e_1; e_2$) allows us to listen to multiple channels if the messages on those channels have the same type. `never`{ τ } is for when you don't

listen to any channel. This can technically be left out so that you're always listening to at least one channel for an event, and it won't be part of the language for Assignment 6.

7.1 Statics

$$\frac{\Gamma \vdash_{\Sigma} e : \mathbf{event}(\tau)}{\Gamma \vdash_{\Sigma} \mathbf{sync}(e) \dot{\sim} \tau} \quad \frac{}{\Gamma \vdash_{\Sigma, a \sim \tau} \mathbf{rcv}[a] : \mathbf{event}(\tau)} \quad \frac{}{\Gamma \vdash_{\Sigma} \mathbf{never}\{\tau\} : \mathbf{event}(\tau)}$$

$$\frac{\Gamma \vdash_{\Sigma} e_1 : \mathbf{event}(\tau) \quad \Gamma \vdash_{\Sigma} e_2 : \mathbf{event}(\tau)}{\Gamma \vdash_{\Sigma} \mathbf{or}(e_1; e_2)}$$

Note that the event type for $\mathbf{rcv}[a]$ depends on the type of the channel a in the signature.

7.2 Dynamics

$$\frac{}{\mathbf{rcv}[a] \mathbf{val}_{\Sigma, a \sim \tau}} \quad \frac{}{\mathbf{never}\{\tau\} \mathbf{val}_{\Sigma}} \quad \frac{e_1 \mathbf{val}_{\Sigma} \quad e_2 \mathbf{val}_{\Sigma}}{\mathbf{or}(e_1; e_2) \mathbf{val}_{\Sigma}} \quad \frac{e_1 \xrightarrow{\Sigma} e'_1}{\mathbf{or}(e_1; e_2) \xrightarrow{\Sigma} \mathbf{or}(e'_1; e_2)}$$

Note that for $\mathbf{rcv}[a]$ to be a value, the channel a must be in the signature.

$$\frac{e \mathbf{val}_{\Sigma, a \sim \tau} \quad \vdash_{\Sigma, a \sim \tau} e : \tau}{\mathbf{sync}(\mathbf{rcv}[a]) \xrightarrow{\Sigma} \mathbf{ret}(e)} \quad \frac{\mathbf{sync}(e_1) \xrightarrow{\Sigma} m_1}{\mathbf{sync}(\mathbf{or}(e_1; e_2)) \xrightarrow{\Sigma} m_1} \quad \frac{\mathbf{sync}(e_2) \xrightarrow{\Sigma} m_2}{\mathbf{sync}(\mathbf{or}(e_1; e_2)) \xrightarrow{\Sigma} m_2}$$

The first rule tells us that an acceptance on channel a may synchronize only with messages classified by a . We shouldn't just receive any message on any channel. The synchronization steps only to expressions tagged with the class a (i.e. expressions sent on the channel a).

The second and third rules tells us that either event between the two choices can lead to an action. The synchronization is satisfied if we receive a message in either of the two choices.