

Egon Börger (Pisa) & Alexander Raschke (Ulm)

Communicating ASMs

illustrated by modeling monitoring network runs

Università di Pisa, Dipartimento di Informatica, boerger@di.unipi.it
Universität Ulm, Abteilung Informatik, alexander.raschke@uni-ulm.de

See Ch. 3.3 of Modeling Companion
<http://modelingbook.informatik.uni-ulm.de>

Goal of the lecture

- *define communicating ASMs*, i.e. multi-agent ASMs without shared memory whose actions are
 - either **local actions**, affecting only each agent's local state
 - or **inter-process communication** actions, i.e. sending/receiving msgs
- *illustrate their use* to model monitoring network runs (i.e. runs of concurrent ASMs which communicate with their neighbors) for
 - concurrent leader election
 - GRAPHLEADELECT (Exl.1)
 - termination detection of diffusing system runs
 - TERMINATIONDETECTOR(\mathcal{D}) (Exl.2)
 - TERMINATIONDETECTOR(GRAPHLEADELECTDIFFUSE)
 - concurrent (asynchronous) emulation of synchronous process runs
 - CONCURSYNCEMULATOR(*Process*, *Edge*) (Exl.3)

Definition of communicating ASMs

A system of *communicating ASMs* is defined as multi-agent ASM of components $p = (ag(p), pgm(p), mailbox(p))$ for $p \in Process$ where:

- the signatures are pairwise disjoint so that each agent has its own private state, also called *internal state* or *local state*,
- each agent is enriched by a *mailbox* for incoming messages,
- each program $pgm(p)$ may contain, besides the usual ASM constructs, the abstract communication actions $SEND(message)$ and $CONSUME(message)$ and the *Received(message)* predicate.

Usually *Process* is assumed to be finite (unless otherwise stated).

The notion of run is that of concurrent ASM runs. Due to the absence of shared locations (besides *mailbox* which p shares with the communication medium), in such asynchronous runs assume wlog that

each step is an atomic ASM **READ&WRITE** step

including independent actions $SENDTo/RECEIVEFrom$ a *mailbox*.

Keeping communication actions abstract

We use communication constructs which deliberately abstract from communication channels. Their intended interpretation is as follows:

- **SEND**(m , **to** q) means to transfer the message (m , **from self, to** q) to the communication medium whose job is to deliver it to $mailbox(q)$
- **Received**(msg) **iff** $msg \in mailbox(\mathbf{self})$
 - i.e. msg has been delivered to its destination by an interaction bw communication medium and receiver or some delegate, etc.
 - leaving to specify when/how to retrieve msgs from the mailbox
- **CONSUME**(msg) = **DELETE**(msg , $mailbox(\mathbf{self})$)
- $mailbox$ (or $inbox$, $outbox$) treated as set, unless otherwise stated
 - e.g. as multiset, FIFO-queue, priority queue, etc.

The components of a $msg = (m, \mathbf{from} p, \mathbf{to} q)$ are extracted by functions $payload(msg)$ (msg content), $sender(msg)$, $receiver(msg)$.

Some variations of delivery

Depending on properties of the communication medium:

- *Immediate reliable communication*: every message sent in one 'step' (e. g. in a synchronous round) is in the receiver's mailbox at the beginning of the next step (e. g. in the next synchronous round).
- *Reliable communication*: every message sent in one 'step' eventually arrives in the receiver's mailbox.
- *Eventually reliable communication* (asynchronous computation model) where either no message is lost or where multiple delivery attempts (message repetition) are performed, assuming that at least one of them eventually succeeds.
- *Lossy uncorrupted communication* (in the asynchronous model), where messages can get lost but not corrupted.
- *Lossy corrupted communication* (in the asynchronous computation model), where messages can get lost or be delivered corrupted.

A notational abbreviation

if $Received(msg(params))$ **then** $M(msg(params))$

abbreviates:

if there is some $msg(params) \in mailbox$ **then**

choose $m \in \{msg(params) \mid msg(params) \in mailbox\}$ **in** $M(m)$

When the order of messages is relevant, we still abstract notationally from the order-reflecting *next* function, which retrieves the next message from the *mailbox*. In that case

if $Received(msg(params))$ **then** $M(msg(params))$

stands for the following rule:

if $m = next(mailbox)$ **and for some** $params$ $m = msg(params)$
then $M(m)$

Disjoint signature assumption

- Communicating ASMs $(a, M, mailbox(a))$, $(b, N, mailbox(b))$ may have the same program $M = N$.
 - To guarantee disjoint locations (read: local states) we assume all function symbols f as implicitly parameterized ('instantiated') by the executing agents a, b in the form f_a, f_b .
 - This parameterization to partition states is used in full generality by what we call ambient ASMs (see Ch.4).
- Communicating ASMs may have input or output locations, but those are used only for providing input or output from/to the environment (if any) and not for inter-process communication (in case the environment is not seen as an additional process).

Ex1.1: Concurrent Leader Election Requirements

Typical *LeaderElectionRequirements*:

PlantReq. Consider a network of finitely many linearly ordered *Processes* without shared memory, located at the nodes of a directed connected graph and communicating asynchronously with their neighbors (only).

FunctionalReq. Design and verify a distributed algorithm whose execution lets every process know the leader.

Corresponding signature elements and constraints:

- Finite connected graph $(Process, Edge)$ (static) with sets $Neighb(p)$
 - set of neighbors q linked to $p \in Process$ by $(p, q) \in Edge$
- linear order $<$ (static) of *Processes* (NB. processes used as nodes)
- $mailbox(p)$, shared with communication medium, for each $p \in Process$; no other shared locations
- $SEND(msg, \mathbf{from} p, \mathbf{to} q)$ implies $q \in Neighb(p)$
- $msg \in mailbox(p)$ implies $p \in Neighb(sender(msg))$

Design idea for LeaderElectionReq

A typical idea to design such an algorithm (see Lynch 1996, Sect.15.2):


- *StepReq*. Every process maintains a record, say *cand*, of the greatest process it has seen so far, initially its own. It alternates between:
 - sending *cand* to all its neighbors
 - updating *cand* in case its knowledge has been improved by receiving a message (say *curMsg*) with larger value from some neighbor.

Corresponding additional signature elements for each $p \in Process$, all three controlled by p :

- $mode \in \{send, check\}$
- $cand \in Process$
- $curMsg$

A system GRAPHLEADELECT of communicating ASMs

We use process names $p \in Process$ as agent name and equip each p with an instance of the communicating ASM program **LEADELECT** defined by the following control state ASM which implements the above design idea:¹



$$\mathbf{GRAPHLEADELECT} = (M_p)_{p \in Process}$$

$$\mathbf{where} \ M_p = (p, \mathbf{LEADELECT}_p, mailbox_p)$$

¹ Figure © 2003 Springer-Verlag Berlin Heidelberg, reprinted with permission

Components and predicates of LEADSELECT

Initialization by:

$mode = send$ **and** $cand = self$ **and** $mailbox = \emptyset$

PROPOSE = **forall** $q \in Neighb$ SEND($cand$, **to** q)

RECEIVEMSG =

choose $msg \in mailbox$

$curMsg := payload(msg)$ -- NB. msg payload is a process

CONSUME(msg)

$KnowledgeImproved$ iff $curMsg > self$

UPDATEKNOWL = ($cand := curMsg$) -- 'increase' of $cand$

NB. $mode$, $cand$, $mailbox$, $Neighb$, $curMsg$ are instantiated (read: implicitly parameterized) by the executing process **self**.

Termination property of GRAPHLEADELECT

Correctness Lemma. In every properly initialized concurrent GRAPHLEADELECT run with reliable communication, fair and not infinitely lazy components—i. e. every enabled process will eventually make a move and every $msg \in mailbox$ will eventually be chosen by RECEIVEMSG—eventually for every $p \in Process$ holds:

- $cand = \max(Process)$ w. r. t. $<$ (everybody ‘knows’ the leader).
- $mailbox = \emptyset$ (there is no more communication)
- $mode = check$

Proof. Follow in the given run R the propagation of $cand = Max$

- which holds in the initial state S_0 of R for $Max = \max(Process)$ via PROPOSE-steps to *Neighbors* along paths through which any given p is reachable from Max .

Formally we proceed by induction on the minimal path-length n connecting Max to p .

GRAPHLEADELECT correctness proof

Decompose R into initial segments $InitSegm_0 = [S_0]$ and $InitSegm_n = [S_0, S_n]$ (of minimal length for $n > 0$) such that:

- for each $n > 0$ and each $p \neq Max$ that is reachable from Max by a path of minimal length n , before reaching S_n process p did:
 - have a msg with $payload(msg) = Max$ in its mailbox
 - choose a msg with payload Max to RECEIVEMSG, UPDATEKNOWL by Max and PROPOSE Max to its Neighbors

Lemma. S_n is well-defined for each n and eventually $S_n = S_{n+1}$.

Proof. For $n = 1$ holds $p \in Neighb(Max)$ so that eventually Max will PROPOSE to p a msg with $payload(msg) = Max$ so that p eventually chooses msg to RECEIVEMSG, to UPDATEKNOWL by Max and to PROPOSE Max to its Neighbors. For $n > 1$, for p (if there is some, otherwise $S_n = S_{n+1}$) apply induction hypothesis to a q that is reachable from Max by a path of minimal length n with $p \in Neighb(q)$.

GRAPHLEADELECT correctness proof (Cont'd)

Corollary 1. For all $n > 0$, Max and every p that is reachable from Max by a path of length $\leq n$:

- has $can_d = Max$ in state S_n and maintains it in the rest of R
- SENDs in R no msg neither in nor after state S_n
- in or in states after S_n is in $mode = check$ when its $mailbox = \emptyset$

Proof 1. Use that since PROPOSE is unguarded, in R every $p \in Process$ always returns eventually to $mode = check$.

Corollary 2. For some k , the algorithm reaches S_k from where eventually it reaches a final state S of R in which every $p \in Process$ is in $mode = check$ and has an empty $mailbox$.

Proof 2. Follows from the finiteness of the graph so that for some $k > 0$ every $p \neq Max$ is reachable from Max by a path of length $\leq k$.

On the meaning of 'let every process know the leader'

Problem with the interpretation of *FunctionalRequirement* by:

eventually every GRAPHLEADELECT-run terminates in a state where for every $p \in Process$ $cand = Max$ holds.

- every $p \neq Max$ can recognize eventually that it is not the leader, namely when its $cand$ assumes a value $cand \neq p$
- no process knows when the run terminates, so no p (not even Max) can recognize when its $cand$ has the correct leader value $cand = Max$

A solution:

- we refine GRAPHLEADELECT to a 'diffusing' concurrent ASM GRAPHLEADELECTDIFFUSE
- we define for every diffusing system \mathcal{M} TERMINATIONDETECTOR rules which permit to monitor without any central control every \mathcal{M} -run to recognize when this run terminates

Exl.2: Termination Detector for Diffusing Computations

A **diffusing system** (Dijkstra/Scholten 1980) is a communicating ASM for which each of its runs satisfies the following:

- the run starts in a state where all system components are *quiescent* (read: no step is enabled so that no local action can happen), say in *mode = idle*
- the run is started by the environment which enables exactly one system component by sending a message to this component to become the *master* of the run—only once until the run terminates (if it terminates at all)
- every other system component can be enabled in the run only by receiving a message from some system component
- the run terminates if it reaches a state in which all components are again quiescent, say in *mode = idle*

NB. An ASM M is quiescent in a state if the disjunction of all rule guards of M is false in this state.

Termination detection idea for diffusing system runs

Let $\mathcal{M} = (m)_{m \in Machine}$ be a diffusing communicating ASM.

Idea: Since components are assumed to be enabled only by receiving some message (assuming reliable msg passing), namely

- either a ‘monitor’ msg, say *Start*, from the environment
 - sent by a `BEGINENDSHELL(\mathcal{M})` env program
- or an ‘ \mathcal{M} -internal’ msg from another component

it suffices to **monitor acknowledgements of \mathcal{M} -internal msgs:**

- require for every sent \mathcal{M} -internal msg an acknowledgement
- check, when a component becomes quiescent, that for every sent \mathcal{M} -internal msg an acknowledgement has been received

Then the *master*’s role is to monitor entering and exiting the first of the resulting spanning trees.

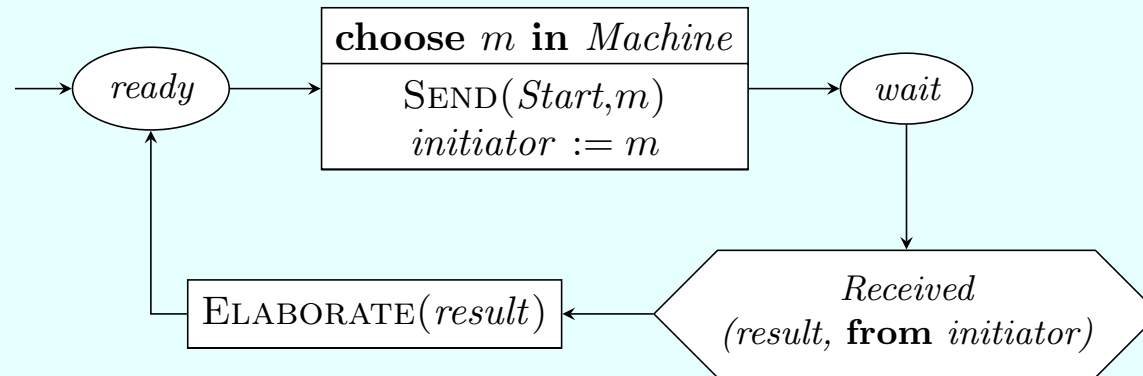
Environment program $\text{BEGINENDSHELL}(\mathcal{M})$

Let $\mathcal{M} = (m)_{m \in \text{Machine}}$ be a diffusing communicating ASM. We extend it to a machine $\text{TERMINATIONDETECTOR}(\mathcal{M})$ by

- adding an env program $\text{BEGINENDSHELL}(\mathcal{M})$ to trigger \mathcal{M} -runs
- extending each $\text{pgm}(m)$ to $\text{pgm}(m)^*$ by components
 - which react to an env trigger and monitor send/receive actions

Define $\text{BEGINENDSHELL}(\mathcal{M})$ as the following control state ASM:

- located at a distinguished graph node without incoming edges²



² Figure © 2018 Springer-Verlag Berlin Heidelberg, reprinted with permission

Master Start/Termination rules

START *Diffuse* =

if *status*(**self**) = *idle* **and** *Received*(*Start*, **from** *env*) **then**

status(**self**) := *master*

START(**self**) -- assumed to enable **self**

CONSUME(*Start*, **from** *env*)

TERMINATE *Diffuse* =

if *status*(**self**) = *master* **and** *quiescent*(**self**)

and *AllAcksArrived*(**self**) **then**

status(**self**) := *idle*

SEND(*computationResult*, *env*) -- payload of termination msg

NB. Initially each component is assumed to be in *status* = *idle*, with empty *mailbox* and without msgs in the communication medium.

Spanning tree requirements

- *SpanningTreeBuildReq*. When a non-master component becomes enabled, namely by receiving in *status = idle* an \mathcal{M} -internal (a non-ack) message, it designates the message sender—a neighbor—as its parent in a new spanning tree (`ENTERSPANNINGTREE`). Any further received \mathcal{M} -internal message is immediately acknowledged (also by the master).
- *ConvergeCastReq*. When a non-master component becomes quiescent and has received an ack for every \mathcal{M} -internal message it had sent out, it will `EXITSPANNINGTREE` and send an ack to its parent (thus acknowledging the msg which triggered building this spanning tree).
- *TerminationReportReq*. When the master becomes quiescent and all the \mathcal{M} -internal messages it had sent—some of which enabled other components—have been acknowledged, then it will report to the environment that the computation did terminate.

NB. Components may be enabled & disabled multiple times during a run.

Rules to monitor sending \mathcal{M} -internal messages

Each m has for each other m' a counter $msgToBeAckBy(m')$ of the number of \mathcal{M} -internal msgs sent by m to m' but not yet acknowledged by m' .

- $msgToBeAckBy(m')$ is increased at each **SEND**(msg , **to** m')

MONITORSENTMSG(msg , **to** receiver) = -- for \mathcal{M} -internal msg
 INCREASE($msgToBeAckBy(receiver)$)

- $msgToBeAckBy(m')$ is decreased when **Received**(ack , **from** m')

MONITORACKMSG = -- decrease counter of expected acks
if **Received**(ack , **from** receiver) **then**
 DECREASE($msgToBeAckBy(receiver)$)
 CONSUME(ack , **from** receiver)

Spanning tree rules to monitor msg acknowledgement

```
MONITORRECEIVEDMSG(msg, from sender) = --  $\mathcal{M}$ -internal msg  
if status(self) = idle  
  then ENTERSPANNINGTREE(sender) -- get enabled  
  else SEND(ack, sender) -- immediate ack if already woken up  
  CONSUME(msg, from sender)  
  where ENTERSPANNINGTREE(sender) = -- NB. no ack sent  
    status(self) := treeNode      parent(self) := sender  
EXITSPANNINGTREE = if status(self) = treeNode and  
quiescent(self) and AllAcksArrived(self) then  
  SEND(ack, parent(self)) -- Ack msg which created spanning tree  
  parent(self) := undef      status(self) := idle  
  where AllAcksArrived(self) iff  
    forall  $m \in \text{Machine}$  msgToBeAckByself( $m$ ) = 0
```

TERMINATIONDETECTOR(\mathcal{M}) for diffusing $\mathcal{M} = (m)_{m \in Machine}$

BEGINENDSHELL(\mathcal{M})
 $(pgm(m)^*)_{m \in Machine}$

where $pgm(m)^*$ is obtained by adding to $pgm(m)$ the following rules:

START *Diffuse* -- just once if triggered by the environment
MONITORSENTMSG(*msg*, **to** *receiver*) -- record expected acks
in parallel to any occurrence of SEND(*msg*, **to** *receiver*) in $pgm(m)$
MONITORACKMSG -- when ack arrives decrease expected acks
MONITORRECEIVEDMSG(*msg*, **from** *sender*) -- wake up or do ack
in parallel to each occurrence of RECEIVE(*msg*, **from** *sender*)
or action triggered by *Received*(*msg*, **from** *sender*) in $pgm(m)$
EXITSPANNINGTREE -- only quiescent tree nodes if *AllAcksArrived*
TERMINATE *Diffuse* -- only master if quiescent & *AllAcksArrived*

TERMINATIONDETECTOR(\mathcal{M}) Lemma

In every TERMINATIONDETECTOR(\mathcal{M}) run every state satisfies:

- For every machine $m \neq initiator$: $status = idle$ iff $parent = \mathbf{undef}$ and in that case (read: if m is not a node of the spanning tree) m has not to wait for any message to be acknowledged (formally expressed $msgToBeAckBy(m')_m = 0$ for each $m' \in Machine$).
 - The machines linked by a *parent* path to the root *initiator* form a spanning tree of all machines with $status \neq idle$.
-

The **initialization** is defined s.t. for every $m \in Machine$:

- m quiescent in $status(m) = idle$ with empty $mailbox(m)$
- no sent but not yet received message in the communication medium and no to-be-acknowledged msg
 - i. e. for each $m' \in Machine$ holds $msgToBeAckBy(m')_m = 0$)
- no parent defined ($parent(m) = \mathbf{undef}$)
- initial mode is *ready*

Applying TERMINATIONDETECTOR to GRAPHLEADELECT

To make GRAPHLEADELECT diffusing, every p which—by having received a *Start* msg (initially from the environment)—is enabled, namely by updating $mode(p) = terminated$ to $mode = send$, will also STARTNEIGHBORS and PROPAGATESTART

- so that eventually the leader *Max* enters $mode = send$.

Therefore *Start* is treated as GRAPHLEADELECT-internal msg so that all *Neighbors* of p MONITORRECEIVEDMSG(*Start*, **from** p).

- in $START_{Diffuse}$ refine **START****self** to:

$mode := send$ -- enabling **self**

STARTNEIGHB -- defined as **forall** $q \in Neighb$ SEND(*Start*, **to** q)

- to LEADELECT add **PROPAGATESTART**, defined by

if *Received*(*Start*, **from** p) **and** $p \in Process$ **then**

if $mode = terminated$ **then** START(**self**) -- enable **self**

CONSUME(*Start*, **from** p)

Refinement to GRAPHLEADELECTDIFFUSE

Refine $\text{TERMINATE}_{\text{Diffuse}}$ to trigger a final round which resets each component to $\text{mode} = \text{terminated}$ (for the next diffusing run).

- $\text{TERMINATE}_{\text{Diffuse}} =$
 - if** $\text{status} = \text{master}$ **and** quiescent **and** AllAcksArrived **then**
 - if** $\text{mode} = \text{check}$ **then** TERMINATE -- launch termination round
 - if** $\text{mode} = \text{terminated}$ **then**
 - $\text{status} := \text{idle}$ $\text{SEND}(\text{computationResult}, \text{to } \text{env})$
 - where** $\text{TERMINATE} =$
 - $\text{TERMINATE}_{\text{NEIGHB}}$ -- **forall** $q \in \text{Neighb}$ $\text{SEND}(\text{Stop}, \text{to } q)$
 - $\text{mode} := \text{terminated}$ -- initialize mode for next diffusing run
- Stop (like Start) is treated as GRAPHLEADELECT-internal msg.
- To LEADELECT add **PROPAGATE TERMINATION**, defined by:
 - if** $\text{Received}(\text{Stop}, \text{from } p)$ **then**
 - if** $\text{mode} \neq \text{terminated}$ **then** TERMINATE
 - $\text{CONSUME}(\text{Stop}, \text{from } p)$

Termination detection lemma for GRAPHLEADELECTDIFFUSE

Let $T = \text{TERMINATIONDETECTOR}(\text{GRAPHLEADELECTDIFFUSE})$.
Every diffusing run of T eventually terminates and does
 $\text{ELABORATE}(Max)$ (refining *computationResult* to *cand*).

More precisely: Let R be any diffusing run of T (with reliable message passing and without infinitely lazy components).

R eventually terminates and does $\text{ELABORATE}(Max)$ when the system of GRAPHLEADELECTDIFFUSE machines, started by an *initiator* in a quiescent state, after the second $\text{TERMINATE}_{Diffuse}$ step of the *initiator* enters again a quiescent state, where all components have *status = idle* and *mode = terminated*.

Proof. Follows from the termination of GRAPHLEADELECT runs and from the above explained behavior of the monitoring components of TERMINATIONDETECTOR. For details see ModelingBook pg.125.

Exl.3: Concurrent emulation of synchronous processes

Goal. Emulation of synchronous runs of a network $(Process, Edge)$ of communicating processes by concurrent runs of a network

$$\text{CONCURSYNCEMULATOR}(Process, Edge) = (Process^*, Edge^*)$$

(also called LOCALSYNCTRANSFORMER) of communicating processes, assuming immediate reliable communication bw neighbors.

=====

A *round* in a synchronous run is characterized by each $p \in Process$

- reading its mailbox only at the beginning of the round
- sending messages only at the end of the round, msgs which in the next round are in the receivers' mailbox
- performing otherwise only non-communication local actions

Wlog we abstract from the possible sequence of local actions, treating them as one atomic step, and assume that in each round, each process sends exactly one (possibly empty) message to each of its neighbors.

Characterization of synchronous *Process* runs

Let $(Process, Edge)$ be a network of communicating ASMs p with agent $ag(p)$ executing $pgm(p)$ using a $mailbox(p)$ for immediate reliable communication with the *Neighbors* of p (defined via $Edge$).

The synchronous runs of $(Process, Edge)$ can be described as the runs of the following (highly parallel) ASM. In each step it performs one step (communication and local actions) of each of its components.

SYNCNET $(Process, Edge) =$

forall $p \in Process$

$pgm(p)$

INCREASEROUND

-- i. e. $curRound := curRound + 1$

Each round corresponds to one step of $SYNCNET(Process, Edge)$ so that $curRound$ works as step counter.


Concurrent emulation of synchronous processes

Idea (Lynch 1996, Sect.16.2, here generalized from interleaved processes to concurrent communicating ASMs):

- associate each process $p \in Process$ with a $synchronizer(p)$ which for each round r synchronizes
 - each round- r -step of p with a round- r -step of all its *Neighbors*
 - the *Process*-internal communication in round r , i.e. between p and its *Neighbors* (via *Edge*)
- replace $pgm(p)$ by SYNC SHELL(p) which
 - simulates one step of $pgm(p)$ when $ReadyForNextRound_p$
 - SUSPENDS p until it has *ReceivedAllMsgsFor* round $r + 1$ and the $synchronizer(p)$ sends a *resume* msg (after p and all its *Neighbors* *MadeOneStep* in round r)

Graph $Edge^*$ of the concurrent synchronization emulator

→ indicates a communication line between neighbors.³



NB. $pgm(p)$ is changed to $SYNCSHELL(p)$ whose round- r -steps are synchronized, by an agent with pgm $SYNCHRONIZER(p)$, with those of (the $SYNCSHELL$ instances of) all *Neighbors* of p (in $Edge$).

³ Figure © 2003 Springer-Verlag Berlin Heidelberg, reprinted with permission

Process^{*} of the concurrent synchronization emulator

For each process p

- its $pgm(p)$ is replaced by a program $SYNCSHELL(p)$ which refines the communication actions of p
- a synchronization controller $SYNCTL(p)$ is added which synchronizes the round actions of all processes (including their communication)

Process^{*} =

$Process(pgm(p)/SYNCSHELL(p))$ -- process refinement

$\cup SYNCTL$ -- adding synchronization controller

where

$SYNCTL = (SYNCTL(p))_{p \in Process}$

$SYNCTL(p) =$

$(synchronizer(p), SYNCHRONIZER(p), mailbox(synchronizer(p)))$

SYNCSHELL(p) for the concurrent synchronization emulator

if *ReadyForNextRound* _{p} **then**

pgm(p)* -- simulate one step of *pgm*(p)

SUSPEND(p)

else

if *Received*(*resume*, **from** *synchronizer*(p)) **then**

RESUME(p)

where

ReadyForNextRound _{p} **iff**

Resumed(p) **and** *ReceivedAllMsgsFor*(*curRound* _{p} , p)

SUSPEND and RESUME actions

$\text{SUSPEND}(p)$ sets $\text{WaitingForNextRoundTick}(p)$ to true, informs the $\text{synchronizer}(p)$ that p has performed its curRound_p step and prepares itself for the next round (by an $\text{INCREASE}(\text{curRound}_p)$).

SUSPEND(p) = -- making p not *ReadyForNextRound*

$\text{WaitingForNextRoundTick}(p) := \text{true}$ -- reset by **RESUME**

INFORMABOUTSTEP(p) -- step performed by $\text{pgm}(p)^*$

INCREASE(curRound_p) -- preparing for next round $\text{curRound}_p + 1$

RESUME(p) = -- upon receiving *resume* msg

$\text{WaitingForNextRoundTick}(p) := \text{false}$

CONSUME((*resume*, **from** $\text{synchronizer}(p)$))

where

INFORMABOUTSTEP(p) =

SEND($\text{stepInfo}(p, \text{curRound}_p)$, **to** $\text{synchronizer}(p)$)

$\text{Resumed}(p)$ **iff** $\text{WaitingForNextRoundTick}(p) = \text{false}$

Simulation of $pgm(p)$ -steps by $pgm(p)^*$ -steps

Simulating a $pgm(p)$ -step consists in performing this step except for SENDING *Process*-internal msgs (analogously for receiving)

- together with the $curRound_p$ information
- not to their receiver $q \in Neighb(p)$ but to the $synchronizer(p)$

where

- SYNCHRONIZER(p) does FORWARDMSGSENTBY($p, curRound_p$) to $synchronizer(q)$
- SYNCHRONIZER(q) will PASSMSGSENTTO(q, r) to q

$pgm(p)^* = pgm(p)$ replacing

SEND(m , **to** q) -- communication with round info via *synchronizer*
by SEND($(m, curRound_p, \mathbf{to} q), \mathbf{to} synchronizer(p)$)

Received(m , **from** q) by

Received($(m, curRound_p - 1, \mathbf{from} q), \mathbf{from} synchronizer(p)$)

-- NB. m received in round r has been sent in round $r - 1$

The two SYNCHRONIZER roles

- FORWARDMSGSENTBY process p in a round- r -step, namely to the $synchronizer(q)$ of each $Neighbor\ q$ of p . These msgs can be:
 - monitor msgs: $stepInfo$ and $resume$ msg
 - a $Process$ -internal $msg \in ProcessMsg$, exchanged between processes via a $SEND(msg)$ in some $pgm(p)$ with $p \in Process$
- Check when to CLOSEROUND r for p by
 - passing to p in one blow all $ProcessMsgs$ which have been sent in round r to p — $Received$ by $synchronizer(p)$ from $synchronizer(q)$ of some $q \in Neighb(p)$ —to be $Received$ by p in round $r + 1$
 - waking up p (by a $resume$ msg) and proceeding to the next round

SYNCHRONIZER(p) =

let $r = curRound(\mathbf{self})$

FORWARDMSGSENTBY(p, r)

CLOSEROUND(r, p)

CLOSEROUND(r, p) rule of SYNCHRONIZER(p)

if $Received(stepInfo(p, r), \text{from } p)$ -- p made a round- r -step
and (**forall** $q \in Neighb(p)$ $MadeOneStep(q, r)$)
and $ReceivedAllMsgsToPassTo(p, r)$ **then**
 PASSMSGSENTTO(p, r) -- msgs to-be-received in round $r + 1$
 WAKEUP(p) -- i.e. SEND(*resume*, **to** $syncShell(p)$)

where

$MadeOneStep(q, r)$ **iff** -- NB. $r = curRound_q$
 $Received(stepInfo(q, r), \text{from } synchronizer(q))$
 $ReceivedAllMsgsToPassTo(p, r)$ **iff**
 forall $q \in Neighb(p)$ -- q has sent a msg to p in round r
 forsome m $Received((m, r, \text{to } p), \text{from } synchronizer(q))$

NB. p assumed to send per round to each neighbor exactly one msg.

Components of $\text{CLOSEROUND}(r, p)$

$\text{PASSMSGSENTTO}(p, r) =$ -- all msgs p receives in round $r + 1$
forall $q \in \text{Neighb}(p)$ **forall** -- all msgs of type ProcessMsg
 $((m, r, \mathbf{to} p), \mathbf{from} \text{synchronizer}(q)) \in \text{mailbox} \cap \text{ProcessMsg}$
 $\text{SEND}((m, r, \mathbf{from} q), \mathbf{to} p)$
 $\text{CONSUME}(m, r, \mathbf{from} q)$

This rule is executed when $\text{ReceivedAllMsgsToPassTo}(p, r)$ is true.
 Correspondingly we can define for $\text{ReadyForNextRound}_p$:

$\text{ReceivedAllMsgsFor}(r + 1, p)$ **iff**
forall $q \in \text{Neighb}(p)$ **forsome** m -- q has sent some msg in round r
 $\text{Received}((m, r, \mathbf{from} q), \mathbf{from} \text{synchronizer}(p))$
 $\text{ReceivedAllMsgsFor}(0, p) = \text{true}$ -- by initialization
 -- If $\text{Received}(m)$ in round $r + 1$, m has been sent in round r

Initial states & runs of CONCURSYNCEMULATOR($Process, Edge$)

Initial states satisfy:

- $curRound = 0$ and $mailbox = \emptyset$ for each $p \in Process$ (same for $SYNcNET(Process, Edge)$) and for each $SYNcCTL(p)$
- $ReadyForNextRound_p = ReceivedAllMsgsToPassTo(p) = true$ and $ReceivedAllMsgsFor(0, p) = true$ for each $p \in Process$.

The concurrent emulation refines each p -step of one $SYNcNET(Process, Edge)$ round- r -step into:

- a first $SYNcSHELL(p)$ step—when $ReadyForNextRound_p$ holds for r
- followed by message forwarding steps performed by $synchronizer(p)$ and the $synchronizers$ of its neighbors q
- followed by a $CLOSEROUND(r, p)$ step—when also all neighbors of p made a round- r -step and the messages sent in round r to p have been received by the $synchronizer(p)$
- concluded by the second $SYNcSHELL(p)$ round- r -step

CONCURSYNCEMULATOR(*Process*, *Edge*) **correctness property**

Let R_s be any run of $\text{SYNCSNET}(\text{Process}, \text{Edge})$ and R_c any concurrent $\text{CONCURSYNCEMULATOR}(\text{Process}, \text{Edge})$ run, both properly initialized and started with equal values in same-named locations.

Then for every process p the following holds for every round number $r = \text{curRound}_{\text{SYNCSNET}(\text{Process}, \text{Edge})} = \text{curRound}_p$:

- when $ag(p)$ starts its round r in R_s resp. in R_c , the runs are in corresponding states $state_{R_s}(r)$, $state_{R_c}(p, r)$ with same values
 - in same-named p -locations
 - for the payload and destination of corresponding *ProcessMsgs*
 - $(m, \mathbf{to} q)$ resp. $(m, \text{curRound}_p, \mathbf{to} q)$
sent by $ag(p)$ in round r to (similarly *Received* from) its neighbor q .

NB. Same-value property in corresponding states means

$$state_{R_s}(r) \downarrow \Sigma_p = state_{R_c}(p, r) \downarrow \Sigma_p$$

Proof: induction on r (see ModelingBook pg. 131 for details).

References

- E. Börger and A. Raschke: Modeling Companion for Software Practitioners. Springer 2018
<http://modelingbook.informatik.uni-ulm.de>
The book bibliography provides exact references to the literature from where the three examples are taken.
- E. Börger and R. Stärk: Abstract State Machines. Springer 2003.

Copyright Notice

It is permitted to (re-) use these slides under the CC-BY-NC-SA licence

<https://creativecommons.org/licenses/by-nc-sa/4.0/>

i.e. in particular under the condition that

- the original authors are mentioned
- modified slides are made available under the same licence
- the (re-) use is not commercial