

Abschnitt I

Asynchronous System Model

Inhalt: I/O automata · traces and executions · operations on automata: composition + hiding · fairness · properties and proof methods · safety and liveness

Literatur: The material is taken from [Lyn96, Chapter 8].

Overview

- big step: from *synchronous* to *asynchronous*¹ model
- more complex
 - more “nondeterminism”
 - more uncertainty, due to the relative speed of the parallel components
 - fairness, liveness
- machine model:

I/O automata

¹the word “asynchronous” is used differently sometimes elsewhere. *Asynchronous* communication/message passing is often meant as *buffered* communication.

I/O automaton

- simple machine/automaton model for **processes** in an **asynchronous** distributed network
- named “**actions**” for the transitions (internal or external = input or output)
- \Rightarrow notion of **interface/signature**
- **composable**
- general/unspecific enough for
 - **shared memory** concurrent systems (cf. Chapter 9 and following)
 - **message passing** (“network”) systems (cf. Chapter 14 and following)
- Examples:

1. process P_i of Figure 8.1
2. fifo channel $C_{i,j}$ of Figure 8.2²

²One sees that the model is so unspecific, that channels are not built in, but have to be programmed.

Actions and signature

- **actions** = will be “labels” on transition
- **signature** $S = (in(S), out(S), int(S))$: three **disjoint** set of actions:
 1. $int(S)$: **internal**
 2. $out(S)$: **output**
 3. $in(S)$: **input**
- furthermore
 - **external** actions: $ext(S) \triangleq in(S) \dot{\cup} out(S)$
 - **locally controlled** actions $local(S) \triangleq int(S) \dot{\cup} out(S)$.
 - thus: **external signature/interface** $extsig(S) = (in(S), out(S), \emptyset)$
 - A **closed** if $in(A) = \emptyset$ (“autonomous”)

I/O automaton: Definition

Definition 1. An *I/O automaton* A is given by

1. $\text{sig}(A)$, a *signature*
2. $\text{states}(A)$ (finite or countably infinite)
3. $\text{start}(A)$: subset of initial states
4. $\text{trans}(A) \subseteq \text{states}(A) \times \text{acts}(\text{sig}(A)) \times \text{states}(A)$: *state-transition relation, input-enabled*
5. $\text{tasks}(A)$: *task-partition* = equivalence-relation on $\text{local}(\text{sig}(A))$, at most countably infinite equivalence classes

Enabledness, tasks

- **transition/step**: $(s, \pi, s') \in \text{trans}(A)^3$ (input/output/internal ... transitions)
- action π **enabled** in s , if $s \xrightarrow{\pi} s' \in \text{trans}(A)$.
- **input-enabledness**: every input action must be enabled in every state
 - it's better to consider all possible reaction, otherwise: **error prone** design
 - nicer theory
- **quiescent** state: no actions except input actions are enabled⁴
- **tasks**
 - abstract representation for “tasks/jobs/threads of control”
 - useful specifically for “parallel composition”
 - primarily used later to specify **fairness conditions** to ensure **liveness** properties
 - word “task” $\hat{=}$ task-partition class of the automaton

³I sometimes write also $s \xrightarrow{\pi}_A s'$, or $s \xrightarrow{\pi} s'$, if A is clear from the context.

⁴Cf. the definition of quiescent state for the synchronous model [Lyn96, p. 19].

Examples: channel and process

- cf. Examples 8.1.1 and 8.1.2 in [Lyn96, p. 204]
- written in **precondition/effect**-style, grouping together “analogous” actions (parameter)
- note:
 - **input enabledness**: empty precondition = true
 - *tasks*:
 - * for the process: separation per receiver channel
 - * for the channel: all outputs⁵ in the same class

⁵i.e., all locally controlled actions in this example.

Executions and traces

- straightforward **step semantics**
 - starting from an initial state: do transitions
 - non-deterministic⁶
 - we distinguish between
 - * **internal** steps observable: **executions**
 - * only **actions** of **external** steps observable: **traces** = interface behavior

Definition 2. [Traces and executions] *Given A*

1. an **execution fragment** of A is

- (a) **finite** sequence $s_0\pi_1s_1\pi_2\dots\pi_r s_r$
- (b) or **infinite** sequence $s_0\pi_1s_1\pi_2\dots,$

such that $s_i \xrightarrow{\pi_{i+1}} s_{i+1}$. Note: in the finite case we end in a state.

⁶because the transitions allow this. Parallel composition will add another source of non-determinism.

execution = *execution fragment* starting with an initial state

2. *traces*:

- (a) the *trace* of an execution (fragment) α of A : *subsequence* of α consisting of the *external actions*⁷
- (b) of an *automaton* A : lifted on the set A 's executions

- *notation* for executions and traces: $execs(A)$, $trace(\alpha)$, $traces(A)$.
- a *state* s is *reachable*, if there *exists* a (finite) execution with s and end-state
- *concatenation*: $\alpha \cdot \alpha'$ = "glueing" together (finite) execution fragments⁸
- Example 8.1.3 about the Fifo-channel $C_{i,j}$

⁷no states!

⁸Assuming, that the end state of α equals the first state of α' and of course, not mention the glue state twice in $\alpha \cdot \alpha', \dots$

Composition

- **parallel composition** of larger/more complex system by smaller ones
- hierarchical description
- standard **product** construction:
 - states are **paired**
 - transitions (“interleaving”)
 - * **synchronizing** on **common** actions
 - * **non-synchronizing** on local actions: automaton does **nothing**⁹
- to obtain the desired intuition: certain **restrictions**
 - internal actions should not synchronize: **internal actions disjoint**¹⁰

⁹In other models, one could call this *stuttering* steps.

¹⁰That’s one possible way to formalize the informal, intuitive intention that some actions are considered to be internal.

- only one process **controls** other actions:¹¹ **output actions** disjoint¹²
 - only a *finite* number of **synchronizing, common** actions
- ⇒ definition of **compatible** signatures and automata

¹¹By synchronizing over output actions. It's one of the underlying intuitions in this model: input forces the component to do something and especially the component cannot *refuse* to accept input, but the internal actions and the output actions are under *component control*.

¹²Note that input actions are not required to be disjoint: an “outputter” can trigger many “inputters”. This form of multi-synchronization is allowed.

Compatible

a countable collection $\{S_i\}_{i \in I}$ is **compatible**:

locality $int(S_i) \cap acts(S_j) = \emptyset$

indepedent outputs $out(S_i) \cap out(S_j) = \emptyset$

finite sync. no action is contained in infinitely many $acts(S_i)$

- note: $local(S_i) \cap local(S_j) = \emptyset$
- **compatible** collection of automata: correspondingly

\Rightarrow composition of sig's and automata

Composition: signatures

signatures: given a compatible, countable collection $\{S_i\}_{i \in I}$ of **signatures**

\Rightarrow $S = \prod_{i \in I} S_i$ given by

- $out(S) \triangleq \dot{\bigcup}_{i \in I} out(S_i)$
- $int(S) \triangleq \dot{\bigcup}_{i \in I} int(S_i)$
- $in(S) \triangleq \bigcup_{i \in I} in(S_i) - \dot{\bigcup}_{i \in I} out(S_i)$

Composition: automata

automata : given a compatible, countable collection $\{A_i\}_{i \in I}$ of automata

\Rightarrow $A = \prod_{i \in I} A_i$ given by

- $sig(A) \triangleq \prod_{i \in I} sig(A_i)$
- $states(A) \triangleq \prod_{i \in I} states(A_i)$
- $start(A) \triangleq \prod_{i \in I} start(A_i)$
- $s \xrightarrow{\pi} s' \in trans(A)$ if, for all i
 - if $\pi \in acts(A_i)$: $s_i \xrightarrow{\pi}_i s'_i$
 - if $\pi \notin acts(A_i)$: $s_i = s'_i$
- $tasks(A) \triangleq \bigcup_{i \in I} tasks(A_i)$.

Composition: remarks

- remember: A_i 's are **input enabled** $\Rightarrow A$ is, too
- note: **local** actions cannot be used to synchronize (= more than one process makes a real step) by convention
- for **finite/binary** composition: $A \times B$, or $A_1 \times \dots \times A_n$
- it's to **prove**, that $\prod_{i \in I} A_i$ yields an automaton, same for signatures
- note: composing **output** and input actions —both are **external**— yields: **an output**¹³ \Rightarrow **broadcast** communication can directly be modelled.
- \times is **associative**¹⁴

¹³Other known models make a different plausible choice here: input parallel with an output gives an internal action.

¹⁴Technical remark: associativity hinges, in this formalization of parallel composition on the fact: output + input gives output, not an internal action. This, on the other hand, does not mean to say, that *binary* associative synchronization is impossible. One would have to give up to *force* synchronization on common actions in parallel actions. Such a communication model is, for instance, used in CCS, and similar calculi.

Composition: example

- parallel composition of [processes](#) + [buffers](#): 8.2.1

Local and global executions/traces

- relating executions of a composed automaton \leftrightarrow executions of the components
- given $\alpha = s_0\pi_1s_1\dots$ of A : “projection” $\alpha \downarrow_{A_i} =$ subsequence obtained by
 1. deleting all pairs $\pi_r s_r$ where π_r not an A -action
 2. replace each remaining s_r by the i -th component s_i
- projection is analogously used on traces or arbitrary sequences of actions

Theorem 1. [Decomposition (8.1)] Given $A = \prod_{i \in I} A_i$.

- $\alpha \in \text{execs}(A)$, then $\alpha \downarrow_{A_i} \in \text{execs}(A_i)$
- $\alpha \in \text{traces}(A)$, then $\alpha \downarrow_{A_i} \in \text{traces}(A_i)$

Theorem 2. [Composition from executions (8.2)] Given

- $A = \prod_{i \in I} A_i$
- α_i : an execution of A_i
- β : a sequence of actions in $\text{ext}(A)$ s.t. $\beta \downarrow_{A_i} = \text{trace}(\alpha_i)$

Then there is an *execution* α of A such that $\beta = \text{trace}(\alpha)$ and $\alpha_i = \alpha \downarrow_{A_i}$

Theorem 3. [Composition from traces (8.3)] Given

- $A = \prod_{i \in I} A_i$
- β a sequence of actions in $\text{ext}(A)$
- $\beta \downarrow_{A_i} \in \text{traces}(A_i)$

Then

$$\beta \in \text{traces}(A).$$

Hiding

- **operation** on an I/O-automaton: hiding **output**
- after hiding: operation is **internal** \Rightarrow no further **synchronization** possible

for signatures: given S and $\Phi \subseteq out(S) \Rightarrow hide_{\Phi}(S) = S'$ defined as

- $int(S') \triangleq int(S) \cup \Phi$
- $in(S') \triangleq in(S)$
- $out(S') \triangleq out(S) - \Phi$

note: $\Phi \cap in(S') = \emptyset$ by definition

for automata: simply using the definition for signatures, i.e., given A and $\Phi \subseteq out(A)$:

- $hide_{\Phi}(A)$ defined as A' given by replacing $sig(A)$ by $hide_{\Phi}(A)$

Fairness: informal

- asynchronous model: fairness becomes quite an issue
- informally: fair = “each one get’s his turn” (here based on tasks)
- formally: abstract (and often difficult) notion, considering infinite behavior
- various “flavors” of fairness useful
- abstraction of a scheduler
- here: each task gets infinitely many opportunities to perform one of its actions

Fairness

Definition 3. [Fairness] An *execution fragment* α of A is *fair*, if for each equivalence class C of $tasks(A)$:

1. if α is finite, then C is *not enabled* in the final state of α

2. if α is *infinite*, then α contains

(a) infinitely many events from C , or

(b) infinitely many occurrences of states in which C is *disabled*

- *event* = occurrence of an action in a sequence (execution, trace)

- $fairexecs(A)$: fair executions of A , and $fairtraces(A)$: fair traces of A , where a *fair trace* is a trace of a fair execution¹⁵

¹⁵Remember: trace cf. slide 9.

- example 8.3.1
- example 8.3.2: *discrete clock*

(De-)composition of fair traces

Analogous to the corresponding properties for general traces/executions

Theorem 4. [Fair decomposition (8.4)] Given $A = \prod_{i \in I} A_i$.

- $\alpha \in \text{fairexecs}(A)$, then $\alpha \downarrow_{A_i} \in \text{fairexecs}(A_i)$
- $\alpha \in \text{fairtraces}(A)$, then $\alpha \downarrow_{A_i} \in \text{fairtraces}(A_i)$

Theorem 5. [Composition from fair executions (8.5)] Given

- $A = \prod_{i \in I} A_i$
- α_i : a *fair* execution of A_i
- β : a sequence of actions in $\text{ext}(A)$ s.t. $\beta \downarrow_{A_i} = \text{trace}(\alpha_i)$

Then there is a *fair execution* α of A such that $\beta = \text{trace}(\alpha)$ and $\alpha_i = \alpha \downarrow_{A_i}$

Theorem 6. [Composition from fair traces (8.6)] *Given*

- $A = \prod_{i \in I} A_i$
- β a sequence of actions in $\text{ext}(A)$
- $\beta \downarrow_{A_i} \in \text{fairtraces}(A_i)$

Then

$$\beta \in \text{fairtraces}(A).$$

Fairness example

- Example 8.3.3: 3 processes
 - every sent message is **eventually** received
 - if there is at least one init-event for each i , each processes sends infinitely many messages to each other, and each process send infinitely many decide messages

Fairness and finite prefix

“(Un-)fairness does not show up on finite prefixes of executions”

- fairness is not a **safety** property
- more formally: every **finite** executions can be **extended** into a fair execution (same for traces)

Theorem 7. [Fairness] 1. If α is a **finite** execution of A , then there is a **fair** execution of A that **starts with** α .

2. The same holds analogously for traces.

3. If α is a **finite** execution of A and β is any (finite or infinite) sequence of **input** actions of A , then there is a fair execution $\alpha \cdot \alpha'$ of A such that the sequence of input actions in α' is exactly β .

4. *similar for traces: If β is a **finite** trace of A and β' is any (finite or infinite) sequence of **input** actions of A , then there is a fair execution $\alpha \cdot \alpha'$ of A such that $\text{trace}(\alpha) = \beta$ and s.t. the sequence of input actions in α' is exactly β' .*

Input and output

- to **apply** this general model: make some conventions about **I/O**
- remember: in the synchronous model: designated **state variables** (write-once for outputs)
- in the asynchronous model: we simply use **input** and **output actions**

Invariant assertions

- most basic/simple class of properties: **invariant assertion** or **invariant**

“something which always holds”

- **invariant** of an automaton A : property which holds for all **reachable** states
- typically proven by **induction** on the number of steps¹⁶

¹⁶In the synchronous setting, we used induction on the number of **rounds**.

Trace properties and satisfaction

- remember: **trace** = “visible/observable” actions of an execution

⇒ often, properties of “external” interest are formulated over (fair) traces, i.e., extensionally by a **set of traces**,

Trace property P given by

1. $sig(P)$, signature, containing **no internal** action¹⁷
2. $traces(P) \subseteq acts(sig(P))$

- (at least) **2 interpretations** of A **satisfies** P (“ $A \models P$ ”)

1. $extsig(A) = sig(P)$ and $traces(A) \subseteq traces(P)$
2. $extsig(A) = sig(P)$ and $fairtraces(A) \subseteq traces(P)$

¹⁷One cannot “look inside”

Safety

- important class of trace properties¹⁸
- general slogan

“never something bad happens”

Definition 4. [Safety trace property] A trace property P is a *safety (trace) property*:

1. $\text{traces}(P)$ is not empty
2. $\text{traces}(P)$ is *prefix closed*¹⁹

¹⁸In this setting, safety is phrased in terms of trace properties.

¹⁹That implies the first point already.

3. $\text{traces}(P)$ is *limit closed*

Definition 5. [Limit closure] An set of traces T is *limit closed*: if β_1, β_2, \dots is an infinite sequence of *finite* traces in T , such than for all i :

$$\beta_i \text{ is a prefix of } \beta_{i+1} ,$$

then the *limit* β (the unique sequence β that is the limit of β_i under the successive extension ordering) is also in T .

- intuition: if something bad happens, it happens in a finite amount of time by a *particular event* \Rightarrow limit closure
- note: a *fairness* property is not a safety property

Example 1. [8.5.2] “No *decide* happens without a preceding *init*”.

Lemma 1. If P is a safety property, then the following statements are equivalent

1. $\text{traces}(A) \subseteq \text{traces}(P)$.

2. $\text{fairtraces}(A) \subseteq \text{traces}(P)$.
3. $\text{finitetraces}(A) \subseteq \text{traces}(P)$.

Liveness

- general slogan:

“something good will (indeed) happen”

Definition 6. [Liveness] A trace property P is a *(trace) liveness property*, if every *finite* sequence over $acts(P)$ has some *extension* in P .

Example 2. [8.5.3] For every *init* event, there will be a/infinite many matching *decide later*.

Proving liveness/temporal logic(s)

- proving a liveness property:²⁰ $\text{fairtraces}(P) \subseteq \text{traces}(P) = \text{liveness property}$
- temporal logic:
 - logic(s) targeted towards reasoning about reactive behavior/traces/ ...
 - “temporal” does not (necessarily) mean real time: but ordering of events
 - typical operators: “Always”, “Eventually” (sometimes written \square , \diamond , \bigcirc ...)
- “do-it-yourself-method”: progress functions²¹
 - proving that a particular event will happen
 - mapping from states of the automaton to well-founded set
 - show that actions decrease the value

²⁰Often, fairness assumptions needed to prove liveness.

²¹one can use temp. logic to formalize this method. Anyway, we won't use temporal logic.

Safety and liveness

- safety and liveness: intuitive **dual** classes of properties
- this can be made **formal**/proven (in the following)
- note: the properties are phrased for **traces**, analogous properties holds for **executions**

Theorem 8. [Safe and live (8.8)] *P is a safety **and** a liveness property \Rightarrow*

$$P = acts(P)$$

Theorem 9. [Safety and liveness decomposition] *Given trace property P with $traces(P) \neq \emptyset$. Then there exist a **safety** and a **liveness** property S and L :*

1. $sig(S) = sig(L) = sig(P)$.
2. $traces(P) = traces(S) \cap traces(L)$.

Compositional reasoning

- “divide-and-conquer” approach sometimes helpful in reasoning
- given²² $A = \prod_{i \in I} A_i$ and $P = \prod_{i \in I} P_i$, and **local satisfaction**: $A_i \models P_i$

Theorem 10. [Compositional reasoning] *1. If $\text{extsig}(A_i) = \text{sig}(P_i)$ and $\text{traces}(A_i) \subseteq \text{traces}(P_i)$, then and $\text{extsig}(A) = \text{extsig}(P)$ and $\text{traces}(A) \subseteq \text{traces}(P)$.*

2. If $\text{extsig}(A_i) = \text{sig}(P_i)$ and $\text{traces}(A_i) \subseteq \text{fairtraces}(P_i)$, then and $\text{extsig}(A) = \text{extsig}(P)$ and $\text{traces}(A) \subseteq \text{fairtraces}(P)$.

²²Unless stated otherwise: by writing the products $\prod_{i \in I} A_i$ and $\prod_{i \in I} P_i$, we implicitly assert that those are well defined, especially that the constituents form a **compatible** collection of automata resp. trace properties.

$$\frac{\text{extsig}(A_i) = \text{sig}(P_i) \quad \text{traces}(A_i) \subseteq \text{traces}(P_i)}{\text{traces}(A) \subseteq \text{traces}(P)} \text{COMP-TR}$$

$$\frac{\text{extsig}(A_i) = \text{sig}(P_i) \quad \text{fairtraces}(A_i) \subseteq \text{traces}(P_i)}{\text{fairtraces}(A) \subseteq \text{traces}(P)} \text{COMP-FTR}$$

Further compositional reasoning

- using composition Theorem 3 on page 19 (or the corresponding one for fair traces)
 - “inverting” the previous reasoning
 - drawing conclusion from **projections** to the **common** trace.
- compositional proof of **safety** properties

Compositional reasoning for safety properties

- exploiting that safety is violated in a **finite** prefix
- A **preserves** P : A is not the **first** to violate P
- if **no** A_i is the **first** violator, safety holds

Definition 7. [Preserving] P a **safety** property with $acts(P) \cap int(A) = \emptyset$ and $in(P) \cap out(A) = \emptyset$.

A **preserves** B , if for every **finite** sequence β of actions that does not include **internal** actions of A , and every $\pi \in out(A)$:

*If $\beta \downarrow_{acts(P)} \in traces(P)$ and $\beta\pi \downarrow_A \in traces(A)$, then **important** $\beta\pi \downarrow_{acts(P)} \in traces(P)$.*

The internal actions are not mentioned, of course.

Compositional reasoning & safety

Theorem 11. [Safety and comp. reasoning] Given $A \prod_{i \in I} A_i$. and a *safety* property with $acts(P) \cap int(A) = \emptyset$ and $in(P) \cap out(A) = \emptyset$:

1. If A_i *preserves* P for all i , then A *preserves* P
2. If A is *closed*, A *preserves* P , and $acts(P) \subseteq ext(A)$, then

$$traces(A) \downarrow_{acts(P)} \subseteq traces(P) .$$

3. If A is *closed*, A *preserves* P , and $acts(P) = ext(A)$, then

$$traces(A) \subseteq traces(P) .$$

Hierarchical proofs

- hierarchy: different **level of abstraction**
- from top to bottom:²³ **successive refinements**
- examples: simulation proofs in the material of the synchronous model
- problem: the two systems under comparison are not a strictly **coupled** as before
- needed: generalization of the **simulation** method to the **asynchronous** setting
- anyway: general approach: **directed relationship**:²⁴

“for any execution of the lower-level automaton, there is a “corresponding” execution of the higher-level automaton.”

²³ “Higher” means: more abstract, less details etc. One could for instance introduce more parallelism going top-down.

²⁴ In the synchronous setting, the goal was that the implementation had **the same behavior** than the abstract system., at least under the assumption of determinism.

Simulation relation

Definition 8. [Simulation relation] • A and B with identical *external* interfaces.

- $f \subseteq \text{states}(A) \times \text{states}(B)$.

Then f is a *simulation relation* from A to B , if

Start condition: if $s \in \text{start}(A)$, then $f(s) \cap \text{start}(B) \neq \emptyset$

Step condition: • if s is a *reachable* state of A , and if $u \in f(s)$ where u is a reachable state of B :

- $s \xrightarrow{\pi} s'$, then there is an execution *fragment* α of B , starting in u and ending with some $u' \in f(s')$ such that

$$\text{trace}(\alpha) = \text{trace}(\pi)$$

Simulation proof method

- **important** proof technique, e.g., for trace
- aux. definition: B **simulates** A ($B \succeq A$), if **there exists** a simulation relation **from** A **to** B .

alternative words: A is simulated by B ($A \preceq B$)

- **Caveat:** some people use the words “the other way around”

Theorem 12. [Simulation and trace inclusion] *If there is a **simulation** relation from A to B , the the traces of A are included in the traces of B :*

$$\frac{A \preceq B}{\text{traces}(A) \subseteq \text{traces}(B)} \text{SIM}$$

Simulation and safety property

- simulation: preservation of safety properties²⁵

$$\frac{P \text{ is a safety property} \quad B \preceq A \quad A \models P}{B \models P}$$

- for liveness properties:
 - not that simple/direct

²⁵Remember: $A \models P$ means $\text{traces}(A) \subseteq \text{traces}(P)$.

Simulation and liveness properties

- we need to **strengthen** the coupling

Definition 9. [Correspondance] *Given:*

- *A and B with identical input and output actions*
- *α and β executions of A resp. B*
- *relation $f \subseteq \text{states}(A) \times \text{states}(B)$.*

*Then α and β **correspond wrt. f** (written $\alpha \bowtie_f \beta$), if*

1. *there exists a mapping g from **indices** (occurrences) of states in α to indices of states in β ,*
 - *g is **monotone nondecreasing***

- g *exhausts* all of β ²⁶
- g -corresponding pairs of states are *related by* f
- between *successive* g -corresponding pairs of states, the traces in α and β are *identical*

Theorem 13. [Simulation with correspondence] *If there is a **simulation relation** f from A to B ²⁷, then for every execution α of A , there exists execution β of B such that $\alpha \bowtie_f \beta$.*

²⁶i.e., the supremum of the range of g is the supremum of the indices of states in β .

²⁷i.e., $A \preceq B$

Complexity measures

- **upper** time bound for any subset of equivalence classes in $tasks(A)$ (thus also for any task, a “full” equivalence class) \Rightarrow
- $upper_C \in \mathbb{R}^{>0} + \infty$
- given a **fair execution** α of $A \Rightarrow$ associate a **real-value** time with each event of α .
 1. times are **monotone non-decreasing** in α
 2. if α is infinite, then the times approach ∞ ²⁸
 3. from any point in α , a task C can be **enabled** for time at most $upper_C$, before some action in C **must occur**
- **timed execution** = **fair** execution with times associated as described
- note: for a given set of bounds in $upper_C$: many ways of associating times to the events of α

²⁸sometime called **non Zeno-ness**.

- time until some designated event π in α occurs: **supremum** of times assignable to π in all such timed executions
- time **between** two events: likewise by the **supremum** of differences
- example 8.6.1

Indistinguishable executions

- **indistinguishable** = from the **perspective** of a **subcomponent** \Rightarrow projection

Definition 10. [Indistinguishable] *Given: α, α' : executions of two composed systems each containing automaton A . Then:*

α and α' are **indistinguishable to A** , if their projections onto A are identical:

$$\alpha \sim_A \alpha' \triangleq \alpha \downarrow_A = \alpha' \downarrow_A$$

Randomization

- remember: transition **relation** in the asynchronous model: **nondeterministic**²⁹
- sometimes: “weighted” nondeterminism is useful \Rightarrow **randomization/probabilistic I/O automata**

Definition 11. [Probabilistic I/O automaton] A *probabilistic I/O automaton* is defined the same way as an ordinary I/O automaton (cf. Definition 1 on page 6), except that the successor states are given by a **probability distribution**

$$(s, \pi, P) ,$$

where P is a probability distribution over some subset of the states.

- semantics/execution: series of **2 non-deterministic choices!**

²⁹In the synchronous (network) model, we had some next-state *function*.

1. choose the next **transition** $s \xrightarrow{\pi} P$
 2. choose the **state** according to the distribution
- restriction: choice of the transition should be **fair**
 - “inside” each probabilistic automaton, there is a nondeterministic one, forgetting the distribution $(\mathcal{N}(A))^{30}$

³⁰successor states with probability 0 are not represented in the non-deterministic choice.

Literatur

[Lyn96] Nancy Lynch. *Distributed Algorithms*. Kaufmann Publishers, 1996.