Distributed automata in an assumption-commitment framework*

SWARUP MOHALIK¹ and R RAMANUJAM²

¹Sasken Communications Technology Ltd, Bangalore 560 071, India

²The Institute of Mathematical Sciences, Taramani, Chennai 600 113, India
e-mail: swarup@sasken.com; jam@imsc.ernet.in

Abstract. We propose a class of finite state systems of synchronizing distributed processes, where processes make assumptions at local states about the state of other processes in the system. This constrains the global states of the system to those where assumptions made by a process about another are compatible with the commitments offered by the other at that state. We model examples like reliable bit transmission and sequence transmission protocols in this framework and discuss how assumption-commitment structure facilitates compositional design of such protocols. We prove a decomposition theorem which states that every protocol specified globally as a finite state system can be decomposed into such an assumption compatible system. We also present a syntactic characterization of this class using top level parallel composition.

Keywords. Assumption-commitment; automata theory; concurrency theory; verification; decomposition.

1. The assumption-commitment framework

Compositionality is a desired criterion for verification methodologies, particularly for the development and analysis of large systems. The idea is to decompose a system into smaller subsystems; then the specification for the system is verified (with respect to its implementation) using only the specifications of the subsystems without referring to their internal structure. This idea is suggested by Dijkstra (1965), who discusses hierarchical decomposition and verification of a given program, and formalized by Floyd (1967), where properties of a *sequential* program are derived from the properties of its atomic actions.

Compositional verification of *parallel* programs, on the other hand, adds substantially many complications, mainly because of the complex interaction of independently executing entities. The first proof systems for parallel programs of the form $P_1 \parallel \dots \parallel P_n$ were suggested by Owicki & Gries (1976) and Apt *et al.* In the former, the central idea was that of *interference freedom test* and in the latter it was *cooperation test*.

*The main results here were first reported in an earlier paper (Mohalik & Ramanujam 1998). We thank the anonymous reviewers for detailed comments that helped to improve the presentation

Both these needed to probe into the body of the component programs to verify these tests. In this sense, they were not compositional proof systems. Misra & Chandy (1981) provided the foundation for the much-studied *assumption-commitment* framework (AC-framework) for compositional verification. The main idea is to specify a system as a module such that if some assumptions about the external environment are satisfied then the module commits to some desired behaviour. When one has a number of such modules acting together, then each module is effectively in the environment created by the other modules. If this environment satisfies the assumptions, then the module delivers the right behaviour. Thus for the desired behaviour of the global system, assumptions and commitments of the components must mutually satisfy each other. This facilitates compositional reasoning: we can reason about the behaviour of each component separately, assuming that others maintain relevant properties and reason globally about their compatibility.

In the work by Misra & Chandy (1981), in order to capture the assumptions of the environment and commitments of the modules, one has predicates over *communication histories*. Communication histories encode the kind of interaction a module undergoes with the environment. Assumptions on communication histories are essentially constraints on the environment (in case of systems of modules running in parallel, they are constraints on the communication behaviour of other modules). But the framework itself is very general and can be applied in various ways to prove global system properties (Jones 1983; Barringer *et al* 1984; Pandya & Joseph 1991; Abadi & Lamport 1993, 1995; Quiwen & Mohalik 1997).

In this paper, we consider a kind of reasoning for distributed systems that is somewhat different in spirit from the classical compositional reasoning mentioned above. In the latter, a component is looked upon as a black-box that maintains some invariant when the environment guarantees some properties. Then, one hopes to derive system properties from the compatibility of assumptions and commitments of components, without looking into internal structure.

This black-box approach works very well as long as we are composing safety properties of the system (Owicki & Lamport 1982; Abadi *et al* 1991; Manna & Pnueli 1992; Abadi & Lamport 1993, 1995). On the other hand, it is generally agreed that composing liveness is hard. This is because, in distributed systems, local enabling of actions does not ensure global enabling. It depends crucially on the internal structure of components in the system. Therefore, for liveness it looks as if one needs to *look into* the black-boxes and keep some global information as part of *local structures* (either states or local transitions). If one does this judiciously, then all necessary global information has been distributed so that global behaviour can be obtained by a product of component processes. Our thesis is that this global information can be distributed in the local structures by suitable assumptions about other processes and commitments. Thus, a process does not have to know the detailed structure of other processes in so far as it can make suitable assumptions about others, and rely on the protocol to ensure that in the global execution these assumptions are met by the commitments of other processes.

This kind of distribution through assumption and commitment is not novel. This happens routinely when one develops subsystems without having access to a global view of the system, for example, when different groups develop parts of a large program. For instance, suppose we are designing a *receiver* that receives a bit from a sender and processes it. Then independent of the sender, the receiver can be designed as follows:

⟨Assume there is a bit in the channel⟩ Receive the bit; Process the bit.
--

The internal actions of components change their state or the state of the environment. These effects can be said to be commitments of the components. Thus the components go on making assumptions about environments and make commitments as well. When such components are put together, one gets compatible behaviours where, as intended, the assumptions of a process are met by commitments of the other processes that constitute its environment. Consider the design of the sender too.

Send the bit; ⟨Commit that a bit is in the channel⟩
--

When we put the sender and receiver together, we expect the normal sequence of transfer of the bit. Notice that the receiver cannot receive the bit before it is sent because then its assumption about the bit in the channel can not be met. In this way, causal dependence can be *encoded* by assumptions.

(Sender) Send the bit; (Receiver) Receive the bit; (Receiver) Process the bit.
--

We call this way of reasoning *local reasoning* since each component reasons “locally” about the environment (other processes in the system) to make appropriate assumptions. Observe that this view is different from the compositionality principle since here one looks into the internal design of components. Our concern is to model one aspect of system design that occurs in many situations, notably in distributed algorithms and program development, as we tried to illustrate in the example above.

Note also that *local reasoning* in the sense described above is very different from what researchers call *modular reasoning*, employed in modular model checking (Kupferman & Vardi 1997; Vardi 1997). In modular reasoning there is no constraint at all on the environment of the module. This makes perfect sense because one is concerned about design of modules as open systems, systems that can potentially be embedded in *any* environment. On the other hand, in local reasoning we are interested in closed systems where the environment of a process is the set of other processes in the system and the processes know the protocol of interaction. Thus, the environments here are very much constrained and the processes know a lot about the environment. Therefore, while modular reasoning is hard (Pneuli & Rosner 1990; Kupferman & Vardi 1997; Vardi 1997), one can expect local reasoning to be simpler.

While a number of researchers seem to have studied the AC-framework in the context of programming methodology, process algebras or temporal logics, there seems to have been little effort in formulating it from an automata-theoretic viewpoint. Implicitly, these models assume each process to be a machine of some sort, but studying formally the implications of each process being a finite-state machine is a different exercise altogether. There have been efforts in modular model checking (Kupferman & Vardi 1997; Vardi 1997), but the kind of complex interaction and compatibility that is reflected in the behaviour of parallel systems is not quite transparent.

Why should one look for an automata theoretic account of the AC-framework? An important reason is that these automata can serve as natural models for temporal logics based on local reasoning (for instance the one by Ramanujam (1996b)). Compositional model checking is one of the major goals of computer-aided verification (Alur & Henzinger 1995), and we believe that local reasoning with automata as the component processes (particularly over infinite words) may help.

In order to “locally reason” about the process model, we attach assumptions and commitments to local states and stipulate that only those global states of the system are valid where assumptions and commitments of local states are mutually compatible. A valid behaviour of the system is determined only by compatible global states. In the next section we show that such modelling is natural, by describing systems like the sequence transmission protocol. We then go on to formally define the class of automata and show that the modelling is expressive enough to capture *all regular behaviours*; that is, every (globally specified) finite state system can be decomposed into an assumption compatible system suitably. We refine this observation further to show that the decomposition can be carried out at a *syntactic* level as well. We close the presentation with a discussion.

The paper presents only automata-theoretic results on the assumption compatible systems defined here. While the definition of the model is clearly inspired by compositional verification, this paper does not address the related logical issues, except for a minor discussion.

2. System examples

In order to express assumptions and commitments at local states of processes, alongwith a distributed alphabet of actions in the system, we have an alphabet we call a *commit alphabet*. It is a tuple $\mathcal{C} = ((\mathcal{C}_1, \preceq_1), \dots, (\mathcal{C}_n, \preceq_n))$, where $\mathcal{C}_i$'s are nonempty alphabets called local commit alphabets and $\preceq_i$ is an order on $\mathcal{C}_i$. The intention is to have boolean formulas over some finite set of propositions as commitment alphabet and logical implication as the partial order: $P \preceq_i Q$, iff $Q \Rightarrow P$. Thus we can think of a commitment at a state as saying which propositions hold at that state. Moreover, suppose process i has an assumption c_j about process j at a local state p_i . Since $c_j \in \mathcal{C}_j$ can be seen as a boolean formula, i assumes that j is in some state which satisfies c_j . If actually j is in such a state p_j , then it will be compatible with p_i ; in other words, the commitments at p_j should logically imply c_j for compatibility. The commit alphabet is essentially an abstraction of this idea. We may also think of commit alphabets as denoting subsets of local states with compatibility defined by set-inclusion.

For instance, in a system of two processes, P_1 and P_2 , at a local state s of P_1 , the assumption-commitment pair may be (λ_1, λ_2) and similarly (γ_1, γ_2) for process P_2 in state t . At state s , P_1 commits to maintaining λ_1 assuming that P_2 would commit to λ_2 , and at state t , P_2 commits to maintaining γ_2 assuming that P_1 would commit to γ_1 . Then the global state (s, t) is compatible iff $\lambda_2 \preceq \gamma_2$ and $\gamma_1 \preceq \lambda_1$. We may think of λ_2 as a logical assertion whose invariance is maintained by P_2 -local states in all global states that map to s for P_1 .

Our formulation is partly inspired by *knowledge-based programs* (Fagin *et al* 1995), where atomic statements of a process are of the form $K_i\varphi \rightarrow a$. These statements are called *knowledge statements* and they are read as “if i knows φ then execute a ”. The

usual semantics of K_i is on Kripke models of global states: $K_i\varphi$ is true at a global state s if φ holds at all the global states i considers possible at s . However, for local reasoning, we need a notion of knowledge based on local states. Ramanujam (1996a) gives the semantics of K_i at local states of i : $K_i\varphi$ is true at an i -local state p if φ holds at all the global states i considers possible at p . While the connection is intuitive, we do not have a formal result relating knowledge-based programs with the automata studied here.

There is one relevant observation regarding the commit alphabet to be made here. It is not necessary that the commit alphabet be fixed universally for the system, as we have done above. This is because different processes may have access to different variables of a process i , hence their assumptions about the states of i will vary from each other. We can define each process with its own n -tuple of assumption alphabets and subsequently ensure in the definition of systems that for all $i, j \in Loc$, the j th assumption set of automaton i is contained in the j th commit set of automaton j . But such fine structure plays no technical role and clutters up notation considerably. Hence, we stick with the (more restricted) notation of a globally determined commit alphabet.

A theoretically simple framework for assumption-commitment in automata is when processes make assumptions only about those other processes that they communicate with. This is naturally modelled by having assumption and commitment on synchronization transitions. Here an automaton A_1 may synchronize with another automaton A_2 on an action $(a, \lambda_1, \lambda_2)$, $a \in \Sigma_1 \cap \Sigma_2$ where we see A_1 as committing to λ_1 provided A_2 commits to λ_2 . Symmetrically, for such a synchronization to occur, A_2 must have a transition on a where it commits to λ_2 . (In this case, A_2 may not require the λ_1 commitment from A_1 .)

2.1 A mutual exclusion example

In order to motivate the kind of framework we are leading to, we give an example of a simple two-process mutual exclusion problem. For simplicity, we abstract away the internal computational states and assume that the processors are always requesting for or executing in the critical section. We model this as follows.

Each process i can either be in state W_i (process i waiting to enter the critical section) or in state C_i (process i is in the critical section) state. In order to gain access to the critical section from the wait state, the processes do a joint action c . Actions a and b are actions taken by processes in the critical section. When an internal action is taken, the process in the critical section goes back to the waiting state.

We set the commitment alphabet $\mathcal{C} = \langle C_1 = (\{p_1, np_1\}, \leq_1), C_2 = (\{p_2, np_2\}, \leq_2) \rangle$. Here, $p_i, i = 1, 2$ denotes that process i is permitted access to critical section and np_i denotes that it is not permitted to enter the critical section. The commit alphabet is shown in figure 1.

The design of process 1 can then be as follows: when 1 is in the state W_1 , it stays in the same state if it is not permitted entry to critical section. When it is permitted entry, assuming that process 2 is not permitted entry, it can go to the state C_1 denoting access to critical section. Process 2 is designed in a symmetric way. Figure 1 shows the two processes and also the product showing the global behaviour. See that the assumptions at the local transition $W_1 \xrightarrow{c} C_1$ are (p_1, np_2) and those at $W_2 \xrightarrow{c} C_2$ are (np_1, p_2) . These assumptions are not compatible because $np_2 \not\leq_2 p_2$ and $np_1 \not\leq_1 p_1$. Hence the global transition $(W_1, W_2) \xrightarrow{c} (C_1, C_2)$ is not possible; so that at no point both the processes can be in the critical section, thus satisfying the safety requirement.

Figure 1. Two-processor mutual exclusion.

It is possible to formally define a class of automata where assumptions and commitments are associated with synchronizing transitions like in this example, along the lines followed in the next section. It then turns out that the class of languages accepted by such systems is the same as the so-called class of *regular consistent languages* (Mohalik & Ramanujam 1997). However, we will not do so, but define a larger class of automata from which these can be obtained by suitable restrictions.

In general, a process may make assumptions about other processes in the system even in the absence of any communication from them. This leads us to a type of systems where at a local state a process makes assumptions about the states in which other processes may be, and in the product only mutually compatible states are admissible. We call these systems Assumption-Compatible Systems (ACS).

We now show how the formal model of assumption compatible systems helps in reasoning about typical problems in distributed computing. For this we choose simplified versions of two well-known problems, namely, that of *reliable bit transmission* and *sequence transmission*.

2.2 Bit transmission problem

There are two processes, a *sender* S and a *receiver* R . Assume that they communicate by asynchronous message passing over a possibly faulty channel. Further, we assume that message loss in the channel is the only kind of fault in the system and that the number of such faults is finite (but unbounded) in any execution sequence.

The sender wants to send one bit (0 or 1) to the receiver. Since messages may get lost, there is no guarantee that a message sent by either of the processes will be received. The problem is to ensure that till R receives the bit, S has to go on sending the bit to R .

Finite state solutions for the above problem are simple. The main idea is to let R send back an acknowledgment when it receives the bit. When S gets the acknowledgment it stops sending the bit. We illustrate how the design can be done in an assumption-commitment framework.

The sender: (See figure 2). The alphabet Σ_1 of sender S is $\{s, g\}$, where

- s : S sends the bit to R , and
- g : S receives the acknowledgment.

S has three states. At the initial state (called p_0), it is yet to send the bit and hence it knows that R cannot have received the bit. So it commits that it has not sent the bit and assumes that R is in a state where it has not received the bit. At the second state (called p_1) it has sent a bit and is waiting for acknowledgment. Now that the bit is sent, S does not know whether R has received the bit or not and if R has received the bit then whether it has sent any acknowledgment or not. In short, it can assume “nothing” about the states of R but commits to having sent the bit and not having got any acknowledgment. At the third state (called p_2), it commits that it has got an acknowledgment from the receiver. Further this can happen only with the assumption that R has received the bit and has sent an acknowledgment.

The receiver: (See figure 3) The alphabet Σ_2 of sender R is $\{r, a\}$, where

- r : R receives the bit from S , and
- a : R sends the acknowledgment.

Figure 2. The sender.

Figure 3. The receiver.

The receiver also has three states. At the initial state (called q_0), it is yet to receive the bit, at the second state (called q_1) it has just received the bit and at the third state (called q_2) it has sent acknowledgment to S . At q_0 , R commits (naturally) that it has not received the bit. Also, since it is yet to send back an acknowledgment, it assumes that S is in a state where it has not got any acknowledgment. Note that at q_0 , the receiver cannot assume anything about whether the sender has sent the bit or not. At q_1 , R commits that it has received the bit and has not yet sent the acknowledgment. Clearly, the assumption is that S must already have sent the bit and that the latter has not yet got the acknowledgment. At q_2 , R commits to having received the bit and sent an acknowledgment. While it cannot assume that the sender has received this acknowledgment, R assumes that the bit has been sent (otherwise, R would not have sent any acknowledgment).

The commitment alphabet of S can then be taken as $(C_1, \leq_1)$ where C_1 is the boolean closure of the set of propositions $\{ack_recd, bit_sent\}$, where the meaning of the propositions is obvious. Similarly, the commitment alphabet of S can then be taken as $(C_2, \leq_2)$ where C_2 is the boolean closure of the set of propositions $\{bit_rec, ack_recd\}$, In both cases the orders $\leq_i$ are logical implication.

With this the assumptions and commitments of the sender at its states can be summarized in the following table. ($\neg bit_sent$ means “bit is not sent” etc.)

	Commitment	Assumption
(S) p_0	$\neg bit_sent$ and $\neg ack_recd$ by S	$\neg bit_recd$ by R
p_1	bit_sent and $\neg ack_recd$ by S	R may be in any state.
p_2	bit_sent and ack_recd by S	bit_recd and ack_sent by R

Similarly, the assumptions and commitments of the receiver at its states are given in the following table.

	Commitment	Assumption
(R) q_0	$\neg bit_recd$ and $\neg ack_sent$ by R	$\neg ack_recd$ by S
q_1	bit_recd and ack_not_sent by R	bit_sent and $\neg ack_recd$ by S
q_2	bit_recd and ack_sent by R	bit_sent by S

Product**Figure 4.** ACS for bit transmission.

Product automaton and global behaviour: Figure 4 gives the product automaton of the ACS consisting of S and R . Notice that the global state (p_0, q_1) (which says R has received the bit before it has been sent) is ruled incompatible since the assumption at q_1 about S is $\lambda = \text{bit_sent} \wedge \neg \text{ack_rcvd}$, the commitment at p_0 is $\mu = \neg \text{bit_sent} \wedge \neg \text{ack_rcvd}$ but $\mu \not\approx \lambda$. Similar compatibility considerations rule out global states (p_2, q_0) , (p_2, q_1) and (p_0, q_2) .

Then, we see from the product automaton that with the final state $F = (p_2, q_2)$, the behaviour of the ACS is $s^+ r s^* a (s + a)^* g a^*$, which captures the desired behaviour of the system, namely, S sends the bit till it receives an acknowledgment and then stops; R starts by receiving the bit and then goes on sending acknowledgments.

2.3 Sequence transmission problem

We now study a slightly more involved protocol, namely the *Sequence Transmission Problem* and discuss how it can be modelled naturally in the assumption - commitment framework. The problem is as follows.

As before, there is a sender S and a receiver R . The sender now wants to send a bit-stream to the receiver. For each message bit, the bit transmission protocol is used. Essentially, S goes on sending the i th message bit till it receives an acknowledgment and then it starts sending the $(i + 1)$ th bit and so on. On the other hand, R initially waits till it gets the first bit. After this, for each bit (say it is the i th bit), it goes on sending the acknowledgment till it receives the $(i + 1)$ th bit. There are, of course, some other requirements of the problem that makes the design slightly harder.

- *Totality.* All the bits of the stream are delivered.
- *Sequentiality.* The bits are delivered in the order in which they occur in the stream.
- *Non-duplication.* A particular bit may be delivered many times over the channel because S might not have got any acknowledgment for the bit, but once R receives the bit and sends acknowledgment, it does not receive the same bit from the channel.

From the description above, one can design the protocol as a series of bit transmission protocols for each bit in the stream. For each bit, if we label the actions and also the assumptions

Figure 5. An infinite state protocol for STP.

and commitments with the position of the bit, we get an infinite state protocol which is given in figure 5 with the assumption and commitments from figure 6. (Note that there are now infinite number of elements of the commitment alphabet; in the table a state p_1i denotes the i -th p_1 state. In the figure it is given as the state p_1 with i quotes.)

	Commitment	Assumption
(S)	p_0 $\neg bit_0_sent$ and $\neg ack_0_recd$ by S	$\neg bit_0_recd$ by R
	p_{1i} bit_i_sent and $\neg ack_i_recd$ by S	R may be in any state.
	p_{2i} bit_i_sent and ack_i_recd by S	bit_i_recd and ack_i_sent by R
	$\neg bit_{i+1}_sent$ and $\neg ack_{i+1}_recd$ by S	$\neg bit_{i+1}_recd$ by R
(R)	q_0 $\neg bit_0_recd$ and $\neg ack_0_sent$ by R	$\neg ack_0_recd$ by S
	q_{1i} bit_i_recd and $ack_i_not_sent$ by R	bit_i_sent and $\neg ack_i_recd$ by S
	q_{2i} bit_i_recd and ack_i_sent by R	bit_i_sent by S
	$\neg bit_{i+1}_recd$ and $\neg ack_{i+1}_sent$ by S	$\neg bit_{i+1}_recd$ by R

Figure 6. Infinitely many assumptions and commitments for the infinite state protocol.

The product shows that the requirements are actually satisfied, namely, for every $i \geq 0$, the action r_i (receive bit i) takes place (totality), it occurs only once (non-duplication) and r_j occurs strictly before r_k when $j < k$ (sequentiality).

Now, we want to *fold* this protocol so that the sender and receiver actually have finite number of states. Since the states are determined by their assumptions and commitments and the actions, this folding would have to bound the labels attached to the letters of the commitment alphabet and also the actions. As a first attempt, if we banish the labels altogether, we get a protocol as in figure 7. Immediately we see that this protocol does not satisfy the requirements.

- Since there is no distinction between consecutive messages, S might be sending the i th message even after it gets an acknowledgment from R . But this is a minor difficulty because we can always ensure that S moves to the $(i + 1)$ th bit after it gets an acknowledgment for the i th message. Hence, sequentiality is ensured.
- But one can see that in the product, between a rec_bit by R and rec_ack by S , there are other rec_ack 's, which means duplication takes place.

Figure 7. An incorrect folding of the infinite state protocol for STP.

Figure 8. A correct folding of the infinite state protocol for STP.

- Also, there is loss of message bits in this design for the following reason: *S* receives an *ack* and sends *i*th message, *R* does not receive it and sends another *ack* corresponding to $(i - 1)$ th message, *S* receives this *ack* and assumes that this is an acknowledgment of the *i*th message. Hence it then starts sending the $(i + 1)$ th message. Thus the message *i* is never received by *R*. The problem arises because when *S* sends message $(i + 1)$, it assumes that *R* has already received message *i*, but then the protocol assumptions and commitments do not reflect this.

In order to remedy this deficiency, we have one bit (0 or 1) attached to the actions of sending and receiving message and acknowledgments and also to the assumptions and

	Commitment	Assumption
(S)	p_0 $\neg bit_0_sent$ and $\neg ack_0_recd$ by S	$\neg bit_0_recd$ by R
	p_1 bit_0_sent and $\neg ack_0_recd$ by S	R may be in any state.
	p_2 bit_0_sent and ack_0_recd by S	bit_0_recd and ack_0_sent by R
	$\neg bit_1_sent$ and $\neg ack_1_recd$ by S	$\neg bit_1_recd$ by R
	p'_1 bit_1_sent and $\neg ack_1_recd$ by S	R may be in any state.
	p_2 bit_1_sent and ack_1_recd by S	bit_1_recd and ack_1_sent by R
(R)	$\neg bit_0_recd$ and $\neg ack_0_sent$ by R	$\neg ack_0_recd$ by S
	q_1 bit_0_recd and $\neg ack_0_sent$ by R	bit_0_sent and $\neg ack_0_recd$ by S
	q_2 bit_0_recd and ack_0_sent by R	bit_0_sent by S
	$\neg bit_1_recd$ and $\neg ack_1_sent$ by S	$\neg bit_1_recd$ by R
	q'_1 bit_1_recd and $\neg ack_1_sent$ by R	bit_1_sent and $\neg ack_1_recd$ by S
	q'_2 bit_1_recd and ack_1_sent by R	bit_1_sent by S
	$\neg bit_0_recd$ and $\neg ack_0_sent$ by S	$\neg bit_0_recd$ by R

Figure 9. Assumptions and commitments for a correct finite state protocol for STP.

commitments. This is essentially to distinguish between consecutive messages and acknowledgments. Thus, we get a protocol as in figure 8 with the assumptions and commitments as in the table in figure 9 and the product shows that this satisfies all the requirements of sequence transmission problem. This, in fact, is the Alternating Bit Protocol (ABP) for the sequence transmission problem.

The above analysis does not give any clue as to why only two distinguishing sets of states were sufficient for the sequence transmission problem. In fact it looks a bit like an accident that we hit upon the ABP. But, at the least, this guides us to have more states to capture relevant assumptions and to design the protocol correctly. While it is not argued here that this transformation presented here makes it easier to design the components (it does not), we remark that it illustrates a paradigm: design the components locally assuming an infinite number of assumptions first and then fold it. The structure of ACS's offers the possibility of designing components 'bearing in mind' global requirements, and translating them into local dependencies. In general, this may involve the construction of an infinite state space (which may be conceptually easy) and then folding it to obtain a finite system.

3. Assumption-compatible systems

We now formally define the class of assumption compatible systems that we have discussed intuitively so far.

Notation. Let Σ be a finite and nonempty alphabet. A *transition system* (TS) over Σ is a tuple $M = (Q, \longrightarrow, q^0)$, where Q is a *finite* set of states, $\longrightarrow \subseteq (Q \times \Sigma \times Q)$ is the transition relation and $q^0 \in Q$ is the initial state. The tuple $A = (M, F)$, where M is a TS and $F \subseteq Q$, is called a *finite state automaton* (FA). When $q_0 \xrightarrow{x} q_k$ and $q_k \in F$, we say that the string x is accepted by A (where $\Longrightarrow$ is the transition relation extended to Σ^*). The set of all strings accepted by A (also called the language of A) is denoted as $L(A)$ or $L(M, F)$. Let Reg_Σ denote the class of languages $L(A)$ accepted by finite state automata.

We model the distribution of actions among the processes by a *distributed alphabet* $\tilde{\Sigma}$. It is a tuple $(\Sigma_1, \dots, \Sigma_n)$, where each Σ_i is a finite nonempty set of actions and is called a *local alphabet*. The local alphabets are not required to be disjoint. In fact, when $a \in \Sigma_i \cap \Sigma_j, i \neq j$, we think of it as a potential synchronization action between i and j .

Given a distributed alphabet $\tilde{\Sigma}$, we often speak of the set $\Sigma \stackrel{\text{def}}{=} \Sigma_1 \cup \dots \cup \Sigma_n$ as the alphabet of the system since the overall behaviour of the system is expressed by strings from Σ^* . For any action in Σ , we have the notion of processes participating in this action. Let $loc : \Sigma \rightarrow 2^{\{1, \dots, n\}}$ be defined by $loc(a) \stackrel{\text{def}}{=} \{i \mid a \in \Sigma_i\}$. So $loc(a)$ (called “locations of a ”) gives the set of processes that participate (or, synchronize) in the action a . By definition, for all $a \in \Sigma, loc(a) \neq \emptyset$. When $loc(a) \cap loc(b) = \emptyset$, we think of a and b as *independent* actions, as they can be performed concurrently by disjoint sets of processes.

A *distributed system* over $\tilde{\Sigma}$ is a tuple $\tilde{M} = (M_1, \dots, M_n)$, where for all i in Loc , $M_i = (Q_i, \rightarrow_i, q_i^0)$ is a TS over Σ_i . Global transitions of $\tilde{M}$ are given by a *product transition system* (product TS): let $Q \stackrel{\text{def}}{=} Q_1 \times \dots \times Q_n$. A TS $\hat{M} = (\hat{Q}, \rightarrow, (q_1^0, \dots, q_n^0))$, where $\hat{Q} \subseteq Q, \rightarrow \subseteq (\hat{Q} \times \Sigma \times \hat{Q})$ and $(q_1^0, \dots, q_n^0) \in \hat{Q}$, is called a *product TS* of $\tilde{M}$ on Σ if it satisfies the *asynchrony* condition: $(p_1, \dots, p_n) \xrightarrow{a} (q_1, \dots, q_n)$ iff $\forall i \in loc(a), p_i \xrightarrow{a} q_i$, and $\forall j \notin loc(a), p_j = q_j$. The language accepted by $\tilde{M}$ is $L(\tilde{M}) \stackrel{\text{def}}{=} L(\hat{M}, F)$. We use $\bar{p}, \bar{q}$ etc. to denote elements of Q . For any global state $\bar{s} = (p_1, p_2, \dots, p_n)$, $\bar{s}[i]$ denotes the i -th component p_i .

Fix a distributed alphabet $\tilde{\Sigma} = (\Sigma_1, \dots, \Sigma_n)$. We also have a *commitment alphabet* $\tilde{C} = ((C_1, \leq_1), \dots, (C_n, \leq_n))$. Each C_i is a nonempty set and for all $i \neq j, C_i \cap C_j = \{\perp\}$. $\leq_i$ is a binary relation on C_i such that $\perp \leq_i c$ for each $c \in C_i$. The element $\perp$ is the null assumption (or commitment). We call $C = C_1 \cup \dots \cup C_n$ the *commit set*. Since we work with finite state systems, it suffices to consider finite C_i . Let $\Phi \stackrel{\text{def}}{=} \{\phi : Loc \mapsto C \mid \forall i \in Loc, \phi(i) \in C_i\}$.

In specific systems, the relations $\leq_i$ may have more structure. For example, if the commitment alphabet is constructed from boolean formulae with implication as the ordering, we get a partial order on the alphabets. In fact, in all our figures of commitment alphabets, we treat them as pre-orders (reflexive and transitive relations) for ease of depiction, but such ordering is not mandatory by definition.

DEFINITION 1

Let $i \in \{1, 2, \dots, n\}$. An *AC transition system* (AC-TS) over $(\Sigma_i, \tilde{C})$ is a tuple (M_i, f_i) where $M_i = (Q_i, \rightarrow_i, q_i^0)$ is a TS over Σ_i and $f_i : Q_i \rightarrow \Phi$ is called an *assumption map*.

At a state $p \in Q_i$, if $f_i(p) = \phi$ then $\phi(i)$ is the commitment of M_i at p and $\phi(j), j \neq i$, is the assumption of M_i about M_j at p .

DEFINITION 2

An *Assumption-compatible system* (ACS) over $(\tilde{\Sigma}, \tilde{C})$ is given by a tuple

$$\tilde{M} = (M_1, M_2, \dots, M_n, \langle f_1, f_2, \dots, f_n \rangle, F),$$

where for each $i \in Loc$, $(M_i = (Q_i, \rightarrow_i, q_i^0), f_i)$ is an AC-TS over $(\Sigma_i, \tilde{C})$, and $F \subseteq (Q_1 \times \dots \times Q_n)$.

Global behaviour of $\tilde{M}$ is given below as that of the product automaton $\hat{M}$ associated with the system.

DEFINITION 3

Given an ACS $\tilde{M} = (M_1, M_2, \dots, M_n, < f_1, f_2, \dots, f_n >, F)$, a global state $(p_1, p_2, \dots, p_n) \in Q$ is called *compatible* iff

$$\text{for all } i, j \in Loc : f_i(p_i)(j) \preceq_j f_j(p_j)(j).$$

DEFINITION 4

The product automaton of $\tilde{M}$ is defined to be $(\hat{M}, F)$ where

- (1) $\hat{M} = (\hat{Q}, \longrightarrow, (q_1^0, \dots, q_n^0))$ is a product TS of $\tilde{M}$ over Σ with $\hat{Q}$ as the set of all compatible global states,
- (2) $(q_1^0, \dots, q_n^0) \in \hat{Q}$, and
- (3) $F \subseteq \hat{Q}$.

The class of languages over $\tilde{\Sigma}$ accepted by ACS's is denoted as $\mathcal{L}(ACS_{\tilde{\Sigma}})$.

$$\mathcal{L}(ACS_{\tilde{\Sigma}}) = \{L \subseteq \Sigma^* \mid \exists \tilde{C} \text{ and an ACS } \tilde{M} \text{ over } (\tilde{\Sigma}, \tilde{C}) \text{ s.t. } L = L(\hat{M}, F)\}.$$

Thus the behaviour of ACS's is given by a product construction. Note that the product contains compatible global states, and the rest are eliminated.

It is worth noting here how an apparent circularity between assumptions and commitments is resolved. If we were to consider a commitment as an 'offer', we could consider the following situation. Process 1, at local state s commits to c assuming d to be available from process 2. If, simultaneously, 2 in local state t offers d' assuming c' to be available from 1, we can have circularity or 'deadlock' when c and c' , as well as d and d' are incompatible. However, the way compatible global states have been defined here, a global state consisting of local states s and t would simply be declared incompatible. Operationally, this can be seen as both processes making their offers first, and then checking whether what is on offer meets with the assumptions made.

The use of compatible states adds a great deal of modelling power. For instance, define the relation $\sim$ on Σ^* as: for all $x, y \in \Sigma^*$, $x \sim y \stackrel{\text{def}}{=} x[i = y[i$ for all $i \in Loc$, where $\lceil : (\Sigma^* \times Loc) \rightarrow \Sigma^*$ is the component projection map giving the maximal subsequence over any local alphabet. It is easy to see that $\sim$ is an equivalence, and it can be verified that $x \sim y$ holds exactly when x can be obtained by commuting independent actions in y . This is a relation of crucial interest in concurrency theory (Mazurkiewicz 1989). Languages accepted by product TS's (defined earlier) are closed under $\sim$, but $\mathcal{L}(ACS_{\tilde{\Sigma}})$ need not be closed under $\sim$. As an example, we describe an ACS in figure 10 that accepts the language $(ab)^*$, where $n = 2$ and $\tilde{\Sigma} = (\{a\}, \{b\})$ (hence a and b are independent). Note that this language is not closed under $\sim$.

Example 1. Let $\tilde{C} = (\mathcal{C}_1, \mathcal{C}_2)$ where the commit alphabet $\mathcal{C}_i$ are as shown in figure 10. For ease of reference, we have annotated the local states by the respective $f_i(\cdot)$. We see that of the possible 9 global states, only 6 are compatible. For example, the global state (p_2, q_2) is compatible because at p_2 , agent 1's assumption about agent 2 is $f_1(p_2)(2) = v_4$, at q_2 agent 2's

Figure 10. An example ACS accepting $(ab)^*$ over the alphabet $(\{a\}, \{b\})$.

commitment is $f_2(q_2)(2) = \nu_2$ and $\nu_4 \preceq_2 \nu_2$. Also, at q_2 , agent 2's assumption about agent 1 is $f_2(q_2)(1) = \mu_5$, agent 1's commitment at p_2 is $f_1(p_2)(1) = \mu_2$ and $\mu_5 \preceq_1 \mu_2$. On the other hand the global state (p_1, q_2) is not compatible because $f_1(p_1)(2) = \nu_6 \not\preceq_2 \nu_2 = f_2(q_2)(2)$. The final states are: $\{(p_1, q_1), (p_2, q_2), (p_3, q_3)\}$.

In the design of ACS's, a crucial decision relates to the choice of commit alphabets. In effect, elements of commit alphabets code up sets of local states or regions of transition systems. That is, when we write $f_1(p_2)(2) = \nu_4$ above, it is an assumption by process 1 at state p_2 that process 2 is in a set of states coded by ν_4 . This approach, as opposed to the 'black-box' way of defining assumptions and commitments over the entire system behaviours, makes it possible for us to reason with a limited form of liveness. We can use commit alphabets and the assumption map to code up situations whereby process 1 eventually makes a condition p true, assuming that until then, 2 is in a region of its transition system, where it eventually gets to a state making a condition q true. We can symmetrically build a similar condition for process 2, and then the product eventually reaches a global state where 1 makes p true and 2

makes q true. However, formalizing such reasoning as a liveness composition rule in a formal logic is left to future work.

The following theorem constitutes the central result of the paper:

Theorem 1. $\mathcal{L}(ACS_{\tilde{\Sigma}}) = Reg_{\Sigma}$.

The inclusion $\mathcal{L}(ACS_{\tilde{\Sigma}}) \subseteq Reg_{\Sigma}$ is easy and follows from the fact that the product automata of ACS's are just finite state automata. The other inclusion, showing that every regular language over Σ is in $\mathcal{L}(ACS_{\tilde{\Sigma}})$, is (understandably) complicated because when we want to accept arbitrary regular languages we need the ability to 'force' specific interleavings. For instance, when a and b are independent actions the language $(ab)^*$ specifies that a is always preferred over b ; coming up with product constructions on automata that achieve such forcing systematically is the difficulty. Hence the problem here is different from that in the construction of, say, asynchronous automata (Zielonka 1987) or cellular asynchronous automata (Cori *et al* 1993), where global states are decomposed preserving concurrency.

The proof is in two stages. First, we construct an automaton for the given regular language where the states are distributed but the transitions are globally specified hence it is not locally presented. From this automaton, we compute the assumptions and commitments for the ACS and distribute the transitions as well so that the product of the ACS accepts the same language. The detailed proof is given in appendix A.

4. Behavioural analogue of compatible products

A product operation on finite state automata corresponds to a shuffle operation on regular languages. Thus we can ask, what manner of shuffle corresponds to the compatible product of automata with assumption and commitment on states? We first answer this question below, and then present a distributed version of Kleene's theorem. For this section, fix a distributed alphabet $\tilde{\Sigma} = (\Sigma_1, \dots, \Sigma_n)$ and a finite commit alphabet $\tilde{\Phi}$. Recall that $\Phi \stackrel{\text{def}}{=} \{\phi : Loc \mapsto \mathcal{C} \mid \forall i \in Loc, \phi(i) \in \mathcal{C}_i\}$. Since each $\mathcal{C}_i$ is finite, $|\Phi|$ is finite.

Note that in AC-transition systems local transitions are labelled by letters from the local alphabet, but the product of these TS's crucially depends on the assumptions and commitments assigned to local states through assumption maps. Hence, if we want to have a shuffle operation corresponding to the product, local languages must encode assumptions and commitments. Then the shuffle operation on these languages should generate strings over Σ from local ones by using this information.

In order to bring assumptions and commitments into languages over distributed alphabets, we extend any given alphabet using assumption maps so that we have actions of the form $\langle a, \phi \rangle$ where $\phi \in \Phi$.

DEFINITION 5

Given a distributed alphabet, $\tilde{\Sigma}$ and a commit alphabet $\tilde{\Phi}$, we define extended alphabets as follows:

$$\begin{aligned} \Sigma_i^c &\stackrel{\text{def}}{=} \{\langle a, \phi \rangle \mid a \in \Sigma_i, \phi \in \Phi\}; \\ \tilde{\Sigma}^c &\stackrel{\text{def}}{=} \langle \Sigma_1^c, \dots, \Sigma_n^c \rangle; \quad \Sigma^c \stackrel{\text{def}}{=} \bigcup_{i \in Loc} \Sigma_i^c. \end{aligned}$$

4.1 Compatible shuffle

We want to define n -way shuffle for strings $x_1, \dots, x_n$, where $x_i \in \Sigma_i^{c*}$, $i \in Loc$. By their structure, each x_i has an assumption map at every point essentially denoting the assumptions of the local state corresponding to x_i . Hence, globally, at every point one has n assumption maps, one for each string, denoting a global state. For a valid shuffle of the given strings, the global states that occur at each point have to be compatible. In terms of strings, the collection of local assumption maps have to be compatible.

Let Ξ denote the set of all possible assumption-commitment tuples for all the agents in the system. Formally, $\Xi = \{\xi \mid \xi : Loc \rightarrow \Phi\}$. For any $i, j \in Loc$, $\xi(i)$ is an assumption map, and $\xi(i)(j) \in \mathcal{C}_j$ is i 's assumption about j . We speak of ξ as an *assumption environment*.

DEFINITION 6

Let $\xi \in \Xi$. ξ is said to be *feasible* iff $\forall i, j \in Loc, \xi(i)(j) \preceq_j \xi(j)(j)$.

DEFINITION 7

$\Sigma^\Xi = \{ \langle a, \xi \rangle \mid a \in \Sigma \text{ and } \xi \in \Xi \}$.

We use $\widehat{x}, \widehat{y}, \dots$ to denote strings over Σ^Ξ which, intuitively, stand for sequences of global states. Since our goal is to establish a correspondence between the global strings and the runs of a compatible product automaton, these strings must somehow capture the compatibility condition internally. We call these strings *good*. Formally, we have the following.

DEFINITION 8

Let $\widehat{x} = \langle a_1, \xi_1 \rangle \langle a_2, \xi_2 \rangle, \dots, \langle a_k, \xi_k \rangle \in \Sigma^{\Xi*}$. Then, $\widehat{x}$ is **good** w.r.t. an *initial environment* ξ_0 iff

- (1) for all $1 \leq l \leq k$, for all $j \notin loc(a_l)$, $\xi_{l-1}(j) = \xi_l(j)$, and
- (2) for all $0 \leq l \leq k$, ξ_l is feasible.

As a minor observation, note that ϵ (the null string) is good w.r.t. some initial environment ξ iff ξ is feasible.

Now we define the notion of when strings over local alphabets can generate global strings. For this we need two projection maps. The *commit erasure* map: $\sigma : \Sigma^{c*} \rightarrow \Sigma^*$ is defined as $\sigma(\langle a_1, \phi_1 \rangle \dots \langle a_k, \phi_k \rangle) \stackrel{\text{def}}{=} a_1 \dots a_k$. We use σ as a commit erasure map on strings over Σ^Ξ as well since there is no scope of confusion here.

The second projection map is the *component projection* map: $\widehat{\cdot} : (\Sigma^{\Xi*} \times Loc) \rightarrow \Sigma^{c*}$ defined by:

$$\widehat{x}\widehat{i} = \begin{cases} \epsilon, & \text{if } x = \epsilon, \\ \widehat{y}\widehat{i}, & \text{if } \widehat{x} = \widehat{y} \cdot \langle a, \xi \rangle \text{ and } i \notin loc(a), \text{ and} \\ (\widehat{y}\widehat{i}) \cdot \langle a, \xi(i) \rangle & \text{if } \widehat{x} = \widehat{y} \cdot \langle a, \xi \rangle \text{ and } i \in loc(a). \end{cases}$$

DEFINITION 9

A string $\widehat{x} \in \Sigma^{\Xi*}$ is called a *witness* for $x \in \Sigma^*$ under $\xi \in \Xi$ if $\widehat{x}$ is good w.r.t. ξ and $\sigma(\widehat{x}) = x$.

Let $x_i \in \Sigma_i^{c*}$, $i \in \{1, \dots, n\}$. A string $x \in \Sigma^*$ is said to be *generated* by $(x_1, x_2, \dots, x_n)$ under ξ if there is a witness $\widehat{x}$ for x under ξ such that for all $i \in Loc$, $x_i = \widehat{x}\widehat{i}$.

Notice that when $x \in \Sigma^*$ is generated by $(x_1, x_2, \dots, x_n)$ under ξ , $x \upharpoonright i = \sigma(x_i)$. This is because, $\sigma(x_i) = \sigma(\widehat{x} \upharpoonright i) = \sigma(\widehat{x}) \upharpoonright i = x \upharpoonright i$.

We now define the compatible shuffle of languages over local extended alphabets using the definition of generation.

DEFINITION 10

For all $i \in Loc$, let $L_i \subseteq \Sigma_i^{c*}$, and let $\xi \in \Xi$. We define the n -ary **compatible shuffle** of these languages under the assumption environment ξ by:

$(L_1 \parallel L_2 \parallel \dots \parallel L_n)_\xi \stackrel{\text{def}}{=} \{x \in \Sigma^* \mid x \text{ is generated by a tuple } (x_1, x_2, \dots, x_n) \text{ under } \xi, \text{ where for all } i, x_i \in L_i\}.$

We are interested in the compatible shuffle of regular languages over the local alphabets. So we define the following class of languages.

DEFINITION 11

Let $\mathcal{L}(AC - shuffle)_{\widetilde{\Sigma}}$ denote the least class that includes the set $\{L \subseteq \Sigma^* \mid \text{for some commit alphabet } \widetilde{C} \text{ and } \xi \text{ over } \widetilde{C}, \text{ there exist regular languages } L_i \subseteq \Sigma_i^{c*} \text{ such that } L = (L_1 \parallel \dots \parallel L_n)_\xi\}$ and is closed under union.

A language in $\mathcal{L}(AC - shuffle)_{\widetilde{\Sigma}}$ will be referred to as an AC-shuffle language (with $\widetilde{\Sigma}$ implicit).

4.2 ACS's and commitment structure

There is a fairly obvious association between states and runs of an ACS and good strings over $\Sigma^{\exists}$. We make it explicit in the following. Let $\widehat{M} = (M_1, \dots, M_n, < f_1, \dots, f_n >, F)$ be an ACS and $\widetilde{M} = (\widehat{Q}, \longrightarrow, < q_0^1, \dots, q_0^n >, F)$ be the compatible product of $\widehat{M}$.

DEFINITION 12

Let $(q_1, \dots, q_n) \in \widehat{Q}$. Then, $env(q_1, \dots, q_n) \stackrel{\text{def}}{=} \xi$, where $\xi(i) = f_i(q_i)$, for all $i \in Loc$.

Let $(q_1, \dots, q_n)$ be a compatible state and $\xi = env(q_1, \dots, q_n)$. Then ξ is feasible. This is because of the following reason. Since $(q_1, \dots, q_n)$ is compatible, for all $i, j \in Loc$, $f_i(q_i)(j) \preceq_j f_j(q_j)(j)$. Then by the construction, for all i, j in Loc , $\xi(i)(j) \preceq_j \xi(j)(j)$. This implies feasibility of ξ .

We can associate strings over Σ^{c*} with runs of $\widehat{M}$ in a canonical fashion. Fix $x = a_1 a_2 \dots a_k$, and a run $\rho = (q_1^0, \dots, q_n^0) \xrightarrow{a_1} (q_1^1, \dots, q_n^1) \dots \xrightarrow{a_k} (q_1^k, \dots, q_n^k)$ on x in $\widehat{M}$. Define $c(\rho) = < a_1, \xi_1 > \dots < a_k, \xi_k > \in \Sigma^{\exists*}$ by: for $1 \leq l \leq k$, $\xi_l = env(q_1^l, \dots, q_n^l)$.

Note that when $k = 0$, that is, when $x = \epsilon$, $c(\rho) = \epsilon$. The initial environment associated with the run is defined as $env(q_1^0, \dots, q_n^0)$. Note that $env(q_1^0, \dots, q_n^0)$ is a feasible assumption environment (because $(q_1^0, \dots, q_n^0)$ is compatible.)

PROPOSITION 1

Let ρ be a run in $(q_1, \dots, q_n) \xRightarrow{x} (p_1, \dots, p_n)$. Then, $c(\rho)$ is a witness for x under $env(q_1, \dots, q_n)$.

Proof. By definition, $\sigma(c(\rho)) = x$. Hence it suffices to show that $c(\rho)$ is good w.r.t. $\text{env}(q_1, \dots, q_n)$.

Let the run ρ be $(q_1^0, \dots, q_n^0) \xrightarrow{a_1} (q_1^1, \dots, q_n^1) \dots \xrightarrow{a_k} (q_1^k, \dots, q_n^k)$.

Fix l such that $1 \leq l \leq k$. For all $j \notin \text{loc}(a_l)$, $q_j^{l-1} = q_j^l$ by asynchrony. Hence, by definition of $c(\rho)$, $\xi_{l-1}(j) = f_j(q_j^{l-1}) = f_j(q_j^l) = \xi_l(j)$.

Also, since all the states on the run are compatible, as we have observed before, the assigned assumption environments are feasible. This proves that $c(\rho)$ is good w.r.t. $\text{env}(q_1, \dots, q_n)$. $\square$

PROPOSITION 2

Let ρ denote a run $(q_1, \dots, q_n) \xrightarrow{x} (p_1, \dots, p_n)$ in $\widehat{M}$. Then for all $i \in \text{Loc}$, there exist $x_i \in \Sigma_i^{c*}$ such that $q_i \xrightarrow{x_i}^c p_i$ and x is generated by the tuple $(x_1, x_2, \dots, x_n)$ under $\text{env}(q_1, \dots, q_n)$.

Proof. By the previous proposition $c(\rho)$ is a witness for x under $\text{env}(q_1, \dots, q_n)$. Take $x_i = c(\rho) \upharpoonright i$, for all $i \in \text{Loc}$. Then, x is generated by the tuple $(x_1, x_2, \dots, x_n)$ under $\text{env}(q_1, \dots, q_n)$. From the definition of $\widehat{\upharpoonright}$ and $\xrightarrow{c}_i$, one can carry out an induction argument on the length of x to show that $q_i \xrightarrow{x_i}^c p_i$. $\square$

PROPOSITION 3

Let $\widehat{x} = \langle a_1, \xi_1 \rangle, \langle a_2, \xi_2 \rangle, \dots, \langle a_k, \xi_k \rangle \in \Sigma^{\Xi*}$ be good w.r.t. $\xi_0 = \text{env}(q_1, \dots, q_n)$ such that $q_i \xrightarrow{\widehat{x} \upharpoonright i}^c p_i$. Then, $\xi_k(i) = f_i(p_i)$. (Since ξ_k is feasible this implies that $(p_1, \dots, p_n)$ is compatible.)

Proof. Let l be the last i -action in $\widehat{x}$. If $l = 0$ (meaning $\widehat{x} \upharpoonright i$ is empty), then $p_i = q_i$ and by goodness of $\widehat{x}$, $\xi_k(i) = \xi_0(i) = f_i(p_i)$.

If $l \neq 0$, then there is an $r_i \in Q_i$ such that $q_i \xrightarrow{y}^c r_i \xrightarrow{\langle a_l, \xi_l(i) \rangle}^c p_i$, where $\widehat{x} \upharpoonright i = y \cdot \langle a_l, \xi_l(i) \rangle$. Hence, $f_i(p_i) = \xi_l(i)$. By goodness of $\widehat{x}$, $\xi_l(i) = \xi_k(i)$. Therefore, $f_i(p_i) = \xi_k(i)$ and we are done. $\square$

PROPOSITION 4

Suppose $x \in \Sigma^*$ and for all $i \in \text{Loc}$, there exist $x_i \in \Sigma_i^{c*}$ such that $q_i \xrightarrow{x_i}^c p_i$ and x is generated by the tuple $(x_1, x_2, \dots, x_n)$ under $\text{env}(q_1, \dots, q_n)$. Then $(q_1, \dots, q_n) \xrightarrow{x} (p_1, \dots, p_n)$ in $\widehat{M}$.

Proof. (by induction on length of x) The assumptions of the claim are rephrased as follows: there exist $x_i \in \Sigma_i^{c*}$ such that $q_i \xrightarrow{x_i}^c p_i$ and there exists a witness $\widehat{x} \in \Sigma^{\Xi*}$ of x under $\text{env}(q_1, \dots, q_n)$ such that $\widehat{x} \upharpoonright i = x_i$. We have to show that $(q_1, \dots, q_n) \xrightarrow{x} (p_1, \dots, p_n)$. Note that since $\widehat{x}$ is good w.r.t. $\text{env}(q_1, \dots, q_n)$, $(q_1, \dots, q_n)$ is compatible from the definition of goodness. From Proposition 3 $(p_1, \dots, p_n)$ is also compatible, so both $(q_1, \dots, q_n)$ and $(p_1, \dots, p_n)$ are in $\widehat{M}$.

Let $x = \epsilon$. Then the witness $\widehat{x}$ must be ϵ since $\sigma(\widehat{x}) = x$. This implies, for all $i \in \text{Loc}$ $x_i = \epsilon$. Hence, $q_i = p_i$ for all i . Then, it is obvious that $(q_1, \dots, q_n) \xrightarrow{x} (p_1, \dots, p_n)$.

For the induction step, let $x = ya$. Then the witness $\widehat{x}$ for x must be of the form $\widehat{x} = \widehat{y} \cdot \langle a, \xi \rangle$, where $\sigma(\widehat{y}) = y$. Since $\widehat{x}$ is a witness under $\text{env}(q_1, \dots, q_n)$ it is good w.r.t. $\text{env}(q_1, \dots, q_n)$. Hence, $\widehat{y}$ must also be good w.r.t. $\text{env}(q_1, \dots, q_n)$. Therefore, $\widehat{y}$ is a witness of y under $\text{env}(q_1, \dots, q_n)$. Let $\widehat{y}[i] = y_i$ for all $i \in \text{Loc}$. Then, $\widehat{y}$ is generated by $(y_1, \dots, y_n)$ under $\text{env}(q_1, \dots, q_n)$.

It is given that for all $i \in \text{Loc}$, $\widehat{x}[i] = x_i$. Hence for every $i \notin \text{loc}(a)$, $x_i = \widehat{x}[i] = \widehat{y}[i] = y_i$ and for every $i \in \text{loc}(a)$, $x_i = y_i \cdot \langle a, \xi(i) \rangle$.

It is also given that for all $i \in \text{Loc}$, $q_i \xRightarrow{x_i}_i^c p_i$. Hence, for all $i \notin \text{loc}(a)$, $q_i \xRightarrow{y_i}_i^c p_i$ and for all $i \in \text{loc}(a)$, there is an $r_i \in Q_i$ such that $q_i \xRightarrow{y_i}_i^c r_i \xrightarrow{\langle a, \xi(i) \rangle}_i^c p_i$. Therefore, $f_i(p_i) = \xi(i)$.

Since we have:

(1) for all $i \notin \text{loc}(a)$, $q_i \xRightarrow{y_i}_i^c p_i$ and

(2) for all $i \in \text{loc}(a)$, $q_i \xRightarrow{y_i}_i^c r_i$,

by induction hypothesis, $(q_1, \dots, q_n) \xRightarrow{y}_i^c (s_1, \dots, s_n)$, where for $i \notin \text{loc}(a)$, $s_i = p_i$ and for $i \in \text{loc}(a)$, $s_i = r_i$.

Since $\widehat{M}$ satisfies the asynchrony condition (definition of Product TS), from the fact that for $i \in \text{loc}(a)$, $s_i = r_i \xrightarrow{\langle a, \xi(i) \rangle}_i^c p_i$, it follows that $(s_1, \dots, s_n) \xrightarrow{a}_i^c (p_1, \dots, p_n)$. Therefore, finally, $(q_1, \dots, q_n) \xRightarrow{ya}_i^c (p_1, \dots, p_n)$, as required. $\square$

4.3 Equivalence of AC-shuffle and AC-systems

We now establish a correspondence between AC-shuffle of languages over Σ_i^c and languages accepted by ACS's. Note that the local automata are over Σ_i . How does one relate languages over Σ_i^c and local automata over Σ_i ? This is simple: since the states are annotated with assumption maps from Φ , we take these into account (like in Moore machines) alongwith the labels of transitions so that languages accepted by local automata are actually over Σ_i^c . For example, if $q_0 \xrightarrow{a_1}_i q_1 \xrightarrow{a_2}_i q_2$, q_0 is the initial state and q_2 is the final state, then we say that this is an accepting path for $\langle a_1, f_i(q_1) \rangle \cdot \langle a_2, f_i(q_2) \rangle$.

Since the local automata are FSA's, languages thus accepted by these are regular over Σ_i^c . But, is the reverse true? In other words, for any regular language L over Σ_i^c , is there an AC-automaton over Σ_i that accepts L ? We show in the following that this is indeed the case.

We first define language acceptance (in the above sense) of AC-automata. Let $(M_i = (Q_i, \xrightarrow{\cdot}_i, q_i^0), f_i)$ be an AC-transition system over Σ_i . We transform M_i into a TS $M'_i = (Q'_i, \xrightarrow{\cdot}_i, q_i^{c0})$ over Σ_i^c where, $Q'_i = Q_i$, $q_i^{c0} = q_i^0$ and $\xrightarrow{\cdot}_i^c \subseteq Q_i^c \times \Sigma_i^c \times Q_i^c$ is defined as follows: $p \xrightarrow{\langle a, \phi \rangle}_i^c q$ iff $p \xrightarrow{a}_i q$ and $\phi = f_i(q)$. This one step transition can be extended to $\xRightarrow{\cdot}_i^c$ for strings from Σ_i^{c*} .

Then given an AC-automaton $((M_i, f_i), F_i)$, we define:

$$L_m((M_i, f_i), F_i) \stackrel{\text{def}}{=} L(M'_i, F_i).$$

PROPOSITION 5

$L \in \text{Reg}_{\Sigma_i^c}$ iff there is an AC-automaton (M, F) over Σ_i such that $L = L_m(M, F)$.

Proof. One direction of the proof follows immediately from the definition of L_m .

For the other direction, suppose $L \in \text{Reg}_{\Sigma_i^c}$. Since L is regular over Σ_i^c , there exist some FSA $(A = (Q, \longrightarrow, q^0), F)$ such that $L = L(A, F)$.

Note that in this FSA, two transitions with different assumption maps may be pointing to the same state. So we refine the states so that transitions that point to a state have the same assumption map. This is possible because the number of assumption maps is finite. We do this as follows.

Define the AC-automaton $((M, f), G)$ as follows. Let $M = (P, \Longrightarrow, p^0)$ where,

- $P = \{q^0\} \cup \bigcup_{q \in Q} \{(q, \phi) \mid \text{there is a transition } p \xrightarrow{<a, \phi>} q \text{ in } A\},$
- $p^0 = q^0,$
- $(p, \phi_1) \xrightarrow{a} (q, \phi_2)$ iff $p \xrightarrow{<a, \phi>} q$ and $\phi_2 = \phi$, and
- for all states (p, ϕ) , $f((p, \phi)) = \phi$.

Finally, $G = \{(p, \phi) \mid p \in F\}$.

In order to check that $L_m((M, f), G) = L$, we transform M in to the following TS $(M^c = (Q^c \xrightarrow{c} q^{c0}), F^c)$ over Σ_i^c , where

- $Q^c = P,$
- $q^{c0} = p^0,$
- $(p, \phi_1) \xrightarrow{<a, \phi>^c} (q, \phi_2)$ iff $(p, \phi_1) \xrightarrow{a} (q, \phi_2)$ and $\phi = f((q, \phi_2)) = \phi_2,$
- Finally, $F^c = G$.

By definition, $L_m(M, G) = L(M^c, F^c)$. So it suffices to show that $L(M^c, F^c) = L(A, F)$.

We note that M^c simulates A via the map Θ , where $\Theta((p, \phi)) = p$ and $\Theta(q^0) = q^0$; that is, states of M^c map to states of A so that transitions of M^c on an action map to transitions on the same action, and initial and final states map to initial and final states respectively. It is then routine to show language equality and the proposition follows. $\square$

We prove the following proposition which is used later to show that languages obtained by AC-systems and shuffle coincide.

PROPOSITION 6

Suppose, for all $i \in \text{Loc}$, $L_i = L_m((M_i, f_i), F_i)$. Take an AC-system $\tilde{M} = (M_1, M_2, \dots, M_n, < f_1, \dots, f_n >, \Pi_{i \in \text{Loc}} F_i)$ with the initial state $(q_1^0, \dots, q_n^0)$. Let ξ_0 denote $\text{env}(q_1^0, \dots, q_n^0)$. Then $L(\tilde{M}) = (\| L_i)_{\xi_0}$.

Proof. ($\subseteq$): Let $x \in L(\tilde{M})$. Then there is a final state $(q_1^f, \dots, q_n^f) \in \Pi_{i \in \text{Loc}} F_i$ such that $(q_1^0, \dots, q_n^0) \xrightarrow{x} (q_1^f, \dots, q_n^f)$. By proposition 2, there exist $x_i \in \Sigma_i^{c*}$ such that $q_i^0 \xrightarrow{x_i^c} q_i^f$ and x is generated by the tuple $(x_1, x_2, \dots, x_n)$ under $\xi_0 = \text{env}(q_1^0, \dots, q_n^0)$. Since each q_i^f is in F_i , each x_i is in L_i . Hence by definition of shuffle, $x \in (\| L_i)_{\xi_0}$. Therefore, $L(\tilde{M}) \subseteq (\| L_i)_{\xi_0}$.

($\supseteq$): Let $x \in (\| L_i)_{\xi_0}$. By definition, there exist $x_i \in \Sigma_i^{c*}$ such that $x_i \in L_i$ and x is generated by the tuple $(x_1, x_2, \dots, x_n)$ under ξ_0 .

Since $x_i \in L_i$, for all $i \in \text{Loc}$, there is some $q_i^f \in F_i$ such that $q_i^0 \xrightarrow{x_i^c} q_i^f$. By proposition 4 we have $(q_1^0, \dots, q_n^0) \xrightarrow{x} (q_1^f, \dots, q_n^f) \in \Pi_{i \in \text{Loc}} F_i$. Hence $x \in L(\tilde{M})$. Therefore, $(\| L_i)_{\xi_0} \subseteq L(\tilde{M})$. $\square$

We state the main theorem connecting AC-shuffle and languages accepted by AC-systems.

Theorem 2. $\mathcal{L}(AC\text{-}Shuffle_{\widetilde{\Sigma}}) = Reg_{\Sigma} = \mathcal{L}(ACS_{\widetilde{\Sigma}})$.

Proof. Consider $L \in \mathcal{L}(ACS_{\widetilde{\Sigma}})$. Then for some ACS $\widetilde{M} = (M_1, M_2, \dots, M_n, < f_1, \dots, f_n >, F)$, $L = L(\widetilde{M})$. One can then write L as $L = \bigcup_{q^f \in F} \mathcal{L}(\widetilde{M}_{q^f})$ where $\widetilde{M}_{q^f} = (M_1, \dots, M_n, < f_1, \dots, f_n >, \{q^f\})$. Let $q^f = (q_1^f, q_2^f, \dots, q_n^f)$ and $L_i = L_m((M_i, f_i), \{q_i^f\})$. By the previous proposition there is an assumption environment ξ_0 such that $L(\widetilde{M}_{q^f}) = (\| L_i)_{\xi_0}$. This places $L(\widetilde{M}_{q^f})$ in $\mathcal{L}(AC\text{-}Shuffle_{\widetilde{\Sigma}})$. Then, since $\mathcal{L}(AC\text{-}Shuffle_{\widetilde{\Sigma}})$ is closed under union, L also is placed in it.

Let $L = (\| L_i)_{\xi_0}$. Since L_i 's are regular over Σ_i^c , by proposition 5, there are AC-automata $((M_i, f_i), F_i)$ such that $L_i = L_m((M_i, f_i), F_i)$. We show that $L \in Reg_{\Sigma}$. By proposition 6, the ACS $M = (M_1, \dots, M_n, < f_1, \dots, f_n >, \Pi_{i \in Loc} F_i)$ accepts L . Since M is an FA over Σ , $L \in Reg_{\Sigma}$. Since any language in $\mathcal{L}(AC\text{-}Shuffle_{\widetilde{\Sigma}})$ is either an AC-shuffle language or a union of AC-shuffle languages, $\mathcal{L}(AC\text{-}Shuffle_{\widetilde{\Sigma}}) \subseteq Reg_{\Sigma}$. The theorem then follows from theorem 1. $\square$

5. A Kleene theorem for ACS's

We now consider the question of syntax for languages in $\mathcal{L}(ACS_{\widetilde{\Sigma}})$. The syntax is given in two layers, one for 'local' expressions and another for parallel composition. Fix a distributed alphabet $\widetilde{\Sigma}$, a commit alphabet $\widetilde{C}$, and the associated extended alphabet.

$$\begin{aligned} ACREG_i &::= \emptyset \mid < a, \phi > \in \Sigma_i^c \mid p + q \mid p; q \mid p^* \\ ACREG &::= (\|_{i=0}^n r_i)_{\xi}, r_i \in ACREG_i, \xi \in \Xi \\ &\mid R_1 + R_2, R_i \in ACREG. \end{aligned}$$

The semantics of these expressions is given as follows: for each $i \in Loc$, we have a map $[\cdot]_i : ACREG_i \rightarrow 2^{\Sigma_i^{c*}}$, and globally a map $[\cdot] : ACREG \rightarrow 2^{\Sigma^*}$. These maps are defined by structural induction:

- $[\emptyset]_i = \emptyset$.
- $[< a, \phi >]_i = \{< a, \phi >\}$.
- $[p + q]_i = [p]_i \cup [q]_i$.
- $[p; q]_i = [p]_i \cdot [q]_i$.
- $[p^*]_i = ([p]_i)^*$.
- $[(r_1 \parallel r_2 \parallel \dots \parallel r_n)_{\xi}] = ([r_1]_1 \widehat{\parallel} [r_2]_2 \widehat{\parallel} \dots \widehat{\parallel} [r_n]_n)_{\xi}$.
- $[R_1 + R_2] = [R_1] \cup [R_2]$.

Thus, the $ACREG_i$ expressions give languages over Σ_i^c and then $ACREG$ expressions are given semantics via AC-shuffle of these local languages and their unions.

The class of regular languages generated by the $ACREG$ expressions is denoted as $\mathcal{L}(ACREG)$. Formally, $\mathcal{L}(ACREG)_{\widetilde{\Sigma}} = \{L \subseteq \Sigma^* \mid \text{for some commit alphabet } \widetilde{C}, \text{ there is an } R \in ACREG \text{ over } \widetilde{\Sigma}^c \text{ such that } L = [R]\}$.

Then, from the semantics and theorem 2, we get the following characterization.

Theorem 3. $\mathcal{L}(ACS)_{\Sigma}^{\approx} = \mathcal{L}(ACREG)_{\Sigma}^{\approx} = \text{Reg}_{\Sigma}$.

5.1 A different syntax

The syntax presented above is not entirely satisfactory, since every expression in the language necessarily involves assumptions and commitments. Typically we wish to make assumptions only at *some* control points rather than at all control points. Moreover, the ξ parameter in the parallel composition operator is awkward. As it turns out, these are not serious problems. Consider the modified syntax (where *FACREG* stands for assumptions and commitments occurring ‘free’ within expressions):

$$\begin{aligned} FACREG_i &::= \emptyset \mid a \in \Sigma_i \mid \phi \in \Phi \mid p + q \mid p; q \mid p^* \\ FACREG &::= r_1 \parallel r_2 \parallel \cdots \parallel r_n, r_i \in FACREG_i \mid R_1 + R_2, R_i \in FACREG \end{aligned}$$

We wish to map *FACREG_i* expressions via a function $\langle \rangle_i$ to languages over Σ_i^c so that at the global level, *FACREG* expressions can be given semantics by AC-shuffle of local languages as before. The natural semantics of *FACREG_i* expressions give regular languages over $\Sigma \cup \Phi$ (call the semantic function $[\]_i$). We translate these languages to languages over Σ_i^c in a systematic way, preserving regularity.

We describe a translation scheme that uniformly translates each string of the given language. We illustrate this with a running example to make the basic ideas clear. Let $x = a_1\phi_1\phi_2a_2a_3\phi_3a_4 \in (\Sigma \cup \Phi)^*$. The basic idea is to first convert every string over $(\Sigma \cup \Phi)$ to one where the letters and assumption maps alternate.

- (1) if there are consecutive letters we insert a $\perp$ in between them. Thus we translate x to $a_1\phi_1\phi_2a_2\perp a_3\phi_3a_4$.
- (2) if there are consecutive assumption maps then retain only the last one in the sequence. Thus we get $a_1\phi_2a_2\perp a_3\phi_3a_4$.
- (3) ensure that there is an assumption map at the beginning and at the end. If there are not any, we put $\perp$. Thus, we get $\perp a_1\phi_2a_2\perp a_3\phi_3a_4\perp$.
- (4) from the first letter onwards, pair up consecutive letter and assumption map. Thus, finally, $\perp \cdot < a_1, \phi_2 > \cdot < a_2, \perp > \cdot < a_3, \phi_3 > \cdot < a_4, \perp > \cdot$.

For some more examples, note the following:

- (1) $a_1\phi_1a_2\phi_2$ is translated to $\perp \cdot < a_1, \phi_1 > \cdot < a_2, \phi_2 > \cdot$,
- (2) $a_1a_2\phi_1\phi_2$ is translated to $\perp \cdot < a_1, \perp > \cdot < a_2, \phi_2 > \cdot$, and
- (3) $\phi_0a_1\phi_1a_2$ is translated to $\phi_0 \cdot < a_1, \phi_1 > \cdot < a_2, \perp > \cdot$.

It should be clear that this translation is actually a function, call it h such that $h : (\Sigma_i \cup \Phi)^* \rightarrow \Phi \cdot \Sigma_i^{c*}$. Also following the steps described above, one can show that h can be expressed as a composition of homomorphisms, hence h is itself a homomorphism.

Suppose $L \subseteq (\Sigma \cup \Phi)^*$ is regular. Then, $L' = h(L)$ is a regular language over $\Phi \cup \Sigma_i^c$. Moreover, since for any string in L' , only the first element is in Φ and $|\Phi|$ is finite, L' can be expressed as $\bigcup_{\phi \in \Phi} \phi \cdot L'_\phi$ where L'_ϕ is regular over Σ_i^c .

DEFINITION 13

Let $r \in FACREG_i$. Then $\langle r \rangle_i \stackrel{\text{def}}{=} h([r]_i)$.

Consider the languages $\phi_i \cdot L_i, i \in Loc$. The shuffle of these languages is given as: $(L_1 \hat{\parallel} \cdots \hat{\parallel} L_n)_{(\phi_1, \dots, \phi_n)}$ which is the AC-shuffle of the L_i 's with the initial environment $(\phi_1, \dots, \phi_n)$.

The *FACREG* expressions are now given semantics via AC-shuffle.

DEFINITION 14

Let $R = (r_1 \hat{\parallel}_F \cdots \hat{\parallel}_F r_n) \in FACREG$ and $\langle r_i \rangle_i = \bigcup_{k \in \Gamma_i} \phi_{ik} L_{ik}$. Then,

$$\langle R \rangle \stackrel{\text{def}}{=} \bigcup_{j_l \in \Gamma_l, l \in Loc} (L_{1j_1} \hat{\parallel} \cdots \hat{\parallel} L_{nj_n})_{(\phi_{1j_1}, \dots, \phi_{nj_n})}.$$

By very definition, $\mathcal{L}(FACREG_{\widetilde{\Sigma}}) \subseteq \mathcal{L}(ACREG_{\widetilde{\Sigma}})$.

In order to show that the other inclusion also holds, we need to prove that for any *ACREG* expression r , there is a *FACREG* expression r' such that $\langle r' \rangle = [r]$.

From the way h is described above, it is easy to see that the alphabetic homomorphism $d : \Sigma_i^c \rightarrow (\Sigma_i \cdot \Phi)$ defined as $d(< a, \phi >) = a\phi$ suffices for the proof, because then $h(\phi \cdot d(x)) = \phi \cdot x$. We omit the monotonous technical details and summarize the result below.

Theorem 4. $\mathcal{L}(ACREG_{\widetilde{\Sigma}}) = \mathcal{L}(FACREG_{\widetilde{\Sigma}})$.

6. Discussion

We have proved that ACS's over $\widetilde{\Sigma}$ characterize all regular languages over Σ . Hence, languages accepted by ACS's are closed under complementation, union and concatenation. Indeed, the construction gives a *deterministic* ACS for any regular language, hence we have determinization as well. We can also directly define constructions on ACSs to show closure under boolean operations though we do not as yet have a direct proof of determinization.

A natural extension of the notions here relates to *infinite* behaviour of ACS's. We can indeed show that using appropriate Muller acceptance conditions, the results here extend to ω -regular languages as well (Mohalik 1999). This is of importance in the context of verifying liveness properties of systems, which are specified on infinite system computations.

An interesting consideration on ACSs relates to *sub-regular* languages. It is easily seen that the class of languages obtained by shuffle or synchronized shuffle is but a small subclass of regular languages. A larger class of sub-regular languages are those regular languages closed under the equivalence $\sim$ defined in §3. We can show that these classes can be obtained by simple restrictions on the structure of the commit alphabets of ACS's. Proceeding along these lines, we can study ACSs with a specific commit structure and we find new classes of sub-regular languages (Mohalik 1999).

It should be clear that ACS's have strong relationships with structures studied in concurrency theory. In fact, the techniques that we use while constructing distributed state automata in the process of decomposing regular behaviours arise mainly from partial order models of concurrency. It is also worth noting here that the study of ACSs arose from the initial work (Ramanujam 1995) where we studied such a decomposition theorem in a categorical

framework, obtaining local presentations for a subclass of event structures. We first proved (Mohalik & Ramanujam 1997) a decomposition theorem only for the class of regular trace languages, and only later generalized it to the results here. Thus, it seems evident that ACS's bear a formal relationship to automata over partial orders, but a precise elucidation of such a relationship seems difficult.

These remarks are intended here to emphasize the fact that the automata theory of ACS's does seem to be rich in interesting notions. The study of these, particularly the complexity of decision questions on ACS's, is important for the design and verification methodology that we outline here.

We conclude the paper with some remarks on the relevance of ACS's in the context of compositional verification introduced in §1. Note that each process i in an ACS over alphabet Σ_i is describing a regular subset of Σ_i^* . Consider a linear time temporal logic defined as follows (along the lines of Ramanujam 1996b): Let P_i be a set of atomic propositions, $i \in Loc$. The formulas of Φ_i are defined by:

$$\alpha \in \Phi_i ::= p \in P_i \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid \langle a \rangle \alpha, a \in \Sigma_i \mid \alpha_1 \mathbf{U} \alpha_2 \\ \mid [now](\alpha @ j), j \in Loc, \alpha \in \Phi_j.$$

The semantics of these formulas can be defined easily on runs of ACS's. Note that formulas of Φ_i are interpreted on time instants on the 'local clock' of process i , and $[now]\alpha @ j$ constrains any global state to be such that the local instant for process j satisfies α . Thus, this modality acts like an assumption specification; the commitment made by a process is simply whatever is true at that state. Now let Ψ define the syntax of global formulas:

$$\phi \in \Psi ::= \alpha @ j, j \in Loc, \alpha \in \Phi_j \mid \neg\phi \mid \phi_1 \vee \phi_2.$$

Now, we can associate an ACS with every formula ϕ such that the language it accepts exactly corresponds to the models of the formula ϕ . Such a *formula automaton construction* is at the heart of automata based model checking procedures. We are currently investigating how we can exploit the structure of the constructed formula automaton for compositional model checking.

Appendix A. Proof of completeness

Here we prove the main result, that all regular behaviours on a distributed alphabet can be accepted by Assumption-Compatible systems.

Theorem A1. $Reg_\Sigma \subseteq \mathcal{L}(ACS_{\widetilde{\Sigma}})$.

We first define an automaton for the given regular language where the states are distributed but the transitions are globally specified hence it is not locally presented. From this automaton, we compute the assumptions and commitments for the ACS and distribute the transitions as well so that the product of the ACS accepts the same language.

A.1 Distributed state automaton

A Distributed State Automaton (DSA) on $\widetilde{\Sigma}$ is a tuple $\mathcal{A} = (A_1, \dots, A_n, \longrightarrow_{\mathcal{A}}, F)$, where

- (1) for every $i \in Loc$, $A_i = (Q_i, \longrightarrow_{\mathcal{A}i}, s_i^0)$ is the i -th TS on Σ_i ,

- (2) $\widehat{Q} = \prod_{i \in Loc} Q_i$ is the state space of $\mathcal{A}$,
- (3) $(s_1^0, \dots, s_n^0)$ is the initial state,
- (4) $F \subseteq Q$ is the set of final states, and
- (5) $\longrightarrow_{\mathcal{A}} \subseteq (Q \times \Sigma \times Q)$ satisfies the following condition:
 if $((p_1, p_2, \dots, p_n), a, (q_1, q_2, \dots, q_n)) \in \longrightarrow_{\mathcal{A}}$ then
 - (a) for all $i \in loc(a)$, $(p_i, a, q_i) \in \longrightarrow_{\mathcal{A}i}$, and
 - (b) for all $j \notin loc(a)$, $p_j = q_j$.

Notice that the transition relation $\longrightarrow_{\mathcal{A}}$ satisfies only one half of the asynchrony condition. This is the crucial difference between DSA's and locally presented systems. In DSA's global transitions are given *a priori* while in locally presented systems, global transitions for the product are *constructed* by the asynchrony condition. We emphasize that DSA's are intended only as intermediate representation for a given regular language. We believe that this gives some structure to both the construction of the ultimate locally presented automaton and the proofs needed to show language equivalence. (We will continue refer to condition (5) above as the asynchrony condition.)

When $(\bar{p}, a, \bar{q}) \in \longrightarrow_{\mathcal{A}}$ we write it as $\bar{p} \xrightarrow{a}_{\mathcal{A}} \bar{q}$. $\mathcal{A}$ is *deterministic* if $\longrightarrow_{\mathcal{A}}$ is a deterministic relation.

A.2 Properties of DSA

Fix a *deterministic* DSA $\mathcal{A}$ over $\widetilde{\Sigma}$. Recall the projection operator $\lceil : (\Sigma^* \times Loc) \rightarrow \Sigma^*$.

$$x \lceil i = \begin{cases} \epsilon, & \text{if } x = \epsilon, \\ y \lceil i, & \text{if } x = ya \text{ and } i \notin loc(a), \\ (y \lceil i)a, & \text{if } x = ya \text{ and } i \in loc(a). \end{cases}$$

Recall also that for a non-empty string $x = a_1 \dots a_k$, $last(x) \stackrel{\text{def}}{=} a_k$.

Now we define $\Downarrow : (\Sigma^* \times Loc) \rightarrow \Sigma^*$ as follows.

DEFINITION A1

$$x \Downarrow i \stackrel{\text{def}}{=} \begin{cases} \epsilon & \text{if } x \lceil i = \epsilon \\ y & \text{such that } \exists u \in \Sigma^* : x = yu, \\ & \text{last}(y) \in \Sigma_i \text{ and } u \lceil i = \epsilon \\ & \text{if } x \lceil i \neq \epsilon. \end{cases}$$

$x \Downarrow i$ gives the maximal prefix of x ending in an i -action. For example, let $\widetilde{\Sigma} = \{\{a, c\}, \{b, c\}\}$. Then $ab \Downarrow 1 = a$, $abcaa \Downarrow 2 = abc$ and $abcaac \Downarrow 1 = abcaac$.

We make a couple of simple observations about $\Downarrow$. Both of them follow from the maximality of $x \Downarrow i$ in x .

- (1) For all $x \in \Sigma^*$, $i, j \in Loc$, if $x \Downarrow i$ is a prefix of $x \Downarrow j$ then $x \Downarrow i = (x \Downarrow j) \Downarrow i$.
- (2) For all $x \in \Sigma^*$, $i, j \in Loc$, if $x \Downarrow i = ya$ and $x \Downarrow j = y'a$, then $x \Downarrow j = x \Downarrow i$ and hence $y = y'$.

Let, for all $x \in \Sigma^*$, $(x)_{\mathcal{A}} \stackrel{\text{def}}{=} (q_1, \dots, q_n) \in Q$ such that $(s_1^0, \dots, s_n^0) \xrightarrow{x}_{\mathcal{A}} (q_1, \dots, q_n)$. Since $\mathcal{A}$ is deterministic, $(x)_{\mathcal{A}}$ is well-defined.

DEFINITION A2

The *events* associated with strings over Σ are defined inductively as follows:

$$event(x) \stackrel{\text{def}}{=} \begin{cases} ((\epsilon)_{\mathcal{A}}, \epsilon, (\epsilon)_{\mathcal{A}}) & \text{if } x = \epsilon \\ ((y)_{\mathcal{A}}, a, (ya)_{\mathcal{A}}) & \text{if } x = ya \end{cases}$$

An event, say $event(x)$, is called an i -event if either $x = \epsilon$ or $x = ya$ and $a \in \Sigma_i$. In the DSA $\mathcal{A}$ there is a precedence relation among the events associated with strings. The definition below says when an i -event precedes a j -event.

DEFINITION A3

$event(u) \text{ } i \triangleright\!-\!j \text{ } event(v)$ iff there is an $r \in \Sigma^*$ such that $r \Downarrow i$ is a prefix of $r \Downarrow j$, $event(u) = event(r \Downarrow i)$ and $event(v) = event(r \Downarrow j)$.

It is clear from this and definition of $\Downarrow$ that $event(u) \text{ } i \triangleright\!-\!j \text{ } event(v)$ iff there is an $r \in \Sigma^*$ such that $r = r \Downarrow j$, $event(v) = event(r)$ and $event(u) = event(r \Downarrow i)$.

We prove certain properties of DSA's in terms of the definitions above. An immediate corollary of the asynchrony condition on $\mathcal{A}$ is the following.

PROPOSITION A1

For all $x \in \Sigma^*$, $(x \Downarrow i)_{\mathcal{A}}[i] = (x)_{\mathcal{A}}[i]$.

PROPOSITION A2

Let $i, j \in loc(a)$. $event(xa) \text{ } i \triangleright\!-\!j \text{ } event(ya)$ implies $(x)_{\mathcal{A}} = (y)_{\mathcal{A}}$.

Proof. Suppose $event(xa) \text{ } i \triangleright\!-\!j \text{ } event(ya)$. By definition, there exists an $r \in \Sigma^*$ such that $event(r \Downarrow i) = event(xa)$ and $event(r \Downarrow j) = event(ya)$ and $r \Downarrow i$ is a prefix of $r \Downarrow j$ (see figure A1).

From the definition of $event$, we have

- $((x)_{\mathcal{A}}, a, (xa)_{\mathcal{A}}) = ((r')_{\mathcal{A}}, c, (r'c)_{\mathcal{A}})$, where $r \Downarrow i = r'c$, and
- $((y)_{\mathcal{A}}, a, (ya)_{\mathcal{A}}) = ((r'')_{\mathcal{A}}, d, (r''d)_{\mathcal{A}})$, where $r \Downarrow j = r''d$.

From these conditions, we get $c = d = a$. From the observation about $\Downarrow$, we immediately get that $r \Downarrow i = r \Downarrow j$, or, equivalently, $r'a = r''a$. This means $r' = r''$ and therefore $(x)_{\mathcal{A}} = (r')_{\mathcal{A}} = (r'')_{\mathcal{A}} = (y)_{\mathcal{A}}$. $\square$

DEFINITION A4

A DSA $\mathcal{A}$ is **stutter-free** if for every $\bar{p}, \bar{q} \in \widehat{Q}$,

$$\bar{p} \xrightarrow{a}_{\mathcal{A}} \bar{q} \text{ implies for all } i \in loc(a), \bar{p}[i] \neq \bar{q}[i].$$

Informally, this means that transitions change the local states of all participating agents.

PROPOSITION A3

Suppose $\mathcal{A}$ is stutter-free. Let $u, w \in \Sigma^*$, $a \in \Sigma_i$ and $k \in Loc$. Then $event(ua) \text{ } i \triangleright\!-\!k \text{ } event(w)$ and $event(u \Downarrow i) \text{ } i \triangleright\!-\!k \text{ } event(w)$ can not both be true simultaneously.

Figure A1. $\text{event}(xa) \text{ i-}\triangleright\text{-j } \text{event}(ya)$.

Proof. Suppose that $\text{event}(ua) \text{ i-}\triangleright\text{-k } \text{event}(w)$ and $\text{event}(u \Downarrow i) \text{ i-}\triangleright\text{-k } \text{event}(w)$ simultaneously.

Claim. Let $x, y, z \in \Sigma^*$, $i, k \in \text{Loc}$ and suppose both $\text{event}(x) \text{ i-}\triangleright\text{-k } \text{event}(z)$ and $\text{event}(y) \text{ i-}\triangleright\text{-k } \text{event}(z)$. Then, $(x)_{\mathcal{A}}[i] = (y)_{\mathcal{A}}[i]$.

Assume the claim. Then, $(ua)_{\mathcal{A}}[i] = (u \Downarrow i)_{\mathcal{A}}[i] = (u)_{\mathcal{A}}[i]$, thus violating the stutter-free condition. But $\mathcal{A}$ is given to be stutter-free, hence we get a contradiction, thus proving the proposition.

Proof of claim. Suppose $\text{event}(x) \text{ i-}\triangleright\text{-k } \text{event}(z)$. Then, there is an $r \in \Sigma^*$ such that $r \Downarrow k$, $\text{event}(z) = \text{event}(r)$ and $\text{event}(x) = \text{event}(r \Downarrow i)$. Hence, from definition of event, we have $(r)_{\mathcal{A}} = (z)_{\mathcal{A}}$ and $(x)_{\mathcal{A}} = (r \Downarrow i)_{\mathcal{A}}$.

Therefore, $(x)_{\mathcal{A}}[i] = (r \Downarrow i)_{\mathcal{A}}[i]$ and using Proposition A1 we get $(x)_{\mathcal{A}}[i] = (r)_{\mathcal{A}}[i] = (z)_{\mathcal{A}}[i]$.

Arguing similarly, we have $(y)_{\mathcal{A}}[i] = (z)_{\mathcal{A}}[i]$. Therefore, $(x)_{\mathcal{A}}[i] = (y)_{\mathcal{A}}[i]$. This proves the claim and the proposition. $\square$

DEFINITION A5

$\mathcal{A}$ is *four-alternation-free* if for all $i, k \in \text{Loc}$, for $a, c \in \Sigma_i \setminus \Sigma_k$, $b, d \in \Sigma_k \setminus \Sigma_i$ and $r \in \Sigma^*$, $(r \Downarrow i = r_1 a, r_1 \Downarrow k = r_2 b, r_2 \Downarrow i = r_3 c, r_3 \Downarrow k = r_4 d \text{ and } r_4 \Downarrow i = r_5)$ implies $(r_5)_{\mathcal{A}} \neq (r)_{\mathcal{A}}$.

The *four* in the definition is because i -transitions and k -transitions alternate four times in the loop (see figure A2).

This definition says that a particular kind of “bad” loops are not present in a four-alternation-free DSA. This characteristic of a DSA helps in synthesizing a behaviour preserving ACS as we will see in the next section.¹

¹We really do not have any intuition to offer as to why only “four” works. This seems to be connected to some kind of minimal unfolding of the DSA necessary for the synthesis

Figure A2. When $\mathcal{A}$ is four-alternation-free, such loops are not present in $\mathcal{A}$.

PROPOSITION A4

Suppose DSA $\mathcal{A}$ is four-alternation-free. Let $u, y \in \Sigma^*$, $a \in \Sigma$, $i \in \text{loc}(a)$, $k \notin \text{loc}(a)$ such that

$$(\text{event}(ua) \text{ } i \rightarrow k \text{ } \text{event}(y) \text{ and } \text{event}(y) \text{ } k \rightarrow i \text{ } \text{event}(u \downarrow i)).$$

Then, $(y)_{\mathcal{A}}[k] = (u)_{\mathcal{A}}[k]$.

Proof. By definition of $\rightarrow$, we have

- (1) $\exists r \in \Sigma^*$ such that $\text{event}(ua) = \text{event}(r \downarrow i)$, $\text{event}(y) = \text{event}(r \downarrow k)$ and $r \downarrow i$ is a prefix of $r \downarrow k$. Hence,

$$r \downarrow i = (r \downarrow k) \downarrow i, (ua)_{\mathcal{A}} = (r \downarrow i)_{\mathcal{A}}, \text{ and } (y)_{\mathcal{A}} = (r \downarrow k)_{\mathcal{A}}. \quad (\text{A1})$$

- (2) $\exists r' \in \Sigma^*$ such that $\text{event}(y) = \text{event}(r' \downarrow k)$, $\text{event}(u \downarrow i) = \text{event}(r' \downarrow i)$ and $r' \downarrow k$ is a prefix of $r' \downarrow i$. Hence,

$$r' \downarrow k = (r' \downarrow i) \downarrow k, (y)_{\mathcal{A}} = (r' \downarrow k)_{\mathcal{A}} \text{ and } (u \downarrow i)_{\mathcal{A}} = (r' \downarrow i)_{\mathcal{A}}. \quad (\text{A2})$$

Claim. $u \downarrow k$ is a prefix of $u \downarrow i$.

Figure A3. The case $u \downarrow k \neq (u \downarrow i) \downarrow k$.

Assume the claim. Then, $u \downarrow k = (u \downarrow i) \downarrow k$. (recall the simple observation we made about $\downarrow$). Then,

$$\begin{aligned}
 (u)_{\mathcal{A}}[k] &= (u \downarrow k)_{\mathcal{A}}[k] && \text{(from proposition A1)} \\
 &= ((u \downarrow i) \downarrow k)_{\mathcal{A}}[k] && \text{(by the claim)} \\
 &= (u \downarrow i)_{\mathcal{A}}[k] && \text{(from proposition A1 again)} \\
 &= (r' \downarrow i)_{\mathcal{A}}[k] && \text{(from (A2) above)} \\
 &= ((r' \downarrow i) \downarrow k)_{\mathcal{A}}[k] && \text{(from proposition A1 yet again)} \\
 &= (r' \downarrow k)_{\mathcal{A}}[k] && \text{(from (A2) above)} \\
 &= (y)_{\mathcal{A}}[k] && \text{(from (A2) above),}
 \end{aligned}$$

and thus we prove the proposition.

Proof of claim. Suppose, $u \downarrow k$ is not a prefix of $u \downarrow i$. This implies $(u \downarrow i)$ is a prefix of $u \downarrow k$ or, in other words,

$$u \downarrow i = (u \downarrow k) \downarrow i. \quad (\text{A3})$$

We show that this leads to a contradiction. See figure A.2.

- Because of (A3), u can be written as $x_1 c x_2 d x_3$, for some $x_1 \in \Sigma^*$, x_2 in $(\Sigma \setminus \Sigma_i)^*$, $x_3 \in (\Sigma \setminus (\Sigma_i \cup \Sigma_k))^*$ where $u \downarrow i = x_1 c$, $c \in \Sigma_i$, $d \in \Sigma_k \setminus \Sigma_i$.
- Since $r \downarrow i = (r \downarrow k) \downarrow i$, we can write $r \downarrow k$ as $(r \downarrow i) \cdot x_4 b$, for some x_4 in $(\Sigma \setminus \Sigma_i)^*$ and $b \in \Sigma_k \setminus \Sigma_i$, by (A1) above.
- Finally, since $r' \downarrow k = (r' \downarrow i) \downarrow k$, we can write $r' \downarrow i$ as $(r' \downarrow k) \cdot x_5 c$, for some $x_5 \in (\Sigma \setminus \Sigma_k)^*$ and $c \notin \Sigma_i$, by (A2) above.

Note that

$$\begin{aligned}
 (x_1 c)_{\mathcal{A}} &= (u \downarrow i)_{\mathcal{A}}, && \text{from item 1 above,} \\
 &= (r' \downarrow i)_{\mathcal{A}}, && \text{since } \text{event}(u \downarrow i) = \text{event}(r' \downarrow i), \\
 &= (r' \downarrow k \cdot x_5 c)_{\mathcal{A}}, && \text{from item 3 above.}
 \end{aligned}$$

By (A1) and (A2) above, $(r \Downarrow k)_{\mathcal{A}} = (y \Downarrow k)_{\mathcal{A}} = (r' \Downarrow k)_{\mathcal{A}}$.

Hence, we get

$$\begin{aligned}
 (x_1 c)_{\mathcal{A}} &= (r \Downarrow k \cdot x_5 c)_{\mathcal{A}} && \text{(from preceding observation)} \\
 &= ((r \Downarrow i \cdot x_4 b) \cdot x_5 c)_{\mathcal{A}} && \text{(from item 2)} \\
 &= (((ua) \cdot x_4 b) \cdot x_5 c)_{\mathcal{A}} && \text{(from (A1) above)} \\
 &= (x_1 c \cdot x_2 d \cdot x_3 a \cdot x_4 b \cdot x_5 c)_{\mathcal{A}} && \text{(from item 1).}
 \end{aligned}$$

Let $y_5 = x_1 c$ and $y = x_1 c \cdot x_2 d \cdot x_3 a \cdot x_4 b \cdot x_5 c$. Then what we have got here is

$$(y)_{\mathcal{A}} = (y_5)_{\mathcal{A}}. \quad (\text{A4})$$

Also, one verifies that

- (1) $y \Downarrow i = x_1 c \cdot x_2 d \cdot x_3 a \cdot x_4 b \cdot x_5 c$, since $c \in \Sigma_i$. Let $y_1 = x_1 c \cdot x_2 d \cdot x_3 a \cdot x_4 b \cdot x_5$.
- (2) $y_1 \Downarrow k = x_1 c \cdot x_2 d \cdot x_3 a \cdot x_4 b$, since $b \in \Sigma_k$ and $x_5 \in (\Sigma \setminus \Sigma_k)^*$. Let $y_2 = x_1 c \cdot x_2 d \cdot x_3 a \cdot x_4$.
- (3) $y_2 \Downarrow i = x_1 c \cdot x_2 d \cdot x_3 a$, since $a \in \Sigma_i$ and $x_4 \in (\Sigma \setminus \Sigma_i)^*$. Let $y_3 = x_1 c \cdot x_2 d \cdot x_3$.
- (4) $y_3 \Downarrow k = x_1 c \cdot x_2 d$, since $d \in \Sigma_k$ and $x_3 \in (\Sigma \setminus \Sigma_k)^*$. Let $y_4 = x_1 c \cdot x_2$.
- (5) Finally, $y_4 \Downarrow i = x_1 c = y_5$, since $c \in \Sigma_i$ and $x_2 \in (\Sigma \setminus \Sigma_i)^*$.

But since $\mathcal{A}$ is four-alternation-free and y and y_5 satisfy the assumptions of for four-alternation-freeness, $(y)_{\mathcal{A}} \neq (y_5)_{\mathcal{A}}$. Hence, from (A4) we get a contradiction. Thus we settle the claim and hence the proposition. $\square$

A.3 From DSA to ACS

The idea behind introducing the notions of *stutter-freeness* and *four-alternation-freeness* is that DSA's with these properties facilitate construction of equivalent ACS's. We demonstrate this construction first. Then, given a regular language, we will show how to construct a DSA with the above-mentioned properties. Thus we will get an ACS for the given regular language.

The following theorem establishes the required correspondence between DSA's and ACS's over $\tilde{\Sigma}$.

Theorem A2. *Let $\mathcal{A}$ be a stutter-free and four-alternation-free deterministic DSA on $\tilde{\Sigma}$. Then, there exists an ACS $\tilde{M}$ on $\tilde{\Sigma}$ such that $L(\mathcal{A}) = L(\tilde{M})$.*

Proof. We first define some sets that will be used to construct the commit alphabet and local states of the ACS.

DEFINITION A6

For $i \in Loc$, $S_i \stackrel{\text{def}}{=} \{event(x) \mid x \in \Sigma^*, x \Downarrow i = x\}$.

The commit alphabet $\mathcal{C}$ is defined as follows: for all $i \in Loc$, $C_i = 2^{S_i}$ and $\lambda \preceq_i \mu$ iff $\lambda \supseteq \mu$.

Now we construct an ACS $\tilde{M} = (M_1, \dots, M_n, < f_1, \dots, f_n >, F)$ over $(\tilde{\Sigma}, \tilde{\mathcal{C}})$ from the DSA $\mathcal{A}$ as follows.

For each $i \in Loc$, i -local TS $M_i = (Q_i, \longrightarrow_i, q_i^0)$ where

- $Q_i = S_i$,
- $q_i^0 = event(\epsilon)$, and

- $p \xrightarrow{a}_i q$ iff there is a $u \in \Sigma^*$ such that $p = \text{event}(u \Downarrow i)$ and $q = \text{event}(ua \Downarrow i)$.
- Local assumption maps f_i are defined as follows.
For every $\text{event}(x) \in Q_i$, $f_i(\text{event}(x))(j) = \{\text{event}(y \Downarrow j) \mid \text{event}(y \Downarrow i) = \text{event}(x)\}$.

Lastly, the set of final states is defined to be

$$F = \{(\text{event}(x \Downarrow 1), \dots, \text{event}(x \Downarrow n)) \mid x \in L(\mathcal{A})\}.$$

Let $(\widehat{M} = (\widehat{Q}, \longrightarrow, (q_1^0, \dots, q_n^0)), F)$ be the compatible product automaton of $\widetilde{M}$ where $\widehat{Q}$ is the set of all compatible states. Since $L(\widetilde{M}) \stackrel{\text{def}}{=} L(\widehat{M}, F)$, we show in the following that $L(\mathcal{A}) = L(\widehat{M}, F)$.

We first observe some properties of the assumption maps.

- (1) For all $i \in \text{Loc}$, $f_i(\text{event}(x))(i) = \{\text{event}(x)\}$.

Proof. For $\text{event}(x) \in Q_i$, $x \Downarrow i = x$. Hence, by construction,

$$f_i(\text{event}(x))(i) = \{\text{event}(y \Downarrow i) \mid \text{event}(y \Downarrow i) = \text{event}(x)\} = \{\text{event}(x)\}.$$

□

An immediate consequence is the following.

- (1) For all $i, j \in \text{Loc}$,

$$f_i(\text{event}(x))(j) \leq_j f_j(\text{event}(y))(j) \text{ iff } \{\text{event}(y)\} \subseteq f_i(\text{event}(x))(j).$$

- (2) For all $i, j \in \text{Loc}$ and $q \in Q_i$,

$$f_i(q)(j) = \{p \in Q_j \mid \text{either } p \text{ j} \triangleright \text{-i } q \text{ or } q \text{ i} \triangleright \text{-j } p\}.$$

Proof. Let $s \in f_i(q)(j)$. Then, there is an $x \in \Sigma^*$ such that $s = \text{event}(x \Downarrow j)$ and $q = \text{event}(x \Downarrow i)$. Since either $x \Downarrow j \leq x \Downarrow i$ or vice-versa, we have either $s \text{ j} \triangleright \text{-i } q$ or $q \text{ i} \triangleright \text{-j } s$ and hence $s \in RHS$.

The other inclusion follows from the definition of assumption maps. □

We now observe some properties of the compatible product.

- (1) For all $x \in \Sigma^*$, $(\text{event}(x \Downarrow 1), \dots, \text{event}(x \Downarrow n))$ is compatible.

Proof. We need to show that $f_i(\text{event}(x \Downarrow i))(j) \leq_j f_j(\text{event}(x \Downarrow j))(j)$ for all $i, j \in \text{Loc}$. This is equivalent to showing $\text{event}(x \Downarrow j) \in f_i(\text{event}(x \Downarrow i))(j)$ (by the observation about assumption maps above).

By definition,

$$f_i(\text{event}(x \Downarrow i))(j) = \{\text{event}(y \Downarrow j) \mid \text{event}(y \Downarrow i) = \text{event}(x \Downarrow i)\}.$$

Therefore, the property is proved by setting y to be x . □

- (2) For all $x \in \Sigma^*$,

$$(\text{event}(x \Downarrow 1), \dots, \text{event}(x \Downarrow n)) \xrightarrow{a} (\text{event}(xa \Downarrow 1), \dots, \text{event}(xa \Downarrow n)).$$

Proof. Notice first that the given states are compatible by the preceding result. One just has to check that the asynchrony condition holds for the transition.

(1) For $i \notin \text{loc}(a)$, $x \Downarrow i = xa \Downarrow i$, hence, $\text{event}(x \Downarrow i) = \text{event}(xa \Downarrow i)$.

(2) From the definition of $\longrightarrow_i$, it directly follows that for all $i \in \text{loc}(a)$,

$$\text{event}(x \Downarrow i) \xrightarrow{a}_i \text{event}(xa \Downarrow i), \text{ and we are done.} \quad \square$$

In the light of the above, by a simple induction on the length of strings we get for all $x \in \Sigma^*$, $(\text{event}(\epsilon), \dots, \text{event}(\epsilon)) \xrightarrow{x} (\text{event}(x \Downarrow 1), \dots, \text{event}(x \Downarrow n))$.

These observations immediately give us the following lemma proving one direction of theorem A2.

Lemma A1. $L(\mathcal{A}) \subseteq L(\tilde{M})$.

Proof. Let $x \in L(\mathcal{A})$. Then, $(\text{event}(x \Downarrow 1), \dots, \text{event}(x \Downarrow n)) \in F$ by construction. Also, from the observation above,

$$\text{event}(\epsilon), \dots, \text{event}(\epsilon) \xrightarrow{x} (\text{event}(x \Downarrow 1), \dots, \text{event}(x \Downarrow n)).$$

Hence, $x \in L(\tilde{M})$.

We now prove the more difficult direction, which uses the special properties of the given DSA.

Lemma A2. $L(\tilde{M}) \subseteq L(\mathcal{A})$.

Proof. Let $x \in L(\tilde{M})$. Then there is a $u \in L(\mathcal{A})$, such that there is a path $(\text{event}(\epsilon), \dots, \text{event}(\epsilon)) \xrightarrow{x} (\text{event}(u \Downarrow 1), \dots, \text{event}(u \Downarrow n))$.

Claim. $\forall x \in \Sigma^*$, $(\text{event}(\epsilon), \dots, \text{event}(\epsilon)) \xrightarrow{x} (\text{event}(z_1), \dots, \text{event}(z_n))$ implies for all $i \in \text{Loc}$, $\text{event}(z_i) = \text{event}(x \Downarrow i)$.

Assuming the claim, for all $i \in \text{Loc}$, $\text{event}(x \Downarrow i) = \text{event}(u \Downarrow i)$. This implies for all $i \in \text{Loc}$, $(x \Downarrow i)_{\mathcal{A}} = (u \Downarrow i)_{\mathcal{A}}$, and hence $(x)_{\mathcal{A}}[i] = (u)_{\mathcal{A}}[i]$. So $(x)_{\mathcal{A}} = (u)_{\mathcal{A}}$. This means that both x and u lead to the same state in $\mathcal{A}$. Since $u \in L(\mathcal{A})$, and $\mathcal{A}$ is deterministic, $(u)_{\mathcal{A}} \in F_{\mathcal{A}}$. Therefore, $(x)_{\mathcal{A}} \in F_{\mathcal{A}}$ which means $x \in L(\mathcal{A})$.

Proof of claim. The proof is by induction on $|x|$. The base case is trivial. Assume that the claim holds for all strings of length k . Take $x = ya$.

Suppose $(\text{event}(\epsilon), \dots, \text{event}(\epsilon)) \xrightarrow{ya} (\text{event}(z_1), \dots, \text{event}(z_n))$. Then there is a state $(\text{event}(y_1), \dots, \text{event}(y_n)) \in \hat{Q}$ such that

$$\begin{aligned} (\text{event}(\epsilon), \dots, \text{event}(\epsilon)) &\xrightarrow{y} (\text{event}(y_1), \dots, \text{event}(y_n)), \text{ and} \\ (\text{event}(y_1), \dots, \text{event}(y_n)) &\xrightarrow{a} (\text{event}(z_1), \dots, \text{event}(z_n)). \end{aligned}$$

By induction hypothesis

$$\text{event}(y_i) = \text{event}(y \Downarrow i) \text{ for all } i \in \text{Loc}. \quad (\text{A5})$$

We have to show that $\text{event}(z_i) = \text{event}(ya \Downarrow i)$ for all $i \in \text{Loc}$.

Case 1. $i \notin \text{loc}(a)$. Then, $\text{event}(z_i) = \text{event}(y \Downarrow i)$, by asynchrony and induction hypothesis. Since $\text{event}(y \Downarrow i) = \text{event}(ya \Downarrow i)$ for $i \notin \text{loc}(a)$, we are done.

Case 2. $i \in \text{loc}(a)$. Since $ya \Downarrow i = ya$, we have to show $\text{event}(z_i) = \text{event}(ya)$.

By asynchrony, for all $j \in \text{loc}(a)$, $\text{event}(y_i) \xrightarrow{a}_i \text{event}(z_i)$. Hence, from the construction of $\xrightarrow{i}$, we have, for all $j \in \text{loc}(a)$, a string $u_j \in \Sigma^*$ such that

$$\text{event}(y \Downarrow j) = \text{event}(u_j \Downarrow j), \text{ and} \quad (\text{A6})$$

$$\text{event}(u_j a) = \text{event}(z_j). \quad (\text{A7})$$

So it suffices to show that $\text{event}(u_i a) = \text{event}(ya)$. From the definition of event , it suffices to show that $(u_i)_{\mathcal{A}} = (y)_{\mathcal{A}}$ or equivalently,

$$\boxed{\text{(to show) for all } k \in \text{Loc}, (u_i)_{\mathcal{A}}[k] = (y)_{\mathcal{A}}[k].}$$

We do this separately for two cases: $k \notin \text{loc}(a)$ and $k \in \text{loc}(a)$.

Let $k \in \text{loc}(a)$. Because the state $(\text{event}(z_1), \dots, \text{event}(z_n))$ is compatible, we know that $\text{event}(z_i) \in f_k(\text{event}(z_k))(i)$. Hence, $\text{event}(u_i a) \in f_k(\text{event}(u_k a))(i)$ by (A7).

From the property of assumption maps, it follows that either

$$\text{event}(u_i a) \text{ i-}\triangleright\text{-k event}(u_k a) \text{ or } \text{event}(u_k a) \text{ k-}\triangleright\text{-i event}(u_i a).$$

In both the cases, from proposition A2, $(u_i)_{\mathcal{A}} = (u_k)_{\mathcal{A}}$. Hence $(u_i)_{\mathcal{A}}[k] = (u_k)_{\mathcal{A}}[k]$.

By (A6) and using proposition A1 we get, for all $m \in \text{loc}(a)$, $(u_m)_{\mathcal{A}}[m] = (y)_{\mathcal{A}}[m]$. Therefore, $(u_i)_{\mathcal{A}}[k] = (y)_{\mathcal{A}}[k]$.

Now, let $k \notin \text{loc}(a)$.

- (1) By compatibility of $(\text{event}(z_1), \dots, \text{event}(z_n))$, $\text{event}(z_i) \in f_k(\text{event}(z_k))(i)$.

Therefore, by property of assumption maps,

$$[\text{either } \text{event}(z_i) \text{ i-}\triangleright\text{-k event}(z_k) \text{ or } \text{event}(z_k) \text{ k-}\triangleright\text{-i event}(z_i)].$$

- (2) By compatibility of $(\text{event}(y_1), \dots, \text{event}(y_n))$, $\text{event}(y_i) \in f_k(\text{event}(y_k))(i)$.

Therefore, by property of assumption maps

$$[\text{either } \text{event}(y_i) \text{ i-}\triangleright\text{-k event}(z_k) \text{ or } \text{event}(z_k) \text{ k-}\triangleright\text{-i event}(y_i)].$$

We consider the possible cases. Remember that we are considering the case when $i \in \text{loc}(a)$ and $k \notin \text{loc}(a)$. Hence, in the following, $\text{event}(z_k) = \text{event}(y_k)$.

Case $\text{event}(z_k) \text{ k-}\triangleright\text{-i event}(z_i)$: There is an $r \in \Sigma^*$ such that $r \Downarrow i = r$, $\text{event}(z_i) = \text{event}(r)$ and $\text{event}(z_k) = \text{event}(r \Downarrow k)$.

From this and (A7), we get $\text{event}(u_i a) = \text{event}(z_i) = \text{event}(r)$ and from asynchrony, we get $\text{event}(y_k) = \text{event}(z_k) = \text{event}(r \Downarrow k)$.

Then

$$\begin{aligned} (y)_{\mathcal{A}}[k] &= (y \Downarrow k)_{\mathcal{A}}[k] \quad (\text{by proposition A1}) \\ &= (y_k)_{\mathcal{A}}[k] \quad (\text{by induction hypothesis (A5)}) \\ &= (r \Downarrow k)_{\mathcal{A}}[k] \quad (\text{since } \text{event}(y_k) = \text{event}(r \Downarrow k)) \\ &= (r)_{\mathcal{A}}[k] \quad (\text{by proposition A1}) \\ &= (u_i a)_{\mathcal{A}}[k] \quad (\text{since } \text{event}(u_i a) = \text{event}(r)) \\ &= (u_i)_{\mathcal{A}}[k] \quad (\text{since } k \notin \text{loc}(a)), \end{aligned}$$

which is the required result.

Case $\text{event}(z_i) \text{ i-}\triangleright\text{-k event}(z_k)$ and $\text{event}(y_i) \text{ i-}\triangleright\text{-k event}(z_k)$: We have $\text{event}(z_i) = \text{event}(u_i a)$ by (A7) and $\text{event}(y_i) = \text{event}(u_i \Downarrow i)$ by (A5) and (A6). Hence, for this case,

$$\text{event}(u_i a) \text{ i-}\triangleright\text{-k event}(z_k) \text{ and } \text{event}(u_i \Downarrow i) \text{ i-}\triangleright\text{-k event}(z_k).$$

Then, since we assume that $\mathcal{A}$ is *stutter-free*, by proposition A3 this case is not possible.

Case $event(z_i) \text{ i-}\triangleright\text{-}k \text{ event}(z_k)$ and $event(z_k) \text{ k-}\triangleright\text{-}i \text{ event}(y_i)$: Recall that for the case under consideration (namely, $i \in loc(a)$, $k \notin loc(a)$),

- $event(z_i) = event(u_i a)$ (from (A7)),
- $event(z_k) = event(y_k)$ (by asynchrony) and
- $event(y_i) = event(u_i \Downarrow i)$ (from (A5) and (A6)).

Hence, $event(u_i a) \text{ i-}\triangleright\text{-}k \text{ event}(y_k)$ and $event(y_k) \text{ k-}\triangleright\text{-}i \text{ event}(u_i \Downarrow i)$. Then, since the DSA $\mathcal{A}$ is four-alternation-free, from Proposition A4, we have $(u_i)_{\mathcal{A}}[k] = (y_k)_{\mathcal{A}}[k]$. Then, from (A5), applying Proposition A1, we finally get $(u_i)_{\mathcal{A}}[k] = (y)_{\mathcal{A}}[k]$.

Thus we have proved that for all locations $k \in Loc$, $(u_i)_{\mathcal{A}}[k] = (y)_{\mathcal{A}}[k]$. Thereby we prove the claim and the lemma. $\square$

Lemmas A1 and A2 prove theorem A2, stated below again.

Theorem A2. *Let $\mathcal{A}$ be a stutter-free and four-alternation-free deterministic DSA on $\tilde{\Sigma}$. Then, there exists an ACS $\tilde{M}$ on $\tilde{\Sigma}$ such that $L(\mathcal{A}) = L(\tilde{M})$.*

A.4 Regular languages and DSA

Because of theorem A2, in order to show that ACS's over $\tilde{\Sigma}$ characterize the class of all regular languages over Σ , it suffices to prove the following theorem.

Theorem A3. *Let the class of deterministic DSA's over $\tilde{\Sigma}$ that are stutter-free and four-alternation-free be denoted as FDSA. Then $\mathcal{L}(FDSA_{\tilde{\Sigma}}) = Reg_{\Sigma}$.*

Proof. Since the global automaton of a DSA (hence that of an FDSA) is a finite state automaton over Σ , languages accepted by DSA's are regular over Σ . Thus we get the easy direction.

For the other direction, let $L \in Reg_{\Sigma}$. We construct a deterministic FDSA $\mathcal{A}_{\tilde{\Sigma}}$ such that $L(\mathcal{A}) = L$ and $\mathcal{A}$ has the required properties.

DEFINITION A7 (*A labelling scheme*)

Let $l : \Sigma^* \rightarrow (Loc \times Loc \rightarrow \{0, 1, 2\})$ be a labeling function defined inductively as follows: $l(\epsilon) \stackrel{\text{def}}{=} C$ where $C(i, j) = 0$ for every $i, j \in Loc$. Let $l(x) = C$. Then $l(xa) \stackrel{\text{def}}{=} C'$ where

$$C'(i, j) = \begin{cases} C(i, j) & \forall i \notin loc(a) \\ 0 & \forall i, j \in loc(a), \text{ and} \\ (C(j, i) + 1) \bmod 3 & \forall i \in loc(a), j \notin loc(a). \end{cases}$$

(When for some $x \in \Sigma^*$, $l(x) = C$, we use the notation C_{ij} to denote $C(i, j)$).

Note that, by definition, $l(x)(i, i) = 0$ for all x and $i \in Loc$.

Example A1. Let $\tilde{\Sigma} = \{a, c\}$. Then,

$$\begin{aligned} l(\epsilon) &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, l(a) = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, l(ac) = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, l(ab) = \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}, l(aba) = \begin{pmatrix} 0 & 0 \\ 2 & 0 \end{pmatrix}, \\ l(abab) &= \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}. \end{aligned}$$

PROPOSITION A5

For all $x \in \Sigma^*$, $i, j \in Loc$, $l(x)(i, j) = l(x \Downarrow i)(i, j)$.

Proof. (by induction on $|x|$)

When $x = \epsilon$, $l(x)(i, j) = l(\epsilon)(i, j) = l(x \Downarrow i)(i, j)$. Suppose $x = ya$. By induction hypothesis, for all $i, j \in Loc$, $l(y)(i, j) = l(y \Downarrow i)(i, j)$.

Case 1. $a \in \Sigma_i$. $x = x \Downarrow i$ and hence the proposition holds.

Case 2. $a \notin \Sigma_i$. Then, by definition of l , $l(y)(i, j) = l(ya)(i, j) = l(x)(i, j)$. Since $y \Downarrow i = ya \Downarrow i = x \Downarrow i$, using induction hypothesis we have $l(x)(i, j) = l(y)(i, j) = l(y \Downarrow i)(i, j) = l(x \Downarrow i)(i, j)$. $\square$

PROPOSITION A6

If $x \Downarrow j$ is a proper prefix of $x \Downarrow i$ then $l(x)(i, j) = l(x)(j, i) + 1 \bmod 3$.

Proof. By the preceding result,

$$l(x)(i, j) = l(x \Downarrow i)(i, j) \text{ and } l(x)(j, i) = l(x \Downarrow j)(j, i).$$

By the given condition, $x \Downarrow i = (x \Downarrow j) \cdot ua$, for some $a \in (\Sigma_i \setminus \Sigma_j)$ and $u \in \Sigma^*$ such that $u \upharpoonright j = \epsilon$. Hence, $x \Downarrow j = (x \Downarrow j \cdot u) \Downarrow j$. Then from previous observation we can derive:

$$\begin{aligned} l(x)(i, j) &= l(x \Downarrow i)(i, j), \\ &= l(x \Downarrow j \cdot u)(j, i) + 1 \bmod 3 \\ &= l((x \Downarrow j \cdot u) \Downarrow j)(j, i) + 1 \bmod 3 \\ &= l(x \Downarrow j)(j, i) + 1 \bmod 3 \\ &= l(x)(j, i) + 1 \bmod 3 \end{aligned}$$

$\square$

Let L be a regular language over Σ . We define the following relations on Σ^* . It can be easily checked that these are equivalence relations of finite index.

DEFINITION A8

Let $x \setminus L \stackrel{\text{def}}{=} \{y \in \Sigma^* \mid xy \in L\}$. For all $i \in Loc$, $x \equiv_i y$ iff $l(x \Downarrow i) = l(y \Downarrow i)$ and $(x \Downarrow i) \setminus L = (y \Downarrow i) \setminus L$.

Lemma A3. If (for all $x, y \in \Sigma^*$ and for all $i \in Loc$, $x \equiv_i y$) then $x \setminus L = y \setminus L$.

Proof. Case 1. $x = \epsilon$ and $y = \epsilon$. Lemma holds trivially.

Case 2. $x = \epsilon$ and $y \neq \epsilon$. Let $j \in loc(last(y))$. Then $y \Downarrow j = y$. Also $x \Downarrow j = \epsilon = x$. By the assumption of the lemma, $x \equiv_j y$. By definition of $\equiv_j$, $x \setminus L = y \setminus L$.

Case 3. $x \neq \epsilon$ and $y \neq \epsilon$. Consider the following claim.

Claim. Let x, y be non-empty strings over Σ^* such that for all $k \in Loc$ $x \equiv_k y$. Let $last(x) = a$ and $last(y) = b$. Then $loc(a) \cap loc(b) \neq \emptyset$.

Assuming the claim, there is a $k \in Loc$ such that $k \in loc(last(x)) \cap loc(last(y))$. Hence, $x \Downarrow k = x$ and $y \Downarrow k = y$. Then, since $x \equiv_k y$ (by the antecedent of the lemma), by definition of $\equiv_k$, then, $x \setminus L = y \setminus L$.

Proof of claim. Suppose $loc(a) \cap loc(b) = \emptyset$. Take $i \in loc(a)$ and $j \in loc(b)$. Then, $x \Downarrow i = x$ and $y \Downarrow j = y$. Let $x \Downarrow j = x'$ and $y \Downarrow i = y'$. By our assumption, x' (resp y') is a proper prefix of x (resp. y).

Since $x \equiv_i y$, $l(x) = l(x \Downarrow i) = l(y \Downarrow i) = l(y')$. Let $l(x) = C = l(y')$. Then, by proposition A6, $l(y) = C'$, where $C'_{ji} = C_{ij} + 1 \bmod 3$.

Further, since $x \equiv_j y$, $l(x') = l(x \Downarrow j) = l(y \Downarrow j) = l(y)$. Now $l(x') = C' = l(y)$. Again, by proposition A6, $C_{ij} = C'_{ji} + 1 \bmod 3$.

From the previous paragraph, we then get, $C_{ij} = C_{ij} + 2 \bmod 3$, which is a contradiction. Thus we prove the claim and the lemma. $\square$

The labelling above is so designed as to construct a DSA that is four-alternation-free. In order to ensure that the DSA is also stutter-free, we introduce an obvious notion as below.

DEFINITION A9 (*i*-parity)

Define *i*-parity : $\Sigma^* \rightarrow \{0, 1\}$ as: for all $x \in \Sigma^*$,

$$i\text{-parity}(x) \stackrel{\text{def}}{=} |x[i]| \bmod 2.$$

For example, let $\tilde{\Sigma} = \{\{a, c\}, \{b, c\}\}$. Then, $1\text{-parity}(a) = 1$, $2\text{-parity}(bab) = 0$ and $1\text{-parity}(abab) = 0$.

The following proposition regarding *i*-parity of strings follows easily from the definition of $\Downarrow$.

PROPOSITION A7

For all $x \in \Sigma^*$ and $i \in Loc$, $i\text{-parity}(x) = i\text{-parity}(x \Downarrow i)$. Therefore, for all $i \notin loc(a)$, $i\text{-parity}(x) = i\text{-parity}(xa)$.

Now we define, for each $i \in Loc$, equivalence classes that are to be the local states of the DSA.

DEFINITION A10

For all $x, y \in \Sigma^*$, $i \in Loc$, $x \sim_i y$ iff $x \equiv_i y$ and $i\text{-parity}(x) = i\text{-parity}(y)$.

Let $[x]_i$ denote the equivalence class under $\sim_i$ containing x .

The following results follow directly from definition of $\sim_i$ and lemma A3.

COROLLARY A1

- (1) $[xa]_i = [x]_i$ for all $i \notin loc(a)$.
- (2) If (for all $x, y \in \Sigma^*$ and for all $i \in Loc$, $x \sim_i y$) then $x \setminus L = y \setminus L$.

The local automata A_i of the constructed DSA $\mathcal{A}$ are defined as $(Q_i, \longrightarrow_{A_i}, s_i^0)$, where $Q_i = \{[x]_i \mid x \in \Sigma^*\}$, $s_i^0 = ([\epsilon]_i, 0)$ and $\longrightarrow_{A_i} = \{([x]_i, a, [xa]_i) \mid x \in \Sigma^*\}$. Finally, $\mathcal{A} = (A_1, A_2, \dots, A_n, \longrightarrow_{\mathcal{A}}, F)$, where $\longrightarrow_{\mathcal{A}} = \{([x]_1, [x]_2, \dots, [x]_n), a, ([xa]_1, [xa]_2, \dots, [xa]_n) \mid x \in \Sigma^*\}$ and $F = \{([x]_1, [x]_2, \dots, [x]_n) \mid x \in L\}$.

Example A2. Let $\tilde{\Sigma} = \{\{a\}, \{b\}\}$. In figure A4 we construct a DSA for the language $(ab)^*$. The *i*-equivalences are computed from table 1.

So, $Q_1 = \{[\epsilon]_1\} \cup \{[(ab)^m a]_1 \mid m \leq 5\}$, and $Q_2 = \{[\epsilon]_2\} \cup \{[(ab)^m]_2 \mid m \leq 6\}$.

Table A1. Computation of i -equivalences.

x	$l(x)$	1-Parity(x)	y	$l(y)$	2-Parity(y)
ϵ	$l(\epsilon) = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$	0	ϵ	$l(\epsilon) = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$	0
a	$l(a) = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$	1	ab	$l(ab) = \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}$	1
aba	$l(aba) = \begin{pmatrix} 0 & 0 \\ 2 & 0 \end{pmatrix}$	0	$(ab)^2$	$l((ab)^2) = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$	0
$(ab)^2a$	$l((ab)^2a) = \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}$	1	$(ab)^3$	$l((ab)^3) = \begin{pmatrix} 0 & 2 \\ 0 & 0 \end{pmatrix}$	1
$(ab)^3a$	$l((ab)^3a) = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$	0	$(ab)^4$	$l((ab)^4) = \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}$	0
$(ab)^4a$	$l((ab)^4a) = \begin{pmatrix} 0 & 0 \\ 2 & 0 \end{pmatrix}$	1	$(ab)^5$	$l((ab)^5) = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$	1
$(ab)^5a$	$l((ab)^5a) = \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}$	0	$(ab)^6$	$l((ab)^6) = \begin{pmatrix} 0 & 2 \\ 0 & 0 \end{pmatrix}$	0
$(ab)^6a$	$l((ab)^6a) = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$	1	$(ab)^7$	$l((ab)^7) = \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}$	1

A.5 Properties of the constructed DSA $\mathcal{A}$

- (1) We observe that by construction, $\longrightarrow_{\mathcal{A}}$ is deterministic.
(2) $\mathcal{A}$ is stutter-free.

Proof. Suppose $\bar{p} \xrightarrow{a}_{\mathcal{A}} \bar{q}$. We need to show that for all $i \in \text{loc}(a)$, $\bar{p}[i] \neq \bar{q}[i]$.

By our construction, there is an $x \in \Sigma^*$ such that $\bar{p} = ([x]_1, [x]_2, \dots, [x]_n)$ and $\bar{q} = ([xa]_1, [xa]_2, \dots, [xa]_n)$. For all $i \in \text{loc}(a)$, $i\text{-parity}(x) \neq i\text{-parity}(xa)$, hence $[x]_i \neq [xa]_i$. Therefore, $\bar{p}[i] \neq \bar{q}[i]$. $\square$

- (3) $\mathcal{A}$ is four-alternation-free.

Proof. Let $r \in \Sigma^*$ be such that $r \Downarrow i = r_1a$, $r_1 \Downarrow k = r_2b$, $r_2 \Downarrow i = r_3c$, $r_3 \Downarrow k = r_4d$ and $r_4 \Downarrow i = r_5$. Further, let $a, c \in \Sigma_i \setminus \Sigma_k$ and $b, d \in \Sigma_k \setminus \Sigma_i$. We show that $(r_5)_{\mathcal{A}} \neq (r)_{\mathcal{A}}$ which shows $\mathcal{A}$ is four-alternation-free.

We show that $l(r)(i, k) \neq l(r_5)(i, k)$, which gives us the result. Let $l(r_5)(i, k) = c \bmod 3$. We know that for all $x \in \Sigma^*$, $i, j \in \text{Loc}$, $l(x)(i, j) = l(x \Downarrow i)(i, j)$. We repeatedly use this observation below.

$$\begin{aligned}
l(r_4)(i, k) &= l(r_4 \Downarrow i)(i, k) \\
&= l(r_5)(i, k) && \text{(assumption)} \\
&= c \bmod 3, && \text{(say)} \\
l(r_3)(k, i) &= l(r_3 \Downarrow k)(k, i) \\
&= l(r_4d)(k, i) && \text{(assumption)} \\
&= (l(r_4)(i, k) + 1) \bmod 3 && \text{(since } d \in \Sigma_k \setminus \Sigma_i) \\
&= (c + 1) \bmod 3,
\end{aligned}$$

Figure A4. A DSA for the language $(ab)^*$.

$$\begin{aligned}
 l(r_2)(i, k) &= l(r_2 \Downarrow i)(i, k) \\
 &= l(r_3c)(i, k) && \text{(assumption)} \\
 &= (l(r_3)(k, i) + 1) \bmod 3 && \text{(since } c \in \Sigma_i \setminus \Sigma_k) \\
 &= (c + 2) \bmod 3, \\
 l(r_1)(k, i) &= l(r_1 \Downarrow k)(k, i) \\
 &= l(r_2b)(k, i) && \text{(assumption)} \\
 &= (l(r_2)(i, k) + 1) \bmod 3 && \text{(since } b \in \Sigma_k \setminus \Sigma_i) \\
 &= c \bmod 3, \\
 l(r)(i, k) &= l(r \Downarrow i)(i, k) \\
 &= l(r_1a)(i, k) && \text{(assumption)} \\
 &= (l(r_1)(k, i) + 1) \bmod 3 && \text{(since } a \in \Sigma_i \setminus \Sigma_k) \\
 &= (c + 1) \bmod 3.
 \end{aligned}$$

Hence, $l(r_5)(i, k) \neq l(r)(i, k)$, as required. $\square$

Lemma A4. $\forall x \in \Sigma^* : (([\epsilon]_1, 0), ([\epsilon]_2, 0), \dots, ([\epsilon]_n, 0)) \xRightarrow{x}_{\mathcal{A}} ([x]_1, [x]_2, \dots, [x]_n)$.

Proof. By the construction, $\longrightarrow_{\mathcal{A}}$ is deterministic and by corollary A1(1) $\longrightarrow_{\mathcal{A}}$ satisfies the *asynchrony* property. Hence we get the result by induction on $|x|$. $\square$

PROPOSITION A8

$$L(\mathcal{A}) = L.$$

Proof. ($\supseteq$): Let $x \in L$. Then, by the previous lemma

$$([\epsilon]_1, [\epsilon]_2, \dots, [\epsilon]_n) \xRightarrow{x}_{\mathcal{A}} ([x]_1, [x]_2, \dots, [x]_n) \in F.$$

Hence $x \in L(\mathcal{A})$.

($\subseteq$): Let $x \in L(\mathcal{A})$. Then, $([x]_1, [x]_2, \dots, [x]_n) \in F$. This implies

$$([x]_1, [x]_2, \dots, [x]_n) = ([y]_1, [y]_2, \dots, [y]_n)$$

for some $y \in L$.

By corollary A1(2), $x \setminus L = y \setminus L$. So since $y \in L$, $x \in L$. $\square$

This proposition completes the proof of theorem A3, which, along with theorem A2, gives us theorem A1.

References

- Apt K R, Francez N, de Roever W P 1980 A proof system for communicating sequential processes. *ACM Trans. Program. Lang. Syst.* 2: 359–385
- Alur R, Henzinger T 1995 Local liveness for compositional modelling of fair reactive systems. *Lecture notes in Computer Science* 939 (Berlin: Springer Verlag) pp 166–179
- Abadi M, Lamport L 1993 Composing specifications. *ACM Trans. Program. Lang. Syst.* 15: 73–132
- Abadi M, Lamport L 1995 Conjoining specifications. *ACM Trans. Program. Lang. Syst.* 17: 507–534
- Abadi M, Alpern B, Apt K R, Francez N, Katz S, Lamport L, Schneider F B 1991 Preserving liveness: comments on “Safety and liveness from a methodological point of view”. *Inf. Process. Lett.* 40: 141–142
- Barringer H, Kuiper R, Pnueli A 1984 Now you may compose temporal logic specifications. *Proc. ACM Symp. on Theory of Computing*, pp 51–63
- Cori R, Metivier Y, Zielonka W 1993 Asynchronous mappings and asynchronous cellular automata. *Inf. Comput.* 106: 159–202
- Dijkstra E W 1965 Programming considered as a human activity. *Proc. Int. Fed. Inf. Process.* 65: 213–217
- Fagin R, Halpern J, Moses Y, Vardi M 1995 *Reasoning about knowledge* (Cambridge, MA: MIT Press)
- Floyd W F 1967 Assigning meanings to programs. *Proc. Symp. in Appl. Math. XIX* (Am. Math. Soc.) pp 19–32
- Jones C B 1983 Specification and design of (parallel) programs. *Proc. Int. Fed. Inf. Process.* pp 321–331
- Kupferman O, Vardi M 1997 An automata-theoretic approach to modular model checking. *Proc. Compositionality 97*, Malente-Gremsmühlen, Germany
- Mazurkiewicz A 1989 Basic notions of trace theory. *Lecture notes in Computer Science* 354 (Berlin: Springer Verlag) pp 285–363

- Misra J, Chandy M 1981 Proofs of networks of processes. *IEEE Trans. Software Eng.* 7: 417–426
- Manna Z, Pnueli A 1992 *Temporal logic of reactive systems: Specification* (New York: Springer-Verlag)
- Mohalik S K 1999 *Local presentations for finite state distributed systems*. Ph D thesis, University of Madras, Chennai
- Mohalik S K, Ramanujam R 1997 Assumption-commitment in automata. *Proc. of Foundations of Software Technology and Theoretical Computer Science '97, LNCS 1346* 153–168
- Mohalik S K, Ramanujam R 1998 A presentation of regular languages in the assumption-commitment framework. *Proc. of Concurrency and System Design '98*, Aizu-Wakamatsu, Japan, pp 250–260
- Owicki S, Gries D 1976 Verifying properties of parallel programs: an axiomatic approach. *Commun. ACM* 19: 279–285
- Owicki S, Lamport L 1982 Proving liveness properties of concurrent programs. *ACM Trans. Program. Lang. Syst.* 4: 455–495
- Pandya P K, Joseph M 1991 P-A logic: a compositional proof system for distributed programs. *Distr. Comput.* 5: 37–54
- Pnueli A, Rosner R 1990 Distributed reactive systems are hard to synthesize. *IEEE Symp. on Foundations of Computer Science* (IEEE Press) 746–757
- Quiwen Xu, Mohalik S K 1997 Compositional reasoning using the assumption-commitment paradigm. *Proc. Compositionality 97*, Malente-Greismühlen, Germany
- Ramanujam R 1995 A local presentation of synchronizing systems. In *Structures in Concurrency Theory* (ed.) Jörg Desel (Springer workshops in computing) pp 264–278
- Ramanujam R 1996a Local knowledge assertions in a changing world. In *Proc. Theoretical Aspects of Rationality and Knowledge* (San Francisco, CA: Morgan Kaufmann) pp 1–17
- Ramanujam R 1996b Locally linear time temporal logic. In *Proc. IEEE Logic in Computer Science*, pp 118–127
- Vardi M Y 1997 Verification of open systems. Invited talk in *Proc. of Foundations of Software Technology and Theoretical Computer Science '97, LNCS 1346* pp 250–266
- Zielonka W 1987 Notes on finite asynchronous automata. *RAIRO-Inf. Theor. Appl.* 21: 99–135