

Probabilistic Contracts for Component-based Design^{*}

Dana N. Xu, Gregor Gössler, and Alain Girault

INRIA, France

Abstract. We define a probabilistic contract framework for the construction of component-based embedded systems, based on the theory of Interactive Markov Chains. A contract specifies the assumptions a component makes on its context and the guarantees it provides. Probabilistic transitions allow for uncertainty in the component behavior, e.g. to model observed black-box behavior (internal choice) or reliability. An interaction model specifies how components interact.

We provide the ingredients for a component-based design flow, including (1) contract satisfaction and refinement, (2) parallel composition of contracts over disjoint, interacting components, and (3) conjunction of contracts describing different requirements over the same component. Compositional design is enabled by congruence of refinement.

1 Introduction

Typical embedded and distributed systems often encompass unreliable software or hardware components, as it may be technically or economically impossible to make a system entirely reliable. As a result, system designers have to deal with probabilistic specifications such as “the probability that this component fails at this point of its behavior is less than or equal to 10^{-4} ”. More generally, uncertainty in the observed behavior is introduced by abstraction of black-box — or simply too complex — behavior of components, the environment, or the execution platform. In this paper we introduce a framework for the design of correct systems from probabilistic, interacting components.

Figure 1(a) shows a Link system that transmits data between a Client and a Server. The Link receives a request from the Client and encodes the request before sending it to the Server. The encoding process fails with probability 0.02. After receiving a response from the Server, it decodes the data before delivering it to the Client. To model components, we adopt the discrete time Interactive Markov Chain (IMC) semantic model [8], which combines Labeled Transition System (LTS) and Markov Chain. Figure 1(b) shows an IMC describing the network Link of Figure 1(a). From its initial state l_0 , the Link goes to state l_1 as soon as it receives (*rec*) a request from a Client; the probability that it delivers (*del'*) this request to the Server is 0.98 and the probability that it fails before delivering to the Server is 0.02. The Link goes to state l_4 immediately after

^{*} supported by the European project COMBEST no. 215543.

receiving a response (rec') from the Server; the probability that it delivers (del) the response to the Client is 0.95 and the probability of failing to do so is 0.05. In state l_8 , the Link may still communicate with the Server regarding other services, but will not deliver any response to the Client.

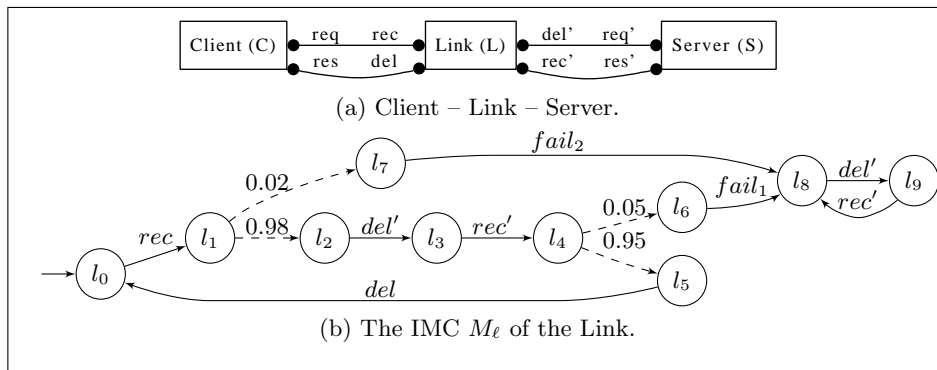


Fig. 1: An example of IMC: a Client-Link-Server.

Components communicate through interactions, that is, synchronized action transitions. Interactions are essential in component frameworks, as they allow the modeling of how components cooperate and communicate. We use the BIP framework [7] to model interactions between components.

Since the deploying context of a component is not known at design time, we use probabilistic *contracts* to specify and reason about correct behaviors of a component. Contracts have first been introduced in [11]. They allow us to specify what a component can expect from its context, what it must guarantee, and explicitly limit the responsibilities of both.

The framework we propose here allows us to model components, their interactions, and uncertainty in their observed behavior (§2). It supports different steps in a design flow: refinement, satisfaction, and abstraction (§3), parallel composition (§4.1), and conjunction (shared refinement) (§4.2). We prove that these operations satisfy the desired properties of *independent implementability* and *congruence* for parallel composition, and *soundness* for conjunction. Thus,

- refinement is compositional, that is, contracts over different components can be refined and implemented independently;
- the parallel composition of two contracts is satisfied by the parallel composition of any two implementations of the contracts; and
- several contracts C_i over the same component may be used to independently specify different requirements, possibly over different subsets of the component interactions. The conjunction is a common refinement of all C_i .

As pointed out in [2], conjunction of probabilistic specifications is non trivial, as a straight-forward approach would introduce spurious behaviors.

2 Components and Contracts

We give a formal definition to the discrete-time Interactive Markov Chains described in [8], used to model the behavior of components.

Definition 1 (Probability distribution) *A probability distribution over a set X is a function $f : X \rightarrow [0, 1]$ such that $\sum_{x \in X} f(x) = 1$.*

Definition 2 (Interactive Markov Chain (IMC)) *An IMC is a tuple $(\mathcal{Q}, \mathcal{A}, \rightarrow, \pi, s_0)$ where:*

- \mathcal{Q} is a nonempty finite set of states, partitioned into \mathcal{Q}^p , the set of probabilistic states, and \mathcal{Q}^a , the set of action states;
- \mathcal{A} is a finite alphabet of actions;
- $\rightarrow \subseteq \mathcal{Q}^a \times \mathcal{A} \times \mathcal{Q}$ is an action transition relation;
- $\pi : \mathcal{Q}^p \rightarrow (\mathcal{Q} \rightarrow [0, 1])$ is a transition probability function such that, for each $s \in \mathcal{Q}^p$, $\pi(s)$ is a probability distribution over \mathcal{Q} ;
- s_0 is the initial state.

IMC may interact with each other by synchronizing on action transitions (details in §4). Each action state in \mathcal{Q}^a has outgoing action transitions like those in an LTS. Each probabilistic state in \mathcal{Q}^p has outgoing probabilistic transitions like those in a Markov Chain. Probability distributions on states are memory-less, i.e., the future of an IMC depends only on the current state, but not on past choices. For example, in Figure 1(b), the probabilistic choice that the Link delivers the response to the Client (i.e., $\pi(l_4)(l_5) = 0.95$) is independent of the probabilistic choice of delivering a request to the Server (i.e., $\pi(l_1)(l_2) = 0.98$).

Notation: For convenience, we sometimes write the transition probability function π as a transition relation $\dashrightarrow \subseteq \mathcal{Q}^p \times [0, 1] \times \mathcal{Q}$ such that

$$\dashrightarrow = \{(s, p, s') \mid s \in \mathcal{Q}^p \wedge s' \in \mathcal{Q} \wedge p = \pi(s)(s')\}$$

We introduce contracts as a finite specification for a possibly infinite number of IMCs. In contrast to IMCs, the probabilistic transitions of a contract are labeled with probability *intervals*, similar to [9,15]. Moreover, a distinct \top state is used to distinguish assumptions on the use of the component from the guarantees it provides.

Definition 3 (Contract) *A contract is a tuple $(\mathcal{Q}, \mathcal{A}, \rightarrow, \sigma, t_0)$ where:*

- \mathcal{Q} is a nonempty finite set of states, partitioned into $\mathcal{Q} = \mathcal{Q}^p \cup \mathcal{Q}^a \cup \{\top\}$, where \mathcal{Q}^p is the set of probabilistic states, \mathcal{Q}^a is the set of action states, and \top is a distinct state without any outgoing transitions;
- \mathcal{A} is a finite alphabet of actions;
- $\rightarrow \subseteq \mathcal{Q}^a \times \mathcal{A} \times \mathcal{Q}$ is the action transition relation;
- $\sigma : \mathcal{Q}^p \rightarrow (\mathcal{Q} \rightarrow 2^{[0,1]})$ is a transition probability predicate, associating with each pair of states $(s, s') \in \mathcal{Q}^p \times \mathcal{Q}$ an interval of probabilities;
- t_0 is the initial state.

Notations: We also write σ as a transition relation $\dashrightarrow \subseteq \mathcal{Q}^p \times 2^{[0,1]} \times \mathcal{Q}$ such that $\dashrightarrow = \{(s, P, s') \mid s \in \mathcal{Q}^p \wedge s' \in \mathcal{Q} \wedge P = \sigma(s)(s')\}$. We write $q \xrightarrow{>0} q'$ if $\exists p > 0 : p \in \sigma(q, q')$, and $\xrightarrow{>0+}$ for the transitive closure of $\xrightarrow{>0}$. We extend arithmetic operations to intervals: $[\ell_1, u_1] + [\ell_2, u_2] = [\ell_1 + \ell_2, u_1 + u_2]$.

The meaning of a contract over a component C is the following:

- a transition $s \xrightarrow{a} \top$ specifies the *assumption* of the component C that an interaction involving action a does not occur in state s ;
- in an action state s , an action a labeling a transition not leading to \top specifies the *guarantee* of the component C that a is enabled in s ; conversely, the absence of any outgoing transition labeled with a specifies the guarantee that an interaction involving a will not occur;
- the \top state represents the fact that the assumption has been violated, and henceforth, the component C can show arbitrary, uncontrollable behavior;
- a transition $s \xrightarrow{[a,b]} t$ specifies an interval of allowed transition probabilities.




Fig. 2: Contract examples.

Example 1. The contract C_s in Figure 2(b) specifies that, after the Server receives a request req' , the probability that it reaches state t_3 is within $[0, 0.1]$; in state t_3 , it *assumes* that the environment does not give req' again; if this occurs, its implementation is not bound by C_s any more; the probability that it reaches t_2 from t_1 is within $[0.9, 1]$; in state t_2 , it *guarantees* to send a response (res'). In §3, we show how to check that M_s (in Figure 2(a)) satisfies C_s .

From the definitions of IMC and contract, we can see that an IMC can be trivially converted to a contract. For this, we define a lifting operator $[\cdot]$ (Figure 3 (c)). For the sake of simplicity, we use the same notation \dashrightarrow to represent both kinds of probabilistic transitions (i.e., those in an IMC and in a contract).

Definition 4 (Delimited contract) A contract $C = (\mathcal{Q}, \mathcal{A}, \rightarrow, \sigma, t_0)$ is delimited [5] iff $\forall s \in \mathcal{Q}^p \forall s' \in \mathcal{Q} \forall p \in \sigma(s)(s') : 1 - p \in \sum_{s'' \in \mathcal{Q} \setminus \{s'\}} \sigma(s)(s'')$.

Example 2. Figure 3 (a) shows a delimited contract: for all $p \in [0, 2, 0.3]$, we can find $p' \in [0.7, 0.8]$ such that $p + p' = 1$ and vice versa. Figure 3 (b) shows a contract that is not delimited. However, we can *cut* [5] the redundant sub-interval $[0.8, 0.9]$ from the interval $[0.7, 0.9]$ to obtain a delimited contract.




Fig. 3: Delimited contract and rules for lifting IMC to contract.

3 Contract Refinement

System synthesis involves refining a contract several times until an implementation is obtained. We therefore define formally the notion of contract refinement.

3.1 Refinement and Satisfaction

We first define contract refinement, and give thereafter some explanations.

Definition 5 (Contract refinement) Let $C_1 = (\mathcal{Q}_1, \mathcal{A}, \rightarrow_1, \sigma_1, s_0)$ and $C_2 = (\mathcal{Q}_2, \mathcal{A}, \rightarrow_2, \sigma_2, t_0)$ be two contracts. C_1 refines C_2 (written $C_1 \leq C_2$) iff $s_0 \leq t_0$, where $\leq \subseteq \mathcal{Q}_1 \times \mathcal{Q}_2$ is the greatest relation s.t. for all $s \leq t$ we have:

1. $s = \top \implies t = \top$;
2. If $(s, t) \in \mathcal{Q}_1^a \times (\mathcal{Q}_2^a \cup \{\top\})$ then
 - (a) $\forall t' \neq \top \in \mathcal{Q}_2, (t \xrightarrow{\alpha}_2 t') \implies (\exists s' \in \mathcal{Q}_1, s \xrightarrow{\alpha}_1 s' \wedge s' \leq t')$;
 - (b) $\forall s' \in \mathcal{Q}_1, (s \xrightarrow{\alpha}_1 s') \implies (t = \top \vee \exists t' \in \mathcal{Q}_2, t \xrightarrow{\alpha}_2 t' \wedge s' \leq t')$.
3. If $(s, t) \in \mathcal{Q}_1^p \times \mathcal{Q}_2^p$ then there exists a function $\delta : \mathcal{Q}_1 \times \mathcal{Q}_2 \rightarrow [0, 1]$, which, for each $s' \in \mathcal{Q}_1$, gives a probability distribution $\delta(s')$ over \mathcal{Q}_2 , such that for every probability distribution f over \mathcal{Q}_1 with $f(s') \in \sigma_1(s)(s')$ and $\forall t' \in \mathcal{Q}_2$,

$$\sum_{s' \in \mathcal{Q}_1} f(s') * \delta(s')(t') \in \sigma_2(t)(t') \quad \text{and} \quad \forall s' \in \mathcal{Q}_1 : (\delta(s')(t') > 0 \implies s' \leq t')$$
4. If $(s, t) \in \mathcal{Q}_1^a \times \mathcal{Q}_2^p$ then $\exists t^a \in \mathcal{Q}_2^a : t \xrightarrow{>0}_2^+ t^a \wedge s \leq t^a$ and $\forall t' \in \mathcal{Q}_2$,

$$(t \xrightarrow{>0}_2 t' \implies s \leq t').$$
5. If $(s, t) \in \mathcal{Q}_1^p \times \mathcal{Q}_2^a$ then $\exists s^a \in \mathcal{Q}_1^a : s \xrightarrow{>0}_1^+ s^a \wedge s^a \leq t$ and $\forall s' \in \mathcal{Q}_1$,

$$(s \xrightarrow{>0}_1 s' \implies s' \leq t).$$

In Definition 5, condition (1) ensures that C_1 makes no stronger assumptions on the context than C_2 . Condition (2a) says that any transition accepted by C_2 must also be accepted by C_1 . Similarly, condition (2b) says that each action

transition of C_1 must also be enabled in C_2 , unless C_2 is in the \top state. Condition (3), adapted from [9], deals with refinement among probabilistic states. Intuitively, $s \leq t$ if there exists a function δ which distributes the probabilities of transitions from s to s' onto the transitions from t to t' , such that the sum of the fractions is in the range $\sigma(t)(t')$, as illustrated in Example 3 below. Condition (4) says that an action state s refines a probabilistic state t if it refines all action states reachable with a path of positive probability from t . Finally, condition (5) is symmetrical to condition (4).

Before giving an example of refinement, we define the satisfaction of a contract by an implementation (an IMC) as the refinement of the contract by the lifted IMC (i.e., written in the form of a contract).

Definition 6 (Contract satisfaction) *An IMC M satisfies a contract C (written $M \models C$) iff $\lfloor M \rfloor \leq C$.*

Example 3. We illustrate how to check $\lfloor M_s \rfloor \leq C_s$, in particular, $s_1 \leq t_1$. It is easy to check $s_3 \leq t_2$, $s_4 \leq t_2$, and $s_2 \leq t_3$. Dashed lines stand for non-negative distributions δ . Condition (3) in Definition 5 states that $s_1 \leq t_1$ if for each successor state of s_1 there is a function δ (i.e., d_1, d_2, d_3) such that, for each tuple (p_2, p_3, p_4) satisfying constraints (1) to (4) in Figure 4, the constraints (5) and (6) are implied. Condition (3) can be checked efficiently by requiring the set inclusion to hold for the bounds of interval $\sigma(s)(s')$, using a linear programming solver. As $\delta(s')$ is a probability distribution, we obtain for our example $d_1 = d_2 = d_3 = 1$. (Note that if we had $s_2 \leq t_2$ as well, say, we had d_4 from s_2 to t_2 , we would have another constraint $d_3 + d_4 = 1$.)




Fig. 4: Left: Contract refinement $s_1 \leq t_1$. Right: Constraints to be checked.

Definition 7 (Models of contracts) *The set of models of a contract C (written $\mathcal{M}(C)$) is the set of IMCs that satisfy C : $\mathcal{M}(C) = \{M \mid M \models C\}$.*

Definition 8 (Semantical equivalence) *Contracts C_1 and C_2 are semantically equivalent (written $C_1 \equiv C_2$) iff $\mathcal{M}(C_1) = \mathcal{M}(C_2)$.*

Lemma 1 (Monotonicity of satisfaction) *For all IMC M and contracts C_1 and C_2 , if $M \models C_1$ and $C_1 \leq C_2$, then $M \models C_2$.*

Lemma 2 (Refinement and model inclusion) *For all contracts C_1 and C_2 , $C_1 \leq C_2 \implies \mathcal{M}(C_1) \subseteq \mathcal{M}(C_2)$.*

The proofs for all lemmas and theorems in this paper can be found in [14].

3.2 Bisimulation

We adapt the usual notion of bisimulation to contracts, and define reduction of a contract with respect to bisimulation.

Definition 9 (Bisimulation \simeq) *Given a contract $C = (\mathcal{Q}, \mathcal{A}, \rightarrow, \dashrightarrow, s_0)$, let $\simeq \subseteq \mathcal{Q} \times \mathcal{Q}$ be the greatest relation such that if $s \simeq t$ then:*

1. $s = \top \iff t = \top$;
2. If $(s, t) \in \mathcal{Q}^a \times \mathcal{Q}^a$ then
 - (a) $\forall \alpha \in \mathcal{A} \quad \forall s' \in \mathcal{Q}, (s \xrightarrow{\alpha} s' \implies \exists t' \in \mathcal{Q}, (t \xrightarrow{\alpha} t' \wedge s' \simeq t'))$
 - (b) $\forall \alpha \in \mathcal{A} \quad \forall t' \in \mathcal{Q}, (t \xrightarrow{\alpha} t' \implies \exists s' \in \mathcal{Q}, (s \xrightarrow{\alpha} s' \wedge s' \simeq t'))$
3. If $(s, t) \in \mathcal{Q}^p \times \mathcal{Q}^p$ then
 - (a) *there is a function $\delta : \mathcal{Q} \times \mathcal{Q} \rightarrow [0, 1]$, which for each $s' \in \mathcal{Q}$ gives a probability distribution $\delta(s')$ on \mathcal{Q} , s.t. for every probability distribution f over \mathcal{Q} with $f(s') \in \sigma(s)(s')$ and $\forall t' \in \mathcal{Q}$*

$$\sum_{s' \in \mathcal{Q}} f(s') * \delta(s', t') \in \sigma(t)(t') \quad \text{and} \quad \forall s' \in \mathcal{Q} : (\delta(s', t') > 0 \implies s' \simeq t')$$

(b) *symmetric to (3a);*

4. If $(s, t) \in \mathcal{Q}^a \times \mathcal{Q}^p$ then $\exists t^a \in \mathcal{Q}^a : t \dashrightarrow^{>0+} t^a \wedge s \simeq t^a$ and $\forall t' \in \mathcal{Q}, t \dashrightarrow^{>0} t' \implies s \simeq t'$;
5. If $(s, t) \in \mathcal{Q}^p \times \mathcal{Q}^a$ then $\exists s^a \in \mathcal{Q}^a : s \dashrightarrow^{>0+} s^a \wedge s^a \simeq t$ and $\forall s' \in \mathcal{Q}, s \dashrightarrow^{>0} s' \implies s' \simeq t$.

In Definition 9, condition (2) is the standard definition for bisimulation. Conditions (3a) and (3b) deal with the probabilistic transitions. Finally, conditions (4) and (5) say that an action state is bisimilar with a probabilistic state if it is bisimilar with all its successors with non-zero probability, and some action state is reachable from this probabilistic state.

Definition 10 (Reduction modulo \simeq) *Let $C = (\mathcal{Q}, \mathcal{A}, \rightarrow, \sigma, s_0)$ be a contract. For all $s \in \mathcal{Q}$, let $\mathcal{C}_s = \{q \in \mathcal{Q} \mid s \simeq q\}$ be the equivalence class of s . Let $\mathcal{C} = \{\mathcal{C}_s \mid s \in \mathcal{Q}\}$. The reduced contract, written \bar{C} , is $(\mathcal{C}, \mathcal{A}, \rightarrow_{\simeq}, \sigma_{\simeq}, \mathcal{C}_{s_0})$ such that, $\forall s = \{s_1, \dots, s_m\}, t = \{t_1, \dots, t_n\} \in \mathcal{C}$, we have: (1) $s \xrightarrow{\alpha}_{\simeq} t$ iff $\exists i, j : s_i \xrightarrow{\alpha} t_j$, and (2) $\sigma_{\simeq}(s, t) = \sum_{1 \leq j \leq n} \sigma(s_1, t_j)$.*

Notice that an equivalence class may contain both action and probabilistic states. By Definition 9, except for probabilistic transitions with probability interval $[0, 0]$, either all transitions leaving an equivalence class are action transition and Definition 9 (2) applies, or they are all probabilistic transitions and Definition 9 (3) applies as follows. For each probabilistic state $s_i \in s$, the probabilities of transitions to states $t_j \in t$ are summed up (it does not matter which of the transitions is taken since all successors t_j are equivalent). This sum is the transition probability from s_i to some state in t . By definition of \simeq , the sum is

the same for all $s_i \in s$, thus we pick $\sigma(s_1, t_j)$. For example, we can reduce the contract C_2 of Figure 7 (right) by combining the bisimilar states t_2 and t_3 into one: $t_1 \xrightarrow{[0.2, 0.6]} \{t_2, t_3\}$.

Lemma 3 (Model equivalence) *For all delimited contracts C , $\overline{C} \equiv C$.*

3.3 Contract Abstraction

The need of abstraction arises naturally in contract frameworks. We abstract actions in $\mathcal{A} \setminus \mathcal{B}$ that we do not care about by renaming them into internal τ actions. The contract over the alphabet $\mathcal{B} \cup \{\tau\}$ is then projected on the sub-alphabet \mathcal{B} by using the standard determinization algorithm (see e.g. [1]).

Definition 11 (Projection) *Let $C = (\mathcal{Q}, \mathcal{A}, \rightarrow_1, \sigma, s_0)$ be a contract and $\mathcal{B} \subseteq \mathcal{A}$. Let $C' = (\mathcal{Q}, \mathcal{B} \cup \{\tau\}, \rightarrow_1, \sigma, s_0)$ be the contract where all transition labels in $\mathcal{A} \setminus \mathcal{B}$ are replaced with τ . The projection of C on \mathcal{B} (written $\pi_{\mathcal{B}}(C)$) is obtained by τ -elimination (determinization) of C' .*

Example 4. In Figure 2, if we do not care how the implementation handles failure cases, we can check that $\pi_{\mathcal{A}_s \setminus \{\text{handle}\}}(M_s) \models C_s$.

4 Contract Composition

We introduce two composition operations for contracts: parallel composition \parallel , parametrized with an interaction set \mathcal{I} , and conjunction \wedge (shared refinement).

4.1 Parallel Composition of Contracts

Parallel composition allows us to build complex models from simpler components in a stepwise and hierarchical manner. In order to reason about the composition of components at the contract level, we introduce parallel composition of contracts. As in the BIP component framework [7], parallel composition is parametrized with a set of interactions, where each interaction is a set of component actions occurring simultaneously. For instance, an interaction set $\{a, a|b, c\}$ says that action a can interleave or synchronize with b ; action b must synchronize with a ; action c is a singleton interaction that always interleaves. The symbol “ $|$ ” is commutative, which means that $a|b$ is identical to $b|a$. In Figure 5, the interactions α and β are of the form c , $a|b$, or $a|b|d$, and so on.

Definition 12 (Parallel composition of contracts) *Let $C_1 = (\mathcal{Q}_1, \mathcal{A}_1, \rightarrow_1, \dashrightarrow_1, s_0)$ and $C_2 = (\mathcal{Q}_2, \mathcal{A}_2, \rightarrow_2, \dashrightarrow_2, t_0)$ be two contracts. The parallel composition of C_1 and C_2 on interaction set \mathcal{I} (written $C_1 \parallel_{\mathcal{I}} C_2$) is the contract $(\mathcal{Q}, \mathcal{I}, \rightarrow, \dashrightarrow, (s_0, t_0))$ where:*

1. $\mathcal{Q} = \mathcal{Q}_1 \times \mathcal{Q}_2$ with $\top = (\mathcal{Q}_1 \times \{\top_2\}) \cup (\{\top_1\} \times \mathcal{Q}_2)$ — that is, \top of $C_1 \parallel_{\mathcal{I}} C_2$ is an aggregate state reached as soon as C_1 or C_2 reaches its \top_i state —, $\mathcal{Q}^a = \mathcal{Q}_1^a \times \mathcal{Q}_2^a$, and $\mathcal{Q}^p = \mathcal{Q} \setminus (\mathcal{Q}^a \cup \top)$;

$$\begin{array}{ccc}
\frac{q_1 \xrightarrow{\alpha} q'_1 \quad \alpha \in \mathcal{I} \quad q_2 \in \mathcal{Q}_2^a}{(q_1, q_2) \xrightarrow{\alpha} (q'_1, q_2)} \quad [\text{R1}] & \frac{q_2 \xrightarrow{\alpha} q'_2 \quad \alpha \in \mathcal{I} \quad q_1 \in \mathcal{Q}_1^a}{(q_1, q_2) \xrightarrow{\alpha} (q_1, q'_2)} \quad [\text{R2}] & \\
\frac{q_1 \xrightarrow{\alpha} q'_1 \quad q_2 \xrightarrow{\beta} q'_2 \quad \alpha|\beta \in \mathcal{I}}{(q_1, q_2) \xrightarrow{\alpha|\beta} (q'_1, q'_2)} \quad [\text{R3}] & \frac{q_1 \xrightarrow{[p^1, p^2]} q'_1 \quad q_2 \xrightarrow{[p^3, p^4]} q'_2}{(q_1, q_2) \xrightarrow{[p^1 * p^3, p^2 * p^4]} (q'_1, q'_2)} \quad [\text{R4}] & \\
\frac{q_1 \xrightarrow{P} q'_1 \quad q_2 \in \mathcal{Q}_2^a}{(q_1, q_2) \xrightarrow{P} (q'_1, q_2)} \quad [\text{R5}] & \frac{q_2 \xrightarrow{P} q'_2 \quad q_1 \in \mathcal{Q}_1^a}{(q_1, q_2) \xrightarrow{P} (q_1, q'_2)} \quad [\text{R6}] &
\end{array}$$

Fig. 5: Rules for the parallel composition of contracts.

2. \rightarrow is the least relation satisfying the rules [R1]–[R3] in Figure 5; and
3. \dashrightarrow is the least relation satisfying the rules [R4]–[R6] in Figure 5.

Rules [R1] to [R3] are the usual parallel composition rules for interactive processes, while the rule [R4] is similar to the typical parallel composition for Markov chains but on probability intervals. Finally, rules [R5] and [R6] state that probabilistic transition, usually modeling hidden internal behavior, have priority over action transitions.

Example 5. Figure 6 illustrates the parallel composition of contracts C_s (from Figure 2(b)) and $C_\ell = [M_\ell]$ (where M_ℓ is given in Figure 1(b)), with $\mathcal{I} = \{rec, del, req' | del', res' | rec', fail_1, fail_2\}$. The composed contract $C_s \parallel_{\mathcal{I}} C_\ell$ states that a failure in the Link component does not prevent it from continuing to deliver the request req' to the Server and receiving response from the Server, but the failure prevents it from delivering the response res' back to the Client.




Fig. 6: Parallel composition of C_s and C_ℓ .

We end the section on parallel composition with two essential properties.

Theorem 1 (Independent implementability) *For all IMCs M, N , contracts C_1, C_2 , and interaction set \mathcal{I} , if $M \models C_1$ and $N \models C_2$, then $M \parallel_{\mathcal{I}} N \models C_1 \parallel_{\mathcal{I}} C_2$.*

Theorem 2 (Congruence of refinement) *For all contracts C_1, C_2, C_3 , and interaction set \mathcal{I} , if $C_1 \leq C_2$, then $C_1 \parallel_{\mathcal{I}} C_3 \leq C_2 \parallel_{\mathcal{I}} C_3$.*

4.2 Conjunction of contracts

A single component may have to satisfy several contracts that are specified independently, each of them specifying different requirements on the component, such as safety, reliability, and quality of service aspects. Therefore, the contracts may use different, possibly overlapping, sub-alphabets of the component. The *conjunction* of contracts computes a common refinement of all contracts. Prior to conjunction, we define *similarity* of contracts as a test whether a common refinement exists.

Definition 13 (Similarity (\sim)) Let $C_1 = (\mathcal{Q}_1, \mathcal{A}_1, \rightarrow_1, \dashrightarrow_1, s_0)$ and $C_2 = (\mathcal{Q}_2, \mathcal{A}_2, \rightarrow_2, \dashrightarrow_2, t_0)$ be two contracts. $\sim \subseteq \mathcal{Q}_1 \times \mathcal{Q}_2$ is the largest relation such that $\forall (s, t) \in \mathcal{Q}_1 \times \mathcal{Q}_2$, $s \sim t$ iff $(s = \top \vee t = \top)$ or conditions (1) to (4) below hold:

1. If $(s, t) \in \mathcal{Q}_1^a \times \mathcal{Q}_2^a$ then
 - (a) for all $s' \in \mathcal{Q}_1$, if $s \xrightarrow{a} s'$, then either $t \xrightarrow{a} t'$ for some $t' \in \mathcal{Q}_2$ and $s' \sim t'$, or $a \notin \mathcal{A}_2$ and $s' \sim t$; and
 - (b) for all $t' \in \mathcal{Q}_2$, if $t \xrightarrow{a} t'$, then either $s \xrightarrow{a} s'$ for some $s' \in \mathcal{Q}_1$ and $s' \sim t'$, or $a \notin \mathcal{A}_1$ and $s \sim t'$;
2. If $(s, t) \in \mathcal{Q}_1^p \times \mathcal{Q}_2^p$ then
 - (a) for all $s' \in \mathcal{Q}_1$, if $s \xrightarrow{P_1} s'$, then $t \xrightarrow{P_2} t'$ for some $t' \in \mathcal{Q}_2$ with $P_1 \cap P_2 \neq \emptyset$ and $(s' \sim t' \vee 0 \in P_1 \cap P_2)$;
 - (b) for all $t' \in \mathcal{Q}_2$, if $t \xrightarrow{P_2} t'$, then $s \xrightarrow{P_1} s'$ for some $s' \in \mathcal{Q}_1$ with $P_1 \cap P_2 \neq \emptyset$ and $(s' \sim t' \vee 0 \in P_1 \cap P_2)$;
3. If $(s, t) \in \mathcal{Q}_1^p \times \mathcal{Q}_2^a$ then for all $s' \in \mathcal{Q}_1$ with $s \xrightarrow{P} s'$, $(s' \sim t \vee 0 \in P)$;
4. If $(s, t) \in \mathcal{Q}_1^a \times \mathcal{Q}_2^p$ then for all $t' \in \mathcal{Q}_2$ with $t \xrightarrow{P} t'$, $(s \sim t' \vee 0 \in P)$.

Finally, C_1 and C_2 are similar, written $C_1 \sim C_2$, iff $s_0 \sim t_0$.

The P_i in Definition 13 refers to a probabilistic interval in the form of $[\ell_i, u_i]$. Any state is similar with the top state (where the contract does not constrain the implementation in any way). Two action states are similar if they agree on the enabled actions in the common alphabet, and the successor states are similar again. Two probabilistic states are similar if the probabilistic transitions can be matched such that the intervals overlap, and the successor states are either similar, or can be made unreachable by refining the probability interval to $[0, 0]$.

Definition 14 (Unambiguous contract) A contract $C = (\mathcal{Q}, \mathcal{A}, \rightarrow, \dashrightarrow, s_0)$ is unambiguous iff for all $r, s, t \in \mathcal{Q}$, if $r \xrightarrow{>0} s \wedge r \xrightarrow{>0} t \wedge s \sim t$, then $s = t$.

A contract is *unambiguous* if the reachable successor states of any probabilistic state are pairwise non-similar. In Figure 7 (left), the contract C_a is not unambiguous: $s_2 \sim s_3$ (highlighted in gray) but $s_2 \neq s_3$.

We are now ready to define the conjunction of contracts.




Fig. 7: Left: ambiguous contract C_a . Middle: C_1 . Right: C_2 where $C_1 \not\sim C_2$.

Definition 15 (Conjunction of contracts (\wedge)) For unambiguous contracts $C_1 = (\mathcal{Q}_1, \mathcal{A}_1, \rightarrow_1, \dashrightarrow_1, s_0)$ and $C_2 = (\mathcal{Q}_2, \mathcal{A}_2, \rightarrow_2, \dashrightarrow_2, t_0)$ such that C_1 and C_2 are similar, let $C_1 \wedge C_2$ be the contract $(\mathcal{Q}_1 \times \mathcal{Q}_2, \mathcal{A}_1 \cup \mathcal{A}_2, \rightarrow, \dashrightarrow, (s_0, t_0))$ where $\top = (\top_1, \top_2)$ and

1. \rightarrow is the least relation satisfying the rules [C1] – [LIFTR] in Figure 8, and
2. \dashrightarrow is the least relation satisfying the rules [C3] – [C4R] in Figure 8 (where for all other probabilistic transitions $(q_1, q_2) \xrightarrow{P} (q'_1, q'_2)$, $P = [0, 0]$).

$$\begin{array}{c}
 \frac{q_1 \xrightarrow{\alpha} q'_1 \quad q_2 \xrightarrow{\alpha} q'_2}{(q_1, q_2) \xrightarrow{\alpha} (q'_1, q'_2)} \quad [\text{C1}] \quad \frac{q_1 \xrightarrow{\alpha} q'_1}{(q_1, \top) \xrightarrow{\alpha} (q'_1, \top)} \quad [\text{C2L}] \quad \frac{q_2 \xrightarrow{\alpha} q'_2}{(\top, q_2) \xrightarrow{\alpha} (\top, q'_2)} \quad [\text{C2R}] \\
 \\
 \frac{q_1 \xrightarrow{\alpha} q'_1 \quad \alpha \notin \mathcal{A}_2 \quad q_2 \in \mathcal{Q}_2^a}{(q_1, q_2) \xrightarrow{\alpha} (q'_1, q_2)} \quad [\text{LIFTL}] \quad \frac{q_2 \xrightarrow{\alpha} q'_2 \quad \alpha \notin \mathcal{A}_1 \quad q_1 \in \mathcal{Q}_1^a}{(q_1, q_2) \xrightarrow{\alpha} (q_1, q'_2)} \quad [\text{LIFTR}] \\
 \\
 \frac{q_1 \xrightarrow{P_1} q'_1 \quad q_2 \xrightarrow{P_2} q'_2 \quad q'_1 \sim q'_2}{(q_1, q_2) \xrightarrow{P_1 \cap P_2} (q'_1, q'_2)} \quad [\text{C3}] \\
 \\
 \frac{q_1 \xrightarrow{P} q'_1 \quad q_2 \in \mathcal{Q}_2^a \cup \{\top\}}{q'_1 \sim q_2} \quad [\text{C4L}] \quad \frac{q_2 \xrightarrow{P} q'_2 \quad q_1 \in \mathcal{Q}_1^a \cup \{\top\}}{q_1 \sim q'_2} \quad [\text{C4R}] \\
 \\
 \frac{q_1 \dashrightarrow^P q'_1 \quad q_2 \dashrightarrow^P q'_2}{(q_1, q_2) \dashrightarrow^P (q'_1, q'_2)}
 \end{array}$$

Fig. 8: Rules for conjunction of contracts.

Rule [C1] requires the contracts to agree on action transitions over the common alphabet. According to rule [C2L] (resp. [C2R]), the conjunction behaves like the first (resp. second) contract as soon as the other contract is in \top . Rules [LIFTL] and [LIFTR] allow the interleaving of action transitions that are not in the common alphabet. Rules [C3] – [C4R] define the probabilistic transitions whose successor states are similar. For non-similar successor states, the probability interval is refined to $[0, 0]$, according to Definition 15.

Example 6. Figure 9 shows three contracts for the Link component: C_{ℓ_1} specifies that the implementation should receive a request from the Client and deliver it to the Server; C_{ℓ_2} specifies that the implementation should receive a response




Fig. 9: Left: C_{ℓ_1} . Middle: C_{ℓ_2} . Right: C_{ℓ_3} .

from the Server and deliver it to the Client; C_{ℓ_3} requires the response (rec') received from the Server to occur after the request (del') delivered to the Server. We can verify that $M_\ell \models (C_{\ell_1} \wedge C_{\ell_3}) \wedge (C_{\ell_2} \wedge C_{\ell_3})$ (where M_ℓ is in Figure 1(b)).




Fig. 10: Example where $M_b \models C_b \wedge C_b$ but $M_b \not\models C_b$.

Example 7. As a contract that is not in reduced form is not unambiguous, contracts should be reduced before performing conjunction. In Figure 7 (left), contract C_2 is non unambiguous, but $t_2 \simeq t_3$. If we reduce C_2 by applying Definition 10, we get $t_1 \xrightarrow{[0.2, 0.6]} \{t_2, t_3\} \xrightarrow{a} \{t_2, t_3\}$. The reduced contract is unambiguous and $s_1 \sim t_1$, such that conjunction yields a common refinement of C_1 and C_2 .

Theorem 3 (Soundness of conjunction) *For any IMC M and unambiguous contracts C_i with alphabet \mathcal{A}_i , $i = 1, 2$ such that C_1 and C_2 are similar, if $M \models C_1 \wedge C_2$ then $\pi_{\mathcal{A}_i}(M) \models C_i$, $i = 1, 2$.*

Example 8. Figure 10 motivates the requirement of conjunction (Definition 15) for unambiguous contracts. The resulting contract $C_b \wedge C_b$ is reduced such that the model relation can be seen easily. The node v_2 denotes the equivalent class $\{(s_1, s_2), (s_2, s_1), (s_2, s_2)\}$; the node v_3 denotes the equivalent class $\{(s_1, s_3), (s_2, s_3), (s_3, s_1), (s_3, s_2), (s_3, s_3)\}$. As $t_1 \sim t_2 \sim t_3$, duplicated intervals lead to an unsound result.

It is interesting to note the similarity of conjunction with discrete controller synthesis [13], in the sense that conjunction is a refinement of both contracts making bad states (i.e., pairs of states where both contracts are contradictory) unreachable. In this analogy, both action transitions and probabilistic transitions with strictly positive intervals amount to uncontrollable transitions, whereas transitions whose probability interval contains 0 amount to controllable transitions that can be refined to $[0, 0]$ so as to make bad states unreachable.

5 Case Study

We study a dependable computing system with time redundancy. The system specification is expressed by the contract C_S of Figure 11 (top left), which specifies that the computation $comp$ should have a success probability of at least 0.999. If the computation fails, then nothing is specified (state \top).

The processor is specified by the contract C_P of Figure 11 (top right). Following an execution request exe , either the processor succeeds and replies with ok (with a probability at least p), or fails and replies with nok (with a probability at most $1 - p$). The failure rates for successive executions are independent. The probability p is a parameter of the contract.

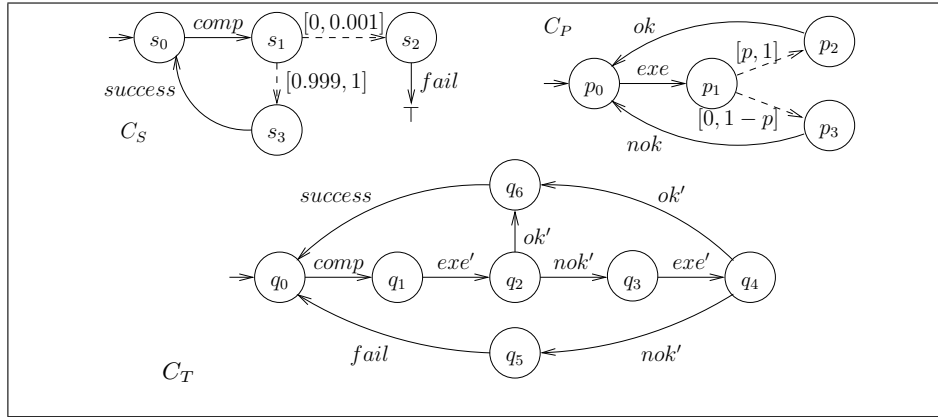


Fig. 11: Specification C_S ; processor contract C_P ; time redundancy contract C_T .

We place ourselves in a setting where the reliability level guaranteed by C_P alone (as expressed by p) cannot fulfill the requirement of C_S (that is, 0.999), and hence some form of redundancy must be used. We propose to use time redundancy, as expressed by the contract C_T of Figure 11 (bottom). Each computation $comp$ is first launched on the processor (exe), either followed by a positive (ok) or negative (nok) answer from the processor. In the latter case, the execution is launched a second time, therefore implementing time redundancy. The contract C_T finally answers with $success$ if *either* execution is followed by ok , or with $fail$ if *both* executions are followed by nok .

In terms of component-based design for reliability, we wonder what is the minimum value of p to guarantee the reliability level of C_S . To compute this minimum value, we first compute the parallel composition $C_T \parallel_{\mathcal{I}} C_P$, with the interaction set $\mathcal{I} = \{comp, exe|exe', ok|ok', nok|nok', success, fail\}$. The reduction modulo bisimulation of this parallel composition is shown in Figure 12 (top), where the interactions $exe|exe'$, $ok|ok'$, and $nok|nok'$ have been replaced for conciseness by **exe**, **ok**, and **nok**, respectively. We call this new contract $C_{T||P}$. We then compute the projection of $C_{T||P}$ onto the set $\mathcal{B} = \{comp, success, fail\}$. The result $C_\pi = \pi_{\mathcal{B}}(C_{T||P})$ is shown in Figure 12 (bottom left).

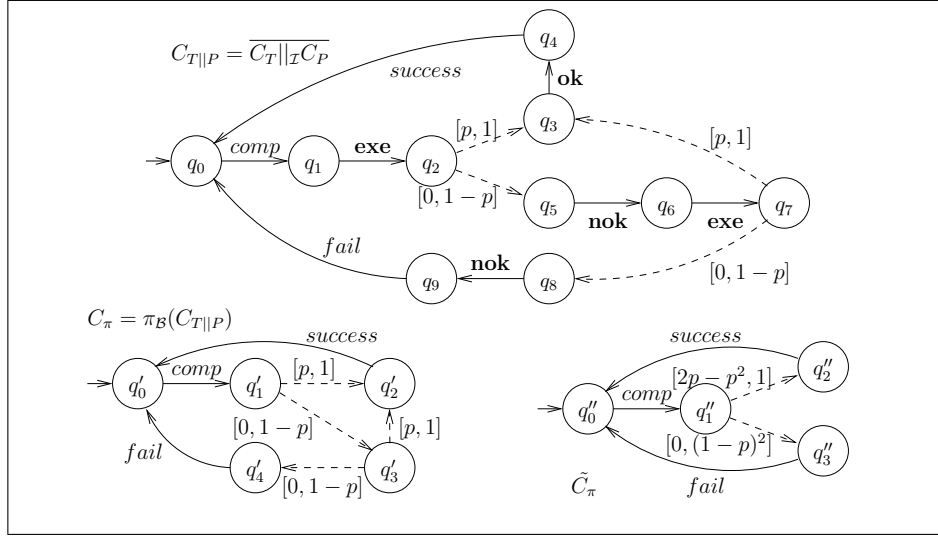


Fig. 12: Parallel composition $C_{T||P}$; projection C_π ; transitive closure \tilde{C}_π .

We are thus faced with a contract C_π having *sequences* of probabilistic transitions; more precisely, since some probabilistic states have several outgoing transitions, we have DAGs of probabilistic transitions. We therefore compute the transitive closure for each such DAG, that is, the equivalent probabilistic transitions from the initial state of the DAG (e.g., q'_1 in C_π) to its final states (e.g., q'_2 and q'_4 in C_π). Without entering into the details of this computation, we show the resulting contract \tilde{C}_π in Figure 12 (bottom right).

The last step involves checking under which condition on p the contract \tilde{C}_π refines the specification C_S . We have $\tilde{C}_\pi \leq C_S \Leftrightarrow (1-p)^2 \leq 0.001$. This means that, with time redundancy and a processor with a reliability level of at least 0.969, we are able to ensure an overall reliability level of 0.999.

6 Discussion

We have introduced a design framework based on probabilistic contracts, and proved essential properties for the use in component-based design. Our definition of contracts is based on the ideas from [9,15,5], although the frameworks in [9,5] do not support interactions.

Shared refinement of interfaces, and conjunction of modal specifications over possibly different alphabets have been defined in [4] and [12]. A framework over modal assume/guarantee-contracts is introduced in [6], for which both parallel composition and conjunction are defined. Probabilistic assume/guarantee-contracts have been introduced in [3] in terms of traces. [10] introduces a compositional framework based on continuous time IMCs, adopting a similar interaction model as done in this paper. This framework supports abstraction, parallel and symmetric composition, but not conjunction. The recently introduced

Constraint Markov Chains (CMC) [2] generalize Markov Chains by introducing constraints on state valuations and transition probability distributions, aiming at a similar goal of providing a probabilistic component-based design framework. Whereas CMCs do not support explicit interactions among components, they allow one to expressively specify constraints on probability distributions. Conjunction is shown to be sound and complete in this framework.

Future work will encompass implementing the framework and carrying out case studies. A particularly interesting application would be the design of adaptive systems where the probabilistic behavior of components may change over time, while the overall system must at any time satisfy a set of safety, reliability, and quality of service contracts.

References

1. A.V. Aho, R. Sethi, and J.D. Ullman. *Compilers – Principles, Techniques, and Tools*. Addison Wesley, 1986.
2. B. Caillaud, B. Delahaye, K.G. Larsen, A. Legay, M. Pedersen, and A. Wasowski. Compositional design methodology with constraint markov chains. Research Report 6993, INRIA, 2009.
3. B. Delahaye and B. Caillaud. A model for probabilistic reasoning on assume/guarantee contracts. Research Report 6719, INRIA, 2008.
4. L. Doyen and T. Petrov T.A. Henzinger, B. Jobstmann. Interface theories with component reuse. In *Proc. EMSOFT’08*, pages 79–88. ACM, 2008.
5. H. Fecher, M. Leucker, and V. Wolf. Don’t know in probabilistic systems. In *Model Checking Software*, LNCS, pages 71–88. Springer, 2006.
6. G. Gössler and J.-B. Ralet. Modal contracts for component-based design. In *Proc. SEFM’09*, pages 295–303. IEEE, 2009.
7. G. Gössler and J. Sifakis. Composition for component-based modeling. *Science of Computer Programming*, 55(1-3):161–183, 3 2005.
8. H. Hermanns. *Interactive Markov Chains: The Quest for Quantified Quality*, volume 2428 of *LNCS*. Springer, 2002.
9. B. Jonsson and K.G. Larsen. Specification and refinement of probabilistic processes. In *LICS*, pages 266–277. IEEE Computer Society, 1991.
10. J.-P. Katoen, D. Klink, and M.R. Neuhäuser. Compositional abstraction for stochastic systems. In *Proc. FORMATS’09*, pages 195–211, 2009.
11. B. Meyer. *Advances in Object-Oriented Software Engineering*, chapter Design by Contract, pages 1–50. Prentice Hall, 1991.
12. J.-B. Ralet, E. Badouel, A. Benveniste, B. Caillaud, and R. Passerone. Why modalities are good for interface theories? In *Proc. ACSD’09*. IEEE, 2009.
13. P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete event processes. *SIAM J. Control and Optimization*, 25(1), 1987.
14. D.N. Xu, G. Gössler, and A. Girault. Probabilistic contracts for component-based design. Research Report 7328, INRIA, 2010.
15. W. Yi. Algebraic reasoning for real-time probabilistic processes with uncertain information. In *FTRTFT*, volume 863 of *LNCS*, pages 680–693. Springer, 1994.