A Constraint Oriented Proof Methodology based on Modal Transition Systems

Kim G. Larsen[1]
BRICS[2]
Aalborg Univ., Denmark,
kgl@iesd.auc.dk

Bernhard Steffen
FB Math. und Informatik,
Univ. of Passau, Germany,
steffen@fmi.uni-passau.de

Carsten Weise
LS für Informatik I,
Aachen Univ., Germany
carsten@informatik.rwth-aachen.de

Abstract

We present a constraint-oriented state-based proof methodology for concurrent software systems which exploits compositionality and abstraction for the reduction of the verification problem under investigation. Formal basis for this methodology are Modal Transition Systems allowing loose state-based specifications, which can be refined by successively adding constraints. Key concepts of our method are *projective views*, *separation of proof obligations*, *Skolemization* and *abstraction*. Central to the method is the use of *Parametrized* Modal Transition Systems. The method easily transfers to real-time systems, where the main problem are parameters in timing constraints.

1 Introduction

The use of formal methods and in particular formal verification of concurrent systems, interactive or fully automatic, is still limited to very specific problem classes. For state-based methods this is mainly due to the state explosion problem: the state graph of a concurrent systems grows exponentially with the number of its parallel components – and with the number of clocks in the real-time case –, leading to an unmanageable size for most practically relevant systems. Consequently, several techniques have been developed to tackle this problem. Here we focus on the four main streams and do not discuss the flood of very specific heuristics. Most elegant and ambitious are *compositional* methods (e.g. [ASW94, CLM89, GS90][3]), which due to the nature of parallel compositions are unfortunately rarely applicable. *Partial order* methods try to avoid the state explosion problem by suppressing unnecessary interleavings of actions [GW91, Val93, GP93]. Although extremely successful in special cases, these methods do not work in general. In practice, *Binary Decision Diagram*-based codings of the state graph are successfully applied to an interesting class of systems, see e.g. [Br86, BCMDH90, EFT91]. These codings of the state graph do not explode directly, but they may explode during verification, and it is not yet fully clear when this happens. All these techniques can be accompanied by *abstraction*: depending on the particular property under investigation, systems may be

[1]This author has been partially supported by the European Communities under CONCUR2, BRA 7166.
[2]Basic Research in Computer Science, Centre of the Danish National Research Foundation.
[3]In contrast to the first reference, the subsequent two papers address compositional reduction of systems rather than compositional verification.

dramatically reduced by suppressing details that are irrelevant for verification, see e.g. [CC77, CGL92, GL93]. Summarizing, all these methods cover very specific cases, and there is no hope for a uniform approach. Thus more application specific approaches are required, extending the practicality of formal methods.

We present a constraint-oriented state-based proof methodology for concurrent software systems which exploits compositionality and abstraction for the reduction of the verification problem under investigation. Formal basis for this methodology are Modal Transition Systems (MTS) [LT88] allowing loose state-based specifications, which can be refined by successively adding constraints. In particular, this allows extremely fine-granular specifications, which are characteristic for our approach: each aspect of a system component is specified by a number of independent constraints, one for each parameter configuration. This leads to a usually infinite number of extremely simple constraints which must all be satisfied by a corresponding component implementation. Beside exploiting compositionality in the standard (vertical) fashion, this extreme component decomposition also supports a horizontally compositional approach, which does not only separate proof obligations for subcomponents or subproperties but also for the various parameter instantiations. This is the key for the success of the following three step reduction, which may reduce even a verification problem for infinite state systems to a small number of automatically verifiable problems about finite state systems:

- *Separating the Proof Obligations.* Sections 4 and 5 present a proof principle justifying the separation and specialization of the various proof obligations, which prepare the ground for the subsequent reduction steps.

- *Skolemization.* The separation of the first step leaves us with problems smaller in size but larger in number. Due to the nature of their origin, these problems often fall into a small number of equivalence classes requiring only one prototypical proof each.

- *Abstraction.* After the first two reduction steps there may still be problems with infinite state graphs. However, the extreme specialization of the problem supports the power of abstract interpretation, which finally may reduce all the proof obligations to finite ones.

Our proof methodology is not complete, i.e., there is neither a guarantee for the possibility of a finite state reduction nor a straightforward method for finding the right amount of separation for the success of the succeeding steps or the adequate abstraction for the final verification. Still, as should be clear from the examples in the paper, there is a large class of problems and systems, where the method can be applied quite straightforwardly. Of course, the more complex the system structure the more involved will be the required search of appropriate granularity and abstraction.

Whereas complex data dependencies may exclude any possibility of 'horizontal' decomposition, our approach elegantly extends to real time systems, even over a dense time domain. In fact, this extension does not affect the possibility of a finite state reduction. For the real-time case, the basis are Timed Modal Transition Systems (TMS) [CGL93], where (weak) refinement is decidable. The TMS tool EPSILON (see again [CGL93]) can be used to find the refinements in demand.

However, in this paper *parametrized* timed modal transition systems are used. Parameters may appear either in actions (so-called *parametrized actions*) or in timing constraints.

While parametrized actions do not interfere with decidability questions, the decision procedures used in the EPSILON tool cannot be directly applied to transition systems with parameters in timing constraints.

We demonstrate our methodology by two examples: an extremely simple problem of pipelined buffers, and a specification problem of a Remote Procedure Call (RPC) posed by Broy and Lamport ([BL93]). The method is explained step by step by applying it to the simple example. For every step, we also give details on how to solve the RPC problem, thus demonstrating that the method is applicable to larger problems as well. Both problems have untimed and timed versions, where in the timed cases parameters are present in the timing constraints. For these special problems we are able to give a (simple) solution to the problem arising from parameters in timing constraints.

The next section recalls the basic theory of Modal Transition Systems, which we use for system specification. Thereafter we describe the RPC problem. The following sections explain our method in detail. Section 4 presents our notion of projective views and discusses the first reduction step. The subsequent two sections are devoted to the second and third reduction step, while Section 7 shows how to extend our method to real time systems over a dense time domain. Finally, Section 8 summarizes our conclusion and directions to future work.

2 Modal Transition Systems

In this section we give a brief introduction to the existing theory of modal transition systems. We assume familiarity with CCS. For more elaborate introductions and proofs we refer the reader to [LT88, HL89, Lar90].

When specifying reactive systems by traditional Process Algebras like e.g. CCS [Mil89], one defines the set of action transitions that can be performed (or observed) in a given system state. In this approach, any valid implementation *must* be able to perform the specified actions, which often constrains the set of possible implementations unnecessarily. One way of improving this situation within the framework of operational specification is to allow specifications where one can explicitly distinguish between transitions that are *admissible* (or allowed) and those that are *required*. This distinction allows a much more flexible specification and a much more generous notion of implementation, and therefore improves the practicality of the operational approach. Technically, this is made precise through the following notion of *modal transition systems*:

Definition 2.1 *A modal transition system is a structure* $\mathcal{S} = (\Sigma, A, \longrightarrow_\Box, \longrightarrow_\Diamond)$, *where* Σ *is a set of states,* A *is a set of actions and* $\longrightarrow_\Box$, $\longrightarrow_\Diamond \subseteq \Sigma \times A \times \Sigma$ *are transition relations, satisfying the consistency condition* $\longrightarrow_\Box \subseteq \longrightarrow_\Diamond$. □

Intuitively, the requirement $\longrightarrow_\Box \subseteq \longrightarrow_\Diamond$ expresses that anything which is required should also be allowed hence ensuring the consistency of modal specifications. When the relations $\longrightarrow_\Box$ and $\longrightarrow_\Diamond$ coincide, the above definition reduces to the traditional notion of labelled transition systems.

Syntactically, we represent modal transition systems by means of a slightly extended version of CCS. The only change in the syntax is the introduction of two prefix constructs $a_\Box.P$ and $a_\Diamond.P$ with the following semantics: $a_\Diamond.P \xrightarrow{a}_\Diamond P$, $a_\Box.P \xrightarrow{a}_\Box P$ and $a_\Box.P \xrightarrow{a}_\Diamond P$. The semantics for the other constructs follow the lines of CCS in the sense that each

rule has a version for $\longrightarrow_\Box$ and $\longrightarrow_\Diamond$ respectively. We will call this version of CCS *modal* CCS.

As usual, we consider a design process as a sequence of *refinement steps* reducing the number of possible implementations. Intuitively, our notion of when a specification S refines another (weaker) specification T is based on the following simple observation. Any behavioural aspect *allowed* by a S should also be allowed by T; and dually, any behavioural aspect which is already guaranteed by the weaker specification T must also be guaranteed by S. Using the derivation relations $\longrightarrow_\Box$ and $\longrightarrow_\Diamond$ this may be formalized by the following notion of *refinement*:

Definition 2.2 *A refinement $\mathcal{R}$ is a binary relation on Σ such that whenever $S\,\mathcal{R}\,T$ and $a \in A$ then the following holds:*

1. *Whenever $S \stackrel{a}{\longrightarrow}_\Diamond S'$, then $T \stackrel{a}{\longrightarrow}_\Diamond T'$ for some T' with $S'\,\mathcal{R}\,T'$,*
2. *Whenever $T \stackrel{a}{\longrightarrow}_\Box T'$, then $S \stackrel{a}{\longrightarrow}_\Box S'$ for some S' with $S'\,\mathcal{R}\,T'$.*

S is said to be a refinement of T in case (S,T) is contained in some refinement $\mathcal{R}$. We write $S \lhd T$ in this case. □

Note that when we apply the above definition to traditional labelled transition systems (where $\longrightarrow = \longrightarrow_\Box = \longrightarrow_\Diamond$), we obtain the well–known notion of bisimulation [Par81, Mil89]. Using standard techniques, one straightforwardly establishes that $\lhd$ is a preorder preserved by all modal CCS operators.

$\lhd$ allows *loose* specifications. This important property, which can be best explained by looking at the 'weakest' specification $\mathcal{U}$ constantly allowing any action, but never requiring anything to happen. Operationally, $\mathcal{U}$ is completely defined by $\mathcal{U} \stackrel{a}{\longrightarrow}_\Diamond \mathcal{U}$ for all actions a. It is easily verified that $S \lhd \mathcal{U}$ for any modal specification S.

Intuitively, S and T are *independent* if they are not contradictory, i.e. any action required by one is not constraint by the other. The following formal definition is due to the fact that for S and T to be *independent* all 'simultaneously'ŕeachable processes S' and T' must be indenpendent too:

Definition 2.3 *An independence relation $\mathcal{R}$ is a binary relation on Σ such that whenever $S\,\mathcal{R}\,T$ and $a \in A$ then the following holds:*

1. *Whenever $S \stackrel{a}{\longrightarrow}_\Box S'$, there is a unique T' such that $T \stackrel{a}{\longrightarrow}_\Diamond T'$ and $S'\,\mathcal{R}\,T'$,*
2. *Whenever $T \stackrel{a}{\longrightarrow}_\Box T'$, there is a unique S' such that $S \stackrel{a}{\longrightarrow}_\Diamond S'$ and $S'\,\mathcal{R}\,T'$,*
3. *Whenever $S \stackrel{a}{\longrightarrow}_\Diamond S'$ and $T \stackrel{a}{\longrightarrow}_\Diamond T'$ then $S'\,\mathcal{R}\,T'$.*

S and T are said to be independent *in case (S,T) is contained in some independence relation $\mathcal{R}$.* □

Note in particular that two specifications are independent if none of them requires any actions. Independence is important, as it allows to define conjunction on modal transition systems by:

$$\frac{S \stackrel{a}{\longrightarrow}_\Box S' \quad T \stackrel{a}{\longrightarrow}_\Diamond T'}{S \wedge T \stackrel{a}{\longrightarrow}_\Box S' \wedge T'} \qquad \frac{S \stackrel{a}{\longrightarrow}_\Diamond S' \quad T \stackrel{a}{\longrightarrow}_\Box T'}{S \wedge T \stackrel{a}{\longrightarrow}_\Box S' \wedge T'}$$

$$\frac{S \xrightarrow{a}_{\Diamond} S' \qquad T \xrightarrow{a}_{\Diamond} T'}{S \wedge T \xrightarrow{a}_{\Diamond} S' \wedge T'}$$

Of course, $S \wedge T$ is always a well-defined modal specifications (i.e. any required transition is also allowed), and in fact, for independent arguments S and T it defines their *logical* conjunction:

Theorem 2.4 Let S and T be independent modal specifications. Then $S \wedge T \lhd S$ and $S \wedge T \lhd T$. Moreover, if $R \lhd S$ and $R \lhd T$ then $R \lhd S \wedge T$.

In order to compare specifications at different levels of abstraction, it is important to abstract from transitions resulting from internal communication. The way this is done for modal transition systems follows the lines of traditional labelled transition systems. That is, for a given modal transition system $\mathcal{S} = (\Sigma, A \cup \{\tau\}, \longrightarrow_{\Box}, \longrightarrow_{\Diamond})$ we derive the modal transition system $\mathcal{S}_{\varepsilon} = (\Sigma, A \cup \{\varepsilon\}, \Longrightarrow_{\Box}, \Longrightarrow_{\Diamond})$, where $\overset{\varepsilon}{\Longrightarrow}_{\Box}$ is the reflexive and transitive closure of $\xrightarrow{\tau}_{\Box}$, and where $T \overset{a}{\Longrightarrow}_{\Box} T'$, $a \neq \varepsilon$, means that there exist T'', T''' such that

$$T \overset{\varepsilon}{\Longrightarrow}_{\Box} T'' \xrightarrow{a}_{\Box} T''' \overset{\varepsilon}{\Longrightarrow}_{\Box} T'$$

The relation $\Longrightarrow_{\Diamond}$ is defined in a similar manner.

The notion of *weak refinement* can now be introduced as follows: S weakly refines T in $\mathcal{S}$, $S \trianglelefteq T$, iff there exists a refinement relation on $\mathcal{S}_{\varepsilon}$ containing S and T.

Weak refinement $\trianglelefteq$ essentially enjoys the same pleasant properties as $\lhd$: it is a preorder preserved by all modal CCS operators except + [HL89] (including restriction, relabelling and hiding). Moreover, for ordinary labelled transition systems weak refinement reduces to the usual notion of weak bisimulation ($\approx$).

3 The Remote Procedure Call Problem

We demonstrate our method by applying it to a specification problem given by Broy and Lamport. Due to space limitations we can only present part of the problem.

The original problem consists of a *memory component* and an *RPC mechanism*. The memory component accepts read and writes from several processes, and returns the requested values (none in case of write) or raises an exception. The only exception here is *memory failure*, i.e. the memory could not read from/write to the hardware. A component in which execptions do never occur is called a *reliable memory*.

The processes are connected to the memory component via an RPC (Remote Procedure Call) mechanism. The RPC mechanism simply forwards calls from the processes to the memory, and returns from the memory to the processes. The RPC should be transparent to the user, i.e. the composition of the memory component and the RPC should be an implementation of the memory. This is what we will call the *untimed RPC problem*.

In the real-time case, the time to forward calls and returns by the RPC should be no more than δ. Further an exception should be raised if a call to the RPC does not return within $2\delta + \epsilon$ seconds. We will prove that if all calls to a reliable memory return within ϵ seconds, then the composition of the RPC and the reliable memory is an implementation of the reliable memory. This is the *timed RPC problem*.

The following is an informal specification of the memory component M, concentrating on write calls only. We assume sets `procId` of process identifiers, `memLocs` of memory

locations and `memVals` of memory values, with typical elements `id`, `loc` and `val` resp. We will often use Z as an abbreviation for the product of the three sets, i.e. $Z := \texttt{procId} \times \texttt{memLocs} \times \texttt{memVals}$, with typical element $z \in Z$.

The events occurring in the memory component are described by *parameterized actions*, taking arguments from `procId`, `memLocs` and `memVals`. The actions of M are:

$\texttt{mWr}(\texttt{id}, \texttt{loc}, \texttt{val})$:	write-call from process `id` of value `val` to location `loc`
$\texttt{write}(\texttt{id}, \texttt{loc}, \texttt{val})$:	atomic write of value `val` to location `loc` initiated by process `id`
$\overline{\texttt{mRetWr}}(\texttt{id})$:	send return from a write-request to process `id`
$\overline{\texttt{mFail}}(\texttt{id})$:	signal memory failure to process `id`

The I/O-behaviour of the memory component M then looks like this:

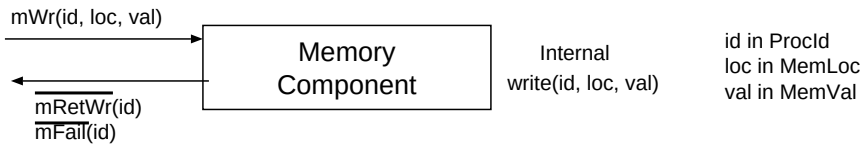

The specification of the (reliable) memory component is a conjunction of the following properties:

$P0$ The memory component engages in actions only when it is called

$P1$ Each write operation (successful or not) performs a sequence of zero or more atomic writes of the correct value to the correct location at some time between the call and return. For a successful write operation, there must be at least one atomic write.

$P2$ A memory failure is never raised.

Clearly, the memory component M is specified by the conjunction of $P0$ and $P1$, while the reliable memory M_R is the conjunction of the M and $P2$.

The RPC R simply hands calls and returns (including the memory failure exception) through. These are the actions of the RPC:

$\texttt{rWr}(\texttt{id}, \texttt{loc}, \texttt{val})$	:	remote write of value `val` to location `loc` issued by process `id`
$\overline{\texttt{rRetWr}}(\texttt{id})$	:	return from remote write issued by process `id`
$\overline{\texttt{rFail}}(\texttt{id})$	:	RPC returns an exception from a call issued by process `id`
$\overline{\texttt{mWr}}(\texttt{id}, \texttt{loc}, \texttt{val})$	:	send a write of value `val` to location `loc` initiated by process `id`
$\texttt{mRetWr}(\texttt{id})$	:	return from a write initiated by process `id`
$\texttt{mFail}(\texttt{id})$	:	memory component raised a memory failure

The I/O-behaviour of the combined components can be depicted as:

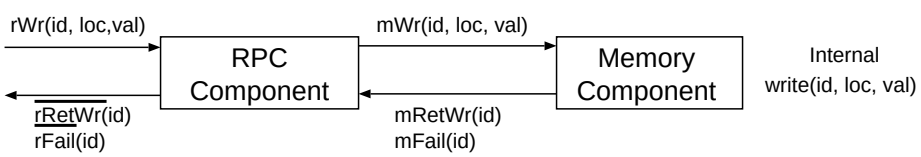

In the next sections, we will explain our method directly using a much simpler example. At the end of each section we show how our method transfers to the RPC problem. We start with the untimed case.

4 Projective Views

In the following, we present, motivate and clarify our proof methodology by means of a minimal example, which is just sufficient to explain the various phenomena.

Consider the parallel system in Figure 1. Here two parameterized, disposable component media (supposed to transmit natural numbers) A and B are composed in parallel yielding a pipeline. Informally, the component A is supposed to input a natural number on port a, then output this number on port b after which it will terminate. The behaviour of B is similar. Using modal transition systems, the parallel system may be expressed as

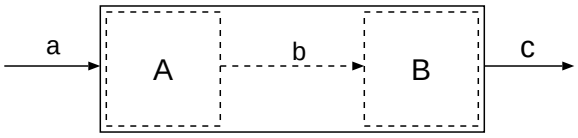

Figure 1: A Pipe Line of Two Disposable Media

follows:

$$\Big(\underbrace{a\Box x.\overline{b}\Box x}_{A} \mid \underbrace{b\Box x.\overline{c}\Box x}_{B}\Big)\backslash\{b\}$$

The behaviour of A and B are given by the two infinite–width transition systems of Figure 2. However, rather than using these direct specifications of A and B we specify the

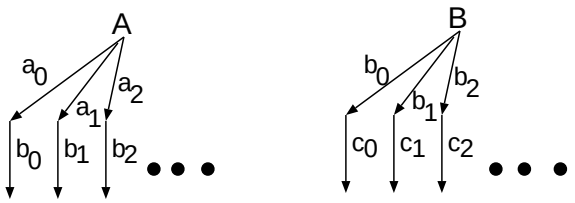

Figure 2: Behaviour of A and B.

two components behaviour using projective views A_n and B_n; one view for each possible natural number n. The projective view A_n specifies the *constraints* on the behaviour of the component A when focusing on transmission of the value n; this constraint can be expressed as the following modal transition system A_n (where we use solid lines for must- and dotted lines for may-transition):

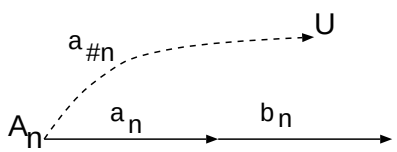

Here $a_{\neq n}$ denotes all labels of the form a_m where $m \neq n$; also $\mathcal{U}$ denotes the universal modal transition system constantly allowing all actions. Note that this 'n-th view' imposes no

constraint on the behaviour of A when transporting values different from n. The complete specification of the component A is the conjunction of all projective views[4] A_n. In fact it is easy to establish the following facts:

$$A \trianglelefteq \bigwedge_n A_n \quad \text{and} \quad \bigwedge_n A_n \trianglelefteq A \tag{1}$$

where A refers to the (infinite) transition system of Figure 2. Obviously, we may obtain similar projective views B_n for component B.

Let us now consider the problem of verifying that the overall system $\Big(A \,|\, B\Big)\backslash\{b\}$ is observationally equivalent to the system $C = a_{\Box}x.\overline{c_{\Box}}x$ (i.e. a slightly different disposable media). As A, B and C are standard transitions systems, i.e., everything allowed is also required, this problem is equivalent to showing

$$\Big(A \,|\, B\Big)\backslash\{b\} \;\trianglelefteq\; C$$

Thus (1), together with the observation that also C may be expressed as a conjunction of an infinite number of constraints C_n, leave us with the following refinement problem:

$$\Big(\bigwedge_n A_n \,|\, \bigwedge_n B_n\Big)\backslash\{b\} \;\trianglelefteq\; \bigwedge_n C_n \tag{2}$$

4.1 Application to the RPC problem

We give modal transition systems for the specification of properties $P0, P1$ and $P2$ of the memory component. Therefore we split $P1$ into two properties $P1a, P1b$ meaning

$P1a$ A write-call from process `id` cannot return unless an atomic write is performed.

$P1b$ As long as a write-call from process `id` has not returned, no atomic write with a wrong location or value occurs

The labels in the following transition systems are sets of actions (called *abstracted actions*). A single action is a shorthand for the set containing this and only this action. For the other sets, we use the usual set-theoretic connectives, and a dot-notation, where a parametrized action with dots as parameters means "the set of all actions where the dotted position is replaced by all legal values for the parameter", e.g. for a fixed $\texttt{id} \in \texttt{procId}$, $\texttt{mWr(id}, ., .)$ is the set $\{\texttt{mWr(id, loc, val)} \,|\, \texttt{loc} \in \texttt{memLocs}, \texttt{val} \in \texttt{memVals}\}$.

Our specification assumes that calls from different processes are handled concurrently. As calls from different processes do not interfere, no actions parametrized with an identifier other than `id` is constrained in the specifications of calls from process `id`. This is modelled by allowing all actions with an identifier different from the fixed `id` in any state. Instead of adding to each state a loop where all these actions are allowed, we draw boxes meaning "a state with a loop for all non-`id` actions". By this the conjunction of the specifications for all processes is the same as their parallel composition.

The following modal transition systems specify the properties for a fixed value `id`:

[4] Note that all the projective views of A are pairwise independent.

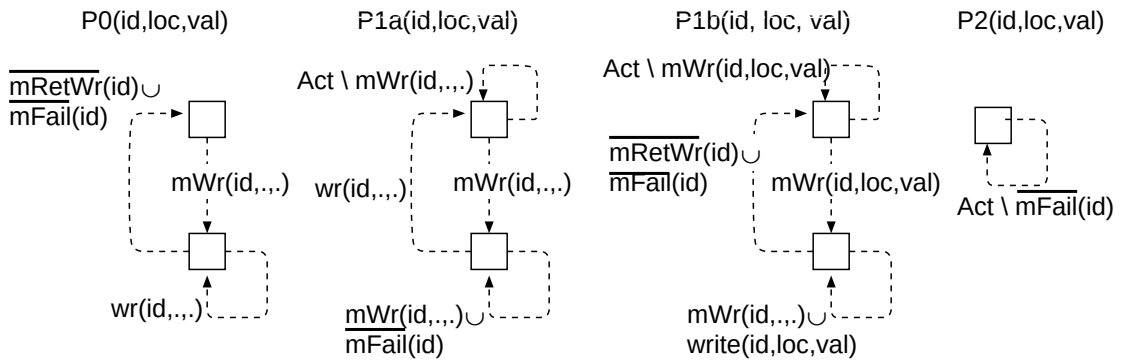

Note that only $P1b$ really depends on `loc` and `val`, and that the properties $P0, P1a, P1b$ and $P2$ are the conjunctions of the above modal specifications over all $z \in Z$.

Let $M(z)$ be the conjunction $P0(z) \wedge P1a(z) \wedge P1b(z)$, and $M_R(z) = M(z) \wedge P2(z)$. The memory component M is the conjunction of $M(z)$ over all $z \in Z$.

The transition systems for $P1a$ and $P1b$ are direct translation of their logical specifications using CTL with modalities. Transition systems resulting from this translation have may-transitions only, therefore conjunction of constraints is defined. CTL with modalities and its usefulness for our approach cannot be demonstrated here due to space limitations.

Let `Act` be the set of all actions. For two sets $R \subseteq$ `Act` (*return set*) and $T \subseteq$ `Act` (*tolerance set*), a state s and actions $a_1, \ldots, a_m \in$ `Act`. Then we use the following *macro state* for the specification of the RPC:

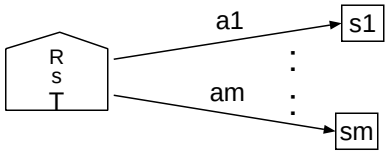

Here the edges leaving the "macro state" can be either may- or must-transition.

For a given transition system with start state s_0 and an auxiliary state s' not already in the transition system, this is meant to expand to

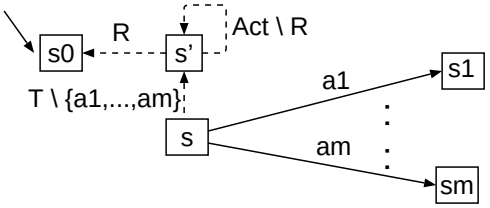

i.e. state s tolerates any action from T. If the behaviour of a tolerated action is already specified by an outgoing edge, nothing new happens. In the other case, the system goes to the auxiliary state s', where it accepts any action until a return action (from R) occurs. Return actions take the system back to the start state. Note that this corresponds to the idea of specifying procedures, where we specify each branch separately, and take care of possible other branches using the set T.

There are two main projective views of the RPC. In the first view, a write is handed through and a return received from the memory. In the second view, instead of a return

a memory failure is received. These two views $R_1(\mathtt{id},\mathtt{loc},\mathtt{val})$ and $R_2(\mathtt{id},\mathtt{loc},\mathtt{val})$ are given in the following picture:

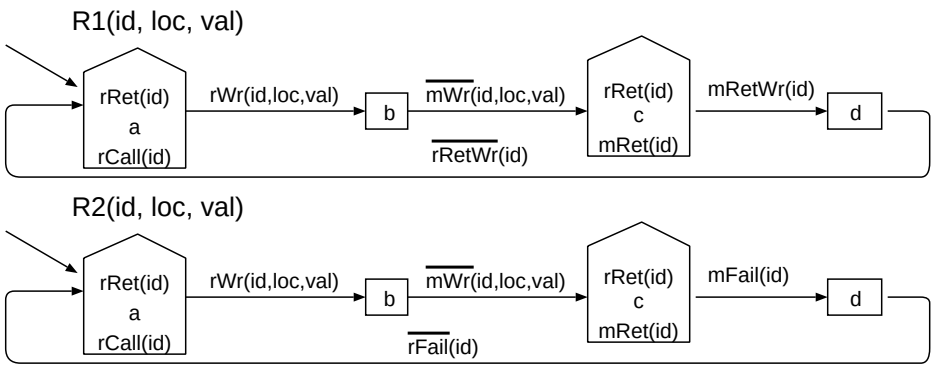

While it is natural to use must-transitions in the specification here, due to the fact that the memory component has no must-transitions at all it follows that must-transitions of the above specification can be replaced by may-transitions without affecting the correctness of the proof. This however makes all our specifications independent, so conjunction is defined. The sets in the macro states are defined as follows:

$$\begin{array}{rcl} \mathtt{rCall(id)} & := & \mathtt{rWr(id},.,.) \\ \mathtt{rRet(id)} & := & \overline{\mathtt{rRetWr}}\mathtt{(id)} \cup \overline{\mathtt{rFail}}\mathtt{(id)} \\ \mathtt{mRet(id)} & := & \mathtt{mRetWr(id)} \cup \mathtt{mFail(id)} \end{array}$$

Let $R(z) := R_1(z) \wedge R_2(z)$. The untimed specification of the RPC R is the conjunction of $R(z)$ over all z.

Let f be a relabelling mapping all actions of the RPC to the appropriate actions of the memory component, and $A := \mathtt{Act}_M$ and $H := \mathtt{write}(.,.,.)$. Then the untimed verification problem is

$$\Big(R \,|\, M/H \Big) \backslash A[f] \;\trianglelefteq\; M/H \tag{3}$$

where the internal actions of the memory (i.e. the atomic writes) are hidden.

5 Sufficient Proof Condition

Due to the properties of conjunction (cf. Lemma 2.4) the proof of (2) can obviously be reduced to the verification of

$$\Big(\bigwedge_n A_n \,|\, \bigwedge_n B_n \Big) \backslash \{b\} \;\trianglelefteq\; C_n$$

for each natural n. Thus due to Lemma 2.4 and the fact that $\trianglelefteq$ is preserved by parallel composition and restriction, it suffices to prove

$$\forall n \in N.\; \Big(A_n \,|\, B_n \Big) \backslash \{b\} \;\trianglelefteq\; C_n \tag{4}$$

There is a general proof principle behind this reduction: in order to conclude:

$$\Big(\bigwedge_{i \in I_1} A_i^1 \,|\, \ldots \,|\, \bigwedge_{i \in I_k} A_i^k \Big) \backslash L \;\trianglelefteq\; \bigwedge_{j \in I} C_j$$

it suffices for each $j \in I$ to establish:

$$\Big(\bigwedge_{i \in I_{1,j}} A_i^1 \,|\, \ldots \,|\, \bigwedge_{i \in I_{k,j}} A_i^k \Big) \backslash L \;\trianglelefteq\; C_j$$

where $I_{\ell,j} \subseteq I_\ell$ for each $\ell = 1 \ldots k$.

Of course, in general the power of this proof principle strongly depends on a good choice of the $I_{\ell,j}$, which was trivial in our example.

5.1 Application to the RPC Problem

With the same argumentation, to prove (3) it is sufficient to show

$$\forall z \in Z.\, \Big(R(z) \,|\, M(z)/H \Big) \backslash A[f] \;\trianglelefteq\; M(z)/H \tag{5}$$

6 Skolemization and Abstraction

So far we have reduced the overall verification problem of (2) to that of (4). At first sight this doesn't seem much of a reduction as (4) requires a refinement proof to be established for each natural number. Fortunately, these proofs are not really sensitive to the actual value of the natural number n. Letting k be an arbitrary natural number (or a Skolem constant) it suffices to prove:

$$\Big(A_k \,|\, B_k \Big) \backslash \{b\} \;\trianglelefteq\; C_k \tag{6}$$

in order to infer (4). Thus we are now left with the problem of establishing a *single* refinement. But still, though finite *state* the specifications A_k and B_k both have infinitely many *transitions* (as $a_{\neq k}$ is an inifinite label set). This problem can be overcome using *abstraction* (or *factorization*) with respect to an appropriate equivalence relation.

Definition 6.1 *Let $\mathcal{S} = (\Sigma, S, \longrightarrow_\Box, \longrightarrow_\Diamond)$ be a modal transition system, let $\equiv_\Sigma$ and $\equiv_S$ be equivalence relations on Σ and S, and let $\Sigma^\equiv$ and $S^\equiv$ be the sets of equivalence classes. Then* the factorization of $\mathcal{S}$ *is the modal transition system $\mathcal{S}^\equiv = (\Sigma^\equiv, S^\equiv, \longrightarrow'_\Box, \longrightarrow'_\Diamond)$, where $\longrightarrow'_\Box, \longrightarrow'_\Diamond$ are defined as follows:*

$$\frac{s \xrightarrow{a}_\Box s'}{[s]^\equiv \xrightarrow[[a]^\equiv]{}'_\Box [s']^\equiv} \qquad \frac{s \xrightarrow{a}_\Diamond s'}{[s]^\equiv \xrightarrow[[a]^\equiv]{}'_\Diamond [s']^\equiv}$$

Equivalence relations $\equiv_\Sigma$ and $\equiv_S$ are called compatible *with the modal transition system $\mathcal{S}$ iff for all $a \equiv_\Sigma b, s \equiv_S t, s' \equiv_S t'$:*

$$s \xrightarrow{a}_\Box s' \quad \textit{iff} \quad t \xrightarrow{b}_\Box t' \qquad \textit{and} \qquad s \xrightarrow{a}_\Diamond s' \quad \textit{iff} \quad t \xrightarrow{b}_\Diamond t'$$

For compatible equivalence relations, the following reduction lemma is straightforward:

Lemma 6.2 *Let $S^\equiv$ and $T^\equiv$ be processes in the factorization $\mathcal{S}^\equiv$ of $\mathcal{S}$ with respect to compatible equivalence relations $\equiv_\Sigma$ and $\equiv_S$. Then we have for arbitrary representatives S of $S^\equiv$ and T of $T^\equiv$:*

$$S^\equiv \trianglelefteq T^\equiv \quad \textit{implies} \quad S \trianglelefteq T$$

This Lemma allows us to reduce verification problems for infinite systems to problems for finite systems, as soon as an appropriate factorization can be found.

For our example, let us consider the equivalence relation $\equiv$ defined by $x_k \equiv x_k$ and $x_i \equiv x_j$ whenever $i, j \neq k$, where x ranges over $\{a, b, c\}$. Obviously, $\equiv$ is compatible with the underlying transition system. Thus the verification of (2) can further be reduced to the refinement proof between the finite $\equiv$–abstracted versions of A_k, B_k and C_k

$$\left({A_k}^{\equiv} \mid {B_k}^{\equiv} \right) \backslash \{b\} \;\trianglelefteq\; {C_k}^{\equiv} \tag{7}$$

which can easily be done by means of the automatic verification tool **Epsilon**.

6.1 Application to the RPC Problem

The same is true for the RPC problem: instead of proving (5) for all z, a proof for a prototypical z is sufficient. Most of the abstraction is already carried out by using abstracted actions. Note however that the abstracted actions are in general *not* the required equivalence classes. For the RPC problem e.g. $\texttt{write}(z)$ is an equivalence class of its own, and the set $\texttt{write}(\texttt{id}, ., .) \setminus \texttt{write}(z)$ is another equivalence classes. This specific partitioning of the atomic write actions reflects the fact that we must distinguish between a write of the correct value to the correct location and all other writes from the same process.

From the diagrams in Sect. 4.1 it is easy to see that the resulting transition systems are of a manageable size, and the proof can be carried out using the **Epsilon** tool.

7 Specifications with Time

The above example can be extended to deal with real time. For the specification we use Wang Yi's Timed CCS (see [Yi91]) together with modal specifications. For details on these so called *Timed Modal Specifications* see [CGL93]. This method can be used with any totally ordered time domain, while in the following we will assume the positive real numbers.

The passing of time is modelled by a delay action $\varepsilon(d)$, where d is a positive real number. The intuitive meaning of such a delay is that d time units pass until the end of this action. Normal actions are enabled immediately, and can be taken at any time. As an example, the process $a_{\Box}x.\varepsilon(2).\overline{b_{\Box}}x$ can execute $a_{\Box}x$ at any time. Thereafter it must delay for at least two time units before it can engage in $b_{\Box}x$.

Further we assume *maximal progress*, i.e. a communication must be performed as soon as possible. Putting $a_{\Box}x.\varepsilon(2).\overline{b_{\Box}}x$ in parallel with $\overline{a_{\Box}}x.\varepsilon(3).b_{\Box}x$ would force the communication via channel a to take place immediately, and the communication via channel b to happen after exactly three time units.

For our specification, the macro $a[l, u]$ is convenient, where a is an action and l, u are real numbers with $l < u$. The intuition is that a process $a[l, u].P$ *may* enable a after l time units and *must* enable a after u time units. In other words, communication via a may be possible after at least l time units, and will be possible at any time after u time units. This macro is defined as $a[l, u].P = (\varepsilon(l).a_{\Diamond} + \varepsilon(u).a_{\Box}).P$.

In our examples, the lower bound is always zero. The graphical presentation we use for $a[0, u].P$ is:

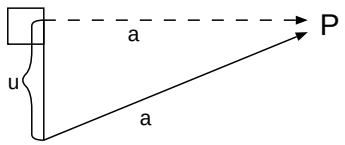

Let d be a fixed real number. Then we specify a timed process $A(d)$, which reads port a and subsequently outputs its input onto port b within d time units, by $a_{\Box}x.\overline{b}x[0,d]$.Note that this is a timed version of process A. The same construction gives timed versions $B(d)$ and $C(d)$ of B and C.

We are now going to establish that a 'pipeline' with two components with delay d should not be slower than one component with delay $2d$, i.e.

$$\Big(A(d) \mid B(d)\Big) \setminus \{b\} \;\trianglelefteq\; C(2d)$$

The same method as in the untimed case reduces the situation to

$$\Big(A_k{}^{\equiv}(d) \mid B_k{}^{\equiv}(d)\Big) \setminus \{b\} \;\trianglelefteq\; C_k{}^{\equiv}(2d)$$

for a Skolem constant k and the equivalence relation of the previous section. Now, given a specific value for d this proof can be carried out using the EPSILON tool, which treats real valued timer domains by means of the clock region automaton technique (see [AD94] for details). This technique relies on integer values for all explicit timer constants in the specification, which can be achieved by multiplication with an appropriate constant in most applications. As all timer constants are multiplied by the same constant, this does not affect the principle behaviour of the system. In our example, the obvious choice for this constant is $1/d$, leaving us with the following refinement problem

$$\Big(A_k{}^{\equiv}(1) \mid B_k{}^{\equiv}(1)\Big) \setminus \{b\} \;\trianglelefteq\; C_k{}^{\equiv}(2)$$

which can be solved using EPSILON.

Note that this proof indeed covers the statement for any d. Thus even in the presence of real time, the original verification problem is reduced to a very simple, automatically solvable problem.

7.1 Application to the RPC Problem

The following is a timed version of R_1, where passing through the calls and returns takes not more than δ seconds:

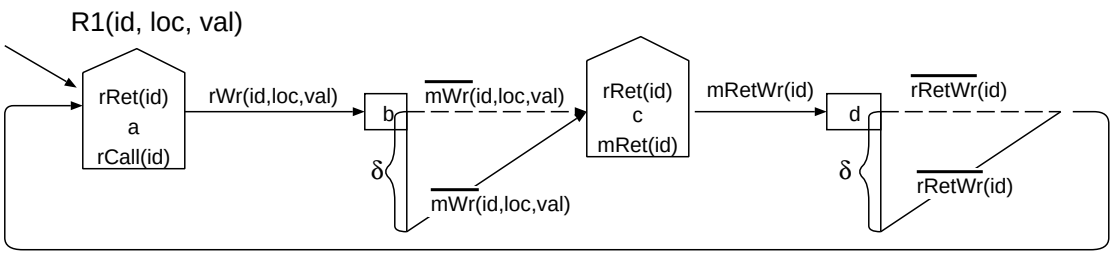

Note that actions without a timing constraint are enabled at any time. The timed version of R_2 is defined analogously (although unnecessary for the reliable memory). Call the timed RPC R^{δ}.

In the same way as the RPC we specify a demon which signals a failure if a call to the RPC does not return within $2\delta + \epsilon$ seconds. The actions of the demon are the same as those of the RPC, only the prefix `r` is replaced by a `d`. Timeout is modelled by a τ-transition. The specification of the demon $D_1(z)$ is

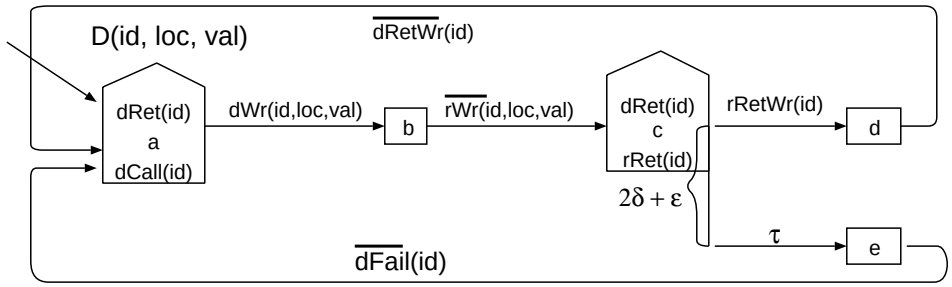

To define a timed reliable memory, we only need to alter property $P0$ by requiring the return to occur within ϵ seconds. This is done by the following:

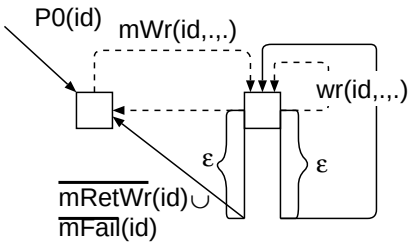

We call the resulting timed specification of the reliable memory M_R^ϵ. The timed verification problem then is

$$\Big(D^{2\delta+\epsilon} \mid R^\delta \mid M_R^\epsilon/H \Big) \backslash A[f] \;\trianglelefteq\; M_R/H$$

Note that the memory on the right hand side is the "untimed" M_R, where we interprete all actions to be enabled all the time. Further the set A and the relabelling f have to be adjusted.

This problem can once again be reduced by our method to a problem where we only need to look at a prototypical z. The involved transition systems are of small size. The problem however are the parameters occuring in the timing constraints. With two parameters δ and ϵ present the multiplication trick does not work as before.

There is a simple workaround. Computing $R^\delta \mid M_R^\epsilon/H$ by hand one finds a transiton system where the expression $2\delta + \epsilon$ occurs in all timing constraints. Thus $2\delta + \epsilon$ can be replaced by d, leaving a system with only d as parameter in timing constraints. Now the previously used multiplication trick is applicable, so that the problem can be verified using EPSILON.

8 Conclusion and Future Work

We have introduced a new constraint-oriented method for the (automated) verification of concurrent systems. Key concepts of our 'divide and conquer' method are *projective views*, *separation of proof obligations*, *Skolemization* and *abstraction*, which together support a drastic reduction of the complexity of the relevant subproblems. Of course, our

proof methodology does neither guarantee the possibility of a finite state reduction nor a straightforward method for finding the right amount of separation or the adequate abstraction. Still, there is a large class of problems and systems, where the method can be applied quite straightforwardly. Typical examples are systems with limited data dependence. Whereas involved data dependencies may exclude any possibility of 'horizontal' decomposition, our approach elegantly extends to real time systems, even over a dense time domain. In fact, the resulting finite state problems can be automatically verified using the EPSILON verification system. All this has been illustrated using a simple example of pipelined buffers. Our experience indicates that our method scales up to practically relevant problems, as demonstrated by the problem of the transparent RPC. Due to space limitations, we did not show how to do a logical specification in CTL with modalities. Such specifications are very natural problem descriptions, and they are in general easy to translate into Modal Transition Systems.

Beside further case studies and the search for good heuristics for proof obligation separation and abstraction, we are investigating the limits of tool support during the construction of constraint based specifications and the application of the three reduction steps. Whereas support by graphical interfaces and interactive editors is obvious and partly implemented in the META-Framework, a management system for synthesis, analysis and verification currently developed at the university of Passau, the limits of consistency checking and tool supported search for adequate separation and abstraction are still an interesting open research topic.

As pointed out, one major problem are parameters in the timing constraints. We are currently investigating methods – similar to the approach presented for parametrized timed automaton in [AHV93] – for checking bisimulation and (weak) refinement for *parametrized modal transition systems*.

References

[ASW94] H. Andersen, C. Stirling, G. Winskel. A Compositional Proof System for the Modal Mu-Calculus. in: Proceedings LICS, 1994.

[AD94] R. Alur, D.L. Dill. A Theory of Timed Automata. in: Theoretical Computer Science Vol. 126, No. 2, April 1994, pp. 183-236.

[AHV93] R. Alur, T.A. Henzinger, M.Y. Vardi. *Parametric real-time reasoning.* Proc. 25th STOC, ACM Press 1993, pp. 592–601.

[BL93] M. Broy, L. Lamport. Specification Problem. Case study for the Dagstuhl Seminar 9439, 1994.

[Br86] R. Bryant. Graph-Based Algorithms for Boolean Function Manipulation. in: IEEE Transactions on Computation, 35 (8). 1986.

[BCMDH90] J. Burch, E. Clarke, K. McMillan, D. Dill, L. Hwang. Symbolic Model Checking: 10^{20} States and Beyond. in: Proceedings LICS'90.

[CGL93] K. Čerāns, J.C. Godesken, K.G. Larsen. Timed Modal Specification - Theory and Tools. in: C. Courcoubetis (Ed.), Proc. 5th Int. Conf. on Computer Aided Verification (CAV '93), Elounda, Greece, June/July 1993. LNCS 697, Springer Berlin 1993, pp. 253–267.

[CGL92] E. Clarke, O. Grumber, D. Long. Model Checking and Abstraction. in: Proceedings XIX POPL'92.

[CLM89] E. Clarke, D. Long, K. McMillan. Compositional Model Checking. in: Proceedings LICS'89.

[CC77] R. Bryant. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction and Approximation of Fixpoints. in: Proceedings POPL'77.

[EFT91] R. Enders, T. Filkorn, D. Taubner. Generating BDDs for Symbolic Model Checking in CCS. in: Proceedings CAV'91, LNCS 575, 1991, pp. 203–213

[GW91] P. Godefroid, P. Wolper. Using Partial Orders for the Efficient Verification of Deadlock Freedom and Safety Properties. in: Proceedings CAV'91, LNCS 575, 1991, pp. 332–342.

[GP93] P. Godefroid, D. Pirottin. Refining Dependencies Improves Partial-Order Verification Methods. in: Proceedings CAV'93, LNCS 697, 1991, pp. 438–449.

[GL93] S. Graf, C. Loiseaux. Program Verification using Compositional Abstraction. in: Proceedings FASE/TAPSOFT'93.

[GS90] S. Graf, B. Steffen. Using Interface Specifications for Compositional Minimization of Finite State Systems. in: Proceedings CAV'90.

[HL89] H. Hüttel and K. Larsen. The use of static constructs in a modal process logic. Proceedings of Logic at Botik'89. LNCS 363, 1989.

[Lar90] K.G. Larsen. Modal specifications. In: Proceedings of Workshop on Automatic Verification Methods for Finite State Systems LNCS 407, 1990.

[LT88] K. Larsen and B. Thomsen. A modal process logic. In: Proceedings LICS'88, 1988.

[Mil89] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.

[Par81] D. Park. Concurrency and automata on infinite sequences. In P. Deussen (ed.), 5th GI Conference, LNCS 104, pp. 167–183, 1981.

[Val93] A. Valmari. On-The-Fly Verification with Stubborn Sets. in: C. Courcoubetis (Ed.), Proc. 5th Int. Conf. on Computer Aided Verification (CAV '93), Elounda, Greece, June/July 1993. LNCS 697, Springer Berlin 1993, pp. 397–408.

[Yi91] W. Yi. CCS + Time = an Interleaving Model for Real-Time Systems, Proc.18th Int. Coll. on Automata, Languages and Programming (ICALP), Madrid, July 1991. LNCS 510, Springer New York 1991, pp. 217-228.