

FORMAL VERIFICATION OF THE BINARY EXPONENTIAL BACKOFF PROTOCOL

Jozef HOOMAN

Department of Mathematics and Computing Science, Eindhoven University of Technology,
P.O. Box 513, NL 5600 MB Eindhoven, The Netherlands; e-mail: hooman@win.tue.nl

Received 12 February 1998

Abstract. We present a formal framework for the specification and verification of distributed real-time systems. To obtain mechanical support, this framework has been defined in the language of the proof checker Prototype Verification System. Intermediate stages of the design are represented by mixed terms where specifications and programming constructs can be combined. Compositional proof rules allow the verification of design steps. Here we focus on the rules for parallel composition and hiding. Their use during protocol verification is illustrated by a part of the Hypertext Transfer Protocol, the binary exponential backoff protocol.

Key words: specification and verification, distributed real-time systems, compositional proof rules, verification of design steps.

1. INTRODUCTION

Distributed real-time systems usually have a large number of possible executions, making exhaustive testing impossible. To increase confidence in the correctness of such systems, which are often safety-critical, we investigate the use of formal methods. Formal verification could be done after the complete program for the system has been written, but this will be extremely complex. Hence, we aim at the verification of design steps during the process of system development. This verify-while-design paradigm allows an early detection of design errors, ideally leading to the systems that are correct by construction.

Clearly, the formal verification of a realistic system requires suitable tool support. Usually there is a large number of proof obligations. Although many of

them are often almost trivial, they have to be checked carefully (not checking such details is a source of many errors). A tool is needed that can discharge trivial proof obligations automatically. Moreover, mechanical support is needed to keep track of the dependencies between definitions, lemmas and theorems, especially since specifications and designs are frequently changed during the development.

Formal specification of real-time systems requires a formalism in which the timed occurrence of events can be expressed. A large number of formalisms has been devised, often extending non-real-time frameworks, such as temporal logic, process algebra, automata, and Hoare logic (see, e.g., [1,2]). Here we consider a basic framework in which a component is simply specified by $spec(Eset, A)$, where $Eset$ is the *alphabet* of the component (the set of events that form the static interface), and A is an *assertion* expressing its timed behaviour. Timing is expressed in terms of a conceptual global clock of an external observer of the system.

Mixed terms are used to represent intermediate stages during system design, where parts of the system are already implemented and other parts are only specified; they allow a free mixture of specifications and programming constructs. Let $Spec_1 \Rightarrow Spec_2$ denote that $Spec_1$ refines (i.e. implements) $Spec_2$. Then, a design can be represented by a chain of subsequent refinements of mixed terms. For instance,

$$\begin{aligned} Spec &\Leftarrow (Spec_1 || Specs_2) \Leftarrow ((Spec_{11}; Spec_{12}) || Specs_2) \\ &\Leftarrow ((x := 5; Spec_{12}) || Specs_2), \text{ etc.} \end{aligned}$$

Formal verification of refinement steps is based on *compositional* proof rules, that is, for each compound programming construct there is a rule in which a specification for the construct can be derived using only the *specifications* of the components, without knowing their implementation. Hence, components can be considered as black boxes, which can be implemented later according to their specification. As an example, consider the parallel composition rule which is especially important during the early stages of system development. We will show that we can obtain a sound rule of the following form, expressing that parallel composition corresponds to the conjunction of specifications, provided some conditions on the specifications hold.

$$\frac{\text{conditions on specifications}}{spec(Eset_1, A_1) || spec(Eset_2, A_2) \Rightarrow spec(Eset_1 \cup Eset_2, A_1 \wedge A_2)}$$

Moreover, the refinement chain above shows that it should be possible to perform refinements in a certain context. For instance, from $Spec_1 \Leftarrow (Spec_{11}; Spec_{12})$ we should be able to derive $(Spec_1 || Specs_2) \Leftarrow ((Spec_{11}; Spec_{12}) || Specs_2)$. This is formalized in the so-called monotonicity rules. For instance, for parallel composition we have

$$\frac{comp_1 \Rightarrow comp_3, \quad comp_2 \Rightarrow comp_4}{comp_1 || comp_2 \Rightarrow comp_3 || comp_4}$$

Similar rules can be given for other programming constructs. Soundness of these proof rules is based on a denotational semantics, which expresses the timed behaviour of programming constructs. In this paper, we focus on parallel composition and hiding. The hiding construct can be used to encapsulate internal events.

As an application of our formal framework, we present a general approach to verify distributed real-time protocols. This is illustrated by the verification of a Binary Exponential Backoff algorithm, which is part of the Hypertext Transfer Protocol (HTTP) [3].

Tool support for our mixed framework is obtained by formulating it in the language of the verification system PVS (Prototype Verification System) [4,5]. The language of PVS is based on typed higher-order logic and contains a large number of pre-defined types. Typechecking is not decidable, but might generate proof obligations (so-called Type Check Conditions). Moreover, PVS contains an interactive theorem prover, which contains powerful decision procedures to prove simple properties automatically. In general, the user can prove a theorem by applying commands that simplify the goal until it can be proved automatically.

We do not only use PVS for the verification of concrete applications, but also for the development of a theoretical framework, for instance, to obtain mechanically checked soundness proofs of the verification rules. Further, we frequently use the “putative” theorems, i.e., we try to prove properties that ought to hold. A failure to prove such a property reveals errors, whereas a successful proof increases confidence in the formalization. All theorems and lemmas presented here have been proved in PVS.

The remainder of this paper is structured as follows. Primitives to reason about real-time properties are defined in Section 2. Section 3 contains the basic semantic framework to express the timing behaviour of components. The timed semantics of parallel composition and hiding is defined in Section 4. Verification rules for these constructs can be found in Section 5. They are used in a general approach to protocol verification, as formulated in Section 6. In Section 7, this approach is applied to the Binary Exponential Backoff Protocol. Finally, Section 8 contains a few concluding remarks.

2. REAL-TIME FRAMEWORK

The verification system PVS provides a very general type higher-order logic and allows structuring and modularization by means of parameterized theories. To be able to deal with a particular class of applications, the users have to define their own framework in this logic. Since we want to describe the behaviour of real-time systems, first, a number of timing primitives are introduced in a general theory *TimePrim*.

This theory has three parameters, a type *Time* and two orders on this type, using the pre-defined predicates *strict_order?* and *partial_order?* The theory contains an assumption about the usual relation between the two orders. If a theory imports *TimePrim* with a particular time domain and certain orders, PVS generates Type Check Conditions, requiring that the orders are indeed strict and total, respectively, and property *leq_less* holds. Naturally, we could also define \leq by $<$, but we have chosen for the assuming clause to be able to exploit the pre-defined types and orders of PVS, and thus make optimal use of the associated decision procedures of PVS.

The standard boolean connectives are overloaded for predicates on *Time*, i.e. functions from *Time* to bool. Intervals are considered as a set of time points, which are in PVS also represented as functions from *Time* to bool. As indicated by the dots "...", we only show some of the definitions.

Next, we define when a predicate holds *inside* or *during* an interval. As an example, we show a simple lemma named *dur_inside* concerning the relation between these primitives. It can be proved automatically by the PVS proof checker.

```
TimePrim[Time : TYPE, < : (strict_order?[Time]), ≤ :
    (partial_order?[Time])] : THEORY
```

```
BEGIN
ASSUMING
  leq_less : ASSUMPTION  $\forall (t_1, t_2 : \text{Time}) : t_1 \leq t_2 \Leftrightarrow t_1 < t_2 \vee t_1 = t_2$ 
ENDASSUMING
```

```
t, t_0, t_1, t_2 : VAR Time
P, Q : VAR pred[Time]
```

```
 $\neg P$  : pred[Time] =  $\lambda t : \neg P(t)$ ;
P  $\wedge$  Q : pred[Time] =  $\lambda t : P(t) \wedge Q(t)$ ;
...
[t_0, t_1] : setof[Time] = {t | t_0 ≤ t  $\wedge$  t ≤ t_1}
[t_0, t_1) : setof[Time] = {t | t_0 ≤ t  $\wedge$  t < t_1}
...
I : VAR setof[Time]
```

```
P in I : bool =  $\exists t : t \in I \wedge P(t)$ 
```

```
P during I : bool =  $\forall t : t \in I \Rightarrow P(t)$ 
```

```
dur_inside : LEMMA  $\neg(P \text{ in } I) \Leftrightarrow (\neg P) \text{ during } I$ 
END TimePrim
```

Henceforth, we do not repeat declarations of variables, using for instance, t, t_0, t_1, \dots as variables over *Time*. Moreover, we omit the structure of theory

names and imported theories, but focus on the main ideas. For more details about the formal framework we refer to [6].

3. SEMANTIC PRIMITIVES

Program semantics is defined in terms of the events that can be observed at any point of time. For simplicity, to emphasize the main concepts, we do not consider the local state of a component here. More details on the incorporation of a local state can be found in [7].

Events are represented by a non-empty type; particular events can later be defined as constants of this type. Basic primitive is an *observation function*, which assigns to each point in time the set of events that occur at that time. Operations on sets are overloaded to operations on observation functions.

Events : NONEMPTY_TYPE

ObsFuncs : TYPE = [Time \rightarrow setof[Events]]

o, o_1, o_2 : VAR ObsFuncs

Eset : VAR setof[Events]

$o_1 \cup o_2$: ObsFuncs = $\lambda t : o_1(t) \cup o_2(t)$

$o \cap$ Eset : ObsFuncs = $\lambda t : o(t) \cap$ Eset

$o \setminus$ Eset : ObsFuncs = $\lambda t : o(t) \setminus$ Eset

$o \subseteq$ Eset : bool = $\forall t : o(t) \subseteq$ Eset

The basic structure to describe program components is given by type *CompInfo*, consisting of records with two fields: α represents the alphabet of the component and *obs* describes the possible behaviours by a predicate on observation functions. The type *Components* requires that an observation of a component contains only events of its alphabet.

CompInfo : TYPE = [# α : setof[Events],
obs : pred[ObsFuncs] #]

ci : VAR CompInfo

CompProp(ci) : bool = $\forall o : \text{obs}(ci)(o) \Rightarrow o \subseteq \alpha(ci)$

Components : TYPE = {ci | CompProp(ci)}

Component comp1 *refines* component comp2 , denoted by $\text{comp1} \Rightarrow \text{comp2}$, if the alphabet of “specification” comp2 is contained in that of “implementation” comp1 and any behaviour of comp1 is also one of comp2 . The refinement relation is reflexive and transitive.

$\text{comp}, \text{comp0}, \text{comp1}, \text{comp2}, \text{comp3} : \text{VAR Components}$

$\text{comp1} \Rightarrow \text{comp2} : \text{bool} = \alpha(\text{comp2}) \subseteq \alpha(\text{comp1}) \wedge \text{obs}(\text{comp1}) \subseteq \text{obs}(\text{comp2})$

RefRef : THEOREM $\text{comp} \Rightarrow \text{comp}$

RefTrans : THEOREM $(\text{comp0} \Rightarrow \text{comp2}) \Leftrightarrow$
 $(\exists \text{comp1} : (\text{comp0} \Rightarrow \text{comp1}) \wedge (\text{comp1} \Rightarrow \text{comp2}))$

Specifications consist of an alphabet and an assertion, i.e., a predicate on observation functions. To obtain a framework of mixed terms, a specification is also of type Components. An observation of a specification should satisfy the assertion and additionally contain only events of the alphabet (to obtain CompProp). For simplicity, specifications contain a single assertion here, but the framework can be extended easily to a framework with, for instance, pre- and post-conditions.

Assertion : TYPE = $\text{pred}[\text{ObsFuncts}]$

$A, A_1, A_2 : \text{VAR Assertion}$

$\text{Valid}(A) : \text{bool} = \forall o : A(o)$

$\text{spec}(\text{Eset}, A) : \text{Components} = (\# \alpha := \text{Eset},$
 $\text{obs} := \lambda o : o \subseteq \text{Eset} \wedge A(o) \#)$

Observe that $o(t)(\text{read})$ expresses that event *read* occurs at time t in observation o . It is often convenient to be able to write $o(\text{read})(t)$ and then, using the primitives of TimePrim, also $o(\text{read})$ in $[^2,7]$. This is achieved by means of a conversion.

$E : \text{VAR Events}$

$\text{At}(o)(E)(t) : \text{bool} = o(t)(E)$

CONVERSION At

This means that an occurrence of $o(E)(t)$ is interpreted as $\text{At}(o)(E)(t)$, i.e., $o(t)(E)$.

4. SEMANTICS OF PARALLEL COMPOSITION AND HIDING

In this section, a denotational semantics of parallel composition and hiding is defined.

4.1. Semantics of parallel composition

The denotational semantics of the parallel composition of two components is defined in terms of semantics of these components. For the alphabet, we simply take the union of the alphabets of the components. Hiding of internal events is considered a separate operation, defined in the next section. Inspired by the trace-based untimed semantics of parallel composition [8], we define the timed behaviour by means of so-called projections, here represented by intersection. The main requirement is that the projection of an observation of the parallel composition onto the alphabet of one component, should lead to an observation of this component. Additionally, observations should only contain events of the alphabet.

$$\begin{aligned} //(\text{comp1}, \text{comp2}) : \text{Comps} = \\ (\# \alpha := \alpha(\text{comp1}) \cup \alpha(\text{comp2}), \\ \text{obs} := \lambda o : (\exists o_1, o_2 : \text{obs}(\text{comp1})(o_1) \wedge \text{obs}(\text{comp2})(o_2) \wedge \\ o \cap \alpha(\text{comp1}) = o_1 \wedge o \cap \alpha(\text{comp2}) = o_2 \wedge \\ o \subseteq \alpha(\text{comp1}) \cup \alpha(\text{comp2})) \#) \end{aligned}$$

The definition above is convenient for the soundness proof of the verification rule that will be presented in the next section; other equivalent versions, e.g., using the intersection of behaviours, can be found in [6]. We have proved that parallel composition is commutative and associative.

$$\text{ParComm} : \text{LEMMA } \text{comp1} // \text{comp2} = \text{comp2} // \text{comp1}$$

ParAssoc :

$$\text{LEMMA } (\text{comp1} // \text{comp2}) // \text{comp3} = \text{comp1} // (\text{comp2} // \text{comp3})$$

4.2. Semantics of hiding

The hiding (or encapsulation) construct $\text{comp} - \text{Eset}$ hides the events of set Eset from component comp . Clearly, this means that these events are removed from the alphabet. Moreover, the events from Eset are removed from the observations of comp .

$$\begin{aligned} - (\text{comp}, \text{Eset}) : \text{Comps} = \\ (\# \alpha := \alpha(\text{comp}) \setminus \text{Eset}, \\ \text{obs} := \lambda o : (\exists o_1 : \text{obs}(\text{comp})(o_1) \wedge o = o_1 \setminus \text{Eset}) \#) \end{aligned}$$

5. VERIFICATION RULES

We present the consequence rule and the rules for parallel composition and hiding. The soundness of these rules has been proved in PVS on the basis of the semantics of the previous section.

5.1. Consequence rule

The consequence rule allows us to weaken assertions.

ConsRule : THEOREM $\text{Eset0} = \text{Eset} \wedge \text{Valid}(A_0 \Rightarrow A) \Rightarrow$
 $(\text{spec}(\text{Eset0}, A_0) \Rightarrow \text{spec}(\text{Eset}, A))$

5.2. Rules for parallel composition

It is rather easy to show the next monotonicity property, which makes it possible to perform refinements in a parallel context.

MonoPar : THEOREM $(\text{comp1} \Rightarrow \text{comp3}) \wedge (\text{comp2} \Rightarrow \text{comp4}) \Rightarrow$
 $(\text{comp1} \parallel \text{comp2} \Rightarrow \text{comp3} \parallel \text{comp4})$

The next rule for parallel composition essentially expresses that the parallel composition of specifications corresponds to the conjunction of assertions. To achieve a sound rule, it is required that the validity of the assertion of a component depends only on its alphabet, as expressed by predicate *OnlyDepEve*.

OnlyDepEve(A, Eset) : bool = $\forall o : A(o) \Leftrightarrow A(o \cap \text{Eset})$

ParCompRule : THEOREM $\text{OnlyDepEve}(A_1, \text{Eset1}) \wedge$
 $\text{OnlyDepEve}(A_2, \text{Eset2}) \Rightarrow$
 $\text{spec}(\text{Eset1}, A_1) \parallel \text{spec}(\text{Eset2}, A_2) \Rightarrow$
 $\text{spec}(\text{Eset1} \cup \text{Eset2}, A_1 \wedge A_2)$

The soundness of the parallel composition rule, i.e., theorem *ParCompRule*, has been proved along the following lines. Assume $\text{OnlyDepEve}(A_i, \text{Eset}_i)$, for $i = 1, 2$.

First observe that the alphabets of the components on both sides of the refinement are equal. Next, consider an observation $o \in \text{obs}(\text{spec}(\text{Eset1}, A_1) \parallel \text{spec}(\text{Eset2}, A_2))$. Let $i \in \{1, 2\}$. By the definition of the semantics, there exists o_i such that $o \cap \text{Eset}_i = o_i$ and $\text{obs}(\text{spec}(\text{Eset}_i, A_i))(o_i)$, thus $A_i(o_i)$. Hence, $A_i(o \cap \text{Eset}_i)$. Then *OnlyDepEve* leads to $A_i(o)$. Since $o \subseteq \text{Eset1} \cup \text{Eset2}$, this leads to $o \in \text{obs}(\text{spec}(\text{Eset1} \cup \text{Eset2}, A_1 \wedge A_2))$.

5.3. Proof rules for hiding

For the hiding construct, we also show a monotonicity property expressed by the theorem *HideMono*. Next, we prove a rule which expresses that a set of events can be hidden by simply removing it from the alphabet, provided the assertion does not depend on the events removed, i.e., it only depends on the resulting events.

HideMono : THEOREM $(\text{comp1} \Rightarrow \text{comp2}) \Rightarrow (\text{comp1} - \text{Eset} \Rightarrow \text{comp2} - \text{Eset})$

HideRule : THEOREM $\text{OnlyDepEve}(A, \text{Eset} \setminus \text{Eset0}) \Rightarrow$
 $\text{spec}(\text{Eset}, A) - \text{Eset0} \Rightarrow \text{spec}(\text{Eset} \setminus \text{Eset0}, A)$

The soundness of this rule is proved as follows. Assume $\text{OnlyDepEve}(A, \text{Eset} \setminus \text{Eset0})$. First observe that $\alpha(\text{spec}(\text{Eset}, A) - \text{Eset0}) = \alpha(\text{spec}(\text{Eset} \setminus \text{Eset0}, A))$. Next consider $o \in \text{obs}(\text{spec}(\text{Eset}, A) - \text{Eset0})$. Hence there exists an o_1 such that $o = o_1 \setminus \text{Eset0}$ and $o_1 \in \text{obs}(\text{spec}(\text{Eset}, A))$, thus $o_1 \subseteq \text{Eset}$ and $A(o_1)$. By $\text{OnlyDepEve}(A, \text{Eset} \setminus \text{Eset0})$ we obtain $A(o_1 \cap \text{Eset} \setminus \text{Eset0})$. Since $o_1 \subseteq \text{Eset}$ this leads to $A(o_1 \setminus \text{Eset0})$ and hence $A(o)$. Since $o \subseteq \text{Eset} \setminus \text{Eset0}$, we obtain $o \in \text{obs}(\text{spec}(\text{Eset} \setminus \text{Eset0}, A))$.

6. APPLICATION OF PROTOCOL VERIFICATION

The formal framework is rather general, intended for a wide range of applications. It is convenient to have some guidelines for a particular class of applications. Here we consider the verification of distributed real-time protocols and formulate a few simple steps for a network of nodes.

1. Model the application domain, i.e., describe the main primitives that are needed to express the service specification.
2. Specify the service to be provided by the network.
3. Specify the communication mechanism between nodes.
4. Specify the protocol performed by each node, in terms of its alphabet only.
5. Verify the protocol, using the parallel composition rule (and the consequence rule), i.e., prove that the specifications of the nodes (point 4) and the communication mechanism (point 3) lead to the required service (point 2). Internal events can be removed by the hiding rule.


7. BINARY EXPONENTIAL BACKOFF PROTOCOL

As an example of the approach in the previous section, we consider the verification of the Binary Exponential Backoff Protocol, a small part of the HyperText Transfer Protocol HTTP/1.1. This protocol is a proposed standard, described in RFC 2068 [3]. It organizes the transfer of data between a client (e.g., a network browser) and a server. The client and the server communicate by means of a TCP connection.

Document RFC 2068 describes only the client part of the protocol, along the following lines.

1. Initiate a new connection to the server.
2. Transmit request headers.
3. Initialize R to the round-trip time (if known) or to five seconds.
4. Compute $T = R^N$, with N , the number of retries.
5. Wait for an error or T seconds.
6. If no error has been received after T seconds, start transmitting data.
7. If the connection is aborted, repeat step 1.

Given this description, the figure shows the events that are relevant for the Binary Exponential Backoff Protocol.



Client and server of the HTTP.

RFC 2068 is unclear about the purpose and the context of this protocol, but a domain expert provided some useful hints. The idea is that the client wants to send data to the server and requests permission to do this. The request, however, might be invalid, e.g., because the authorization code is not correct, and the server sends an error message (instead of a normal response). If the error is delayed, the client decides to start sending the data already (step 6 above). When the server receives too much unwanted data, it aborts the connection, and then also the error message might get lost. Consequently, the client will try again, with a new time-out value. The protocol is intended to guarantee that the client eventually receives the error.

Our aim is to investigate the essential functioning of the protocol on an abstract level to achieve better understanding. This includes relations between the timing of components and, for instance, the value of the time-out of the client.

7.1. Model the application domain

As a time domain, we take the reals with the standard orderings, which are already available in PVS. Our real-time framework is incorporated by importing theory TimePrim. Moreover, the framework of the previous sections is imported, although this is not shown here. Messages and nodes are defined as enumeration types, which implies that all elements are different, a fact that is used by the decision procedures of PVS. A *send* and a *receive* of a message by a node and an *abort* are declared as events. Moreover, assume given a predicate that expresses which messages are erroneous.

Time : TYPE = real

NonNegTime : TYPE = $\{t : \text{Time} \mid t \geq 0\}$

IMPORTING TimePrim[Time, <, ≤], ...

Messages : TYPE = {request, data, error}

m, m_1, m_2 : VAR Messages

Nodes : TYPE = {client, server}

node, node1, node2 : VAR Nodes

send(node, m) : Events

rec(node, m) : Events

abort : Events

erroneous(m) : bool

7.2. Specify the required service

The service specification expresses that if the client sends an erroneous request, it receives an error between lower bound L and upper bound U . These time bounds have been added to be able to investigate the timing relations between components.

$L, U : \text{Time}$

$\text{ETL} : \text{setof}[\text{Events}] = \{E \mid E = \text{send}(\text{client}, \text{request}) \vee E = \text{rec}(\text{client}, \text{error})\}$

$\text{ATL} : \text{Assertion} = \lambda o : \forall t : o(\text{send}(\text{client}, \text{request}))(t) \wedge \text{erroneous}(\text{request}) \Rightarrow$
 $o(\text{rec}(\text{client}, \text{error})) \text{ in } [t + L, t + U]$

$\text{TLSpec} : \text{Components} = \text{spec}(\text{ETL}, \text{ATL})$

7.3. Specify the communication mechanism

Next, we axiomatize the relation between send, receive, and abort events. Only two properties of the underlying TCP protocol are needed. First, we specify that a message sent will be received between certain time bounds, provided no abort event takes place.

$\text{TransDelayL}, \text{TransDelayU} : \text{NonNegTime}$

$\text{SendRec} : \text{AXIOM}$

$o(\text{send}(\text{node1}, m))(t) \wedge \neg o(\text{abort}) \text{ during } [t, t + \text{TransDelayU}] \Rightarrow$
 $\forall (\text{node2} \mid \text{node2} \neq \text{node1}) :$
 $o(\text{rec}(\text{node2}, m)) \text{ in } [t + \text{TransDelayL}, t + \text{TransDelayU}]$

Moreover, we express that a message is only received if it has been sent between certain time bounds.

$\text{RecSend} : \text{AXIOM } o(\text{rec}(\text{node1}, m))(t) \Rightarrow$

$\exists t_0, \text{node2} : \text{node2} \neq \text{node1} \wedge$

$t_0 \in [t - \text{TransDelayU}, t - \text{TransDelayL}] \wedge$

$o(\text{send}(\text{node2}, m))(t_0)$

The time bounds mentioned here are not present in the TCP protocol, but have been added to be able to reason about the transmission speed.

7.4. Specify the protocol performed by the nodes

Next, the protocol performed by the nodes is specified. Here we have to specify the server and the client.

7.4.1. Specify protocol of the server

The server should send an error message between certain time bounds if an erroneous request is received.

ErrL, ErrU : NonNegTime

AS1 : Assertion = $\lambda o : \forall t : o(\text{rec}(\text{server}, \text{request}))(t) \wedge \text{erroneous}(\text{request}) \Rightarrow$
 $o(\text{send}(\text{server}, \text{error})) \text{ in } [t + \text{ErrL}, t + \text{ErrU}]$

The main complication of the protocol is that the server might close the connection by an abort when it receives too much unwanted data. For the correctness it is, however, important that this is the only reason for the server to abort the connection. So when it does an abort, some data must have been received recently, i.e., between certain time bounds.

AbortL, AbortU : NonNegTime

AS2 : Assertion = $\lambda o : \forall t : o(\text{abort})(t) \Rightarrow$
 $\exists t_0 : t_0 \in [t - \text{AbortU}, t - \text{AbortL}] \wedge$
 $o(\text{rec}(\text{server}, \text{data}))(t_0)$

AS : Assertion = AS1 \wedge AS2

ES : setof[Events] = {E | E = rec(server, request) \vee E = rec(server, data) \vee
E = send(server, error) \vee E = abort}

Server : Components = spec(ES, AS)

7.4.2. Specify protocol of the client

To get an insight into the main principles of the protocol, we abstract from the algorithm of the client that dynamically computes the distance between a request and subsequent data. It is essential only that there is sufficient time between a request and subsequent data to allow the error to reach the client. Hence, as a first attempt, the client has been specified by the following formula.

$o(\text{send}(\text{client}, \text{request}))(t) \Rightarrow$
 $\neg o(\text{send}(\text{client}, \text{data})) \text{ during } [t, t + \text{NoDataAfterPeriod}]$

An attempt to verify this protocol revealed that this assertion is too weak. A request might be processed very fast (by TCP and server), whereas an old data message, sent before the request, might be processed very slowly, generating a disturbing abort. Hence, the specification is modified to include also a period before the request.

NoDataBeforePeriod, NoDataAfterPeriod : NonNegTime

AC : Assertion = $\lambda o : \forall t : o(\text{send}(\text{client}, \text{request}))(t) \Rightarrow$
 $\neg o(\text{send}(\text{client}, \text{data})) \text{ during } [t - \text{NoDataBeforePeriod},$
 $t + \text{NoDataAfterPeriod}]$

EC : setof[Events] = $\{E \mid E = \text{send}(\text{client}, \text{request}) \vee$
 $E = \text{send}(\text{client}, \text{data}) \vee E = \text{rec}(\text{client}, \text{error})\}$

Client : Components = spec(EC, AC)

7.5. Verification of the protocol

To verify the protocol, we first prove that a client request implies that no abort occurs in the next D time units, provided certain conditions hold (as explained below).

ReqAbort : LEMMA NoDataBeforePeriod \geq TransDelayU + AbortU \wedge
 NoDataAfterPeriod $\geq D - \text{TransDelayL} - \text{AbortL} \wedge$
 AS2(o) \wedge AC(o) \Rightarrow
 $(\forall t : o(\text{send}(\text{client}, \text{request}))(t) \Rightarrow$
 $\neg o(\text{abort}) \text{ during } [t, t + D])$

Lemma *ReqAbort* requires that *NoDataBeforePeriod* is greater than the slowest processed data, i.e., a maximal TCP delay *TransDelayU* plus the upper bound of the server on performing an abort, *AbortU*. Similarly, *NoDataAfterPeriod* should be greater than the required upper bound D minus the fastest transmission delay *TransDelayL* and the fastest response by the server, *AbortL*.

For the correctness of the protocol, it is required that no abort is generated during the maximal time needed to transmit an error, which equals $2 \times \text{TransDelayU} + \text{ErrU}$. Hence, D in the lemma above is replaced by this expression. Moreover, this expression determines the upper bound U of the service specification. Similarly, lower bound L is determined by the fastest transmission. This leads to the following timing constraints which allow us to prove that the assertions of the server and the client lead to the required specification, as expressed in lemma *ATLLem*.

TimingConstraints : bool =
 $L \leq \text{ErrL} + 2 \times \text{TransDelayL} \wedge$
 $U \geq \text{ErrU} + 2 \times \text{TransDelayU} \wedge$
 NoDataBeforePeriod $\geq \text{TransDelayU} + \text{AbortU} \wedge$
 NoDataAfterPeriod $\geq \text{ErrU} + 2 \times \text{TransDelayU} - \text{TransDelayL} - \text{AbortL}$

ATLLem : LEMMA TimingConstraints \Rightarrow Valid(AS \wedge AC \Rightarrow ATL)

To apply the parallel composition rule, we first prove that the specifications of server and client only depend on their alphabet. Further, observe that we do not yet obtain the alphabet of the service specification, but still have some additional events, represented by *IntEve*.

ASEveLem : FACT OnlyDepEve(AS, ES)

ACEveLem : FACT OnlyDepEve(AC, EC)

IntEve : setof[Events] = {E | E = send(client, data) ∨
E = rec(server, request) ∨ ...}

TLPar : THEOREM TimingConstraints ⇒
(Client || Server ⇒ spec(ETL ∪ IntEve, ATL))

Next, the additional events of *IntEve* can be removed by the hiding rule, provided they are different from the other events. This is expressed by the axiom *CommEventsDiffer*.

ATLEveLem : FACT OnlyDepEve(ATL, ETL)

CommEventsDiffer : AXIOM send(node1, m₁) ≠ rec(node2, m₂) ∧
send(node, m) ≠ abort ∧ ...

TLCor : THEOREM TimingConstraints ⇒
((Client || Server) – IntEve ⇒ TLSpec)

8. CONCLUDING REMARKS

A general framework for the formal specification and mechanical verification of distributed real-time systems has been presented. It can be considered as an extension and a modification of mixed frameworks for untimed systems [9–11]. Alternatives for the semantics and an application to hybrid systems can be found in [6]. The treatment of parallel composition for components with a local state has been studied in [7].

The formalism has been applied to a distributed real-time protocol, the Binary Exponential Backoff Protocol. It has been verified on an abstract level, abstracting for instance, from the algorithm which is used to compute the distance in time between the request and subsequent data dynamically. This computation could be considered in a continuation of this work, where the server and the client can be

refined and implemented. (Note that the correctness of the client protocol described in Section 7 requires $R > 1$.) A related refinement can be found in the work on a distributed real-time arbitration protocol, where, first, the protocol is verified on an abstract level, and next, the nodes are implemented in isolation according to their specification [12].

Other applications of our approach to protocol verification in PVS concern part of the ACCESS.bus protocol [13] and a membership protocol, with a dynamically changing network and local clocks [14].

In [15], an alternative approach has been applied to the specification and verification of the link layer of the serial bus protocol P1394. Since the informal specification is based on communicating state machines, this framework has been formalized in PVS. The intention was to stay close to the informal text, motivated by the importance of the step from the informal to the formal specification. This is also a topic of current work on requirements engineering.

ACKNOWLEDGEMENTS

The author is grateful to Koen Holtman for his suggestion to verify the Binary Exponential Backoff Protocol and to his explanation of context and purpose of the protocol. Thanks are due to Leendert Pieter van Drimmelen, Arnaud Gouder de Beauregard, Ronald Marcelis, and Alex Vrijksen for their work on this protocol as an exercise for their postgraduate course on PVS. Although their verification attempts were different from the solution presented here, their insights certainly contributed to the solution presented here.

REFERENCES

1. de Bakker, I. W., Huizing, C., de Roever, W.-P., and Rozenberg, G. (eds.). *Proc. REX Workshop on Real-Time: Theory in Practice, LNCS 600*. Springer, Berlin, 1992.
2. Langmaack, H., de Roever, W.-P., and Vytupil, J. (eds.). *Formal Techniques in Real-Time and Fault-Tolerant Systems, LNCS 863*. Springer, Berlin, 1994.
3. Fielding, R., Irvine, U., Gettys, J., Mogul, J., Frysky, H., and Berners-Lee, T. *Hypertext Transfer Protocol* – <http://1.1>. Request for Comments (RFC) 2068, <http://ds.internic.net/ds/rfc-index.html>, 1997.
4. Owre, S., Rushby, J., and Shankar, N. PVS: A prototype verification system. In *11th Conf. on Automated Deduction. Lecture Notes in Artificial Intelligence, 607*. Springer, Berlin, 1992, 748–752.
5. Owre, S., Rushby, J., Shankar, N., and von Henke, F. Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, 1995, **21**, 2, 107–125.
6. Hooman, J. Compositional verification of real-time applications. In *Proc. COMPOS '97, Compositionality – The Significant Difference, LNCS*. Springer, Berlin, 1998 (to be published).

7. Hooman, J. Developing proof rules for distributed real-time systems with PVS. In *Proc. the Workshop on Tool Support for System Development and Verification*, 1998 (to be published).
8. Hoare, C. A. R. *Communicating Sequential Processes*. Prentice Hall, Englewood Cliffs, N. J., 1985.
9. Olderog, E.-R. Process theory: Semantics, specification and verification. In *Current Trends in Concurrency*, LNCS 224. Springer, Berlin, 1985, 442–509.
10. Olderog, E.-R. *Nets, Terms and Formulas*, *Cambridge Tracts in Theoretical Computer Science*, 23. Cambridge University Press, Cambridge, 1991.
11. Zwiers, J. *Compositionality, Concurrency and Partial Correctness*, LNCS 321. Springer, Berlin, 1989.
12. Hooman, J. Compositional verification of a distributed real-time arbitration protocol. *Real-Time Systems*, 1994, 6, 2, 173–205.
13. Hooman, J. Verifying part of the ACCESS.bus protocol using PVS. In *Proc. 15th Conf. on the Foundations of Software Technology and Theoretical Computer Science*, 6, 2., LNCS 1026. Springer, Berlin, 1995, 96–110.
14. Hooman, J. Verification of distributed real-time and fault-tolerant protocols. In *Algebraic Methodology and Software Technology (AMAST'97)*, LNCS 1349. Springer, Berlin, 1997, 261–275.
15. Kühne, L., Hooman, J., and de Roever, W.-P. Towards mechanical verification of parts of the IEEE P1394 serial bus. In *2nd International Workshop on Applied Formal Methods in System Design* (Lovrek, I., ed.). University of Zagreb, Faculty of Electrical Engineering and Computing, 1997, 73–85.

BINAAR-EKSPONENTSIAALSE TAGASIVÕTMISPROTOKOLLI FORMAALNE VERIFITSEERIMINE

Jozef HOOMAN

On esitatud formalism reaalaaja hajussüsteemide spetsifitseerimiseks ja verifitseerimiseks. Automaatse tõestamise hõlbustamiseks on formalism automaat-tõestaja PVS keeles. Projekteerimise vahestaadiume on kirjeldatud segatermide keeles, milles termideks võivad olla nii spetsifikatsiooni kui ka programmi konstruktsioonid. Kompositsioonilised tuletusreeglid võimaldavad verifitseerida projekteerimissammude korrektsust. Töös on kasutatud paralleelkompositsioonile ja varjamisoperaatorile vastavaid tuletusreegleid. Nende reeglite kasutamist illustreerib HTTP binaar-eksponentsiaalse tagasivõtmisprotokolli formaalne verifitseerimine.