

Mechanical Verification of a Two-Way Sliding Window Protocol

Bahareh Badban^{a,1}, Wan Fokkink^b and Jaco van de Pol^c

^a *University of Konstanz, Department of Computer and Information Science*

badban@inf.uni-konstanz.de

^b *Vrije Universiteit Amsterdam, Department of Computer Science*

wanf@cs.vu.nl

^c *University of Twente, Department of EEMCS, Formal Methods and Tools*

j.c.vandepol@ewi.utwente.nl

Abstract. We prove the correctness of a two-way sliding window protocol with piggybacking, where the acknowledgments of the latest received data are attached to the next data transmitted back into the channel. The window size of both parties are considered to be finite, though they can be of different sizes. We show that this protocol is equivalent (branching bisimilar) to a pair of FIFO queues of finite capacities. The protocol is first modeled and manually proved for its correctness in the process algebraic language of μCRL . We use the theorem prover PVS to formalize and to mechanically prove the correctness. This implies both safety and liveness (under the assumption of fairness).

Keywords. two-way sliding window protocol, specification in μCRL , verification with PVS, a pair of FIFO queues

Introduction

A sliding window protocol [7] (SWP) ensures successful transmission of messages from a sender to a receiver through a medium in which messages may get lost. Its main characteristic is that the sender does not wait for incoming acknowledgments before sending next messages, for optimal use of bandwidth. Many data communication systems include a SWP, in one of its many variations.

In SWPs, both the sender and the receiver maintain a buffer. We consider a *two-way* SWP, in which both parties can both send and receive data elements from each other. One way of achieving full-duplex data transmission is to have two separate communication channels and use each one for simplex data traffic (in different directions). Then there are two separate physical circuits, each with a forward channel (for data) and a reverse channel (for acknowledgments). In both cases the bandwidth of the reverse channel is almost entirely wasted. In effect, the user is paying for two circuits but using the capacity of one. A better idea is to use the same circuit in both directions. Each party maintains two buffers, for storing the two opposite data streams. In this two-way version of the SWP, an acknowledgment that is sent from one party to the other may get a free ride by attaching it to a data element. This method for efficiently passing acknowledgments and data elements through a channel in the same direction, which is known as *piggybacking*, is used broadly in transmission control protocols, see [39]. The main advantage of piggybacking is a better use of available bandwidth. The

¹Corresponding Author: Bahareh Badban, University of Konstanz, Department of Computer and Information Science, Universitätsstr. 10, P.O. Box D67, 78457 Konstanz, Germany. Tel.: +49 (0)7531 88 4079; Fax: +49 (0)7531 88 3577; E-mail: badban@inf.uni-konstanz.de.

extra acknowledgment field in the data frame costs only a few bits, whereas a separate acknowledgment would need a header and a checksum. In addition, fewer frames sent means fewer ‘frame arrived’ interrupts.

The current paper builds on a verification of a one-way version of the SWP in [1, 9]. The protocol is specified in μCRL [15], which is a language based on process algebra and abstract data types. The verification is formalized in the theorem prover PVS [29]. The correctness proof is based on the so-called *cones and foci* method [10, 18], which is a symbolic approach towards establishing a branching bisimulation relation. The starting point of the cones and foci method are two μCRL specifications, expressing the implementation and the desired external behavior of a system. A *state mapping* ϕ relates each state of the implementation to a state of the desired external behavior. Furthermore, the user must declare which states in the implementation are *focus points*, whereby each reachable state of the implementation should be able to get to a focus point by a sequence of hidden transitions, carrying the label τ . If a number of *matching criteria* are met, consisting of equations between data objects, then states s and $\phi(s)$ are branching bisimilar. Roughly, the matching criteria are: (1) if $s \xrightarrow{\tau} s'$ then $\phi(s) = \phi(s')$, (2) each transition $s \xrightarrow{a} s'$ with $a \neq \tau$ must be matched by a transition $\phi(s) \xrightarrow{a} \phi(s')$, and (3) if s is a focus point, then each transition of $\phi(s)$ must be matched by a transition of s .

The crux of the cones and foci method is that the matching criteria are formulated syntactically, in terms of relations between data terms. Thus, one obtains clear proof obligations, which can be verified with a theorem prover. The cones and foci method provides a general verification approach, which can be applied to a wide range of communication protocols and distributed algorithms.

The main motivations for the current research is to provide a mechanized correctness proof of the most complicated version of the SWP in [39], including the piggybacking mechanism. Here we model buffers (more realistically) as ordered lists, without multiple occurrences of the same index. Therefore two buffers are equal only if they are identical. That is, any swapping or repetition of elements results in a different buffer. It was mainly this shift to ordered lists without duplications (i.e. each buffer is uniquely represented with no more than once occurrence of each index), that made this verification exercise hard work. Proving that each reachable state can get to a focus point by a sequence of τ -transitions appeared to be considerably hard (mainly because communication steps of the two data streams can happen simultaneously).

The medium between the sender and the receiver is modeled as a lossy queue of capacity one. With buffers of sizes $2n_1$ and $2n_2$, and windows of sizes n_1 and n_2 , respectively, we manually (paper-and-pencil) prove that the external behavior of this protocol is branching bisimilar [43] to a pair of FIFO queues of capacity $2n_1$ and $2n_2$. This implies both safety and liveness of the protocol (the latter under the assumption of fairness, which intuitively states that no message gets lost infinitely often).

The structure of the proof is as follows. First, we linearize the specification, meaning that we get rid of parallel operators. Moreover, communication actions are stripped from their data parameters. Then we eliminate modulo arithmetic, using an idea from Schoone [35]. Finally, we apply the cones and foci technique, to prove that the linear specification without modulo arithmetic is branching bisimilar to a pair of FIFO queues of capacity $2n_1$ and $2n_2$. The lemmas for the data types, the invariants, the transformations and the matching criteria have all been checked using PVS 2.3. The PVS files are available via <http://www.cs.utwente.nl/~vdpol/piggybacking.html>.

The remainder of this paper is set up as follows. In Section 1 the μCRL language is explained. In Section 2 the data types needed for specifying the protocol are presented. Section 3 features the μCRL specifications of the two-way SWP with piggybacking, and its external

behavior. In Section 4, three consecutive transformations are applied to the specification of the SWP, to linearize the specification, eliminate arguments of communication actions, and get rid of modulo arithmetic. In Section 5, properties of the data types and invariants of the transformed specification are formulated; their proofs are in [2]. In Section 6, it is proved that the three transformations preserve branching bisimilarity, and that the transformed specification behaves as a pair of FIFO queues. In Section 7, we present the formalization of the verification of the SWP in PVS. We conclude the paper in Section 8.

Related Work

Sliding window protocols have attracted considerable interest from the formal verification community. In this section we present an overview. Many of these verifications deal with unbounded sequence numbers, in which case modulo arithmetic is avoided, or with a fixed finite buffer and window size at the sender and the receiver. Case studies that do treat arbitrary finite buffer and window sizes mostly restrict to safety properties.

Unbounded sequence numbers Stenning [38] studied a SWP with unbounded sequence numbers and an infinite window size, in which messages can be lost, duplicated or reordered. A timeout mechanism is used to trigger retransmission. Stenning gave informal manual proofs of some safety properties. Knuth [25] examined more general principles behind Stenning's protocol, and manually verified some safety properties. Hailpern [19] used temporal logic to formulate safety and liveness properties for Stenning's protocol, and established their validity by informal reasoning. Jonsson [22] also verified safety and liveness properties of the protocol, using temporal logic and a manual compositional verification technique. Rusu [34] used the theorem prover PVS to verify safety and liveness properties for a SWP with unbounded sequence numbers.

Fixed finite window size Vaandrager [40], Groenvelde [12], van Wamel [44] and Bezem and Groote [4] manually verified in process algebra a SWP with window size one. Richier *et al.* [32] specified a SWP in a process algebra based language Estelle/R, and verified safety properties for window size up to eight using the model checker Xesar. Madelaine and Vergamini [28] specified a SWP in Lotos, with the help of the simulation environment Lite, and proved some safety properties for window size six. Holzmann [20, 21] used the Spin model checker to verify safety and liveness properties of a SWP with sequence numbers up to five. Kaivola [24] verified safety and liveness properties using model checking for a SWP with window size up to seven. Godefroid and Long [11] specified a full duplex SWP in a guarded command language, and verified the protocol for window size two using a model checker based on Queue BDDs. Stahl *et al.* [37] used a combination of abstraction, data independence, compositional reasoning and model checking to verify safety and liveness properties for a SWP with window size up to sixteen. The protocol was specified in Promela, the input language for the Spin model checker. Smith and Klarlund [36] specified a SWP in the high-level language IOA, and used the theorem prover MONA to verify a safety property for unbounded sequence numbers with window size up to 256. Jonsson and Nilsson [23] used an automated reachability analysis to verify safety properties for a SWP with a receiving window of size one. Latvala [26] modeled a SWP using Coloured Petri nets. A liveness property was model checked with fairness constraints for window size up to eleven.

Arbitrary finite window size Cardell-Oliver [6] specified a SWP using higher order logic, and manually proved and mechanically checked safety properties using HOL. (Van de Snepscheut [41] noted that what Cardell-Oliver claims to be a liveness property is in fact a safety property.) Schoone [35] manually proved safety properties for several SWPs using assertional verification. Van de Snepscheut [41] gave a correctness proof of a SWP as a sequence of correctness preserving transformations of a sequential program. Paliwoda and Sanders [30]

specified a reduced version of what they call a SWP (but which is in fact very similar to the bakery protocol from [13]) in the process algebra CSP, and verified a safety property modulo trace semantics. Röckl and Esparza [33] verified the correctness of this bakery protocol modulo weak bisimilarity using Isabelle/HOL, by explicitly checking a bisimulation relation. Chkhaev *et al.* [8] used a timed state machine in PVS to specify a SWP with a timeout mechanism and proved some safety properties with the mechanical support of PVS; correctness is based on the timeout mechanism, which allows messages in the mediums to be reordered.

1. μ CRL

μ CRL [15] (see also [17]) is a language for specifying distributed systems and protocols in an algebraic style. It is based on the process algebra ACP [3] extended with equational abstract data types [27]. We will use \approx for equality between process terms and $=$ for equality between data terms.

A μ CRL specification of data types consists of two parts: a signature of function symbols from which one can build data terms, and axioms that induce an equality relation on data terms of the same type. They provide a loose semantics, meaning that it is allowed to have multiple models. The data types needed for our μ CRL specification of a SWP are presented in Section 2. In particular we have the data sort of booleans *Bool* with constants `true` and `false`, and the usual connectives \wedge , \vee , \neg , \rightarrow and \leftrightarrow . For a boolean b , we abbreviate $b = \text{true}$ to b and $b = \text{false}$ to $\neg b$.

The process part of μ CRL is specified using a number of pre-defined process algebraic operators, which we will present below. From these operators one can build process terms, which describe the order in which the atomic actions from a set \mathcal{A} may happen. A process term consists of actions and recursion variables combined by the process algebraic operators. Actions and recursion variables may carry data parameters. There are two predefined actions outside \mathcal{A} : δ represents deadlock, and τ a hidden action. These two actions never carry data parameters.

Two elementary operators to construct processes are *sequential composition*, written $p \cdot q$, and *alternative composition*, written $p + q$. The process $p \cdot q$ first executes p , until p terminates, and then continues with executing q . The process $p + q$ non-deterministically behaves as either p or q . *Summation* $\sum_{d:D} p(d)$ provides the possibly infinite non-deterministic choice over a data type D . For example, $\sum_{n:\mathbb{N}} a(n)$ can perform the action $a(n)$ for all natural numbers n . The *conditional* construct $p \triangleleft b \triangleright q$, with b a data term of sort *Bool*, behaves as p if b and as q if $\neg b$. *Parallel composition* $p \parallel q$ performs the processes p and q in parallel; in other words, it consists of the arbitrary interleaving of actions of the processes p and q . For example, if there is no communication possible between actions a and b , then $a \parallel b$ behaves as $(a \cdot b) + (b \cdot a)$. Moreover, actions from p and q may also synchronize to a communication action, when this is explicitly allowed by a predefined *communication function*; two actions can only synchronize if their data parameters are equal. *Encapsulation* $\partial_{\mathcal{H}}(p)$, which renames all occurrences in p of actions from the set \mathcal{H} into δ , can be used to force actions into communication. For example, if actions a and b communicate to c , then $\partial_{\{a,b\}}(a \parallel b) \approx c$. *Hiding* $\tau_{\mathcal{I}}(p)$ renames all occurrences in p of actions from the set \mathcal{I} into τ . Finally, processes can be specified by means of recursive equations

$$X(d_1:D_1, \dots, d_n:D_n) \approx p$$

where X is a recursion variable, d_i a data parameter of type D_i for $i = 1, \dots, n$, and p a process term (possibly containing recursion variables and the parameters d_i). For example, let $X(n:\mathbb{N}) \approx a(n) \cdot X(n+1)$; then $X(0)$ can execute the infinite sequence of actions $a(0) \cdot a(1) \cdot a(2) \cdot \dots$.

Definition 1 (Linear process equation) A recursive specification is a linear process equation (LPE) if it is of the form

$$X(d:D) \approx \sum_{j \in J} \sum_{e_j : E_j} a_j(f_j(d, e_j)) \cdot X(g_j(d, e_j)) \triangleleft h_j(d, e_j) \triangleright \delta$$

with J a finite index set, $f_j : D \times E_j \rightarrow D_j$, $g_j : D \times E_j \rightarrow D$, and $h_j : D \times E_j \rightarrow \text{Bool}$.

Note that an LPE does not contain parallel composition, encapsulation and hiding, and uses only one recursion variable. Groote, Ponse and Usenko [16] presented a linearization algorithm that transforms μCRL specifications into LPEs.

To each μCRL specification belongs a directed graph, called a labeled transition system. In this labeled transition system, the states are process terms, and the edges are labeled with parameterized actions. For example, given the μCRL specification $X(n:\mathbb{N}) \approx a(n) \cdot X(n+1)$, we have transitions $X(n) \xrightarrow{a(n)} X(n+1)$. Branching bisimilarity \xleftrightarrow{b} [43] and strong bisimilarity $\xleftrightarrow{\quad}$ [31] are two well-established equivalence relations on states in labeled transition systems.¹ Conveniently, strong bisimilarity implies branching bisimilarity. The proof theory of μCRL from [14] is sound with respect to branching bisimilarity, meaning that if $p \approx q$ can be derived from it then $p \xleftrightarrow{b} q$.

Definition 2 (Branching bisimulation) Given a labeled transition system. A strong bisimulation relation \mathcal{B} is a symmetric binary relation on states such that if $s \mathcal{B} t$ and $s \xrightarrow{\ell} s'$, then there exists t' such that $t \xrightarrow{\ell} t'$ and $s' \mathcal{B} t'$. Two states s and t are strongly bisimilar, denoted by $s \xleftrightarrow{\quad} t$, if there is a strong bisimulation relation \mathcal{B} such that $s \mathcal{B} t$.

A strong and branching bisimulation relation \mathcal{B} is a symmetric binary relation on states such that if $s \mathcal{B} t$ and $s \xrightarrow{\ell} s'$, then

- either $\ell = \tau$ and $s' \mathcal{B} t$;
- or there is a sequence of (zero or more) τ -transitions $t \xrightarrow{\tau} \dots \xrightarrow{\tau} \hat{t}$ such that $s \mathcal{B} \hat{t}$ and $\hat{t} \xrightarrow{\ell} t'$ with $s' \mathcal{B} t'$.

Two states s and t are branching bisimilar, denoted by $s \xleftrightarrow{b} t$, if there is a branching bisimulation relation \mathcal{B} such that $s \mathcal{B} t$.

See [42] for a lucid exposition on why branching bisimilarity constitutes a sensible equivalence relation for concurrent processes.

The goal of this section is to prove that the initial state of the forthcoming μCRL specification of a two-way SWP is branching bisimilar to a pair of FIFO queues. In the proof of this fact, in Section 6, we will use three proof techniques to derive that two μCRL specifications are branching (or even strongly) bisimilar: invariants, bisimulation criteria, and cones and foci. An *invariant* $I : D \rightarrow \text{Bool}$ [5] characterizes the set of reachable states of an LPE $X(d:D)$. That is, if $I(d) = \text{true}$ and X can evolve from d to d' in zero or more transitions, then $I(d') = \text{true}$.

Definition 3 (Invariant) $I : D \rightarrow \text{Bool}$ is an invariant for an LPE in Definition 1 if for all $d:D$, $j \in J$ and $e_j:E_j$. $(I(d) \wedge h_j(d, e_j)) \rightarrow I(g_j(d, e_j))$.

If I holds in a state d and $X(d)$ can perform a transition, meaning that $h_j(d, e_j) = \text{true}$ for some $e_j:E$, then it is ensured by the definition above that I holds in the resulting state $g_j(d, e_j)$.

¹The definitions of these relations often take into account a special predicate on states to denote successful termination. This predicate is missing here, as successful termination does not play a role in our SWP specification.

Bisimulation criteria rephrase the question whether $X(d)$ and $Y(d')$ are strongly bisimilar in terms of data equalities, where $X(d:D)$ and $Y(d':D')$ are LPEs. A *state mapping* ϕ relates each state in $X(d)$ to a state in $Y(d')$. If a number of bisimulation criteria are satisfied, then ϕ establishes a strong bisimulation relation between terms $X(d)$ and $Y(\phi(d))$.

Definition 4 (Bisimulation criteria) *Given two LPEs,*

$$\begin{aligned} X(d:D) &\approx \sum_{j \in J} \sum_{e_j: E_j} a_j(f_j(d, e_j)) \cdot X(g_j(d, e_j)) \triangleleft h_j(d, e_j) \triangleright \delta \\ Y(d':D') &\approx \sum_{j \in J} \sum_{e'_j: E'_j} a'_j(f'_j(d', e'_j)) \cdot Y(g'_j(d', e'_j)) \triangleleft h'_j(d', e'_j) \triangleright \delta \end{aligned}$$

and an invariant $I : D \rightarrow \text{Bool}$ for X . A state mapping $\phi : D \rightarrow D'$ and local mappings $\psi_j : E_j \rightarrow E'_j$ for $j \in J$ satisfy the bisimulation criteria if for all states $d \in D$ in which invariant I holds:

- I $\forall j \in J \forall e_j: E_j (h_j(d, e_j) \leftrightarrow h'_j(\phi(d), \psi_j(e_j))),$
- II $\forall j \in J \forall e_j: E_j (h_j(d, e_j) \wedge I(d) \rightarrow (a_j(f_j(d, e_j)) = a'_j(f'_j(\phi(d), \psi_j(e_j))))),$
- III $\forall j \in J \forall e_j: E_j (h_j(d, e_j) \wedge I(d) \rightarrow (\phi(g_j(d, e_j)) = g'_j(\phi(d), \psi_j(e_j)))).$

Criterion I expresses that at each summand i , the corresponding guard of X holds if and only if the corresponding guard of Y holds with parameters $(\phi(d), \psi_j(e_j))$. Criterion II (III) states that at any summand i , the corresponding action (next state, after applying ϕ on it) of X could be equated to the corresponding action (next state) of Y with parameters $(\phi(d), \psi_j(e_j))$.

Theorem 5 (Bisimulation criteria) *Given two LPEs $X(d:D)$ and $Y(d':D')$ written as in Definition 4, and $I : D \rightarrow \text{Bool}$ an invariant for X . Let $\phi : D \rightarrow D'$ and $\psi_j : E_j \rightarrow E'_j$ for $j \in J$ satisfy the bisimulation criteria in Definition 4. Then $X(d) \leftrightarrow Y(\phi(d))$ for all $d \in D$ in which I holds.*

This theorem has been proved in PVS. The proof is available at <http://www.cs.utwente.nl/~vdpol/piggybacking.html>.

The *cones and foci* method from [10, 18] rephrases the question whether $\tau_{\mathcal{I}}(X(d))$ and $Y(d')$ are branching bisimilar in terms of data equalities, where $X(d:D)$ and $Y(d':D')$ are LPEs, and the latter LPE does not contain actions from some set \mathcal{I} of internal actions. A *state mapping* ϕ relates each state in $X(d)$ to a state in $Y(d')$. Furthermore, some $d:D$ are declared to be *focus points*. The *cone* of a focus point consists of the states in $X(d)$ that can reach this focus point by a string of actions from \mathcal{I} . It is required that each reachable state in $X(d)$ is in the cone of a focus point. If a number of *matching criteria* are satisfied, then ϕ establishes a branching bisimulation relation between terms $\tau_{\mathcal{I}}(X(d))$ and $Y(\phi(d))$.

Definition 6 (Matching criteria) *Given two LPEs:*

$$\begin{aligned} X(d:D) &\approx \sum_{j \in J} \sum_{e_j: E_j} a_j(f_j(d, e_j)) \cdot X(g_j(d, e_j)) \triangleleft h_j(d, e_j) \triangleright \delta \\ Y(d':D') &\approx \sum_{\{j \in J \mid a_j \notin \mathcal{I}\}} \sum_{e_j: E_j} a_j(f'_j(d', e_j)) \cdot Y(g'_j(d', e_j)) \triangleleft h'_j(d', e_j) \triangleright \delta \end{aligned}$$

Let $FC: D \rightarrow \text{Bool}$ be a predicate which designates the focus points, and $\mathcal{I} \subset \{a_j \mid j \in J\}$. A state mapping $\phi : D \rightarrow D'$ satisfies the matching criteria for $d:D$ if for all $j \in J$ with $a_j \notin \mathcal{I}$ and all $k \in J$ with $a_k \in \mathcal{I}$:

- I $\forall e_k: E_k (h_k(d, e_k) \rightarrow \phi(d) = \phi(g_k(d, e_k)));$
- II $\forall e_j: E_j (h_j(d, e_j) \rightarrow h'_j(\phi(d), e_j));$
- III $FC(d) \rightarrow \forall e_j: E_j (h'_j(\phi(d), e_j) \rightarrow h_j(d, e_j));$
- IV $\forall e_j: E_j (h_j(d, e_j) \rightarrow f_j(d, e_j) = f'_j(\phi(d), e_j));$
- V $\forall e_j: E_j (h_j(d, e_j) \rightarrow \phi(g_j(d, e_j)) = g'_j(\phi(d), e_j)).$

Matching criterion I requires that the internal transitions at d are inert, meaning that d and $g_k(d, e_k)$ are branching bisimilar. Criteria II, IV and V express that each external transition of d can be simulated by $\phi(d)$. Finally, criterion III expresses that if d is a focus point, then each external transition of $\phi(d)$ can be simulated by d .

Theorem 7 (Cones and foci) *Given LPEs $X(d:D)$ and $Y(d':D')$ written as in Definition 6. Let $I : D \rightarrow Bool$ be an invariant for X . Suppose that for all $d:D$ with $I(d)$:*

1. $\phi : D \rightarrow D'$ satisfies the matching criteria for d ; and
2. there is a $\hat{d}:D$ such that $FC(\hat{d})$ and X can perform transitions $d \xrightarrow{c_1} \dots \xrightarrow{c_k} \hat{d}$ with $c_1, \dots, c_k \in \mathcal{I}$.

Then for all $d:D$ with $I(d)$, $\tau_{\mathcal{I}}(X(d)) \leftrightarrow_b Y(\phi(d))$.

PVS proof of this is in [10]. For example, consider the LPEs $X(b:Bool) \approx a \cdot X(b) \triangleleft b \triangleright \delta + c \cdot X(\neg b) \triangleleft \neg b \triangleright \delta$ and $Y(d':D') \approx a \cdot Y(d')$, with $\mathcal{I} = \{c\}$ and focus point `true`. Moreover, $X(\text{false}) \xrightarrow{c} X(\text{true})$, i.e., `false` can reach the focus point in a single c -transition. For any $d':D'$, the state mapping $\phi(b) = d'$ for $b:Bool$ satisfies the matching criteria.

Given an invariant I , only $d:D$ with $I(d) = \text{true}$ need to be in the cone of a focus point, and we only need to satisfy the matching criteria for $d:D$ with $I(d) = \text{true}$.

2. Data Types

In this section, the data types used in the μCRL specification of the two-way SWP are presented: booleans, natural numbers supplied with modulo arithmetic, buffers, and lists. Furthermore, basic properties are given for the operations defined on these data types. The μCRL specification of the data types, and of the process part are presented in here.

Booleans. We introduce constant functions `true`, `false` of type $Bool$. \wedge and \vee both of type $Bool \times Bool \rightarrow Bool$ represent conjunction and disjunction operators, also \rightarrow and \leftrightarrow of the same exact type, denote implication and bi-implication, and $\neg : Bool \rightarrow Bool$ denotes negation. For any given sort D we consider a function $if : Bool \times D \times D \rightarrow D$ which functions an If-Then-Else operation, and also a mapping $eq : D \times D \rightarrow Bool$ such that $eq(d, e)$ holds if and only if $d = e$. For notational convenience we take the liberty to write $d = e$ instead of $eq(d, e)$.

Natural Numbers. $0 : \rightarrow \mathbb{N}$ denotes zero and $S : \mathbb{N} \rightarrow \mathbb{N}$ the successor function. The infix operations $+$, $-$ and \cdot of type $\mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ represent addition, monus (also called cut-off subtraction) and multiplication, respectively. The infix operations \leq , $<$, \geq and $>$ of type $\mathbb{N} \times \mathbb{N} \rightarrow Bool$ are the less-than(-or-equal) and greater-than(-or-equal) operations. $|$, div of type $\mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ are modulo (some natural number) and dividing functions respectively. The rewrite rules applied over this data type, are explained in detail in Section 2 in [2].

Since the buffers at the sender and the receiver in the SWP are of finite size, modulo calculations will play an important role. $i|_n$ denotes i modulo n , while $i \text{ div } n$ denotes i integer divided by n . In the proofs we will take notational liberties like omitting the sign for multiplication, and abbreviating $\neg(i = j)$ to $i \neq j$, $(k < \ell) \wedge (\ell < m)$ to $k < \ell < m$, $S(0)$ to 1, and $S(S(0))$ to 2. We will also use the standard induction rule to prove some properties.

Buffers. Each party in the two-way SWP will both maintain two buffers containing the sending and the receiving window (outside these windows both buffers will be empty).

$$\begin{aligned}
& [] : \rightarrow Buf; \quad inb, add : \Delta \times \mathbb{N} \times Buf \rightarrow Buf; \\
& |, || : Buf \times \mathbb{N} \rightarrow Buf; \\
& smaller, test : \mathbb{N} \times Buf \rightarrow Bool; \quad sorted : Buf \rightarrow Bool; \\
& retrieve : \mathbb{N} \times Buf \rightarrow \Delta; \quad remove : \mathbb{N} \times Buf \rightarrow Buf;
\end{aligned}$$

$release, release|_n : \mathbb{N} \times \mathbb{N} \times Buf \rightarrow Buf$;
 $next-empty, next-empty|_n : \mathbb{N} \times Buf \rightarrow \mathbb{N}$;
 $in-window : \mathbb{N} \times \mathbb{N} \times \mathbb{N} \rightarrow Bool$ and
 $max : Buf \rightarrow \mathbb{N}$

are the functions we use for buffers. And, the rewrite rules are:

$$\begin{aligned}
add(d, i, []) &= inb(d, i, []) \\
add(d, i, inb(e, j, q)) &= if(i > j, inb(e, j, add(d, i, q)), \\
&\quad inb(d, i, remove(i, inb(e, j, q)))) \\
[]|_n &= [] \text{ and } inb(d, i, q)|_n = inb(d, i|_n, q|_n) \\
[]|_n &= [] \text{ and } inb(d, i, q)|_n = add(d, i|_n, q|_n) \\
smaller(i, []) &= true \text{ and } smaller(i, inb(d, j, q)) = i < j \wedge smaller(i, q) \\
sorted([]) &= true \text{ and } sorted(inb(d, j, q)) = smaller(j, q) \wedge sorted(q) \\
test(i, []) &= false \text{ and } test(i, inb(d, j, q)) = i=j \vee test(i, q) \\
retrieve(i, inb(d, j, q)) &= if(i=j, d, retrieve(i, q)) \\
remove(i, []) &= [] \\
remove(i, inb(d, j, q)) &= if(i=j, remove(i, q), inb(d, j, remove(i, q))) \\
release(i, j, q) &= if(i \geq j, q, release(S(i), j, remove(i, q))) \\
release|_n(i, j, q) &= if(i|_n=j|_n, q, release|_n(S(i), j, remove(i|_n, q))) \\
next-empty(i, q) &= if(test(i, q), next-empty(S(i), q), i) \\
next-empty|_n(i, q) &= if(next-empty(i|_n, q) < n, next-empty(i|_n, q), \\
&\quad if(next-empty(0, q) < n, next-empty(0, q), n)) \\
in-window(i, j, k) &= i \leq j < k \vee k < i \leq j \vee j < k < i \\
max([]) &= 0 \text{ and } max(inb(d, i, q)) = if(i \geq max(q), i, max(q))
\end{aligned}$$

More explanation on this is in [2] Section 2.

Δ represents the set of data elements that can be communicated between the two parties. The buffers are modeled as a list of pairs (d, i) with $d: \Delta$ and $i: \mathbb{N}$, representing that cell (or sequence number) i of the buffer is occupied by datum d ; cells for which no datum is specified are empty. The empty buffer is denoted by $[],$ and $inb(d, i, q)$ is the buffer that is obtained from q by simply putting (d, i) on top of the buffer $q.$

add inserts data into the queue, while keeping it sorted (if the queue itself is so) and avoiding duplications. $q|_n$ is taking the sequence numbers in q of modulo $n,$ and $q|_n$ the resulting buffer is further sorted out. $sorted$ announces whether or not a buffer is sorted. $smaller$ makes sure that the first data in the queue is having the smallest index number.

$test(i, q)$ is true if and only if the i th location in q is occupied. $retrieve(i, q)$ reveals q 's i th element ² $remove(i, q)$ wipes the i th element out. $release(i, j, q)$ empties i th to j th locations, where $release|_n(i, j, q)$ does the analogous modulo $n.$ $next-empty(i, q)$ reveals the first empty cell in q as of $i,$ where $next-empty|_n(i, q)$ operates the same modulo $n.$ $in-window(i, j, k)$ is true if and only if $i \leq j \leq k - 1,$ modulo $n.$ Finally, $max(q)$ reports the greatest occupied place in $q.$

Lists. $List$ is used for the specification of the external behavior of the protocol. $\langle \rangle : \rightarrow List, inl : \Delta \times List \rightarrow List, length : List \rightarrow \mathbb{N}, top : List \rightarrow \Delta, tail : List \rightarrow List, append : \Delta \times List \rightarrow List, ++ : List \times List \rightarrow List$ and $\lambda, \lambda' : List$ represent the functions, where $\langle \rangle$ denotes the empty list, and $inl(d, \lambda)$ adds datum d at the top of list $\lambda.$ A special datum d_0 is specified to serve as a dummy value for data parameters. $length(\lambda)$ denotes the length of $\lambda,$ $top(\lambda)$ produces the datum that resides at the top of $\lambda,$ $tail(\lambda)$ is obtained by removing

²Note that $retrieve(i, [])$ is undefined. One could choose to equate it to a default value in $\Delta,$ or to a fresh error element in $\Delta.$ However, with the first approach an occurrence of $retrieve(i, [])$ might remain undetected, and the second approach would needlessly complicate the data type $\Delta.$ We prefer to work with an under-specified version of $retrieve,$ which is allowed in $\mu CRL,$ since data types have a loose semantics. All operations in μCRL data models, however, are total; partial operations lead to the existence of multiple models.

the top position in λ , $append(d, \lambda)$ adds datum d at the end of λ , and $\lambda++\lambda'$ represents list concatenation. Finally, $q[i..j]$ is the list containing the elements in buffer q at positions i up to but not including j . The rewrite rules which are being used are:

$$\begin{aligned}
 length(\langle \rangle) &= 0 \text{ and } length(inl(d, \lambda)) = S(length(\lambda)) \\
 top(inl(d, \lambda)) &= d, \quad tail(inl(d, \lambda)) = \lambda \\
 append(d, \langle \rangle) &= inl(d, \langle \rangle) \\
 append(d, inl(e, \lambda)) &= inl(e, append(d, \lambda)) \\
 \langle \rangle ++ \lambda &= \lambda \text{ and } inl(d, \lambda) ++ \lambda' = inl(d, \lambda ++ \lambda') \\
 q[i..j] &= if(i \geq j, \langle \rangle, inl(retrieve(i, q), q[S(i)..j]))
 \end{aligned}$$

Detailed description on this data type is written in [2] Section 2.

3. Two-Way SWP with Piggybacking

This section contains the specification of the protocol in μ CRL. Figure 1 illustrates the the protocol we work on (i.e. a two-way SWP with piggybacking). In this protocol sender (S/R) stores data elements that it receives via channel A in a buffer of size $2n$, in the order in which they are received. It can send a datum, together with its sequence number in the buffer, to a receiver R/S via a medium that behaves as a lossy queue of capacity one, represented by the medium K and the channels B and C. Upon receipt, the receiver may store the datum in its buffer, where its position in the buffer is dictated by the attached sequence number. In order to avoid a possible overlap between the sequence numbers of different data elements in the buffers of sender and receiver, no more than one half of each of these two buffers may be occupied at any time; these halves are called the sending and the receiving window, respectively. The receiver can pass on a datum that is located at the first cell in its window via channel D; in that case the receiving window slides forward by one cell. Furthermore, the receiver can send the sequence number of the first empty cell in (or just outside) its window as an acknowledgment to the sender via a medium that behaves as a lossy queue of capacity one, represented by the medium L and the channels E and F. If the sender receives this acknowledgment, its window slides forward accordingly. In a two-way SWP, data streams




Figure 1. A two sided Sliding window protocol

are in both directions, meaning that S/R and R/S both act as sender and receiver at the same time. In addition to this, in our protocol when a datum arrives, the receiver may either send an acknowledgment back to the channel or it might instead wait until the network layer passes on the next datum. In latter case, once this new datum is to be sent into the channel, the awaited acknowledgment can be attached to it, and hence get a free ride. This technique is known as *piggybacking*.

3.1. Specification

The sender/receiver \mathbf{S}/\mathbf{R} is modeled by the process $\mathbf{S}/\mathbf{R}(\ell, m, q, q'_2, \ell'_2)$, where q is its sending buffer of size $2n$, ℓ is the first cell in the window of q , and m the first empty cell in (or just outside) this window. Furthermore, q'_2 is the receiving buffer of size $2n_2$, and ℓ'_2 is the first cell in the window of q_2 .

The μCRL specification of \mathbf{S}/\mathbf{R} consists of seven clauses. The first clause of the specification expresses that \mathbf{S}/\mathbf{R} can receive a datum via channel A and place it in its sending window, under the condition that this window is not yet full. The next two clauses specify that \mathbf{S}/\mathbf{R} can receive a datum/acknowledgment pair via channel F; the data part is either added to q_2 if it is within the receiving window (second clause), or ignored if it is outside this window (third clause). In both clauses, q is emptied from ℓ up to but not including the received acknowledgment. The fourth clause specifies the reception of a single (i.e., non-piggybacked) acknowledgment. According to the fifth clause, data elements for transmission via channel B are taken (at random) from the filled part of the sending window; the first empty position in (or just outside) the receiving window is attached to this datum as an acknowledgment. In the sixth clause, \mathbf{S}/\mathbf{R} sends a single acknowledgment. Finally, clause seven expresses that if the first cell in the receiving window is occupied, then \mathbf{S}/\mathbf{R} can send this datum into channel A, after which the cell is emptied.

$$\begin{aligned}
& \mathbf{S}/\mathbf{R}(\ell:\mathbb{N}, m:\mathbb{N}, n:\mathbb{N}, n_2:\mathbb{N}, q:\text{Buf}, q'_2:\text{Buf}, \ell'_2:\mathbb{N}) \\
& \approx \sum_{d:\Delta} r_A(d) \cdot \mathbf{S}/\mathbf{R}(\ell, S(m)|_{2n}, \text{add}(d, m, q), q'_2, \ell'_2) \triangleleft \text{in-window}(\ell, m, (\ell + n)|_{2n}) \triangleright \delta \\
& + \sum_{d:\Delta} \sum_{i:\mathbb{N}} \sum_{k:\mathbb{N}} r_F(d, i, k) \cdot \mathbf{S}/\mathbf{R}(k, m, \text{release}|_{2n}(\ell, k, q), \text{add}(d, i, q'_2), \ell'_2) \\
& \quad \triangleleft \text{in-window}(\ell'_2, i, (\ell'_2 + n_2)|_{2n_2}) \triangleright \delta \\
& + \sum_{d:\Delta} \sum_{i:\mathbb{N}} \sum_{k:\mathbb{N}} r_F(d, i, k) \cdot \mathbf{S}/\mathbf{R}(k, m, \text{release}|_{2n}(\ell, k, q), q'_2, \ell'_2) \\
& \quad \triangleleft \neg \text{in-window}(\ell'_2, i, (\ell'_2 + n_2)|_{2n_2}) \triangleright \delta \\
& + \sum_{k:\mathbb{N}} r_F(k) \cdot \mathbf{S}/\mathbf{R}(k, m, \text{release}|_{2n}(\ell, k, q), q'_2, \ell'_2) \\
& + \sum_{k:\mathbb{N}} s_B(\text{retrieve}(k, q), k, \text{next-empty}|_{2n_2}(\ell'_2, q'_2)) \cdot \mathbf{S}/\mathbf{R}(\ell, m, q, q'_2, \ell'_2) \triangleleft \text{test}(k, q) \triangleright \delta \\
& + s_B(\text{next-empty}|_{2n_2}(\ell'_2, q'_2)) \cdot \mathbf{S}/\mathbf{R}(\ell, m, q, q'_2, \ell'_2) \\
& + s_A(\text{retrieve}(\ell'_2, q'_2)) \cdot \mathbf{S}/\mathbf{R}(\ell, m, q, \text{remove}(\ell'_2, q'_2), S(\ell'_2)|_{2n_2}) \triangleleft \text{test}(\ell'_2, q'_2) \triangleright \delta
\end{aligned}$$

The μCRL specification of \mathbf{R}/\mathbf{S} (in [2] Appendix A) is symmetrical to the one of \mathbf{S}/\mathbf{R} . In the process $\mathbf{R}/\mathbf{S}(\ell_2, m_2, q_2, q', \ell')$, q' is the receiving buffer of size $2n$, and ℓ' is the first position in the window of q . Furthermore, q_2 is the sending buffer of size $2n_2$, ℓ_2 is the first position in the window of q_2 , and m_2 the first empty position in (or just outside) this window.

Mediums \mathbf{K} and \mathbf{L} , introduced below, are of capacity one. These mediums are specified in a way that they may lose frames or acknowledgments:

$$\begin{aligned}
\mathbf{K} & \approx \sum_{d:\Delta} \sum_{k:\mathbb{N}} \sum_{i:\mathbb{N}} r_B(d, k, i) \cdot (j \cdot s_C(d, k, i) + j) \cdot \mathbf{K} + \sum_{i:\mathbb{N}} r_B(i) \cdot (j \cdot s_C(i) + j) \cdot \mathbf{K} \\
\mathbf{L} & \approx \sum_{d:\Delta} \sum_{k:\mathbb{N}} \sum_{i:\mathbb{N}} r_E(d, k, i) \cdot (j \cdot s_F(d, k, i) + j) \cdot \mathbf{L} + \sum_{i:\mathbb{N}} r_E(i) \cdot (j \cdot s_F(i) + j) \cdot \mathbf{L}.
\end{aligned}$$

For each channel $i \in \{\mathbf{B}, \mathbf{C}, \mathbf{E}, \mathbf{F}\}$, actions s_i and r_i can communicate, resulting in the action c_i . The initial state of the SWP is expressed by $\tau_{\mathcal{I}}(\partial_{\mathcal{H}}(\mathbf{S}/\mathbf{R}(0, 0, [], [], 0) \parallel \mathbf{R}/\mathbf{S}(0, 0, [], [], 0) \parallel \mathbf{K} \parallel \mathbf{L}))$ where the set \mathcal{H} consists of the read and send actions over the internal channels \mathbf{B} , \mathbf{C} , \mathbf{E} , and \mathbf{F} , namely $\mathcal{H} = \{s_B, r_B, s_C, r_C, s_E, r_E, s_F, r_F\}$ while the set \mathcal{I} consists of the communication actions over these internal channels together with j , namely $\mathcal{I} = \{c_B, c_C, c_E, c_F, j\}$.

3.2. External Behavior

Data elements that are read from channel A should be sent into channel D in the same order, and vice versa data elements that are read from channel D should be sent into channel A in the same order. No data elements should be lost. In other words, the SWP is intended to be a solution for the following linear μCRL specification, representing a pair of FIFO queues of capacity $2n$ and $2n_2$.

$$\begin{aligned}
\mathbf{Z}(\lambda_1:List, \lambda_2:List) &\approx \sum_{d:\Delta} r_A(d) \cdot \mathbf{Z}(\text{append}(d, \lambda_1), \lambda_2) \triangleleft \text{length}(\lambda_1) < 2n \triangleright \delta \\
&+ s_D(\text{top}(\lambda_1)) \cdot \mathbf{Z}(\text{tail}(\lambda_1), \lambda_2) \triangleleft \text{length}(\lambda_1) > 0 \triangleright \delta \\
&+ \sum_{d:\Delta} r_D(d) \cdot \mathbf{Z}(\lambda_1, \text{append}(d, \lambda_2)) \triangleleft \text{length}(\lambda_2) < 2n_2 \triangleright \delta \\
&+ s_A(\text{top}(\lambda_2)) \cdot \mathbf{Z}(\lambda_1, \text{tail}(\lambda_2)) \triangleleft \text{length}(\lambda_2) > 0 \triangleright \delta
\end{aligned}$$

Note that $r_A(d)$ can be performed until the list λ_1 contains $2n$ elements, because in that situation the sending window of \mathbf{S}/\mathbf{R} and the receiving window of \mathbf{R}/\mathbf{S} will be filled. Furthermore, $s_D(\text{top}(\lambda_1))$ can only be performed if λ_1 is not empty. Likewise, $r_D(d)$ can be performed until the list λ_2 contains $2n_2$ elements, and $s_A(\text{top}(\lambda_2))$ can only be performed if λ_2 is not empty.

4. Modifying the Specification

This section witnesses three transformations, one to eliminate parallel operators, one to eliminate arguments of communication actions, and one to eliminate modulo arithmetic.

Linearization. The starting point of our correctness proof is a linear specification \mathbf{M}_{mod} , in which no parallel composition, encapsulation and hiding operators occur. \mathbf{M}_{mod} can be obtained from the μCRL specification of the SWP without the hiding operator, i.e., $\partial_{\mathcal{H}}(\mathbf{S}/\mathbf{R}(0, 0, [], [], 0) \parallel \mathbf{R}/\mathbf{S}(0, 0, [], [], 0) \parallel \mathbf{K} \parallel \mathbf{L})$ by means of the linearization algorithm presented in [16]; and according to [16], the following result can be obtained:

$$\text{Proposition 8 } \partial_{\mathcal{H}}(\mathbf{S}/\mathbf{R}(0, 0, [], [], 0) \parallel \mathbf{R}/\mathbf{S}(0, 0, [], [], 0) \parallel \mathbf{K} \parallel \mathbf{L}) \Leftrightarrow \mathbf{M}_{mod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0).$$

\mathbf{M}_{mod} contains eight extra parameters: $e, e_2:D$ and $g, g', h, h', h_2, h'_2:\mathbb{N}$. Intuitively, g is 5 when medium \mathbf{K} is inactive, is 4 or 2 when \mathbf{K} just received a data frame or a single acknowledgment, respectively, and is 3 or 1 when \mathbf{K} has decided to pass on this data frame or acknowledgment, respectively. The parameters e, h and h'_2 represent the memory of \mathbf{K} , meaning that they can store the datum that is being sent from \mathbf{S}/\mathbf{R} to \mathbf{R}/\mathbf{S} , the position of this datum in q , and the first empty position in the window of q , respectively. Initially, or when medium \mathbf{K} is inactive, g, e, h and h'_2 have the values 5, $d_0, 0$ and 0. Likewise, g' captures the five states of medium \mathbf{L} , and e_2, h_2 and h' represent the memory of \mathbf{L} .

The linear specification \mathbf{M}_{mod} of the SWP, with encapsulation but without hiding, is written below. For the sake of presentation, in states that results after a transition we only present parameters whose values have changed. In this specification

- The first summand describes that a datum d can be received by \mathbf{S}/\mathbf{R} through channel \mathbf{A} , if q 's window is not full ($\text{in-window}(\ell, m, (\ell + n)|_{2n})$). This datum is then placed in the first empty cell of q 's window ($q := \text{add}(d, m, q)$), and the next cell becomes the first empty cell of this window ($m := S(m)|_{2n}$).
- By the 2nd summand, a frame ($\text{retrieve}(k, q), k, \text{next-empty}|_{2n_2}(\ell'_2, q'_2)$) can be communicated to \mathbf{K} , if cell k in q 's window is occupied ($\text{test}(k, q)$). And by the 19th summand, an acknowledgment $\text{next-empty}|_{2n_2}(\ell'_2, q'_2)$ can be communicated to \mathbf{K} .
- The fifth and third summand describe that medium \mathbf{K} decides to pass on a frame or acknowledgment, respectively. The fourth summand describes that \mathbf{K} decides to lose this frame or acknowledgment.
- The sixth and seventh summand describe that the frame in medium \mathbf{K} is communicated to \mathbf{R}/\mathbf{S} . In the sixth summand the frame is within the window of q' ($\text{in-window}(\ell', h, (\ell' + n)|_{2n})$), so it is included ($q' := \text{add}(e, h, q')$). In the seventh summand the frame is outside the window of q' , so it is omitted. In both cases, the first cell of the window of q' is moved forward to h'_2 ($\ell_2 := h'_2$), and the cells before h'_2 are emptied ($q_2 := \text{release}|_{2n_2}(\ell_2, h'_2, q_2)$).

- The twentieth and last summand describes that the acknowledgment in medium \mathbf{K} is communicated to \mathbf{R}/\mathbf{S} . Then the first cell of the window of q is moved forward to h'_2 , and the cells before h'_2 are emptied.
- By the eighth summand, \mathbf{R}/\mathbf{S} can send the datum at the first cell in the window of q' (*retrieve*(ℓ' , q')) through channel \mathbf{D} , if this cell is occupied (*test*(ℓ' , q')). This cell is then emptied ($q' := \text{remove}(\ell', q')$), and the first cell of the window of q is moved forward by one ($\ell' := S(\ell')|_{2n}$).
- Other summands are symmetric counterparts to the ones described above.

$$\begin{aligned} & \mathbf{M}_{mod}(\ell, m, q, q'_2, \ell'_2, g, h, e, h'_2, g', h_2, e_2, h', \ell_2, m_2, q_2, q', \ell') \\ & \approx \sum_{d:\Delta} r_A(d) \cdot \mathbf{M}_{mod}(m := S(m)|_{2n}, q := \text{add}(d, m, q)) \triangleleft \text{in-window}(\ell, m, (\ell + n)|_{2n}) \triangleright \delta \end{aligned} \quad (A1)$$

$$\begin{aligned} & + \sum_{k:\mathbb{N}} c_B(\text{retrieve}(k, q), k, \text{next-empty}|_{2n_2}(\ell'_2, q'_2)) \cdot \mathbf{M}_{mod}(g := 4, e := \text{retrieve}(k, q), h := k, \\ & \quad h'_2 := \text{next-empty}|_{2n_2}(\ell'_2, q'_2)) \triangleleft \text{test}(k, q) \wedge g = 5 \triangleright \delta \end{aligned} \quad (B1)$$

$$+ j \cdot \mathbf{M}_{mod}(g := 1, e := d_0, h := 0) \triangleleft g = 2 \triangleright \delta \quad (C1)$$

$$+ j \cdot \mathbf{M}_{mod}(g := 5, e := d_0, h := 0, h_2 := 0) \triangleleft g = 2 \vee g = 4 \triangleright \delta \quad (D1)$$

$$+ j \cdot \mathbf{M}_{mod}(g := 3) \triangleleft g = 4 \triangleright \delta \quad (E1)$$

$$\begin{aligned} & + c_C(e, h, h'_2) \cdot \mathbf{M}_{mod}(\ell_2 := h'_2, q' := \text{add}(e, h, q'), g := 5, e := d_0, h := 0, h'_2 := 0, \\ & \quad q_2 := \text{release}|_{2n_2}(\ell_2, h'_2, q_2)) \triangleleft \text{in-window}(\ell', h, (\ell' + n)|_{2n}) \wedge g = 3 \triangleright \delta \end{aligned} \quad (F1)$$

$$\begin{aligned} & + c_C(e, h, h'_2) \cdot \mathbf{M}_{mod}(\ell_2 := h'_2, g := 5, e := d_0, h := 0, h'_2 := 0, q_2 := \text{release}|_{2n_2}(\ell_2, h'_2, q_2)) \\ & \quad \triangleleft \neg \text{in-window}(\ell', h, (\ell' + n)|_{2n}) \wedge g = 3 \triangleright \delta \end{aligned} \quad (G1)$$

$$+ s_D(\text{retrieve}(\ell', q')) \cdot \mathbf{M}_{mod}(\ell' := S(\ell')|_{2n}, q' := \text{remove}(\ell', q')) \triangleleft \text{test}(\ell', q') \triangleright \delta \quad (H1)$$

$$+ c_E(\text{next-empty}|_{2n}(\ell', q')) \cdot \mathbf{M}_{mod}(g' := 2, h_2 := 0, h' := \text{next-empty}|_{2n}(\ell', q')) \triangleleft g' = 5 \triangleright \delta \quad (I1)$$

$$+ j \cdot \mathbf{M}_{mod}(g' := 1, e_2 := d_0, h_2 := 0) \triangleleft g' = 2 \triangleright \delta \quad (J1)$$

$$+ j \cdot \mathbf{M}_{mod}(g' := 5, h_2 := 0, e_2 := d_0, h' := 0) \triangleleft g' = 2 \vee g' = 4 \triangleright \delta \quad (K1)$$

$$+ j \cdot \mathbf{M}_{mod}(g' := 3) \triangleleft g' = 4 \triangleright \delta \quad (L1)$$

$$+ c_F(h') \cdot \mathbf{M}_{mod}(\ell := h', q := \text{release}|_{2n}(\ell, h', q), g' := 5, h_2 := 0, e_2 := d_0, h' := 0) \triangleleft g' = 1 \triangleright \delta \quad (M1)$$

$$\begin{aligned} & + \sum_{d:\Delta} r_D(d) \cdot \mathbf{M}_{mod}(m_2 := S(m_2)|_{2n}, q_2 := \text{add}(d, m_2, q_2)) \\ & \quad \triangleleft \text{in-window}(\ell_2, m_2, (\ell_2 + n_2)|_{2n_2}) \triangleright \delta \end{aligned} \quad (N1)$$

$$\begin{aligned} & + \sum_{k:\mathbb{N}} c_E(\text{retrieve}(k, q_2), k, \text{next-empty}|_{2n}(\ell', q')) \cdot \mathbf{M}_{mod}(g' := 4, e_2 := \text{retrieve}(k, q_2), \\ & \quad h_2 := k, h' := \text{next-empty}|_{2n}(\ell', q')) \triangleleft \text{test}(k, q_2) \wedge g' = 5 \triangleright \delta \end{aligned} \quad (O1)$$

$$\begin{aligned} & + c_F(e_2, h_2, h') \cdot \mathbf{M}_{mod}(\ell := h', q'_2 := \text{add}(e_2, h_2, q'_2), g' := 5, e_2 := d_0, h_2 := 0, h' := 0, \\ & \quad q := \text{release}|_{2n}(\ell, h', q)) \triangleleft \text{in-window}(\ell'_2, h_2, (\ell'_2 + n_2)|_{2n_2}) \wedge g' = 3 \triangleright \delta \end{aligned} \quad (P1)$$

$$\begin{aligned} & + c_F(e_2, h_2, h') \cdot \mathbf{M}_{mod}(\ell := h', g' := 5, e_2 := d_0, h_2 := 0, h' := 0, q := \text{release}|_{2n}(\ell, h', q)) \\ & \quad \triangleleft \neg \text{in-window}(\ell'_2, h_2, (\ell'_2 + n_2)|_{2n_2}) \wedge g' = 3 \triangleright \delta \end{aligned} \quad (Q1)$$

$$+ s_A(\text{retrieve}(\ell'_2, q'_2)) \cdot \mathbf{M}_{mod}(\ell'_2 := S(\ell'_2)|_{2n_2}, q'_2 := \text{remove}(\ell'_2, q'_2)) \triangleleft \text{test}(\ell'_2, q'_2) \triangleright \delta \quad (R1)$$

$$+ c_B(\text{next-empty}|_{2n_2}(\ell'_2, q'_2)) \cdot \mathbf{M}_{mod}(g := 2, h := 0, h'_2 := \text{next-empty}|_{2n_2}(\ell'_2, q'_2)) \triangleleft g = 5 \triangleright \delta \quad (S1)$$

$$+ c_C(h'_2) \cdot \mathbf{M}_{mod}(\ell_2 := h'_2, q_2 := \text{release}|_{2n_2}(\ell_2, h'_2, q_2), g := 5, h := 0, e := d_0, h'_2 := 0) \triangleleft g = 1 \triangleright \delta \quad (T1)$$

\mathbf{N}_{mod} : No Communication Action's Arguments.

The linear specification \mathbf{N}_{mod} (Written in [2] Appendix A) is obtained from \mathbf{M}_{mod} by renaming all arguments from communication actions (e.g. $c_F(e_2, h_2, h')$) to a fresh action c . Since we want to show that the “external” behavior of this protocol is branching bisimilar to a pair of FIFO queues (of capacity $2n$ and $2n_2$), the internal actions can be removed. The following proposition is then a trivial result of this renaming:

Proposition 9 $\tau_{\mathcal{I}}(\mathbf{M}_{mod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)) \Leftrightarrow \tau_{\{c,j\}}(\mathbf{N}_{mod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)).$

\mathbf{N}_{nonmod} : **No Modulo Arithmetic.**

The specification of \mathbf{N}_{nonmod} is obtained by eliminating all occurrences of $|_{2n}$ (resp. $|_{2n_2}$) from \mathbf{N}_{mod} , and replacing all guards of the form $in\text{-}window(i, j, (i+k)|_{2n})$ (respectively $in\text{-}window(i, j, (i+k)|_{2n_2})$) with $i \leq j < i+n$ (respectively $i \leq j < i+n_2$). According to what just mentioned, only $A1, F1, G1, N1, P1$ and $Q1$ whose guards are of this form, will be subjected to change. We name each new clause after its corresponding one by removing the index 1 from it, that is e.g. $A1$ will become A , and so forth. As an example we show this clause below, the whole specification of \mathbf{N}_{nonmod} is in [2] Appendix A.

$$\sum_{d:\Delta} r_A(d) \cdot \mathbf{N}_{nonmod}(m := S(m), q := add(d, m, q)) \triangleleft l < m < \ell + n \triangleright \delta \quad (A)$$

In Section 6.1, we will prove that \mathbf{N}_{nonmod} and \mathbf{N}_{mod} are strongly bisimilar. In order to demonstrate the correctness of \mathbf{N}_{nonmod} (see Section 6.2) there will be a number of properties on the Data Types which should be investigated first. In the next section we list these properties, and thereafter, in its following section, we will prove the correctness.

5. Properties of Data Types

This section presents some properties of the data types and the *ordered* buffers, also some invariants of the final specification of the system; all proofs are in [2] Appendix B.

5.1. Basic Properties

These properties contain some mathematical reasoning over the functions in our specification of the system, with/without modulo arithmetic. One of them for example is: $test(k, q) \rightarrow add(retrieve(k, q), k, q)[i..j] = q[i..j]$. The entire list is in [2] Appendix B.1.

5.2. Ordered Buffers

Lemma 10 *Some properties on $add(., .)$ function:*

1. $test(i, q) \rightarrow test(i, add(d, j, q))$
2. $next\text{-}empty(i, add(d, j, q)) \geq next\text{-}empty(i, q)$
3. $test(i, add(d, j, q)) = (i=j \vee test(i, q))$
4. $retrieve(i, add(d, j, q)) = if(i=j, d, retrieve(i, q))$
5. $remove(i, add(d, i, q)) = remove(i, q)$
6. $j \neq next\text{-}empty(i, q) \rightarrow next\text{-}empty(i, add(d, j, q)) = next\text{-}empty(i, q)$
7. $next\text{-}empty(i, add(d, next\text{-}empty(i, q), q)) = next\text{-}empty(S(next\text{-}empty(i, q)), q)$
8. $i < j \rightarrow remove(i, add(d, j, q)) = add(d, j, remove(i, q))$
9. $i \neq j \rightarrow add(e, i, add(d, j, q)) = add(d, j, add(e, i, q))$

Lemma 11 *Ordered buffers maintain the following properties:*

1. $smaller(i, q) \rightarrow smaller(i, remove(j, q))$
2. $i < j \wedge smaller(i, q) \rightarrow smaller(i, add(d, j, q))$
3. $smaller(i, q) \rightarrow remove(i, q) = q$
4. $i < j \wedge smaller(j, q) \rightarrow smaller(i, q)$
5. $sorted(q) \rightarrow sorted(add(d, i, q))$
6. $smaller(i, q) \rightarrow add(d, i, q) = inb(d, i, q)$
7. $sorted(q) \wedge j < i \rightarrow remove(i, add(d, j, q)) = add(d, j, remove(i, q))$
8. $sorted(q) \rightarrow add(d, i, q) = add(d, i, remove(i, q))$

Lemma 12 For $n > 0$, the following results hold on $q|_n$.

1. $sorted(q|_n)$
2. $test(i, q|_n) = test(i, q|_n)$
3. $retrieve(i|_n, q|_n) = retrieve(i|_n, q|_n)$
4. $j \neq i \rightarrow remove(i, add(d, j, q|_n)) = add(d, j, remove(i, q|_n))$
5. $\forall j: \mathbb{N}(test(j, q) \rightarrow i \leq j < i + n) \wedge i \leq k \leq i + n \rightarrow$
 $next-empty|_{2n}(k|_{2n}, q|_{2n}) = next-empty|_{2n}(k|_{2n}, q|_{2n})$
6. $\forall j: \mathbb{N}(test(j, q) \rightarrow i \leq j < i + n) \wedge i \leq k \leq i + n \rightarrow remove(k, q)|_{2n} = remove(k|_{2n}, q|_{2n})$
7. $\forall j: \mathbb{N}(test(j, q) \rightarrow i \leq j < i + n) \wedge i \leq k \leq i + n \rightarrow release(i, k, q)|_{2n} =$
 $release|_{2n}(i|_{2n}, k|_{2n}, q|_{2n})$
8. $\forall j: \mathbb{N}(test(j, q) \rightarrow i \leq j < i + n) \wedge i \leq k \leq i + n \rightarrow add(d, k, q)|_{2n} = add(d, k|_{2n}, q|_{2n})$

All the abovementioned lemmas are proved in detail in [2] Appendix B.2.

5.3. Invariants

Invariants of a system are properties of data that are satisfied throughout the reachable state space of the system (see Definition 3). Lemma 13 collects 19 invariants of \mathbb{N}_{nonmod} (and their symmetric counterparts). Occurrences of variables $i, j: \mathbb{N}$ in an invariant are always implicitly universally quantified at the outside of the invariant.

Invariants 6, 8, 15 and 17 are only needed in the derivation of other invariants. We provide some intuition for the (first of each pair of) invariants that will be used in the correctness proofs in Section 6 and in the derivations of the data lemmas. Invariants 4, 11, 12, 13 express that the sending window of **S/R** is filled from ℓ up to but not including m , and that it has size n . Invariants 7, 10 express that the receiving window of **R/S** starts at ℓ' and stops at $\ell' + n$. Invariant 2 expresses that **S/R** cannot receive acknowledgments beyond $next-empty(\ell', q')$, and Invariant 9 that **R/S** cannot receive frames beyond $m - 1$. Invariants 16, 18, 19 are based on the fact that the sending window of **S/R**, the receiving window of **R/S**, and **K** (when active) coincide on occupied cells and frames with the same sequence number. Invariants 1, 3, 5 and 14 give bounds on the parameters h and h' of mediums **K** and **L**.

Lemma 13 $\mathbb{N}_{nonmod}(\ell, m, q, q'_2, \ell'_2, g, h, e, h'_2, g', h_2, e_2, h', \ell_2, m_2, q_2, q', \ell')$ satisfies the following invariants.

1. $h' \leq next-empty(\ell', q')$ and $h'_2 \leq next-empty(\ell'_2, q'_2)$
2. $\ell \leq next-empty(\ell', q')$ and $\ell_2 \leq next-empty(\ell'_2, q'_2)$
3. $g' \neq 5 \rightarrow \ell \leq h'$ and $g \neq 5 \rightarrow \ell_2 \leq h'_2$
4. $test(i, q) \rightarrow i < m$ and $test(i, q_2) \rightarrow i < m_2$
5. $(g = 3 \vee g = 4) \rightarrow h < m$ and $(g' = 3 \vee g' = 4) \rightarrow h_2 < m_2$
6. $test(i, q') \rightarrow i < m$ and $test(i, q'_2) \rightarrow i < m_2$
7. $test(i, q') \rightarrow \ell' \leq i < \ell' + n$ and $test(i, q'_2) \rightarrow \ell'_2 \leq i < \ell'_2 + n_2$
8. $\ell' \leq m$ and $\ell'_2 \leq m_2$
9. $next-empty(\ell', q') \leq m$ and $next-empty(\ell'_2, q'_2) \leq m_2$
10. $next-empty(\ell', q') \leq \ell' + n$ and $next-empty(\ell'_2, q'_2) \leq \ell'_2 + n_2$
11. $test(i, q) \rightarrow \ell \leq i$ and $test(i, q_2) \rightarrow \ell_2 \leq i$
12. $\ell \leq i < m \rightarrow test(i, q)$ and $\ell_2 \leq i < m_2 \rightarrow test(i, q_2)$
13. $m \leq \ell + n$ and $m_2 \leq \ell_2 + n_2$
14. $(g = 3 \vee g = 4) \rightarrow next-empty(\ell', q') \leq h + n$ and
 $(g' = 3 \vee g' = 4) \rightarrow next-empty(\ell'_2, q'_2) \leq h_2 + n_2$
15. $\ell' \leq i < h' \rightarrow test(i, q')$ and $\ell'_2 \leq i < h'_2 \rightarrow test(i, q'_2)$
16. $(g = 3 \vee g = 4) \wedge test(h, q) \rightarrow retrieve(h, q) = e$ and
 $(g' = 3 \vee g' = 4) \wedge test(h_2, q_2) \rightarrow retrieve(h_2, q_2) = e_2$

17. $(test(i, q) \wedge test(i, q')) \rightarrow retrieve(i, q) = retrieve(i, q')$ and
 $(test(i, q_2) \wedge test(i, q'_2)) \rightarrow retrieve(i, q_2) = retrieve(i, q'_2)$
18. $((g = 3 \vee g = 4) \wedge test(h, q')) \rightarrow retrieve(h, q') = e$ and
 $((g' = 3 \vee g' = 4) \wedge test(h_2, q'_2)) \rightarrow retrieve(h_2, q'_2) = e_2$
19. $(\ell \leq i \wedge j \leq next_empty(i, q')) \rightarrow q[i..j] = q'[i..j]$ and
 $(\ell_2 \leq i \wedge j \leq next_empty(i, q'_2)) \rightarrow q_2[i..j] = q'_2[i..j]$

In the initial state $\mathbf{N}_{nonmod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)$ all these invariants are satisfied. Also, all invariants are preserved by all summands. So they are satisfied in all reachable states of \mathbf{N}_{nonmod} . For a proof of this lemma see [2] Appendix B.3.

6. Correctness of \mathbf{N}_{mod}

In Section 6.1, we establish the strong bisimilarity of \mathbf{N}_{mod} and \mathbf{N}_{nonmod} . In order to prove this, we show that the bisimulation criteria in Definition 4 hold. Then according to Theorem 5, proof is complete. Section 6.2 demonstrates that \mathbf{N}_{nonmod} behaves like a pair of FIFO queues. Finally, the correctness of the two-way SWP is established in Section 6.3.

6.1. Equality of \mathbf{N}_{mod} and \mathbf{N}_{nonmod}

Proposition 14 $\mathbf{N}_{nonmod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0) \leftrightarrow$
 $\mathbf{N}_{mod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0).$

Proof. By Theorem 5, it suffices to define a state mapping ϕ and local mappings ψ_j for $j = 1, 2, \dots, 20$ that satisfy the bisimulation criteria in Definition 4, with respect to the invariants in Lemma 13.

Let Ξ abbreviate $\mathbb{N} \times \mathbb{N} \times Buf \times Buf \times \mathbb{N} \times \mathbb{N} \times \mathbb{N} \times \Delta \times \mathbb{N} \times \mathbb{N} \times \mathbb{N} \times \Delta \times \mathbb{N} \times \mathbb{N} \times \mathbb{N} \times Buf \times Buf \times \mathbb{N}$. We use $\xi : \Xi$ to abbreviate $(\ell, m, q, q'_2, \ell'_2, g, h, e, h'_2, g', h_2, e_2, h', \ell_2, m_2, q_2, q', \ell')$, then we define $\phi : \Xi \rightarrow \Xi$ by:

$$\phi(\xi) = (\ell|_{2n}, m|_{2n}, q|_{2n}, q'_2|_{2n_2}, \ell'_2|_{2n_2}, g, h|_{2n}, e, h'_2|_{2n_2}, g', h_2|_{2n_2}, e_2, h'|_{2n}, \ell_2|_{2n_2}, m_2|_{2n_2}, q_2|_{2n_2}, q'|_{2n}, \ell'|_{2n})$$

Furthermore, $\psi_2 : \mathbb{N} \rightarrow \mathbb{N}$ maps k to $k|_{2n}$, and $\psi_{15} : \mathbb{N} \rightarrow \mathbb{N}$ maps k to $k|_{2n_2}$; the other 18 local mappings are simply the identity. We show that ϕ and the ψ_j satisfy the bisimulation criteria. For each summand, we list (and prove) the non-trivial bisimulation criteria that it induces. For a detailed proof, see [2] Appendix C. \blacksquare

6.2. Correctness of \mathbf{N}_{nonmod}

We prove that \mathbf{N}_{nonmod} is branching bisimilar to the pair of FIFO queues \mathbf{Z} (see Section 3.2), using cones and foci (see Theorem 7)

The state mapping $\phi : \Xi \rightarrow List \times List$, which maps states of \mathbf{N}_{nonmod} to states of \mathbf{Z} , is defined by:

$$\phi(\xi) = (\phi_1(m, q, \ell', q'), \phi_2(m_2, q_2, \ell'_2, q'_2))$$

where

$$\begin{aligned} \phi_1(m, q, \ell', q') &= q'[\ell'..next_empty(\ell', q')] ++ q[next_empty(\ell', q')..m] \\ \phi_2(m_2, q_2, \ell'_2, q'_2) &= q'_2[\ell'_2..next_empty(\ell'_2, q'_2)] ++ q_2[next_empty(\ell'_2, q'_2)..m_2] \end{aligned}$$

Intuitively, ϕ_1 collects data elements in the sending window of **S/R** and the receiving window of **R/S**, starting at the first cell in the receiving window (i.e., ℓ') until the first empty cell in this window, and then continuing in the sending window until the first empty cell in that window (i.e., m). Likewise, ϕ_2 collects data elements in the sending window of **R/S** and the receiving window of **S/R**.

The focus points are states where in the direction from **S/R** to **R/S**, either the sending window of **S/R** is empty (meaning that $\ell = m$), or the receiving window from **R/S** is full and all data elements in this receiving window have been acknowledged (meaning that $\ell = \ell' + n$). Likewise for the direction from **R/S** to **S/R**. That is, the focus condition reads

$$FC(\xi) := (\ell = m \vee \ell = \ell' + n) \wedge (\ell_2 = m_2 \vee \ell_2 = \ell'_2 + n_2)$$

Lemma 15 For each $\xi:\Xi$ with $\mathbf{N}_{nonmod}(\xi)$ reachable from the initial state, there is a $\hat{\xi}:\Xi$ with $FC(\hat{\xi})$ such that $\mathbf{N}_{nonmod}(\xi) \xrightarrow{c_1} \dots \xrightarrow{c_n} \mathbf{N}_{nonmod}(\hat{\xi})$, where $c_1, \dots, c_n \in \mathcal{I}$.

Proof. We prove (see [2] Appendix C) that for each $\xi:\Xi$ where the invariants in Lemma 13 hold, there is a finite sequence of internal actions which ends in a state where $(\ell = m \vee \ell = \ell' + n) \wedge (\ell_2 = m_2 \vee \ell_2 = \ell'_2 + n_2)$. ■

Proposition 16 $\tau_{\{c,j\}}(\mathbf{N}_{nonmod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)) \xleftrightarrow{b} \mathbf{Z}(\langle \rangle, \langle \rangle)$.

Proof. We prove this using cones and foci method. See [2] Appendix C. ■

6.3. Correctness of the Two-Way Sliding Window Protocol

Finally, we can prove the main result of our specification which is:

Theorem 17 (Correctness)

$$\tau_{\mathcal{I}}(\partial_{\mathcal{H}}(\mathbf{S/R}(0, 0, [], [], 0) \parallel \mathbf{R/S}(0, 0, [], [], 0) \parallel \mathbf{K} \parallel \mathbf{L})) \xleftrightarrow{b} \mathbf{Z}(\langle \rangle, \langle \rangle)$$

Proof. We combine the equivalences that have been obtained so far:

$$\begin{aligned} & \tau_{\mathcal{I}}(\partial_{\mathcal{H}}(\mathbf{S/R}(0, 0, [], [], 0) \parallel \mathbf{K} \parallel \mathbf{R/S}(0, 0, [], [], 0) \parallel \mathbf{L})) \\ \Leftrightarrow & \tau_{\mathcal{I}}(\mathbf{M}_{mod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)) && \text{(Proposition 8)} \\ \Leftrightarrow & \tau_{\{c,j\}}(\mathbf{N}_{mod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)) && \text{(Proposition 9)} \\ \Leftrightarrow & \tau_{\{c,j\}}(\mathbf{N}_{nonmod}(0, 0, [], [], 0, 5, 0, d_0, 0, 5, 0, d_0, 0, 0, 0, [], [], 0)) && \text{(Proposition 14)} \\ \xleftrightarrow{b} & \mathbf{Z}(\langle \rangle, \langle \rangle) && \text{(Proposition 16)} \end{aligned}$$

■

7. Formalization in PVS

In this section we show the formalization and verification of the correctness proof of the SWP with piggybacking in PVS [29].

The PVS specification language is based on simply typed higher-order logic. Its type system contains basic types such as *boolean*, *nat*, *integer*, *real*, etc. and type constructors such as *set*, *tuple*, *record*, and *function*. Tuple types have the form $[T_1, \dots, T_n]$, where T_i are type expressions. A record is a finite list of fields of the form $R:\text{TYPE}=[\# E_1:T_1, \dots, E_n:T_n \#]$, where E_i are *record accessor* functions. A function type constructor has the form $F:\text{TYPE}=[T_1, \dots, T_n \rightarrow R]$, where F is a function with domain $D=T_1 \times \dots \times T_n$ and range R .

A PVS specification can be structured through a hierarchy of *theories*. Each theory consists of a *signature* for the type names and constants introduced in the theory, and a number of axioms, definitions and theorems associated with the signature. A PVS theory can be parametric in certain specified types and values, which are placed between $[]$ after the theory name.

In μCRL , the semantics of a data specification is the set of all its models. Incomplete data specifications may have multiple models. Even worse, it is possible to have inconsis-

tent data specifications for which no models exist. Here the necessity of specification with PVS emerges, because of this probable incompleteness and inconsistency which exists when working with μCRL . Moreover, PVS was used to search for omissions and errors in the manual μCRL proof of the SWP with piggybacking.

In Section 7.1 we show examples of the original specification of some data functions, then we introduce the modified forms of them. Moreover, we show how measure functions are used to detect the termination of recursive definitions. In Section 7.2 and 7.3 we represent the LPEs and invariants of the SWP with piggybacking in PVS. Section 7.4 presents the equality of μCRL specification of the SWP with piggybacking with and without modulo arithmetic. Section 7.5 explains how the cones and foci method is used to formalize the main theorem, that is the μCRL specification of the SWP with piggybacking is branching bisimilar to a FIFO queue of size $2n$. Finally, Section 7.6 is dedicated to some remarks on the verification in PVS.

7.1. Data Specifications in PVS

In PVS, all the definitions are first type checked, which generates some *proof obligations*. Proving all these obligations ascertains that our data specification is complete and consistent.

To achieve this, having total definitions is required. So in the first place, partially defined functions need to be extended to total ones. Below there are some examples of partial definitions in the original data specification of the SWP with piggybacking, which we changed into total ones. Second, to guarantee totality of recursive definitions, PVS requires the user to define a so-called *measure function*. Doing this usually requires time and effort, but the advantage is that recursive definitions are guaranteed to be well-founded. PVS enabled us to find non-terminating definitions in the original data specification of the SWP with piggybacking, which were not detected within the framework of μCRL . After finding these non-terminating definitions with PVS, we searched for new definition which can express the operation we look for. Then we replaced the old definitions with new terminating ones in our μCRL framework. Below we show some of the most interesting examples.

Example 18 We defined a function *next-empty* which seeks for the first empty position in q from a given position i . This function is identified as:

$$\text{next-empty}(i, q) = \text{if}(\text{test}(i, q), \text{next-empty}(S(i), q), i).$$

We also need to have $\text{next-empty}|_n(i, q)$ as a function which produces the first empty position in q modulo n , from position i . It looked reasonable to define it as:

$$\text{next-empty}|_n(i, q) = \text{if}(\text{test}(i, q), \text{next-empty}|_n(S(i)|_n, q), i)$$

Although the definition looks total and well-founded, this was one of the undetected potential errors that PVS detected during the type checking process. Below we bring an example to show what happens. Let $q = [(d_0, 0), (d_1, 1), (d_2, 2), (d_3, 3), (d_5, 5)]$, $n = 4$, $i = 5$ then

$$\begin{aligned} \text{next-empty}|_4(5, q) &= \text{next-empty}|_4(6|_4, q) = \text{next-empty}|_4(2, q) = \text{next-empty}|_4(3, q) \\ &= \text{next-empty}|_4(0, q) = \text{next-empty}|_4(1, q) = \text{next-empty}|_4(2, q) = \dots \end{aligned}$$

which will never terminate. The problem is that modulo n all the places in q are occupied, and since $0 \leq i|_n < n$ hence $\text{test}(i, q)$ will always be true. Hence each position will call for its immediate next position and so on. Therefore the calls will never stop.

At the end we replaced it with the following definition, which is terminating and operates the way as we expect.

$$\begin{aligned} \text{next-empty}|_n(i, q) &= \text{if}(\text{next-empty}(i|_n, q) < n, \text{next-empty}(i|_n, q), \\ &\quad \text{if}(\text{next-empty}(0, q) < n, \text{next-empty}(0, q), n)) \end{aligned}$$

```

...
D:nonempty_type
Buf:type=list[[D,nat]]
x,i,j,k,l,n: VAR nat
...
dm(i,j,n): nat =
    IF mod(i,n)<=mod(j,n)
    THEN mod(j,n)-mod(i,n)
    ELSE n+mod(j,n)-mod(i,n)
    ENDIF
...
release(n)(i,j,q): RECURSIVE Buf=
    IF mod(i,n)=mod(j,n) THEN q
    ELSE release(n)(mod(i+1,n),j,remove(mod(i,n),q))
    ENDIF
    measure dm(i,j,n)
...

```

Figure 2. An example of data specification in PVS

This function first checks whether there is any empty place after $i|_n$ (incl. $i|_n$ itself). If this is the case then that position would be the result, otherwise using $\text{next-empty}(0, q)$ it will check if there is any empty position in the buffer modulo n . If so then that position would be the value of the function since $\text{next-empty}(i|_n, q)$ will reach it. If all the buffer modulo n is full then n would be the result, because n is bigger than all the possible values for the function (i.e. $i|_n$ at most) and moreover it indicates that the buffer is full modulo n .

In [2] Appendix D there are similar examples for $\text{release}(i, j, q)$ and $\text{release}|_n(i, j, q)$, detected errors by PVS, and also our ultimate solutions for them.

We represented the μCRL abstract data types directly by PVS types. This enables us to reuse the PVS library for definitions and theorems of “standard” data types. Figure 2 illustrates part of a PVS theory defining $\text{release}|_n$. There D is an unspecified but non-empty type which represents the set of all data elements that can be communicated between the sender and the receiver. Buf is list of pairs of type $D \times \mathbb{N}$ defined as $\text{list}[[D, \text{nat}]]$. Here we used list to identify the type of lists, which is defined in the prelude in PVS. Therefore we simply use it without any need to define it explicitly. This figure also represents $\text{release}|_n(i, j, q)$ in PVS. Since it is defined recursively, in order to establish its termination (or totality), it is required by PVS to have a measure function. We define a measure function called dm which is decreasing and non-recursive. Here, PVS uses its type-checker to check the validity of dm . It generates two type-check proof obligations: if $i|_n < j|_n$ then $j|_n - i|_n \geq 0$ and if $i|_n \geq j|_n$ then $n + j|_n - i|_n \geq 0$. The first proof obligation is proved in one trivial step. The second one is proved using Lemma 19.

In [2] Appendix D, we also list the extra data lemmas which had to be proved in PVS while they are considered to be trivial in the manual proof.

7.2. Representing LPEs

We now reuse [10] to show how the μCRL specification of the SWP with piggybacking (an LPE) can be represented in PVS. The main distinction will be that we have assumed so far that LPEs are *clustered*. This means that each action label occurs in at most one summand, so that the set of summands could be indexed by the set of action labels. This is no limitation, because any LPE can be transformed in clustered form, basically by replacing $+$ by \sum over

```

LPE[Act,State,Local:TYPE,n:nat]: THEORY BEGIN
  SUMMAND:TYPE= [State,Local-> [#act:Act,guard:bool,next:State#] ]
  LPE:TYPE= [#init:State,sums:[below(n)->SUMMAND]#]
END LPE

```

Figure 3. Definition of LPE in PVS

```

...
l,m,l12,g,h,h12,g1,h2,h1,l2,m2,l1: var nat
q,q1,q2,q12 : var Buf
e,e2: var D
...
inv(l,m,q,q12,l12,g,h,e,h12,g1,h2,e2,h1,l2,m2,q2,q1,l1): bool= next_empty(l1,q1)<=m
...

```

Figure 4. An example of representing invariants in PVS

finite types. Clustered LPEs enable a notationally smoother presentation of the theory. However, when working with concrete LPEs this restriction is not convenient, so we avoid it in the PVS framework: an arbitrarily sized index set $\{0, \dots, n-1\}$ will be used, represented by the PVS type `below(n)`. A second deviation is that we will assume from now on that every summand has the same set of local variables. Again this is no limitation, because void summations can always be added (i.e. $p = \sum_{d:D} p$, when d does not occur in p). This restriction is needed to avoid the use of polymorphism, which does not exist in PVS. The third deviation is that we do not distinguish action labels from action data parameters. We simply work with one type of expressions for actions. This allows that a summand can generate transitions with various labels. This generalization makes the formalization a bit smoother, but was not really exploited.

So an LPE is parameterized by sets of actions (`Act`), global parameters (`State`) and local variables (`Local`), and by the size of its index set (n). Note that the guard, action and next-state of a summand depend on the global parameters $d : State$ and on local variables $e : Local$. This dependency is represented in the definition `SUMMAND` by a PVS function type. In Figure 3 an LPE consists of an initial state and a list of summands indexed by `below(n)`.

A concrete LPE by a fragment of the linear specification \mathbb{N}_{mod} of SWP with piggybacking in PVS (see Figure 6 in Appendix D in [2]) is introduced as an lpe of a set of actions: `Nnonmod_act`, `states: State`, `local variables: Local`, and a `digit: 20` referring to the number of summands. The LPE is identified as a pair, called `init` and `sums`, where `init` is introducing the initial state of \mathbb{N}_{mod} and `sums` the summands. The first LAMBDA maps each number to the corresponding summand in \mathbb{N}_{mod} . The second LAMBDA is representing the summands as functions over `State` and `Local`. Here, `State` is the set of states and `Local` is the data type $D \times \mathbb{N}$ of all pairs (d, k) of the summation variables, which is considered as a global variable regarding the property: $p = \sum_{(d,k):local} p$, which is mentioned before.

7.3. Representing Invariants

Invariants are boolean functions over the set of states. In Figure 4, we explain how to represent an invariant of the μ CRL specification, in PVS. This figure illustrates the (first part of the) Invariant 13.9 from Section 5.3

```

...
state_f(l,m,q,q12,l12,g,h,e,h12,g1,h2,e2,h1,l2,m2,q2,q1,l1): State=
  (mod(l,2*n),mod(m,2*n),modulo2(q,2*n),modulo2(q12,2*n2),mod(l12,2*n2),
    g,mod(h,2*n),e,mod(h1,2*n2),g1,mod(h2,2*n2),e2,mod(h1,2*n),
    mod(l2,2*n2),mod(m2,2*n2),modulo2(q2,2*n2),modulo2(q1,2*n),
    mod(l1,2*n)),
local_f(l:Local,i:below(20)): Local=
  LET (e,k)=1 IN
  IF i=4 THEN (e,mod(k,2*n)) ELSE (IF i=9 THEN (e,mod(k,2*n2)) ELSE(e,k)) ENDIF
...
Propsimilaosition_6_22: proposition bisimilar (lpe2lts(Nnonmod),lpe2lts(Nmod))
...

```

Figure 5. Equality of N_{mod} and N_{nonmod} in PVS

7.4. Equality of N_{mod} and N_{nonmod}

Strong bisimilarity of N_{mod} and N_{nonmod} (Proposition 14) is depicted in Figure 5. `state_f` and `local_f` are introduced to construct the state mapping between N_{nonmod} and N_{mod} . In PVS we introduce the state mapping (`state_f`, `local_f`) from the set of states and local variables of N_{nonmod} to those of N_{mod} . Then we use the corresponding relation to this state mapping, and we show that this relation is a bisimulation between N_{nonmod} and N_{mod} .

In PVS we defined an LPE as a list of summands (not as a recursive equation), equipped with the standard LTS semantics. It could be proved directly that state mappings preserve strong bisimulation.

By contrast, the manual proof that N_{mod} and N_{nonmod} are strongly bisimilar is based on the proof principle CL-RSP [5], which states that each LPE has a unique solution, modulo strong bisimilarity. An advantage of this approach is that by using algebraic principles only, the stated equivalence also holds in non-standard models for process algebra + CL-RSP. We did not formalize CL-RSP in PVS because it depends on recursive process equations; this would have required a laborious embedding of μ CRL in PVS, which would complicate the formalization too much.

7.5. Correctness of N_{mod}

The branching bisimilarity verification of N_{mod} and Z (Theorem 17) is pictured in Figure 6. The function `fc(l,m,q,q12,l12,g,h,e,h12,g1,h2,e2,h1,l2,m2,q2,q1,l1)` defines the focus condition for $N_{nonmod}(\ell, m, q, q'_2, \ell'_2, g, h, e, h'_2, g', h_2, e_2, h', \ell_2, m_2, q_2, q', \ell')$ as a boolean function on set of states. `qlist(q,i,j)` is used to describe the function $q[i..j]$, which is defined as an application on triples. The state mapping `h` maps states of N_{nonmod} to states of Z , which is called $\phi: \Xi \rightarrow List \times List$ in Section 6.2. `k` is a Boolean function which is used to match each external action of N_{nonmod} to the corresponding one of Z . This is done by corresponding the number of each summand of N_{nonmod} to one of Z . As PVS requires, this function must be total, therefore without loss of generality we map all the summands with an internal action, from N_{nonmod} 's specification, to the second summand of Z 's specification.

According to cones and foci proof method [10], to derive that N_{nonmod} and N_{mod} are branching bisimilar, it is enough to check the matching criteria and the reachability of focus points. The two conditions of the cones and foci proof method are represented by `mc` and `WN`, namely matching criteria and the reachability of focus points, respectively. `mc` establishes that all the matching criteria (see Section 1) hold for every reachable state `d` in N_{nonmod} , with the aforementioned `h`, `k` and `fc` functions. `WN` represents the fact that from all reachable states `S`

```

...
fc(l,m,q,q12,l12,g,h,e,h12,g1,h2,e2,h1,l2,m2,q2,q1,l1): bool =
    (l=m OR l=11+n) AND (l2=m2 OR l2=112+n2)
k(i): below(2)= IF i=18 THEN 0 ELSE
    IF i=10 THEN 1 ELSE
    IF i=11 THEN 2 ELSE 3 ENDIF ENDIF ENDIF
h(l,m,q,q12,l12,g,h,e,h12,g1,h2,e2,h1,l2,m2,q2,q1,l1): [List_,List_]=
    (concat(qlist(q1,l1,next_empty(l1,q1)),qlist(q,next_empty(l1,q1),m)),
    concat(qlist(q12,l12,next_empty(l12,q12)),qlist(q2,next_empty(l12,q12),m2)))
mc: THEOREM FORALL d: reachable(Nnonmod)(d) IMPLIES MC(Nnonmod,Z,k,h,fc)(d)
WN: LEMMA FORALL S: reachable(Nnonmod)(S) IMPLIES WN(Nnonmod,fc)(S)
main: THEOREM brbisimilar(lpe2lts(Nmod),lpe2lts(Z))
...

```

Figure 6. Correctness of N_{mod} in PVS

in N_{nonmod} , a focus point can be reached by a finite series of internal actions. The function `lpe2lts` provides the Labeled Transition System semantics of an LPE (see [10]).

7.6. Remarks on the Verification in PVS

We used PVS to find the omissions and undetected potential errors that have been ignored in the manual μ CRL proofs; some of them have been shown as examples in Section 7.1. PVS guided us to find some important invariants. We affirmed the termination of recursive definitions by means of various measure functions. We represented LPEs in PVS and then introduced N_{mod} and N_{nonmod} as LPEs. We verified the bisimulation of N_{nonmod} and N_{mod} . Finally we used the cones and foci proof method [10], to prove that N_{mod} and the external behavior of the SWP with piggybacking, represented by Z , are branching bisimilar.

8. Conclusions

In this paper we verify a two-sided sliding window protocol which has the acknowledgments piggybacked on data. This way acknowledgments take a free ride in the channel. As a result the available bandwidth is used better. We present a specification of sliding window protocol with piggybacking in μ CRL, and then verify the specification with the PVS theorem prover.

An important aim of this paper is to show how one can incrementally extend a PVS verification effort, in this case the one described in [1]. PVS verification can be reused to check modifications of the SWP nearly automatically. We benefited from the PVS formalizations and lemmas in [1], e.g. properties of data types and those invariants which are not directly working with the internal structure of buffers (i.e. ordered lists). These are also mentioned in [2]. Note that a large part of the complete formalization consists of developing the meta theory. This part is split in generic PVS files with proofs. This generic part can be reused for the correctness proof of many other protocols. In particular, the generic part consists of the definition of an LTS, various forms of bisimulation (with proofs that they form equivalence relations), the definition of LPEs, their operational semantics, the notions of state mappings between LPEs, the notion of an invariant of an LPE (and its relation with reachable states), the proof rules for tau-reachability (with a soundness proof), and the matching criteria (including the proof of the theorem, that from the cones and foci method one may conclude branching bisimilarity).

For a specific protocol verification one must formalize the used data types (or find them in PVS's prelude), define LPEs for the specification and implementation, list the invariants,

the focus conditions and the state mapping. From this, all proof obligations (like invariants and matching criteria) are generated automatically. Most obligations can be discharged automatically, but still many must be proven manually. Also, tau-reachability must typically be proven manually, using the predefined proof rules. However, some steps remain protocol-specific, such as the transition from modulo to full arithmetic in the case of the Sliding Window Protocol.

Here, we model the medium between the sending and receiving window as a queue of capacity one. So a possible extension of this work would be to verify this protocol with mediums of unbounded size, i.e. we can define the mediums as lists of pairs (d, i) by:

$$\begin{aligned} \text{cons} &: [] : \rightarrow \text{Medium} \\ \text{func} &: \text{add} : \Delta \times \mathbb{N} \times \text{Medium} \rightarrow \text{Medium} \end{aligned}$$

Acknowledgements

We would like to thank Jun Pang for the helpful discussions and for his μCRL files on the one-way SWP.

References

- [1] B. Badban, W. Fokkink, J. Groote, J. Pang, and J. van de Pol. Verifying a sliding window protocol in μCRL and PVS. *Formal Aspects of Computing*, 17(3):342–388, 2005.
- [2] B. Badban, W. Fokkink, and J. van de Pol. Mechanical verification of a two-way sliding window protocol (extended version), 2008. <http://www.inf.uni-konstanz.de/~badban/piggybacking.pdf>.
- [3] J. Bergstra and J. Klop. Process algebra for synchronous communication. *Information and Control*, 60(1-3):109–137, 1984.
- [4] M. Bezem and J. Groote. A correctness proof of a one bit sliding window protocol in μCRL . *The Computer Journal*, 37(4):289–307, 1994.
- [5] M. Bezem and J. Groote. Invariants in process algebra with data. In B. Jonsson and J. Parrow, editors, *Proc. 5th Conference on Concurrency Theory*, LNCS 836, pages 401–416, 1994.
- [6] R. Cardell-Oliver. Using higher order logic for modeling real-time protocols. In J. Diaz and F. Orejas, editors, *Proc. 4th Joint Conference on Theory and Practice of Software Development*, Lecture Notes in Computer Science 494, pages 259–282, 1991.
- [7] V. Cerf and R. Kahn. A protocol for packet network intercommunication. *IEEE Transactions on Communications*, COM-22:637–648, 1974.
- [8] D. Chkhaev, J. Hooman, and E. de Vink. Verification and improvement of the sliding window protocol. In H. Garavel and J. Hatcliff, editors, *Proc. 9th Conference on Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science 2619, pages 113–127, 2003.
- [9] W. Fokkink, J. Groote, J. Pang, B. Badban, and J. van de Pol. Verifying a sliding window protocol in μCRL . In S. M. C. Rattray and C. Shankland, editors, *Proc. 10th Conference on Algebraic Methodology and Software Technology*, Lecture Notes in Computer Science 3116, pages 148–163, 2004.
- [10] W. Fokkink, J. Pang, and J. van de Pol. Cones and foci: A mechanical framework for protocol verification. *Formal Methods in System Design*, 29(1):1–31, 2006.
- [11] P. Godefroid and D. Long. Symbolic protocol verification with Queue BDDs. *Formal Methods and System Design*, 14(3):257–271, 1999.
- [12] R. Groenvelde. Verification of a sliding window protocol by means of process algebra. Technical Report P8701, University of Amsterdam, 1987.
- [13] J. Groote and H. Korver. Correctness proof of the bakery protocol in μCRL . In A. Ponse, C. Verhoef, and S. v. Vlijmen, editors, *Proc. 1st Workshop on Algebra of Communicating Processes '94*, Workshops in Computing Series, pages 63–86. Springer-Verlag, 1995.
- [14] J. Groote and A. Ponse. Proof theory for μCRL : A language for processes with data. In D. Andrews, J. Groote, and C. Middelburg, editors, *Proc. Workshop on Semantics of Specification Languages*, Workshops in Computing Series, pages 231–250, 1994.
- [15] J. Groote and A. Ponse. The syntax and semantics of μCRL . In A. Ponse, C. Verhoef, and S. v. Vlijmen, editors, *Proc. 1st Workshop on Algebra of Communicating Processes '94*, Workshops in Computing Series, pages 26–62. Springer-Verlag, 1995.

- [16] J. Groote, A. Ponse, and Y. Usenko. Linearization in parallel pCRL. *Journal of Logic and Algebraic Programming*, 48(1-2):39–72, 2001.
- [17] J. Groote and M. Reniers. Algebraic process verification. In J. Bergstra, A. Ponse, and S. Smolka, editors, *Handbook of Process Algebra*, chapter 17, pages 1151–1208. Elsevier, 2001.
- [18] J. Groote and J. Springintveld. Focus points and convergent process operators. a proof strategy for protocol verification. *Journal of Logic and Algebraic Programming*, 49(1-2):31–60, 2001.
- [19] B. Hailpern. *Verifying Concurrent Processes Using Temporal Logic*. LNCS 129. 1982.
- [20] G. Holzmann. *Design and Validation of Computer Protocols*. Prentice Hall, 1991.
- [21] G. Holzmann. The model checker Spin. *IEEE Transactions on Software Engineering*, 23:279–295, 1997.
- [22] B. Jonsson. *Compositional Verification of Distributed Systems*. PhD thesis, Department of Computer Science, Uppsala University, 1987.
- [23] B. Jonsson and M. Nilson. Transitive closures of regular relations for verifying infinite-state systems. In S. Graf and M. Schwartzbach, editors, *Proc. 6th Conference on Tools and Algorithms for Construction and Analysis of Systems*, Lecture Notes in Computer Science 1785, pages 220–234. Springer-Verlag, 2000.
- [24] R. Kaivola. Using compositional preorders in the verification of sliding window protocol. In *Proc. 9th Conference on Computer Aided Verification*, LNCS 1254, pages 48–59, 1997.
- [25] D. Knuth. Verification of link-level protocols. *BIT*, 21:21–36, 1981.
- [26] T. Latvala. Model checking LTL properties of high-level Petri nets with fairness constraints. In J. Colom and M. Koutny, editors, *Proc. 21st Conference on Application and Theory of Petri Nets*, Lecture Notes in Computer Science 2075, pages 242–262. Springer-Verlag, 2001.
- [27] J. Loeckx, H. Ehrlich, and M. Wolf. *Specification of Abstract Data Types*. Wiley/Teubner, 1996.
- [28] E. Madelaine and D. Vergamini. Specification and verification of a sliding window protocol in LOTOS. In E. Knuth and L. Wegner, editors, *Proc. 4th Conference on Formal Description Techniques for Distributed Systems and Communication Protocols*, IFIP Transactions (C-2), pages 495–510. North-Holland, 1991.
- [29] S. Owre, S. Rajan, J. Rushby, N. Shankar, and M. Srivas. PVS: Combining specification, proof checking, and model checking. In R. Alur and T.A. Henzinger, editors, *Proc. 8th Conference on Computer Aided Verification*, Lecture Notes in Computer Science 1102, pages 411–414. Springer-Verlag, 1996.
- [30] K. Paliwoda and J. Sanders. An incremental specification of the sliding-window protocol. *Distributed Computing*, 5(2):83–94, 1991.
- [31] D. Park. Concurrency and automata on infinite sequences. In P. Deussen, editor, *Proc. 5th GI-Conference on Theoretical Computer Science*, LNCS 104, pages 167–183, 1981.
- [32] J. Richier, C. Rodriguez, J. Sifakis, and J. Voiron. Verification in Xesar of the sliding window protocol. In H. Rudin and C. West, editors, *Proc. 7th Conference on Protocol Specification, Testing and Verification*, pages 235–248. North-Holland, 1987.
- [33] C. Röckl and J. Esparza. Proof-checking protocols using bisimulations. In J. Baeten and S. Mauw, editors, *Proc. 10th Conference on Concurrency Theory*, LNCS 1664, pages 525–540, 1999.
- [34] V. Rusu. Verifying a sliding-window protocol using PVS. In M. Kim, B. Chin, S. Kang, and D. Lee, editors, *Proc. 21st Conference on Formal Techniques for Networked and Distributed Systems*, IFIP Conference Proceedings 197, pages 251–268. Kluwer Academic, 2001.
- [35] A. Schoone. *Assertional Verification in Distributed Computing*. PhD thesis, Utrecht University, 1991.
- [36] M. Smith and N. Klarlund. Verification of a sliding window protocol using IOA and MONA. In T. Bolognesi and D. Latella, editors, *Proc. 20th Joint Conference on Formal Description Techniques for Distributed Systems and Communication Protocols*, pages 19–34. Kluwer Academic Publishers, 2000.
- [37] K. Stahl, K. Baukus, Y. Lakhnech, and M. Steffen. Divide, abstract, and model-check. In D. Dams, R. Gerth, S. Leue, and M. Massink, editors, *Proc. 6th SPIN Workshop on Practical Aspects of Model Checking*, Lecture Notes in Computer Science 1680, pages 57–76. Springer-Verlag, 1999.
- [38] N. Stenning. A data transfer protocol. *Computer Networks*, 1:99–110, 1976.
- [39] A. Tanenbaum. *Computer Networks*. Prentice Hall, 1981.
- [40] F. Vaandrager. Verification of two communication protocols by means of process algebra. Technical Report Report CS-R8608, CWI, 1986.
- [41] J. van de Snepscheut. The sliding window protocol revisited. *Formal Aspects of Computing*, 7(1):3–170, 1995.
- [42] R. van Glabbeek. What is branching time and why to use it? *The Concurrency Column, Bulletin of the EATCS*, 53:190–198, 1994.
- [43] R. van Glabbeek and W. Weijland. Branching time and abstraction in bisimulation semantics. *Journal of the ACM*, 43(3):555–600, 1996.
- [44] J. van Wamel. A study of a one bit sliding window protocol in ACP. Technical Report P9212, University of Amsterdam, 1992.