


Protocol Validation with μ CRL



Wan Fokkink
Modelling Distributed Systems
Springer, 2007

- ▶ Formal modeling of real-life protocols
- ▶ Automated analysis of protocols by means of state space exploration (i.e. simulation or model checking)
- ▶ To get an impression of the difficulties in analysis (state space explosion)
- ▶ Symbolic verification of protocols using equational logic and theorem provers

Algebraic Specification

An **algebraic specification** of data types consists of

- ▶ a **signature**, i.e. function symbols from which one can build **terms**
- ▶ **axioms**, i.e. equations between terms, inducing an equality relation on terms (closed under:
(1) **equivalence**, (2) **substitution**, and (3) **context**)

Natural Numbers - Example

The signature of the **natural numbers** consists of constant **0**, unary successor **S**, and binary addition **plus** and multiplication **mul**.

The axioms are:

$$\text{plus}(n, 0) = n$$

$$\text{plus}(n, S(m)) = S(\text{plus}(n, m))$$

$$\text{mul}(n, 0) = 0$$

$$\text{mul}(n, S(m)) = \text{plus}(\text{mul}(n, m), n)$$

The axioms are directed from left to right, and must constitute a *terminating* rewrite system.

Question: Derive $\text{plus}(S(0), S(0)) = S(S(0))$.

Constructors

μ CRL uses algebraic specification of data, with explicit recognition of **constructor** symbols, which cannot be eliminated from data terms.

Example: For the natural numbers, 0 and S are constructors, while $plus$ and mul are not.

```
sort    $Nat$ 
func    $0 : \rightarrow Nat$ 
          $S : Nat \rightarrow Nat$ 
map    $plus, mul : Nat \times Nat \rightarrow Nat$ 
var    $n, m : Nat$ 
rew    $plus(n, 0) = n$ 
          $plus(n, S(m)) = S(plus(n, m))$ 
          $mul(n, 0) = 0$ 
          $mul(n, S(m)) = plus(mul(n, m), n)$ 
```

Question: Specify $power(n, m)$, denoting n^m .

Booleans

'true' and 'false', together with conjunction, disjunction and negation must be declared in each μ CRL specification.

```
sort   Bool
func    $\top, \text{F} : \rightarrow \text{Bool}$ 
map     $\wedge, \vee : \text{Bool} \times \text{Bool} \rightarrow \text{Bool}$ 
          $\neg : \text{Bool} \rightarrow \text{Bool}$ 
var     $b : \text{Bool}$ 
rew     $b \wedge \top = b$ 
          $b \wedge \text{F} = \text{F}$ 
          $b \vee \top = \top$ 
          $b \vee \text{F} = b$ 
          $\neg \top = \text{F}$ 
          $\neg \text{F} = \top$ 
```

Rewriting of data terms is performed according to the **innermost** strategy, meaning that a term $f(d_1, \dots, d_n)$ can only be rewritten if d_1, \dots, d_n are **normal forms**.

A normal form only consists of constructor symbols.

Equality Function

One needs to define an equality function $eq : D \times D \rightarrow Bool$ for data types D , where $eq(d, e) = T$ if and only if $d = e$.

Example:

map $eq : Bool \times Bool \rightarrow Bool$

rew $eq(T, T) = T$

$eq(F, F) = T$

$eq(T, F) = F$

$eq(F, T) = F$

map $eq : Nat \times Nat \rightarrow Bool$

var $n, m : Nat$

rew $eq(0, 0) = T$

$eq(S(n), S(m)) = eq(n, m)$

$eq(0, S(n)) = F$

$eq(S(n), 0) = F$

Equality Function

A shorter specification of the equality function on booleans is:

$$eq(b, b) = T$$

$$eq(T, F) = F$$

$$eq(F, T) = F$$

The following specification of an equality function does *not* work in μ CRL:

$$eq(x, x) = T$$

$$eq(x, y) = F$$

That is, rewrite rules are not always 'executed' from top to bottom.

Induction

One can prove properties of data terms by **induction** on **constructors**.

Example: We prove by induction that $\neg\neg b = b$ for all booleans b .

$$[b \text{ is T}] \neg\neg T = \neg F = T$$

$$[b \text{ is F}] \neg\neg F = \neg T = F$$

Example: We prove by induction that $plus(0, n) = n$ for all natural numbers n .

$$[\text{Base case, } n \text{ is } 0] plus(0, 0) = 0$$

$$[\text{Inductive case, } n \text{ is } S(m)] plus(0, S(m)) = S(plus(0, m)) = S(m)$$

Basic Process Terms

Basic process terms are built from parametrized actions in a set Act , alternative composition and sequential composition.

- ▶ An action name $a \in \text{Act}$ represents indivisible behavior. It can carry data parameters: $a(d_1, \dots, d_n)$.
- ▶ The process term $p + q$ executes the behavior of either p or q .
- ▶ The process term $p \cdot q$ first executes p , and upon termination proceeds to execute q .

Basic Process Terms - Example

$((a + b) \cdot c) \cdot d$ represents the state space

$((a + b) \cdot c) \cdot d$



$c \cdot d$



d



\checkmark

Basic Process Algebra - Axioms

$$x + y = y + x$$

$$(x + y) + z = x + (y + z)$$

$$x + x = x$$


$$(x + y) \cdot z = (x \cdot z) + (y \cdot z)$$

$$(x \cdot y) \cdot z = x \cdot (y \cdot z)$$

No Left Distributivity

$x \cdot (y + z) \neq (x \cdot y) + (x \cdot z)$ does *not* hold.

Example:



Process one reads d , and then decides whether it writes d on disc 1 or 2. Process two makes a choice for disc 1 or 2 before it reads d . If disc 1 crashes, then process one saves datum d on disc 2, while process two may get stuck.

Bisimulation Equivalence

\downarrow is a special predicate on states, expressing **successful termination**.
That is, \checkmark is the only state where \downarrow holds.

Assume a state space. A **bisimulation** is a binary relation \mathcal{B} on states such that:

1. if $s_1 \mathcal{B} s_2$ and $s_1 \xrightarrow{a} s'_1$, then $s_2 \xrightarrow{a} s'_2$ with $s'_1 \mathcal{B} s'_2$
2. if $s_1 \mathcal{B} s_2$ and $s_2 \xrightarrow{a} s'_2$, then $s_1 \xrightarrow{a} s'_1$ with $s'_1 \mathcal{B} s'_2$
3. if $s_1 \mathcal{B} s_2$ and $s_1 \downarrow$, then $s_2 \downarrow$
4. if $s_1 \mathcal{B} s_2$ and $s_2 \downarrow$, then $s_1 \downarrow$

Two states s_1 and s_2 are **bisimilar**, denoted $s_1 \leftrightarrow s_2$, if there is a bisimulation relation \mathcal{B} such that $s_1 \mathcal{B} s_2$.

Example: $a \cdot (b + c) \not\leftrightarrow (a \cdot b) + (a \cdot c)$

Soundness and Completeness of the Axioms

Theorem: For basic process algebra terms p and q :

$$p = q \Leftrightarrow p \underline{\leftrightarrow} q$$

Communication

To specify that two action names can **communicate** (or **synchronize**):

$$\mathbf{comm} \quad a|b = c$$

Communication is supposed to be **commutative** and **associative**.

$$\begin{aligned} a|b &= b|a \\ (a|b)|c &= a|(b|c) \end{aligned}$$

Actions $a(d_1, \dots, d_n)$ and $b(e_1, \dots, e_m)$ can only communicate if they carry exactly the same data parameters.

In μCRL , the equality function only needs to be defined for data types that are used in parameters of actions that can communicate.

Parallelism

The **merge** \parallel executes the two process terms in its arguments in parallel.

For example, if action names a and b do not communicate,

$$a \parallel b = a \cdot b + b \cdot a$$


The merge can also execute a communication between actions of its arguments.

For example, if $a|b = c$,

$$a \parallel b = (a \cdot b + b \cdot a) + c$$

Parallelism - Example

If all communications between action names result to c , then



Question

Does $(x + y) \parallel z = (x \parallel z) + (y \parallel z)$ hold?

Left Merge and Communication Merge

The **left merge** \ll executes an action of its first argument and then behaves as the merge.

The **communication merge** $|$ executes a communication of actions of its two arguments and then behaves as the merge.

Example: If $a|b = c$, then

$$a \ll b = a \cdot b$$

$$a|b = c$$

These operators are needed to axiomatize the merge. In particular,

$$p \parallel q = (p \ll q + q \ll p) + p|q$$

Left Merge and Communication Merge

$(x + y) \parallel z = (x \parallel z) + (y \parallel z)$ holds.

$x \parallel (y + z) = (x \parallel y) + (x \parallel z)$ does **not** hold.

$(x + y) | z = (x | z) + (y | z)$ holds.

$x | (y + z) = (x | y) + (x | z)$ holds.

Parallelism - Axioms

$$x \parallel y = (x \perp\!\!\!\perp y + y \perp\!\!\!\perp x) + x | y$$

$$a(\vec{d}) \perp\!\!\!\perp x = a(\vec{d}) \cdot x \quad (\vec{d} \text{ denotes } d_1, \dots, d_n)$$

$$(a(\vec{d}) \cdot x) \perp\!\!\!\perp y = a(\vec{d}) \cdot (x \parallel y)$$

$$(x + y) \perp\!\!\!\perp z = x \perp\!\!\!\perp z + y \perp\!\!\!\perp z$$

$$a(\vec{d}) | b(\vec{d}) = c(\vec{d}) \quad \text{if } a | b = c$$

$$a(\vec{d}) | b(\vec{e}) = \delta \quad \text{if } \vec{d} \neq \vec{e}$$

$$a(\vec{d}) | b(\vec{e}) = \delta \quad \text{if } a | b \text{ is undefined}$$

$$(a(\vec{d}) \cdot x) | b(\vec{e}) = (a(\vec{d}) | b(\vec{e})) \cdot x$$

$$a(\vec{d}) | (b(\vec{e}) \cdot x) = (a(\vec{d}) | b(\vec{e})) \cdot x$$

$$(a(\vec{d}) \cdot x) | (b(\vec{e}) \cdot y) = (a(\vec{d}) | b(\vec{e})) \cdot (x \parallel y)$$

$$(x + y) | z = x | z + y | z$$

$$x | (y + z) = x | y + x | z$$

Deadlock and Encapsulation

- * The **deadlock** δ does not display any behavior.
- * The **encapsulation** operators ∂_H , for sets of actions H , rename all actions of H in their argument into δ .

Encapsulation operators enable to enforce actions into communication.

Example: Let $s|r = c$.

$$\begin{aligned} s \parallel r &= (s \cdot r + r \cdot s) + c \\ \partial_{\{s,r\}}(s \parallel r) &= c \end{aligned}$$

Deadlock and Encapsulation - Axioms

$$x + \delta = x$$

$$\delta \cdot x = \delta$$

$$\partial_H(\delta) = \delta$$

$$\partial_H(a(\vec{d})) = a(\vec{d}) \quad \text{if } a \notin H$$

$$\partial_H(a(\vec{d})) = \delta \quad \text{if } a \in H$$

$$\partial_H(x + y) = \partial_H(x) + \partial_H(y)$$

$$\partial_H(x \cdot y) = \partial_H(x) \cdot \partial_H(y)$$

$$\delta \perp\!\!\!\perp x = \delta$$

$$\delta | x = \delta$$

$$x | \delta = \delta$$

Example - Bits Through a Channel

A bit 0 or 1 is sent into a channel: $s(0) + s(1)$

The bit is received at the other side of the channel: $r(0) + r(1)$

The communication of s and r is c .

The behavior of the channel is described by

$$\partial_{\{s,r\}}((s(0) + s(1)) \parallel (r(0) + r(1)))$$

The encapsulation operator enforces that $s(d)$ and $r(d)$ can only occur in communication.

The axioms can be used to equate the process term above to

$$c(0) + c(1)$$

Process Declaration



This process can be captured by means of:

$$\begin{aligned} X &= a \cdot Y \\ Y &= b \cdot X \end{aligned}$$

X and Y represent the two states of the process.

A **process declaration** **proc** consists of **recursive equations**



$$X(x_1:D_1, \dots, x_n:D_n) = p$$

where the term p may contain expressions $Y(e_1, \dots, e_m)$.

The **initial declaration** **init** consists of a single expression $X(d_1, \dots, d_n)$, representing the initial state of the specification.

Question

How can one specify



Process Declaration - Example

The process *Clock* repeatedly performs action *tick* or displays the current time.

act *tick*
 display : *Nat*

proc $Clock(n:Nat) = tick \cdot Clock(S(n)) + display(n) \cdot Clock(n)$

init $Clock(0)$

'Unguarded' process declarations such as $X = X$ and $Y = Y \cdot a$ are illegal.

Conditional

The process term $p \triangleleft b \triangleright q$, where p and q are process terms and b is a data term of sort *Bool*, behaves as p if $b = T$ and as q if $b = F$.

$$x \triangleleft T \triangleright y = x$$

$$x \triangleleft F \triangleright y = y$$

Example: The process *Counter* counts the number of *a*-actions that occur, resetting the internal counter after three *a*'s:

act a, reset

proc $\text{Counter}(n:\text{Nat}) =$
 $a \cdot \text{Counter}(S(n)) \triangleleft n < S(S(S(0))) \triangleright \text{reset} \cdot \text{Counter}(0)$

init $\text{Counter}(0)$

Summation over a Data Type

The **sum** operator $\sum_{d:D} P(d)$ behaves as

$$P(d_1) + P(d_2) + \dots$$

i.e. as the (possibly infinite) choice between process terms $P(d)$ for data terms d that can be built from the *constructors* of D .

In μ CRL, the distinction between **func** and **map** is used to build the set of constructor terms for summation over a data type.

Question

Let $send \mid recv = comm$.

What is the state space of

$$\partial_{\{send,recv\}}(send(S(0)) \parallel \sum_{n:Nat} recv(n)) ?$$

Summation - Axioms

$$\sum_{d:D} x = x$$

$$\sum_{d:D} P(d) = \sum_{d:D} P(d) + P(d_0) \quad (d_0 \in D)$$

$$\sum_{d:D} (P(d) + Q(d)) = \sum_{d:D} P(d) + \sum_{d:D} Q(d)$$

$$(\sum_{d:D} P(d)) \cdot x = \sum_{d:D} (P(d) \cdot x)$$

$$(\forall d:D P(d) = Q(d)) \Rightarrow \sum_{d:D} P(d) = \sum_{d:D} Q(d)$$

$$(\sum_{d:D} P(d)) \ll x = \sum_{d:D} (P(d) \ll x)$$


$$(\sum_{d:D} P(d)) | x = \sum_{d:D} (P(d) | x)$$

$$x | (\sum_{d:D} P(d)) = \sum_{d:D} (x | P(d))$$

$$\partial_H(\sum_{d:D} P(d)) = \sum_{d:D} \partial_H(P(d))$$

Example - Bag

We can put elements of sort D into a bag, and collect these elements from the bag in arbitrary order. For example, if D is $\{0, 1\}$:



How could one specify the bag over $D = \{d_1, d_2\}$ in μCRL ?

Example - Bag

If D is $\{d_1, d_2\}$, then a μ CRL specification of the bag is:

```
act    $in, out : D$   
proc  $Y(n: Nat, m: Nat) = in(d_1) \cdot Y(S(n), m)$   
       $+ in(d_2) \cdot Y(n, S(m))$   
       $+ (out(d_1) \cdot Y(P(n), m) \triangleleft n > 0 \triangleright \delta)$   
       $+ (out(d_2) \cdot Y(n, P(m)) \triangleleft m > 0 \triangleright \delta)$   
  
init   $Y(0, 0)$ 
```

where $P(S(n)) = n$.

An alternative μ CRL specification that works for general D is:

```
act    $in, out : D$   
proc  $X = \sum_{d:D} in(d) \cdot (X \parallel out(d))$   
init   $X$ 
```

Hidden Action and Hiding

- * The **hidden action** τ represents an internal computation step.
- * The **hiding** operators τ_I , for $I \subseteq \text{Act}$, rename all actions of I in their argument into τ .


τ does not communicate with any action names.

Hiding operators can make actions **inert**.

Example: $\tau_{\{c\}}(a \cdot c \cdot b) = a \cdot \tau \cdot b = a \cdot b$

Not All Hidden Actions Are Inert

Example: A malfunctioning channel.



$$\tau_{\{c_2, c_3\}}(\partial_{\{s_5\}}(r_1 \cdot (c_2 \cdot s_4 + c_3 \cdot s_5))) = r_1 \cdot (\tau \cdot s_4 + \tau \cdot \delta) \neq r_1 \cdot s_4$$

Which Hidden Actions Are Inert? - Part I

$$a \cdot (b + \tau \cdot \delta) \neq a \cdot b$$

$$\partial_{\{c\}}(a \cdot (b + \tau \cdot c)) \neq \partial_{\{c\}}(a \cdot (b + c))$$

$$a \cdot (b + \tau \cdot c) \neq a \cdot (b + c)$$

Solution: A hidden action is only inert if it **does not lose possible behaviors**.

Example: $a \cdot (b + \tau \cdot (b + c)) = a \cdot (b + c)$



Branching Bisimulation Equivalence

s_1 and s_2 are **branching bisimilar** states, denoted by $s_1 \leftrightarrow_b s_2$, if:

- * if $s_1 \xrightarrow{\tau} s'_1$ is inert, then $s'_1 \leftrightarrow_b s_2$
- * a non-inert transition $s_1 \xrightarrow{a} s'_1$ (or $s_1 \downarrow$) is simulated by s_2 after zero or more inert τ -transitions: $s_2 \xrightarrow{\tau} \dots \xrightarrow{\tau} \hat{s}_2$, where $s_1 \leftrightarrow_b \hat{s}_2$ and $\hat{s}_2 \xrightarrow{a} s'_2$ with $s'_1 \leftrightarrow_b s'_2$ (or $\hat{s}_2 \downarrow$)

and **vice versa**.

Branching Bisimulation - Examples



Which Hidden Actions Are Inert? - Part II

Initial τ 's are not inert.

$$a \cdot (b + \tau \cdot c) \neq a \cdot (b + c)$$

$$\tau \cdot c \neq c$$

Solution: A hidden action is inert if it does not lose possible behaviors **and is not initial**.

Rooted Branching Bisimulation Equivalence

s_1 and s_2 are **rooted branching bisimilar**, denoted $s_1 \leftrightarrow_{rb} s_2$, if:

1. if $s_1 \xrightarrow{a} s'_1$, then $s_2 \xrightarrow{a} s'_2$ with $s'_1 \leftrightarrow_b s'_2$
2. if $s_1 \downarrow$, then $s_2 \downarrow$

and **vice versa**.

Hidden Action and Hiding - Axioms

$$x \cdot \tau = x$$

$$x \cdot (\tau \cdot (y + z) + y) = x \cdot (y + z)$$

$$\tau_I(\delta) = \delta$$

$$\tau_I(\tau) = \tau$$

$$\tau_I(a(\vec{d})) = a(\vec{d}) \quad \text{if } a \notin I$$

$$\tau_I(a(\vec{d})) = \tau \quad \text{if } a \in I$$

$$\tau_I(x + y) = \tau_I(x) + \tau_I(y)$$

$$\tau_I(x \cdot y) = \tau_I(x) \cdot \tau_I(y)$$

$$\tau_I(\sum_{d:D} P(d)) = \sum_{d:D} \tau_I(P(d))$$

Soundness and Completeness of the Axioms

Theorem: For process algebra terms p and q :

$$p = q \Leftrightarrow p \stackrel{\leftrightarrow}{\text{rb}} q$$


Fair Abstraction

τ -loops can be eliminated.

Example: The process X with

$$\begin{aligned} X &= a \cdot Y \\ Y &= \tau \cdot Y + b \end{aligned}$$

is rooted branching bisimilar to $a \cdot b$.



In the black state there is a fixed chance $\alpha > 0$ that the b -transition is taken. So the chance that b is eventually executed is 100%.

Overview

- * data types: (**sort func map var rew**)
- * action declaration: (**act** $a : D_1 \times \dots \times D_n$, **comm** $a \mid b = c$)
- * basic operators: ($+ \cdot$)
- * data-dependent operators: ($\langle b \rangle \sum_{d:D}$)
- * process declaration: (**proc** $X(d_1, \dots, d_n)$)
- * parallel operators: ($\parallel \perp \mid$)
- * deadlock and encapsulation: ($\delta \partial_H$)
- * hidden action and hiding: ($\tau \tau_I$)


In general, the **init** declaration of a μ CRL specification is of the form

$$\tau_I(\partial_H(X_1(d_1, \dots, d_n) \parallel \dots \parallel X_k(e_1, \dots, e_m)))$$

where the recursive equations for X_1, \dots, X_k use only data, actions, basic operators and data-dependent operators.

Example - One-bit Buffers in Sequence

act $r_1, s_2, r_3, s_3, c_3 : D$
comm $r_3 \mid s_3 = c_3$
proc $B_1 = \sum_{d:D} r_1(d) \cdot s_3(d) \cdot B_1$
 $B_2 = \sum_{d:D} r_3(d) \cdot s_2(d) \cdot B_2$
init $\tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel B_2))$



Buffers B_1 and B_2 of capacity one in sequence behave as a buffer of capacity two:

proc $X = \sum_{d:D} r_1(d) \cdot Y(d)$
 $Y(d:D) = \sum_{d':D} r_1(d') \cdot Z(d, d') + s_2(d) \cdot X$
 $Z(d:D, d':D) = s_2(d) \cdot Y(d')$
init X

Symbolic Proof Example: One-bit Buffers in Sequence

$$\begin{aligned} & B_1 \parallel B_2 && \text{(Summations } \Sigma_{d:D} \text{ and data parameters } d \text{ are omitted)} \\ = & B_1 \ll B_2 + B_2 \ll B_1 + B_1 | B_2 \\ = & (r_1 \cdot s_3 \cdot B_1) \ll B_2 + (r_3 \cdot s_2 \cdot B_2) \ll B_1 + (r_1 \cdot s_3 \cdot B_1) | (r_3 \cdot s_2 \cdot B_2) \\ = & r_1 \cdot ((s_3 \cdot B_1) \parallel B_2) + r_3 \cdot ((s_2 \cdot B_2) \parallel B_1) + \delta \cdot ((s_3 \cdot B_1) \parallel (s_2 \cdot B_2)) \\ = & r_1 \cdot ((s_3 \cdot B_1) \parallel B_2) + r_3 \cdot ((s_2 \cdot B_2) \parallel B_1) \end{aligned}$$

$$\begin{aligned} & \partial_{\{s_3, r_3\}}(B_1 \parallel B_2) \\ = & \partial_{\{s_3, r_3\}}(r_1 \cdot ((s_3 \cdot B_1) \parallel B_2) + r_3 \cdot ((s_2 \cdot B_2) \parallel B_1)) \\ = & \partial_{\{s_3, r_3\}}(r_1 \cdot ((s_3 \cdot B_1) \parallel B_2)) + \partial_{\{s_3, r_3\}}(r_3 \cdot ((s_2 \cdot B_2) \parallel B_1)) \\ = & \partial_{\{s_3, r_3\}}(r_1) \cdot \partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2) + \partial_{\{s_3, r_3\}}(r_3) \cdot \partial_{\{s_3, r_3\}}((s_2 \cdot B_2) \parallel B_1) \\ = & r_1 \cdot \partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2) + \delta \cdot \partial_{\{s_3, r_3\}}((s_2 \cdot B_2) \parallel B_1) \\ = & r_1 \cdot \partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2) \end{aligned}$$


One-bit Buffers in Sequence

Likewise we can derive:

$$\partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2) = c_3 \cdot \partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2))$$

$$\partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2)) = s_2 \cdot \partial_{\{s_3, r_3\}}(B_1 \parallel B_2) + r_1 \cdot \partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel (s_2 \cdot B_2))$$

$$\partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel (s_2 \cdot B_2)) = s_2 \cdot \partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2)$$




One-bit Buffers in Sequence

$$\begin{aligned}\tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel B_2)) &= \tau_{\{c_3\}}(r_1 \cdot \partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2)) \\ &= r_1 \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2)) \\ &= r_1 \cdot \tau_{\{c_3\}}(c_3 \cdot \partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2))) \\ &= r_1 \cdot \tau \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2))) \\ &= r_1 \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2)))\end{aligned}$$

Likewise we can derive:

$$\begin{aligned}\tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2))) &= s_2 \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel B_2)) \\ &\quad + r_1 \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel (s_2 \cdot B_2))) \\ \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel (s_2 \cdot B_2))) &= s_2 \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}((s_3 \cdot B_1) \parallel B_2)) \\ &= s_2 \cdot \tau_{\{c_3\}}(\partial_{\{s_3, r_3\}}(B_1 \parallel (s_2 \cdot B_2)))\end{aligned}$$

Alternating Bit Protocol



Data elements are sent from *Sender* to *Receiver* via a corrupted channel. *Sender* alternately attaches bit 0 or 1 to data elements.

If *Receiver* receives a datum, it sends the attached bit to *Sender* via a corrupted channel, to acknowledge reception. If *Receiver* receives an error message, then it resends the preceding acknowledgement.

Sender keeps sending a datum with attached bit b until it receives acknowledgement b . Then it starts sending the next datum with attached bit $1-b$ until it receives acknowledgement $1-b$, etc.

Alternating Bit Protocol - Sender and Receiver

S_b sends a datum with bit b attached:

$$S_b = \sum_{d:\Delta} r_A(d) \cdot s_B(d, b) \cdot T_{db}$$

$$\begin{aligned} T_{db} &= r_F(b) \cdot S_{1-b} \\ &+ (r_F(1-b) + r_F(\perp)) \cdot s_B(d, b) \cdot T_{db} \end{aligned}$$

R_b expects to receive a datum with b attached:

$$\begin{aligned} R_b &= \sum_{d:\Delta} r_C(d, b) \cdot s_D(d) \cdot s_E(b) \cdot R_{1-b} \\ &+ \sum_{d:\Delta} r_C(d, 1-b) \cdot s_E(1-b) \cdot R_b \\ &+ r_C(\perp) \cdot s_E(1-b) \cdot R_b \end{aligned}$$

Alternating Bit Protocol - Channels

K and L represent the corrupted channels, to model *asynchronous* communication. (The action j represents an internal choice.)

$$K = \sum_{d:\Delta} \sum_{b:\{0,1\}} r_B(d, b) \cdot (j \cdot s_C(d, b) + j \cdot s_C(\perp)) \cdot K$$

$$L = \sum_{b:\{0,1\}} r_E(b) \cdot (j \cdot s_F(b) + j \cdot s_F(\perp)) \cdot L$$

A send and a read action of the same message over the same internal channel communicate:

$$s_B \mid r_B = c_B \quad s_C \mid r_C = c_C \quad s_E \mid r_E = c_E \quad s_F \mid r_F = c_F$$

Alternating Bit Protocol - Initial State

The **initial state** of the alternating bit protocol is specified by

$$\tau_I(\partial_H(R_0 \parallel S_0 \parallel K \parallel L))$$

with H the set of read and send actions over channels B, C, E and F, and I the set of communication actions together with j .

$\tau_I(\partial_H(R_0 \parallel S_0 \parallel K \parallel L))$ exhibits the desired **external behavior**

$$X = \sum_{d:\Delta} r_A(d) \cdot s_D(d) \cdot X$$

Question: What is the behavior of $\tau_I(\partial_H(R_1 \parallel S_0 \parallel K \parallel L))$?

Question

What is the state space of

$$\tau_I(\partial_H(R_0 \parallel S_0 \parallel K \parallel L))$$

for $\Delta = \{d_1, d_2\}$, after minimization modulo \leftrightarrow_b ?

Alternating Bit Protocol - Symbolic Proof

A **symbolic** correctness proof of the alternating bit protocol, for any data set Δ , was checked with the theorem prover Coq.

(Bezem & Groote, 1993)

Question


Which assumptions underlying the alternating bit protocol are unrealistic or impractical?

Alternating Bit Protocol - Unrealistic Assumptions

The alternating bit protocol makes three unrealistic assumptions:

- ▶ Unbounded number of retries
- ▶ Messages are never lost (and error messages can be recognized)
- ▶ Poor use of available bandwidth

Sliding Window Protocol



The sliding window protocol has better use of available bandwidth.

A. Tanenbaum, *Computer Networks* (Chapter 4.2), Prentice Hall, 1981

B. Badban, W. Fokkink, J.F. Groote, J. Pang and J. van de Pol
Verification of a sliding window protocol in μ CRL and PVS
Formal Aspects of Computing, 17(3):342-388, 2005

Model Checking Versus Symbolic Correctness Proofs

In *model checking*, the state space is generated, and logical properties are checked automatically.

- ▶ automated, so convenient to use
- ▶ expressive logic for specifying properties
- ▶ entire state space is searched
- ▶ suffers from state explosion
- ▶ works only for fixed data sets and topologies

Symbolic correctness proofs can be supported by a *theorem prover*.

- ▶ laborious
- ▶ no generation of the state space
- ▶ provides general correctness proof

Model Checking Example: A Distributed Lift System



J. Pang, B. Karstens and W. Fokkink

Analyzing the redesign of a distributed lift system in UPPAAL
Proc. ICFEM'03, Singapore, Lecture Notes in Computer Science
2885, Springer, 2003

Linear Process Equation

A **linear process equation (LPE)** is a symbolic representation of a state space.

$$X(d:D) = \sum_{i:I} \sum_{e:E} a_i(f_i(d, e)) \cdot X(g_i(d, e)) \triangleleft h_i(d, e) \triangleright \delta$$

with

- ▶ $a_i \in \text{Act} \cup \{\tau\}$
- ▶ $f_i : D \times E \rightarrow D_i$
- ▶ $g_i : D \times E \rightarrow D$
- ▶ $h_i : D \times E \rightarrow \text{Bool}$

Linearization - Type I and II Equations

Two types of recursive equations $X(d_1:D_1, \dots, d_n:D_n) = p$ are distinguished:

- I p contains only $\cdot, +, \langle b \rangle, \sum_{d:D}$
- II p also contains $\parallel, \partial_H, \tau_I$

The μ CRL **lineariser** requires that all recursion variables with a recursive equation of **type II** can be substituted away from right-hand sides of recursive equations and from the initial declaration.

Linearization

First the type I recursive equations are linearized, in two steps:

- ▶ Turn them into **Greibach Normal Form**, by replacing “non-initial” actions in right-hand sides of recursive equations into fresh recursion variables.
- ▶ Linearize the resulting recursive equations using a **stack**.

(An alternative method uses **pattern matching**.)

Then all (type I and II) recursive equations are transformed into a single LPE, by eliminating parallel, encapsulation, hiding and renaming operators from right-hand sides of recursive equations and from the initial declaration.

Linearization of Type I Equations - Example

First we explain, by an example, the standard linearization method for type I equations (**mcr1**).

Example: $Y = a \cdot Y \cdot b + c$

Y performs k a 's, then a c , and then k b 's, for any $k \geq 0$.

Step 1: Make a **Greibach Normal Form**.

$$\begin{aligned} Y &= a \cdot Y \cdot Z + c \\ Z &= b \end{aligned}$$

Linearization of Type I Equations - Example

Step 2: Linearization using a stack.

Lists can contain *recursion variables* and their *data parameters* (i.e., function symbols, brackets, comma's).

Empty list $[]$ and $in : D \times List \rightarrow List$ are the constructors of *List*.

$empty : List \rightarrow Bool$ and $head : List \rightarrow D$ and $tail : List \rightarrow List$ are standard operations on lists.

Linearization of Type I Equations - Example

$$\begin{aligned} Y &= a \cdot Y \cdot Z + c \\ Z &= b \end{aligned}$$

is transformed into

$$\begin{aligned} X(\lambda:List) &= a \cdot X(in(Y, in(Z, tail(\lambda)))) \triangleleft eq(head(\lambda), Y) \triangleright \delta \\ &+ (c \triangleleft empty(tail(\lambda)) \triangleright c \cdot X(tail(\lambda))) \triangleleft eq(head(\lambda), Y) \triangleright \delta \\ &+ (b \triangleleft empty(tail(\lambda)) \triangleright b \cdot X(tail(\lambda))) \triangleleft eq(head(\lambda), Z) \triangleright \delta \end{aligned}$$

If recursion variables carry data parameters, then function symbols, brackets and comma's are also pushed on the stack.

Here Y, Z do not carry data parameters, so D contains only Y, Z .

Disadvantage: The stack gives a lot of overhead.

Linearization with Pattern Matching - Example

`mcrl -regular` invokes another linearization algorithm for type I equations, which is based on **pattern matching**.

When the state space is finite, `mcrl -regular` usually works better than `mcrl`. But `mcrl -regular` does not always terminate.

Example:

$$Y = a \cdot Z \cdot Y$$

$$Z = b \cdot Z + b$$

Y repeatedly performs an a followed by one or more b 's.

Linearization with Pattern Matching - Example

Step 1: Replace $Z \cdot Y$ by a fresh recursion variable X .

$$\begin{aligned} Y &= a \cdot X \\ Z &= b \cdot Z + b \\ X &= Z \cdot Y \end{aligned}$$

Step 2: Expand Z in the right-hand side of X . (Store that $X = Z \cdot Y$.)

$$\begin{aligned} Y &= a \cdot X \\ Z &= b \cdot Z + b \\ X &= b \cdot Z \cdot Y + b \cdot Y \end{aligned}$$

Step 3: Replace $Z \cdot Y$ by X in the right-hand side of X .

$$\begin{aligned} Y &= a \cdot X \\ Z &= b \cdot Z + b \\ X &= b \cdot X + b \cdot Y \end{aligned}$$

Linearization with Pattern Matching - Example

$$X(n:\text{Nat}) = a(n) \cdot b(S(n)) \cdot c(S(S(n))) \cdot X(S(S(S(n))))$$

$$X(n:\text{Nat}) = a(n) \cdot Y(n)$$

$$Y(n:\text{Nat}) = b(S(n)) \cdot c(S(S(n))) \cdot X(S(S(S(n))))$$

$$X(n:\text{Nat}) = a(n) \cdot Y(n)$$

$$Y(n:\text{Nat}) = b(S(n)) \cdot Z(n)$$

$$Z(n:\text{Nat}) = c(S(S(n))) \cdot X(S(S(S(n))))$$

Linearization with Pattern Matching - Non-termination

`mcr1 -regular` does not always terminate.

Example:

$$\begin{array}{l} Y = a \cdot Y \cdot b + c \\ \hline Y = a \cdot X_1 + c \\ X_1 = Y \cdot b \\ \hline Y = a \cdot X_1 + c \\ X_1 = a \cdot X_1 \cdot b + c \cdot b \\ \hline Y = a \cdot X_1 + c \\ X_1 = a \cdot X_2 + c \cdot Z_1 \\ X_2 = X_1 \cdot b \\ Z_1 = b \\ \hline Y = a \cdot X_1 + c \\ X_1 = a \cdot X_2 + c \cdot Z_1 \\ X_2 = a \cdot X_2 \cdot b + c \cdot Z_1 \cdot b \\ Z_1 = b \\ \hline \vdots \end{array}$$

Linearization of Type II Equations - Example

We show, by an example, how to reduce the parallel composition of LPEs to an LPE.

Example: Let $a|b = c$, and

$$X(n:\text{Nat}) = a(n) \cdot X(S(n)) \triangleleft n < 10 \triangleright \delta + b(n) \cdot X(S(S(n))) \triangleleft n > 5 \triangleright \delta$$

$Y(m:\text{Nat}, n:\text{Nat}) = X(m) \parallel X(n)$ can be linearized to:

$$\begin{aligned} Y(m:\text{Nat}, n:\text{Nat}) = & a(m) \cdot Y(S(m), n) \triangleleft m < 10 \triangleright \delta \\ & + b(m) \cdot Y(S(S(m)), n) \triangleleft m > 5 \triangleright \delta \\ & + a(n) \cdot Y(m, S(n)) \triangleleft n < 10 \triangleright \delta \\ & + b(n) \cdot Y(m, S(S(n))) \triangleleft n > 5 \triangleright \delta \\ & + c(m) \cdot Y(S(m), S(S(n))) \triangleleft m < 10 \wedge n > 5 \wedge \text{eq}(m, n) \triangleright \delta \\ & + c(n) \cdot Y(S(S(m)), S(n)) \triangleleft m > 5 \wedge n < 10 \wedge \text{eq}(m, n) \triangleright \delta \end{aligned}$$

How can an application of τ_I to an LPE be reduced to an LPE?

How can an application of ∂_H to an LPE be reduced to an LPE?

State Space Generation

From an LPE $X(d:D)$ and initial state d_0 , the state space is generated (instantiator). The algorithm below focuses on finding reachable states (i.e., transitions are ignored).

M contains all “explored” states, and L the generated states that still need to be explored.

Initially, $L = \{d_0\}$ and $M = \emptyset$.

while $L \neq \emptyset$ **do**

 select $d \in L$; $L := L \setminus \{d\}$; $M := M \cup \{d\}$

 from LPE X , compute each transition $d \xrightarrow{a} d'$

if $d' \notin L \cup M$ **then** $L := L \cup \{d'\}$

Challenges: Store large state spaces in memory.

 Check efficiently whether $d' \notin L \cup M$.


Hash Tables

A (random) **hash function** h maps a large domain to a small one, to allow fast lookups:

$$h : D \rightarrow \text{Hash values}$$

Problem: Different states may map to the same hash value.

Solution: A **chained** hash table.



When the hash table gets full, blocks of states from the hash table are swapped to disk (e.g. based on “age”).

Bloom Filter

If a generated state d' is *not* in the hash table, the check $d' \notin L \cup M$ requires an expensive disk lookup.

A **Bloom filter** allows an inexpensive check whether $d' \notin L \cup M$, allowing for **false positives**.

For some (smartly chosen) k, m , fix different hash functions $h_1, \dots, h_k : D \rightarrow \{1, \dots, m\}$.

A Bloom filter is a **bit array** of length m .

Initially, all bits are set to 0.

For each generated state d , set the bits in the Bloom filter at positions $h_1(d), \dots, h_k(d)$ to 1.

Bloom Filter

If a state d' is generated, and does *not* occur at entry $h(d')$ in the hash table, then check if positions $h_i(d')$ for $i = 1, \dots, k$ in the Bloom filter all contain 1.

If not, then $d' \notin L \cup M$.

Else, still an expensive disk lookup is required.

Bloom Filter - Analysis

When n elements have been inserted in $L \cup M$, the possibility that a certain position in the Bloom filter contains 0 is

$$\left(\frac{m-1}{m}\right)^{kn}$$

So the probability that k positions in the Bloom filter all contain 1 is

$$\left(1 - \left(\frac{m-1}{m}\right)^{kn}\right)^k$$

For given m, n , the number of false positives are minimal for $k \approx 0.7 \cdot \frac{m}{n}$.
(Typically, $k = 4$ and 256 MB is given to the Bloom filter.)

Bitstate Hashing

In **bitstate hashing**, a non-chained hash table is maintained.

No extra disk space is used.

If two generated states happen to have the same hash value, the old entry is overwritten by the new entry.

Bitstate hashing approximates an *exhaustive* search for small systems, and slowly changes into a *partial* search for large systems.

Distributed Verification

Store the state space on a cluster of computers (e.g. DAS-3).

Let there be a **globally known hash function**.

States are divided over processors on the basis of their hash values.

When a state is generated at a processor, its hash value is calculated, and the state is forwarded to the appropriate processor. There it is determined whether the state was generated before.

Distributed versions exist of:

- ▶ state space generation
- ▶ minimization modulo \leftrightarrow_b
- ▶ model checking

Challenge: Perform these tasks efficiently, with as little communication overhead as possible.

Distributed Verification - Example

We model checked a **cache coherence protocol** for a *distributed shared memory implementation* of Java.

For 2 processors, each with 1 thread, the protocol is correct.

For 2 processors, one with 1 and one with 2 threads, with **distributed model checking**, we detected a **deadlock**.

Namely, while a thread is waiting for the write lock of a region, the home node of the region may migrate to the thread's processor, so that the thread actually accesses the region at home.

J. Pang, W. Fokkink, R. Hofman and R. Veldema

Model checking a cache coherence protocol for a Java DSM implementation

Journal of Logic and Algebraic Programming, 71(1):1-43, 2007

Sorted Lists to Fight State Explosion

Storing the **message buffers** of processes, and **messages in channels**, in a **sorted** list, reduces the number of states considerably.

Impose a **total order** $<$ on the data type D .

Model Checking

We define some basic **modal** logic operators to express properties of states.

$$\phi ::= \mathbf{T} \mid \mathbf{F} \mid \phi \wedge \phi' \mid \phi \vee \phi' \mid \langle \mathbf{a} \rangle \phi \mid [\mathbf{a}] \phi$$

where a ranges over $\text{Act} \cup \{\tau\}$.

- ▶ \mathbf{T} holds in all states, and \mathbf{F} in no state
- ▶ \wedge denotes conjunction, and \vee disjunction
- ▶ $\langle \mathbf{a} \rangle \phi$ holds in state s if there is a transition $s \xrightarrow{a} s'$ such that ϕ holds in state s'
- ▶ $[\mathbf{a}] \phi$ holds in state s if for each transition $s \xrightarrow{a} s'$, ϕ holds in state s'

Question

Does $\langle a \rangle \phi$ imply $[a] \phi$?

Does $[a] \phi$ imply $\langle a \rangle \phi$?

Model Checking

The states s that satisfy a formula ϕ , denoted $s \models \phi$, are defined inductively by:

$$s \models \top$$

$$s \not\models \text{F}$$

$$s \models \phi \wedge \phi' \quad \text{if } s \models \phi \text{ and } s \models \phi'$$

$$s \models \phi \vee \phi' \quad \text{if } s \models \phi \text{ or } s \models \phi'$$

$$s \models \langle a \rangle \phi \quad \text{if for some state } s', s \xrightarrow{a} s' \text{ with } s' \models \phi$$

$$s \models [a] \phi \quad \text{if for all states } s', s \xrightarrow{a} s' \text{ implies } s' \models \phi$$

Example: If $s \xrightarrow{a}$, then $s \models [a] \text{F}$ and $s \not\models \langle a \rangle \top$.

Fixpoints

Let D be a *finite* set with **partial ordering** \leq , with a *least* and a *greatest* element.

$S : D \rightarrow D$ is **monotonic** if $d \leq e$ implies $S(d) \leq S(e)$.

$d \in D$ is a **fixpoint** of $S : D \rightarrow D$ if $S(d) = d$.

If S is monotonic, then it has a **minimal fixpoint** $\mu X.S(X)$ and a **maximal fixpoint** $\nu X.S(X)$.

Question: How can $\mu X.S(X)$ and $\nu X.S(X)$ be computed?

Question

Give an example to show that if D is *infinite*, monotonic mappings $S : D \rightarrow D$ need not have a fixpoint.

The μ -calculus is a temporal logic.

$$\phi ::= \top \mid \text{F} \mid \phi \wedge \phi' \mid \phi \vee \phi' \mid \langle a \rangle \phi \mid [a] \phi \mid X \mid \mu X. \phi \mid \nu X. \phi$$

where the X are recursion variables.

We restrict to closed formulas, meaning that each occurrence of a recursion variable X is within the scope of a μX or νX .

We need to explain how ϕ (with X as only free variable) is interpreted as a (monotonic) mapping from sets of states to sets of states.

Consider a *finite* state space.

(For simplicity we ignore successful termination.)

Let X be the *only free variable* in μ -calculus formula ϕ .

We define the meaning of $\mu X.\phi$ and $\nu X.\phi$.

ϕ maps each set P of states to the set of states that satisfy ϕ , under the assumption that P is the set of states in which X holds.

Example: Consider the state space $s_0 \xrightarrow{a} s_1$.

$\langle a \rangle X$ maps sets containing s_1 to $\{s_0\}$, and all other sets to \emptyset .

$[a] X$ maps sets containing s_1 to $\{s_0, s_1\}$, and all other sets to $\{s_1\}$.

As partial order we take **set inclusion**.

Theorem: For each ϕ with one free variable X ,
the corresponding mapping is **monotonic**.

So the closed formulas $\mu X.\phi$ and $\nu X.\phi$ are **well-defined**.

They are satisfied only by the states in
the minimal and maximal fixpoint of ϕ , respectively.

μ -calculus - Examples

$\mu X.(\langle a \rangle X \vee \langle b \rangle T)$ represents those states that can execute $a^k b$ for some $k \geq 0$.

$\nu X.(\langle a \rangle X \vee \langle b \rangle T)$ represents those states that can execute a^∞ or $a^k b$ for some $k \geq 0$.


$\nu X.(\langle a \rangle X \vee \langle b \rangle X)$ represents those states that can execute an infinite trace of a 's and b 's.

Question: How about $\mu X.(\langle a \rangle X \vee \langle b \rangle X)$?

μ -calculus - Negation Violates Monotonicity

Absence of negation in the μ -calculus is needed for monotonicity.

Example:



$\mu X. \neg \langle a \rangle X$ has no fixpoint.

μ -calculus - Complexity

Worst-case time complexity: $O(|\phi| \cdot m \cdot n^{N(\phi)})$

where $N(\phi)$ is the longest chain of nested fixpoints in ϕ .

Example:

$$s_0 \begin{array}{c} \xrightarrow{a} \\ \xleftarrow{a} \end{array} s_1$$

Consider $\nu X. \langle b \rangle (\mu Y. (\langle a \rangle X \vee \langle a \rangle Y))$.

| Y | X |
|----------------|----------------|
| \emptyset | $\{s_0, s_1\}$ |
| $\{s_0, s_1\}$ | \emptyset |
| \emptyset | \emptyset |

In the second iteration, recomputation of Y *must* start at \emptyset (instead of $\{s_0, s_1\}$).

Conclusion: If a minimal fixpoint μY is within the scope of a maximal fixpoint νX , the successive values of Y must be recomputed starting at \emptyset every time.

Alternation-free μ -calculus

For two nested minimal (or maximal) fixpoints, recomputing a fixpoint is not so expensive.

Example: $s_0 \xrightarrow{a} s_1 \xrightarrow{b} \dots \xrightarrow{a} s_{2n-3} \xrightarrow{b} s_{2n-2} \xrightarrow{a} s_{2n-1} \xrightarrow{b} s_{2n}$

Consider $\nu X. \nu Y. (\langle a \rangle X \vee \langle b \rangle Y)$.

| Y | X |
|----------------------------|----------------------------|
| $\{s_0, \dots, s_{2n}\}$ | $\{s_0, \dots, s_{2n}\}$ |
| $\{s_0, \dots, s_{2n-1}\}$ | $\{s_0, \dots, s_{2n-2}\}$ |
| $\{s_0, \dots, s_{2n-3}\}$ | $\{s_0, \dots, s_{2n-4}\}$ |
| \vdots | \vdots |
| \emptyset | \emptyset |

Note that the successive values of X and Y decrease.

This is always true for two nested maximal fixpoints. Likewise, for two nested minimal fixpoints, the successive values always increase.

Alternation-free μ -calculus

Worst-case time complexity: $O(|\phi| \cdot m \cdot n^{N(\phi)})$

for model checking the μ -calculus, where $N(\phi)$ is the longest chain of nested *alternating* fixpoints in ϕ (i.e., minimal within maximal, or maximal within minimal fixpoint).

Worst-case time complexity: $O(|\phi| \cdot m \cdot n)$

for model checking the **alternation-free** μ -calculus.

Model checking the *full* μ -calculus is in **NP** \cap **co-NP**.

It is an open question whether it is in **P**.

Regular μ -calculus

$\alpha ::= \top \mid a \mid \neg\alpha \mid \alpha \wedge \alpha' \quad (a \in \text{Act} \cup \{\tau\})$

$\beta ::= \alpha \mid \beta \cdot \beta' \mid \beta \mid \beta' \mid \beta^*$

$\phi ::= \top \mid \text{F} \mid \phi \wedge \phi' \mid \phi \vee \phi' \mid \langle \beta \rangle \phi \mid [\beta] \phi \mid X \mid \mu X. \phi \mid \nu X. \phi$

α represents a *set of actions*: \top denotes all actions, a the set $\{a\}$, \neg complement, and \wedge intersection.

β represents a *set of traces*: \cdot is concatenation, \mid union, and $*$ iteration.

Regular μ -calculus - Examples

Deadlock freeness: $[T^*] \langle T \rangle T$

Absence of *error*: $[T^* \cdot \text{error}] F$

After an occurrence of *send*, **fair** reachability of *read* is guaranteed:

$$[T^* \cdot \text{send} \cdot (\neg \text{read})^*] \langle T^* \cdot \text{read} \rangle T$$

Question: Specify the properties:

- there is an execution sequence to a deadlock state
- *read* cannot be executed before an occurrence of *send*

Regular μ -calculus - Examples

There is an infinite execution sequence:

$$\nu X.(\langle T \rangle X)$$

No reachable state exhibits an infinite τ -sequence:

$$[T^*] \mu X. [\tau] X$$

Each *send* is eventually followed by a *read*:

$$[T^* \cdot \text{send}] \mu X. (\langle T \rangle T \wedge [\neg \text{read}] X)$$

CADP supports model checking of **alternation-free**, **regular** μ -calculus formulas.

Classes of actions can be expressed with the use of **UNIX regular expressions**, e.g. `'send(.*)'`.

If a property is violated, an *error trace* is produced.

To analyse the error trace, omit the hiding operator from the initial state before state space generation.

CADP Syntax for Formulas

Examples: `[(not "send(d)")*. "recv(d)"] false`

`[(not 'send(.*)')*. 'recv(.*)']*] false`

`[(true)*. "send(d)"] mu X. ((true) true and [not "read(d)"] X)`

Beware that text between quotes ("`...`") is interpreted literally (even "`a(d,e)`" and "`a(d, e)`" are taken to be syntactically different!).

A small typo in an action name may therefore mean you verify a property for a non-existent action.

Self-loops to Verify State Properties

To take values of variables into account in a model checking analysis, one can include artificial **self-loops** that carry these variables as action variables.

Example: To the process declaration of a recursion variable $X(d_1:D_1, d_2:D_2, d_3:D_3)$ one can add a summand

$$+ \text{test}(d_1, d_3) \cdot X(d_1, d_2, d_3)$$

where *test* is a “fresh” action name.

Such a self-loop does not increase the number of reachable states.

Types of Requirements

- ▶ *Safety*: something bad will never happen.

(E.g., when motor M1 is on, brake B1 is never applied.)

- ▶ *Liveness*: something good will eventually happen.

(E.g., if the system is in uncalibrated mode, the bed is not in the uppermost position, and the Up button is pressed, then the bed must go up.)

Pitfalls for Requirements

Beware not to formulate requirements that are **too general**.

Example: *“the bed can move up, down, left or right”*

(In which states of the system, under which inputs?)

Requirements must be formulated in terms of **external events** (input/output).

Example: *“the controllers must communicate asynchronously”*

(Implementation detail, cannot be verified using model checking.)

Minimization Algorithm

Assume a finite state space.

Let the set S of states be partitioned into $P_1 \cup \dots \cup P_k$, such that
(*) *branching bisimilar states reside in the same set of the partition.*

For $a \in \text{Act} \cup \{\tau\}$, $s_0 \in \text{split}_a(P_i, P_j)$ if

$s_0 \xrightarrow{\tau} \dots \xrightarrow{\tau} s_{n-1} \xrightarrow{a} s_n$ with $s_0, \dots, s_{n-1} \in P_i$ and $s_n \in P_j$.

If $s \in \text{split}_a(P_i, P_j)$ and $s' \in P_i \setminus \text{split}_a(P_i, P_j)$ with $a \neq \tau$ or $i \neq j$,
then $s \not\leftrightarrow_b s'$.

So after performing a split, (*) remains satisfied.

Minimization Algorithm

Initially, S is partitioned in S . (So $(*)$ is trivially satisfied.)

Suppose that at some point S is partitioned in $P_1 \cup \dots \cup P_k$.

If $a \neq \tau$ or $i \neq j$, and

$$\emptyset \subset \text{split}_a(P_i, P_j) \subset P_i$$

then in the partition, P_i can be replaced by $\text{split}_a(P_i, P_j)$ and $P_i \setminus \text{split}_a(P_i, P_j)$.

Splitting continues until no further split is possible.

$(*)$ is satisfied by the final partition.

Minimization Algorithm - Correctness and Complexity

Let $s \mathcal{B} s'$ if s and s' are in the same set of the final partition.

\mathcal{B} is a branching bisimulation relation.

Theorem: Let $P_1 \cup \dots \cup P_k$ denote the final partition of S .




Two states s and s' are in the same set P_i if and only if $s \leftrightarrow_b s'$ in the original state space.

Worst-case time complexity: $O(mn)$, where m is the number of transitions and n the number of states in the original state space.

Namely, calculating a split takes $O(m)$, and there are no more than n splits.

Minimization Algorithm - Example


Consider $(a \cdot \tau + \tau \cdot b) \cdot \delta$. P contains all four states in the state space.



$split_a(P, P)$ separates s_0 from $\{s_1, s_2, s_3\}$.

$split_b(P_2, P_2)$ separates s_2 from $\{s_1, s_3\}$.

P_1 , P_{21} and P_{22} cannot be split any further, so the minimized state space is



Questions

How can we, in the previous example, split on b followed by a split on a ?

How is $(a \cdot \tau + \tau \cdot a) \cdot \delta$ minimized?

How is $a \cdot a \cdot \delta$ minimized?

Bounded Retransmission Protocol

Data **packets** are sent from RC to TV.

The **last** datum of a packet is labeled.

A datum may get **lost**.

T_1 and T_2 send **time-out** messages.


An alternating bit is attached to each datum.

Only **one** kind of acknowledgement.


If RC does not receive an acknowledgement within a certain **time**, it resends the datum.

A datum is resent a **limited** number of times.

If TV does not receive a next datum within a certain **time**, RC has given up transmission.



Bounded Retransmission Protocol - External Behavior



Messages into channel A:

$s_A(I_{OK})$: transmission was successful

$s_A(I_{NOK})$: transmission was unsuccessful

$s_A(I_{DK})$: transmission *may* have been (un)successful

Bounded Retransmission Protocol - Remote Control

Let Λ consist of lists over Δ . (Only lists of length ≥ 2 can be transmitted.)

$$X = \sum_{\lambda:\Lambda} r_A(\lambda) \cdot Y(\lambda, 0, S(0)) \triangleleft \text{length}(\lambda) > S(0) \triangleright \delta$$

$$Y(\lambda:\Lambda, b:\text{Bit}, n:\text{Nat}) =$$

$$(s_B(\text{head}(\lambda), b) \triangleleft \text{length}(\lambda) > S(0) \triangleright s_B(\text{head}(\lambda), b, \text{last})) \cdot Z(\lambda, b, n)$$

$$Z(\lambda:\Lambda, b:\text{Bit}, n:\text{Nat}) =$$

$$r_F(\text{ack}) \cdot (Y(\text{tail}(\lambda), 1-b, S(0)) \triangleleft \text{length}(\lambda) > S(0) \triangleright s_A(I_{OK}) \cdot X)$$

$$+ r_G(\text{to}) \cdot (Y(\lambda, b, S(n)) \triangleleft n < \text{max} \triangleright$$

$$(s_A(I_{NOK}) \triangleleft \text{length}(\lambda) > S(0) \triangleright s_A(I_{DK})) \cdot s_H(\text{to}) \cdot X)$$

Bounded Retransmission Protocol - Television

$$\begin{aligned}V &= \sum_{d:\Delta} r_C(d, 0) \cdot s_D(d, \text{first}) \cdot s_E(\text{ack}) \cdot W(1) \\ &+ \sum_{d:\Delta} (r_C(d, 0, \text{last}) + r_C(d, 1, \text{last})) \cdot s_E(\text{ack}) \cdot V \\ &+ r_H(\text{to}) \cdot V\end{aligned}$$

$$\begin{aligned}W(b:\text{Bit}) &= \sum_{d:\Delta} r_C(d, b) \cdot s_D(d) \cdot s_E(\text{ack}) \cdot W(1-b) \\ &+ \sum_{d:\Delta} r_C(d, b, \text{last}) \cdot s_D(d, \text{last}) \cdot s_E(\text{ack}) \cdot V \\ &+ \sum_{d:\Delta} r_C(d, 1-b) \cdot s_E(\text{ack}) \cdot W(b) \\ &+ r_H(\text{to}) \cdot V\end{aligned}$$

Bounded Retransmission Protocol - Medium

$$K = \sum_{d:\Delta} \sum_{b:\{0,1\}} (r_B(d, b) \cdot (j \cdot s_C(d, b) + j \cdot s_G(to))) \cdot K \\ + r_B(d, b, last) \cdot (j \cdot s_C(d, b, last) + j \cdot s_G(to)) \cdot K$$

$$L = r_E(ack) \cdot (j \cdot s_F(ack) + j \cdot s_G(to)) \cdot L$$

Bounded Retransmission Protocol - Initial State

The **initial state** is specified by


$$\tau_I(\partial_H(V \parallel X \parallel K \parallel L))$$

with H the internal read and send actions,
and I the communication actions and j .

Bounded Retransmission Protocol - External Behavior

$$\Delta = \{d_1, d_2\}$$

Λ consists of lists of length 3



Bounded Retransmission Protocol - Exercises

- ▶ An incomplete specification of the BRP is available at

<http://www.cs.vu.nl/~tcs/pv/brp-scheme>

You need to supply rewrite rules for the functions *smaller*, *head*, *tail* and *length*, and specify the TV.

- ▶ Generate the minimized state space for $\Delta = \{d_1, d_2\}$.
- ▶ Formulate the next properties in regular μ -calculus:
 - Each action $r_A(\ell)$ from the remote control is eventually followed by an action $s_A(I_{OK})$ or $s_A(I_{NOK})$ or $s_A(I_{DK})$ from the remote control.
 - Each action $s_A(I_{OK})$ from the remote control is preceded by an action $s_D(d, last)$ from the television.

Verify these two properties using the CADP model checker.

Bounded Retransmission Protocol - Exercises

- ▶ Suppose that the timer T_1 is absent.

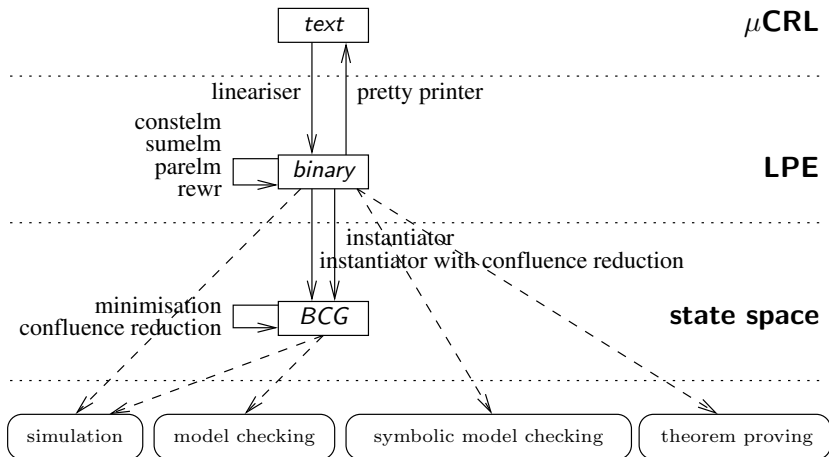
Verify using the CADP toolset that this version of the BRP contains a deadlock.

- ▶ What would go wrong if the timer T_2 were absent from the BRP?

Would the resulting system contain a deadlock?

Analyse the state space of this system using the CADP toolset.

Overview of the μ CRL Toolset



Automated Theorem Prover

A **theorem prover** within the μ CRL toolset provides (semi-)automated support for proving (large) formulas.

This theorem prover is *not* complete!
(Equalities over an abstract data type are in general undecidable.)

If the theorem prover cannot prove validity of a formula, diagnostics are provided. The user can add equations to the data specification.

In some cases, the formula is not valid in all states of the system, but does hold in all reachable states.

The user may supply **invariants**, and formulas can be proved under the assumption of invariants.

Such invariants must be proved separately, which can again be done using the theorem prover.

Download μ CRL, CADP and mCRL2

μ CRL can be downloaded at

`http://homepages.cwi.nl/~mcr1`

CADP can be downloaded at

`http://www.inrialpes.fr/vasy/cadp/`

mCRL2, a successor of μ CRL with **functional data types** that is being developed at Eindhoven University of Technology, can be downloaded at

`http://www.mcr12.org`

Other Useful Links

At the web site of my course *Protocol Validation*

<http://www.cs.vu.nl/~tcs/pv/>

you can find exercises, lab assignments and solutions.

My book *Modelling Distributed Systems* can be downloaded at

<http://www.springer.com>

Basic framework

- ▶ abstract data types
- ▶ process algebra
- ▶ branching bisimilarity

State space

- ▶ linearisation / state space generation
- ▶ minimisation of the state space
- ▶ regular μ -calculus / model checking

Fighting state space explosion

- ▶ store state space on a cluster of computers
- ▶ partial order reduction
- ▶ abstraction
- ▶ symbolic model checking
- ▶ theorem prover

Symbolic proof techniques

- ▶ axioms
- ▶ linear process equations
- ▶ invariants
- ▶ cones and foci