

# Model Checking of Fault-Tolerant Distributed Algorithms

## Part III: Parameterized Model Checking of Fault-tolerant Distributed Algorithms by Abstraction

Annu Gmeiner   Igor Konnov   Ulrich Schmid  
Helmut Veith   Josef Widder



Uni Salzburg, June 2015

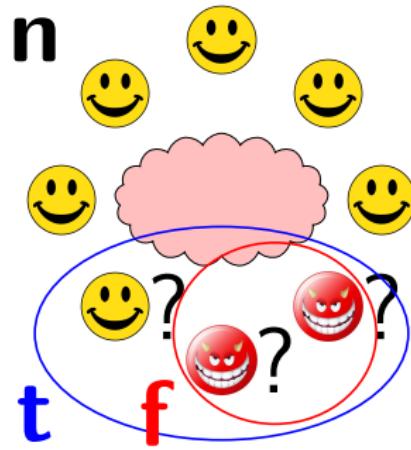
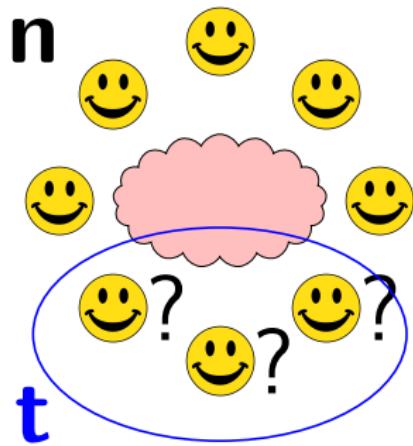
# Fault-tolerant DAs: Model Checking Challenges

- unbounded data types
  - counting how many messages have been received
- parameterization in multiple parameters
  - among  $n$  processes  $f \leq t$  are faulty with  $n > 3t$
- contrast to concurrent programs
  - fault tolerance against adverse environments
- degrees of concurrency
  - many degrees of partial synchrony
- continuous time
  - fault-tolerant clock synchronization

# Model checking problem for fault-tolerant DA algorithms

Parameterized model checking problem:

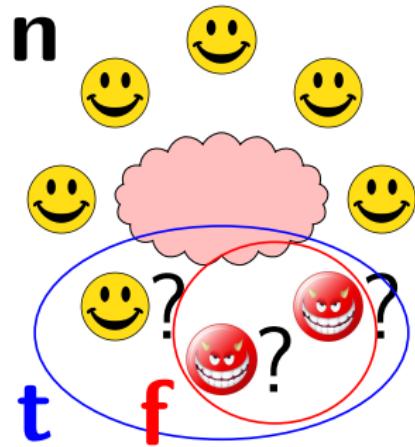
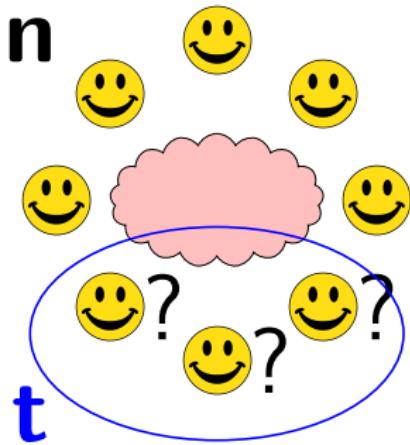
- given a distributed algorithm and spec.  $\varphi$
- show for all  $n$ ,  $t$ , and  $f$  satisfying  $n > 3t \wedge t \geq f \geq 0$   
 $M(n, t, f) \models \varphi$
- every  $M(n, t, f)$  is a system of  $n - f$  correct processes



# Model checking problem for fault-tolerant DA algorithms

Parameterized model checking problem:

- given a distributed algorithm and spec.  $\varphi$
- show for all  $n$ ,  $t$ , and  $f$  satisfying **resilience condition**  
 $M(n, t, f) \models \varphi$
- every  $M(n, t, f)$  is a system of  $N(n, f)$  correct processes



# Properties in Linear Temporal Logic

**Unforgeability (U).** If  $v_i = 0$  for all correct processes  $i$ , then for all correct processes  $j$ ,  $\text{accept}_j$  remains 0 forever.

$$\mathbf{G} \left( \left( \bigwedge_{i=1}^{n-f} v_i = 0 \right) \rightarrow \mathbf{G} \left( \bigwedge_{j=1}^{n-f} \text{accept}_j = 0 \right) \right)$$

**Completeness (C).** If  $v_i = 1$  for all correct processes  $i$ , then there is a correct process  $j$  that eventually sets  $\text{accept}_j$  to 1.

$$\mathbf{G} \left( \left( \bigwedge_{i=1}^{n-f} v_i = 1 \right) \rightarrow \mathbf{F} \left( \bigvee_{j=1}^{n-f} \text{accept}_j = 1 \right) \right)$$

**Relay (R).** If a correct process  $i$  sets  $\text{accept}_i$  to 1, then eventually all correct processes  $j$  set  $\text{accept}_j$  to 1.

$$\mathbf{G} \left( \left( \bigvee_{i=1}^{n-f} \text{accept}_i = 1 \right) \rightarrow \mathbf{F} \left( \bigwedge_{j=1}^{n-f} \text{accept}_j = 1 \right) \right)$$

# Properties in Linear Temporal Logic

Unforgeability (U). If  $v_i = 0$  for all correct processes  $i$ , then for all correct processes  $j$ ,  $\text{accept}_j$  remains 0 forever.

$$\mathbf{G} \left( \left( \bigwedge_{i=1}^{n-f} v_i = 0 \right) \rightarrow \mathbf{G} \left( \bigwedge_{j=1}^{n-f} \text{accept}_j = 0 \right) \right) \quad \text{Safety}$$

Completeness (C). If  $v_i = 1$  for all correct processes  $i$ , then there is a correct process  $j$  that eventually sets  $\text{accept}_j$  to 1.

$$\mathbf{G} \left( \left( \bigwedge_{i=1}^{n-f} v_i = 1 \right) \rightarrow \mathbf{F} \left( \bigvee_{j=1}^{n-f} \text{accept}_j = 1 \right) \right) \quad \text{Liveness}$$

Relay (R). If a correct process  $i$  sets  $\text{accept}_i$  to 1, then eventually all correct processes  $j$  set  $\text{accept}_j$  to 1.

$$\mathbf{G} \left( \left( \bigvee_{i=1}^{n-f} \text{accept}_i = 1 \right) \rightarrow \mathbf{F} \left( \bigwedge_{j=1}^{n-f} \text{accept}_j = 1 \right) \right) \quad \text{Liveness}$$

# Threshold-guarded fault-tolerant distributed algorithms

# Threshold-guarded FTDA

## Fault-free construct: quantified guards ( $t=f=0$ )

- Existential Guard  
if received  $m$  from *some* process then ...
- Universal Guard  
if received  $m$  from *all* processes then ...

These guards allow one to treat the processes in a parameterized way

# Threshold-guarded FTDA

## Fault-free construct: quantified guards ( $t=f=0$ )

- Existential Guard  
if received  $m$  from *some* process then ...
- Universal Guard  
if received  $m$  from *all* processes then ...

These guards allow one to treat the processes in a parameterized way

*what if faults might occur?*



# Threshold-guarded FTDA

## Fault-free construct: quantified guards ( $t=f=0$ )

- Existential Guard  
if received  $m$  from *some* process then ...
- Universal Guard  
if received  $m$  from *all* processes then ...

These guards allow one to treat the processes in a parameterized way

*what if faults might occur?*



## Fault-Tolerant Algorithms: $n$ processes, at most $t$ are Byzantine

- Threshold Guard  
if received  $m$  from  $n - t$  processes then ...
- (the processes **cannot refer to  $f!$** )

# Control Flow Automata

Variables of process  $i$

$v_i : \{0, 1\}$  init with 0 or 1  
 $accept_i : \{0, 1\}$  init with 0

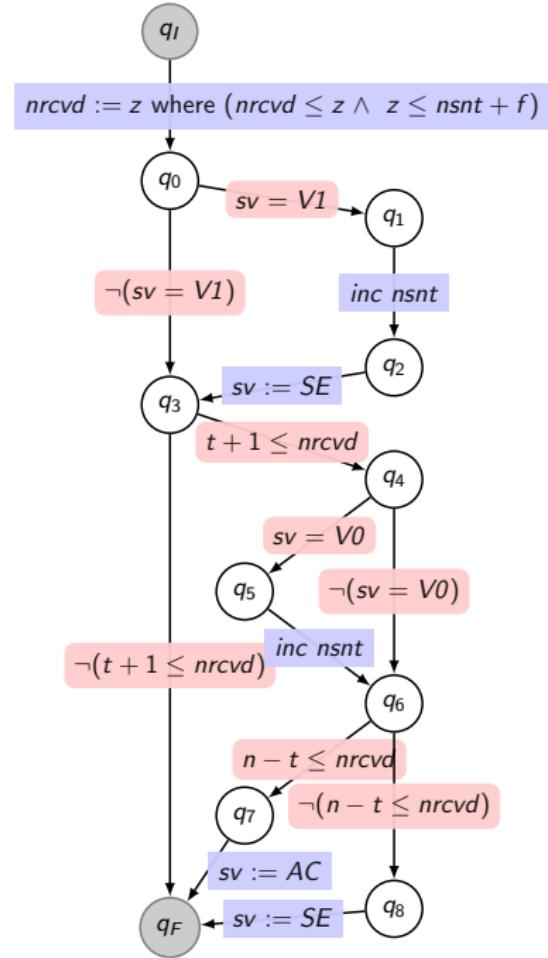
An indivisible step:

**if**  $v_i = 1$   
**then** send (echo) to all;

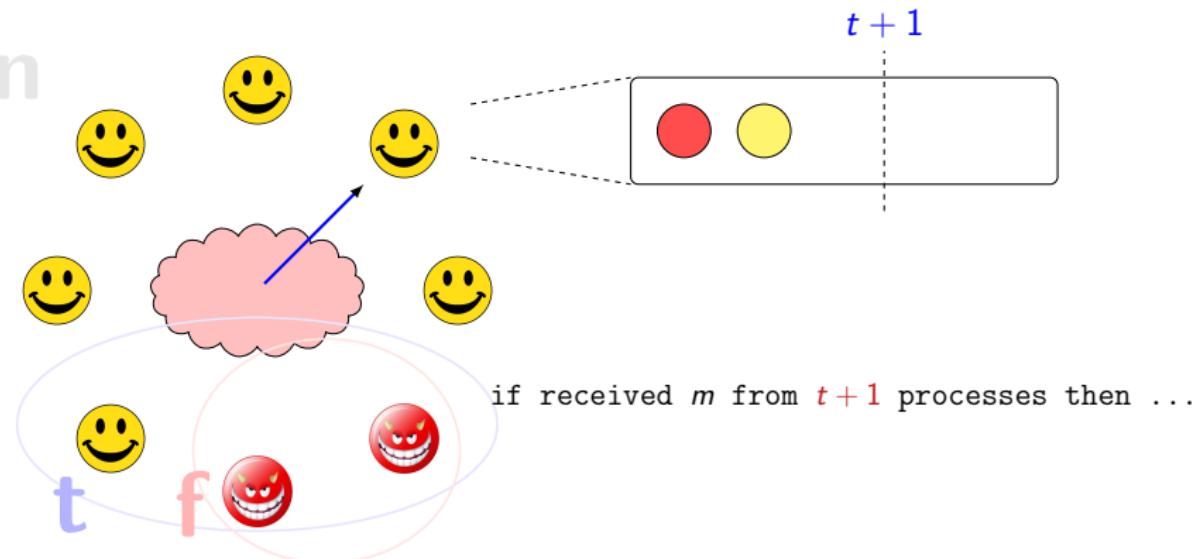
**if** received (echo) from at least  
t + 1 distinct processes  
and not sent (echo) before  
**then** send (echo) to all;

**if** received (echo) from at least  
n - t distinct processes  
**then**  $accept_i := 1$ ;

$n - f$  copies of the process

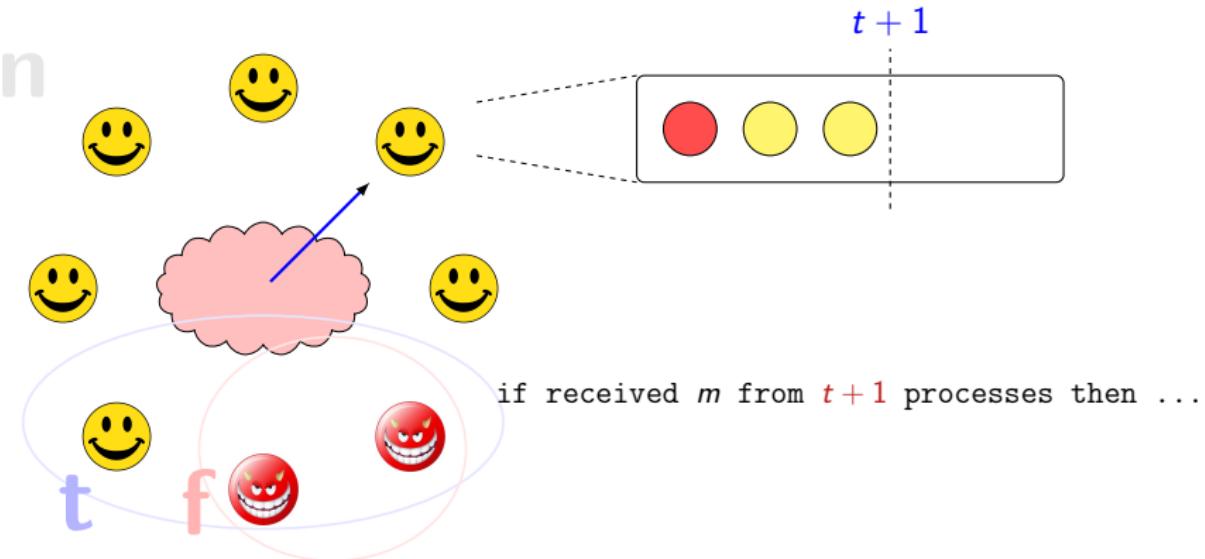


# Counting argument in threshold-guarded algorithms



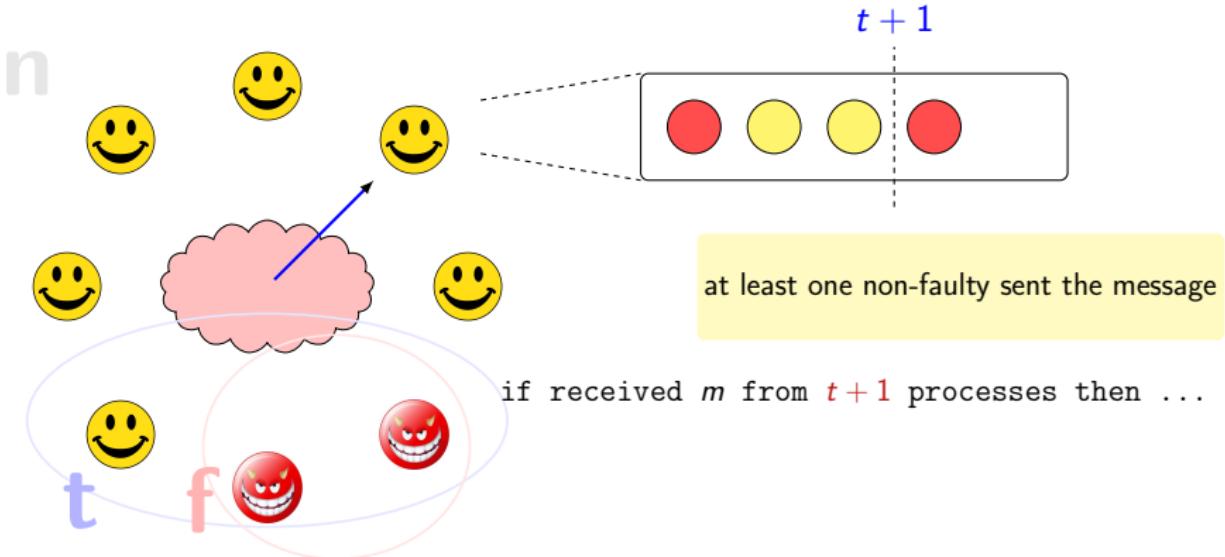
Correct processes count distinct incoming messages

# Counting argument in threshold-guarded algorithms

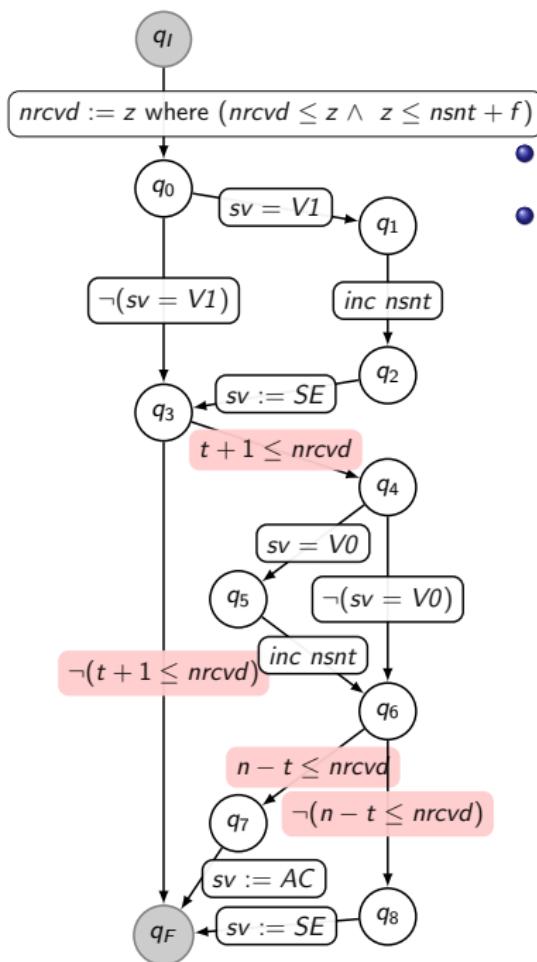


Correct processes count distinct incoming messages

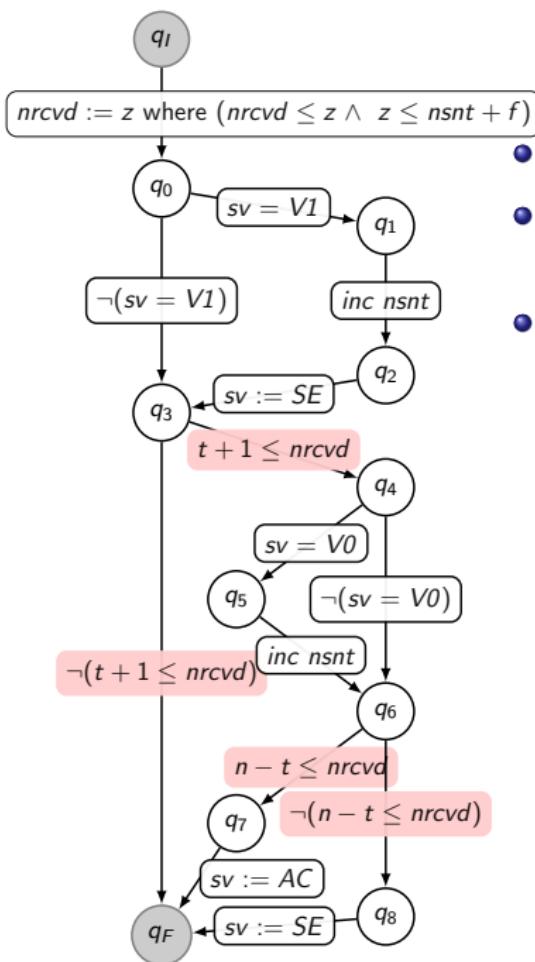
# Counting argument in threshold-guarded algorithms



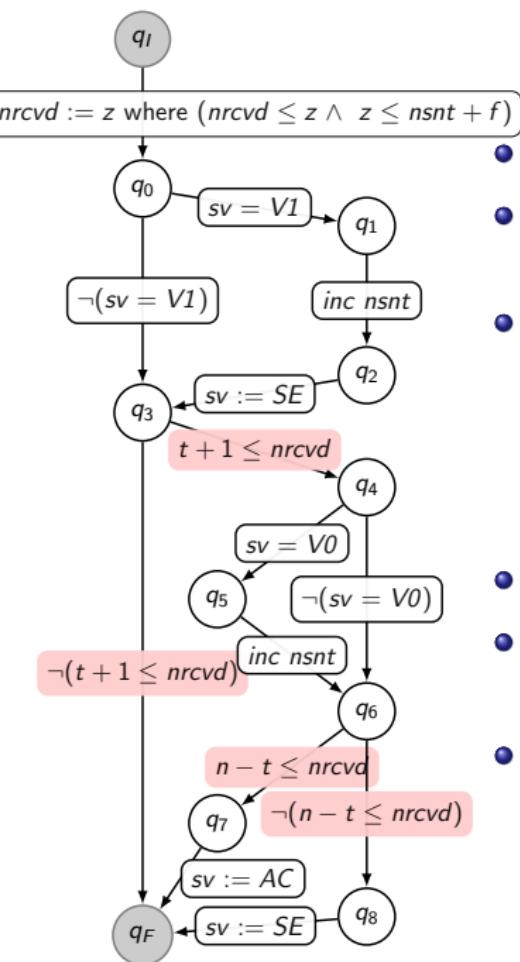
Correct processes count distinct incoming messages



- concrete values are not important
- thresholds are essential:  
 $0, 1, t + 1, n - t$



- concrete values are not important
- thresholds are essential:  
 $0, 1, t + 1, n - t$
- intervals with symbolic boundaries:
  - $I_0 = [0, 1)$
  - $I_1 = [1, t + 1)$
  - $I_{t+1} = [t + 1, n - t)$
  - $I_{n-t} = [n - t, \infty)$



- concrete values are not important
- thresholds are essential:  
 $0, 1, t + 1, n - t$
- intervals with symbolic boundaries:
  - $I_0 = [0, 1)$
  - $I_1 = [1, t + 1)$
  - $I_{t+1} = [t + 1, n - t)$
  - $I_{n-t} = [n - t, \infty)$
- Parameteric Interval Abstraction (PIA)
- Similar to interval abstraction:  
 $[t + 1, n - t)$  rather than  $[4, 10)$ .
- Total order:  $0 < 1 < t + 1 < n - t$  for all parameters satisfying RC:  
 $n > 3t, t \geq f \geq 0$ .

## Technical challenges

We have to reduce the verification of an infinite number of instances where

- ① the process code is parameterized
- ② the number of processes is parameterized

to one finite state model checking instance

# Technical challenges

We have to reduce the verification of an infinite number of instances where

- ① the process code is parameterized
- ② the number of processes is parameterized

to one finite state model checking instance

We do that by:

- ① PIA data abstraction
- ② PIA counter abstraction

# Technical challenges

We have to reduce the verification of an infinite number of instances where

- ① the process code is parameterized
- ② the number of processes is parameterized

to one finite state model checking instance

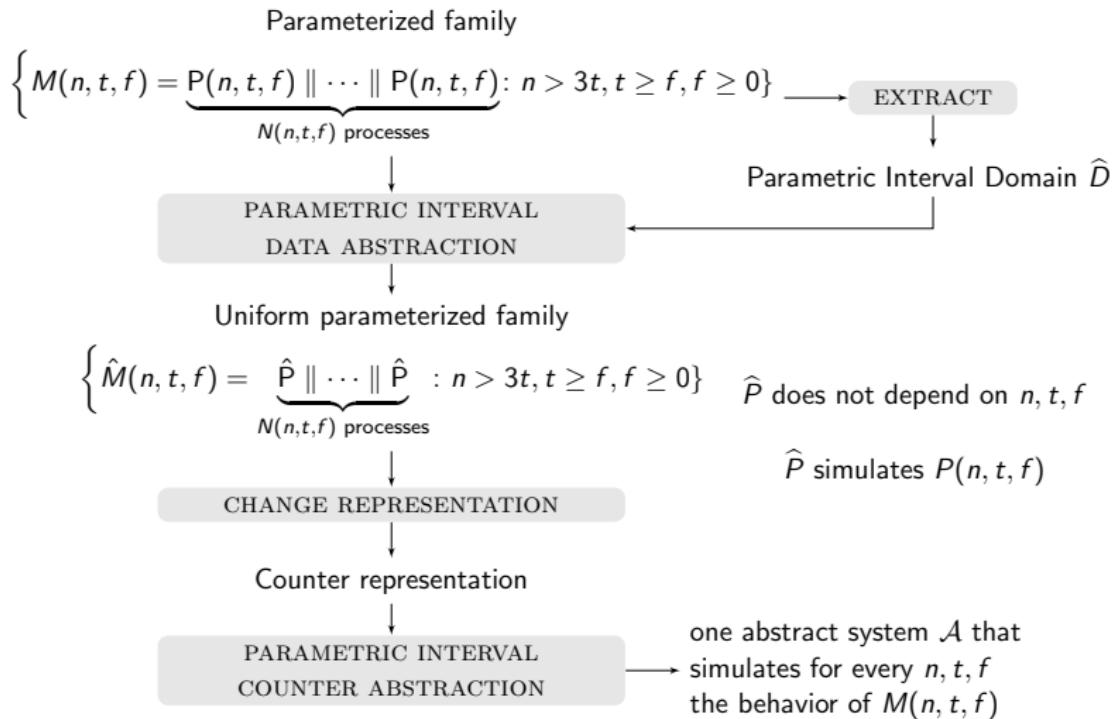
We do that by:

- ① PIA data abstraction
- ② PIA counter abstraction

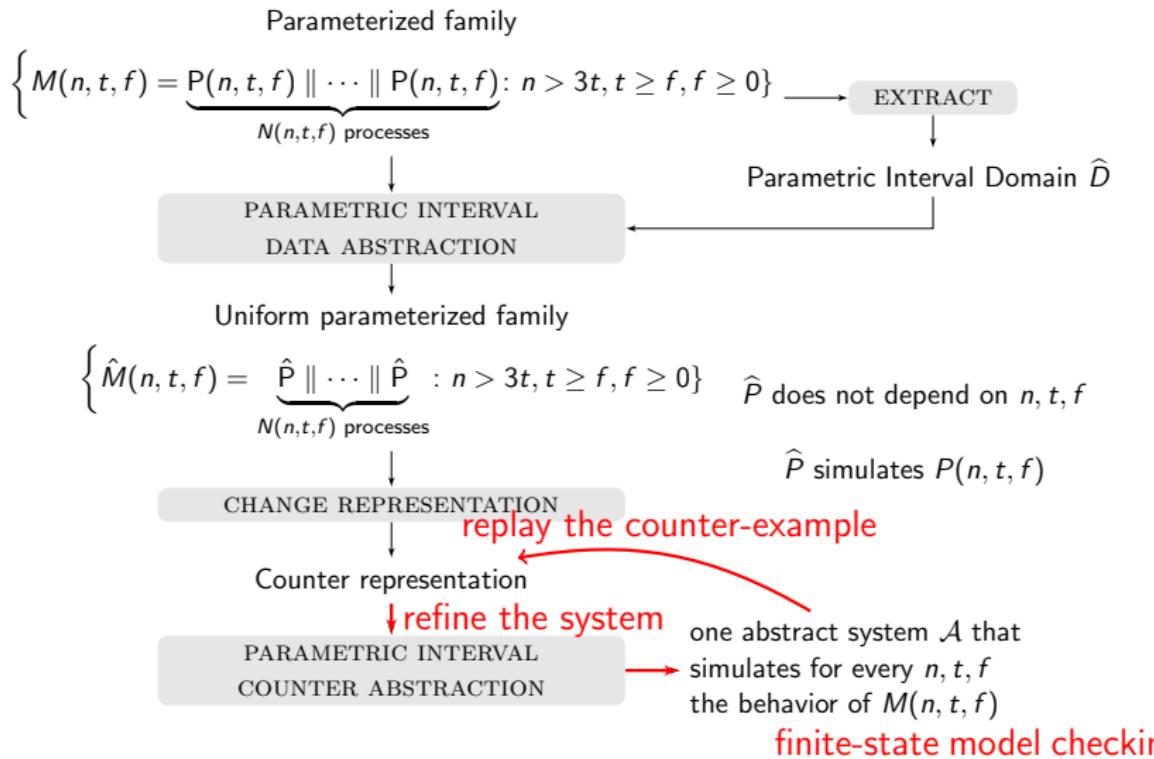
abstraction is an over approximation  $\Rightarrow$  possible abstract behavior that does not correspond to a concrete behavior.

- ③ Refining spurious counter-examples

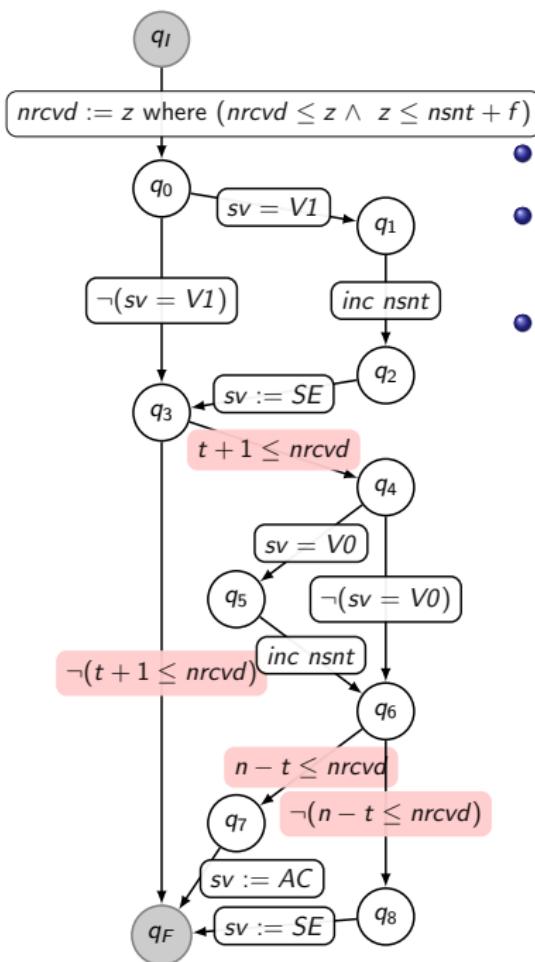
# Abstraction overview



# Abstraction overview

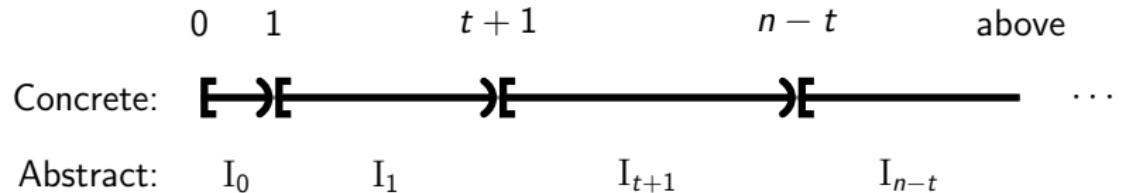


# Data abstraction



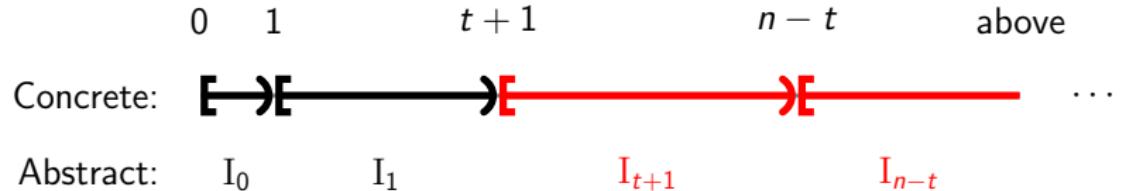
- concrete values are not important
- thresholds are essential:  
 $0, 1, t + 1, n - t$
- intervals with symbolic boundaries:
  - $I_0 = [0, 1)$
  - $I_1 = [1, t + 1)$
  - $I_{t+1} = [t + 1, n - t)$
  - $I_{n-t} = [n - t, \infty)$

# Abstract operations



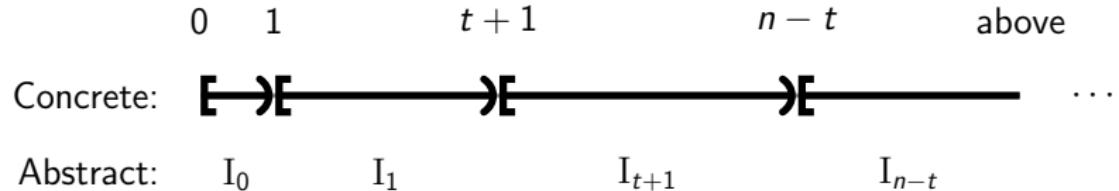
Concrete  $t + 1 \leq x$

# Abstract operations



Concrete  $t + 1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

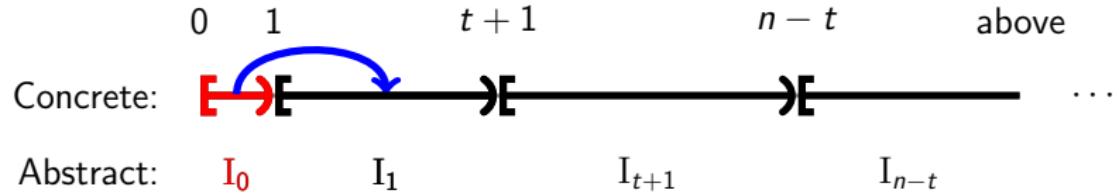
# Abstract operations



Concrete  $t + 1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

Concrete  $x' = x + 1$ ,

# Abstract operations

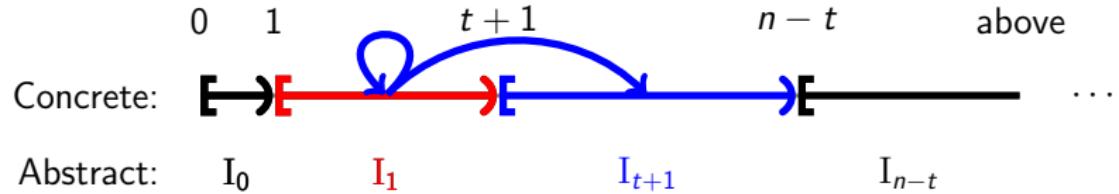


Concrete  $t + 1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

Concrete  $x' = x + 1$ , is abstracted as:

$$x = I_0 \quad \wedge \quad x' = I_1 \dots$$

# Abstract operations



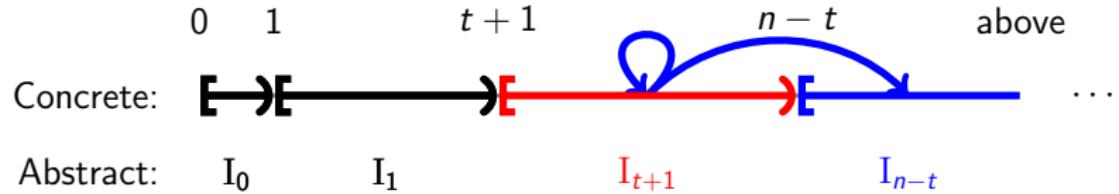
Concrete  $t+1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

Concrete  $x' = x + 1$ , is abstracted as:

$$x = I_0 \quad \wedge \quad x' = I_1$$

$$\forall x = I_1 \quad \wedge \quad (x' = I_1 \quad \vee x' = I_{t+1}) \dots$$

# Abstract operations



Concrete  $t + 1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

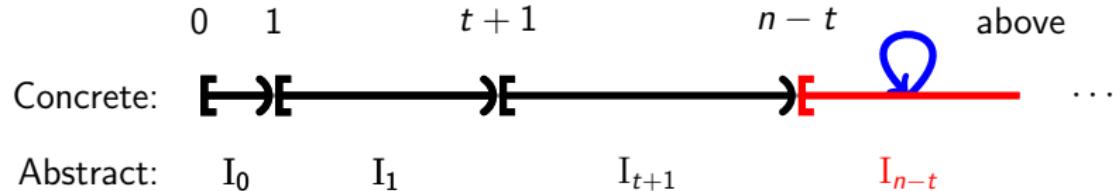
Concrete  $x' = x + 1$ , is abstracted as:

$$x = I_0 \wedge x' = I_1$$

$$\forall x = I_1 \wedge (x' = I_1 \vee x' = I_{t+1})$$

$$\forall x = I_{t+1} \wedge (x' = I_{t+1} \vee x' = I_{n-t}) \dots$$

# Abstract operations

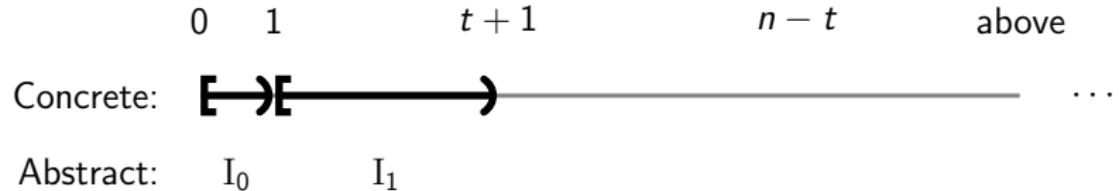


Concrete  $t + 1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

Concrete  $x' = x + 1$ , is abstracted as:

$$\begin{aligned}x &= I_0 \quad \wedge \quad x' = I_1 \\ \forall x &= I_1 \quad \wedge \quad (x' = I_1 \quad \vee \quad x' = I_{t+1}) \\ \forall x &= I_{t+1} \quad \wedge \quad (x' = I_{t+1} \quad \vee \quad x' = I_{n-t}) \\ \forall x &= I_{n-t} \quad \wedge \quad x' = I_{n-t}\end{aligned}$$

# Abstract operations



Concrete  $t + 1 \leq x$  is abstracted as  $x = I_{t+1} \vee x = I_{n-t}$ .

Concrete  $x' = x + 1$ , is abstracted as:

$$x = I_0 \wedge x' = I_1$$

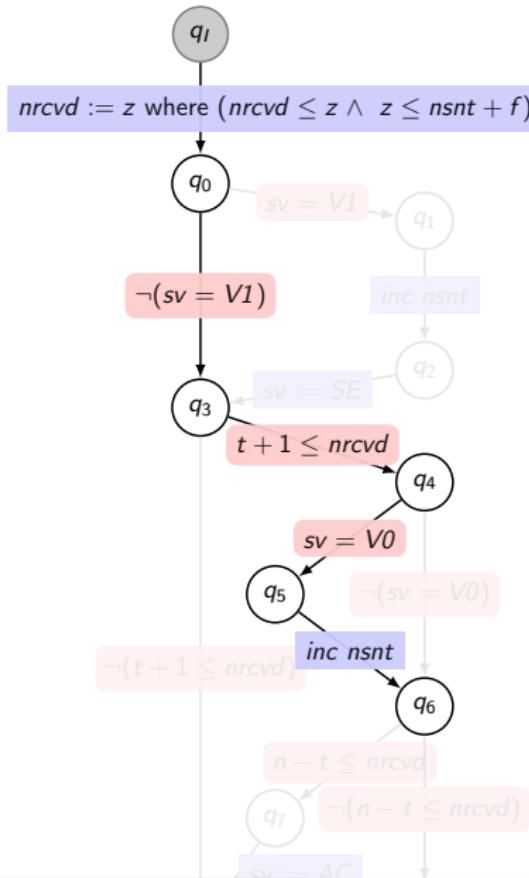
$$\forall x = I_1 \wedge (x' = I_1 \vee x' = I_{t+1})$$

$$\forall x = I_{t+1} \wedge (x' = I_{t+1} \vee x' = I_{n-t})$$

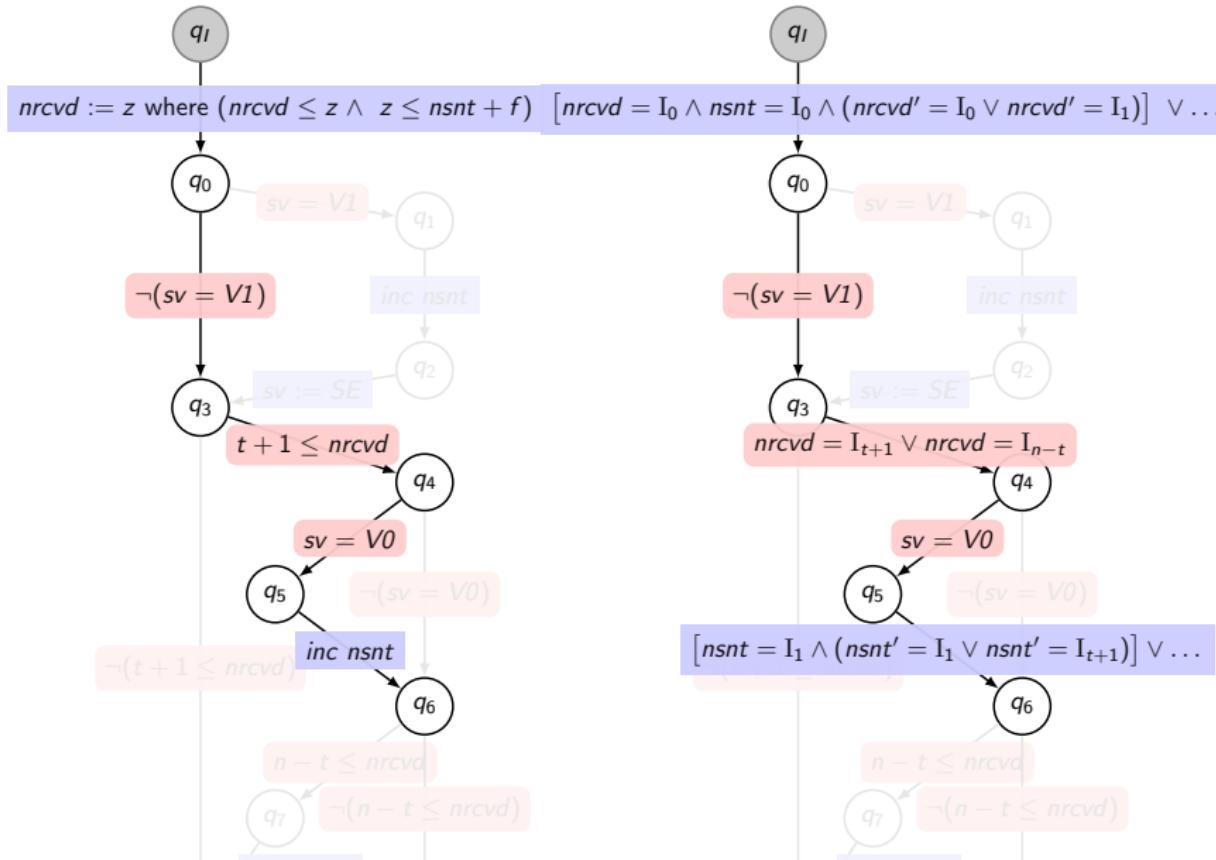
$$\forall x = I_{n-t} \wedge x' = I_{n-t}$$

abstract increase may keep the same value!

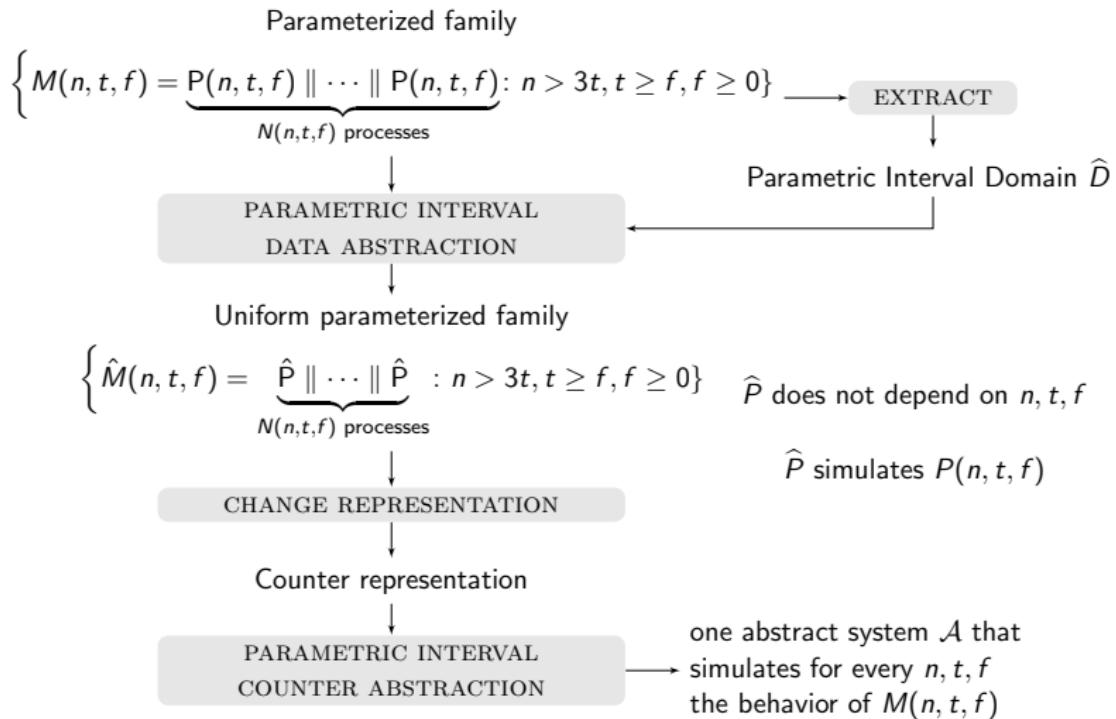
# Abstract CFA



# Abstract CFA



# Abstraction overview



# Counter abstraction

## Classic $(0, 1, \infty)$ -counter abstraction

Pnueli, Xu, and Zuck (2001) introduced  $(0, 1, \infty)$ -counter abstraction:

- finitely many local states,  
e.g.,  $\{N, T, C\}$ .
- based on counter representation:  
for each local states count how many processes are in it

## Classic $(0, 1, \infty)$ -counter abstraction

Pnueli, Xu, and Zuck (2001) introduced  $(0, 1, \infty)$ -counter abstraction:

- finitely many local states,  
e.g.,  $\{N, T, C\}$ .
- based on counter representation:  
for each local states count how many processes are in it
- abstract the number of processes in every state,  
e.g.,  $K : C \mapsto \mathbf{0}, T \mapsto \mathbf{1}, N \mapsto \text{"many"}$ .
- perfectly reflects mutual exclusion properties  
e.g.,  $\mathbf{G}(K(C) = \mathbf{0} \vee K(C) = \mathbf{1})$ .

# Limits of $(0, 1, \infty)$ -counter abstraction

Our parametric data + counter abstraction:

- we require finer counting of processes:
  - $t + 1$  processes in a specific state can force global progress,
  - $t$  processes cannot
- mapping  $t$ ,  $t + 1$ , and  $n - t$  to “**many**” is **too coarse**.

# Limits of $(0, 1, \infty)$ -counter abstraction

Our parametric data + counter abstraction:

- we require finer counting of processes:
    - $t + 1$  processes in a specific state can force global progress,
    - $t$  processes cannot
  - mapping  $t$ ,  $t + 1$ , and  $n - t$  to “**many**” is **too coarse**.
- starting point of our approach...

# Data + counter abstraction over parametric intervals

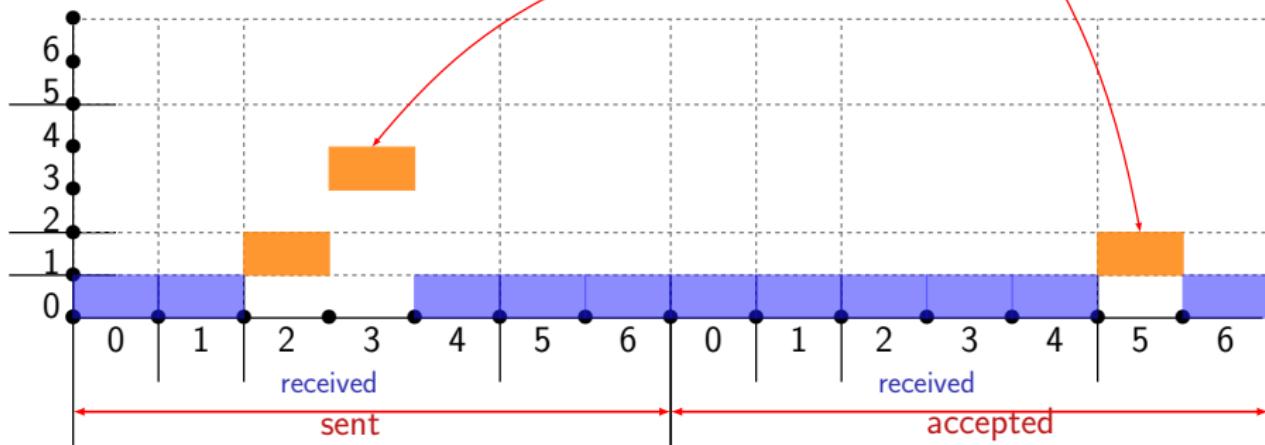
$$n = 6, t = 1, f = 1$$

$$t + 1 = 2, n - t = 5$$

1 process at (accepted, received=5)

nr. processes (counters)

3 processes at (sent, received=3)



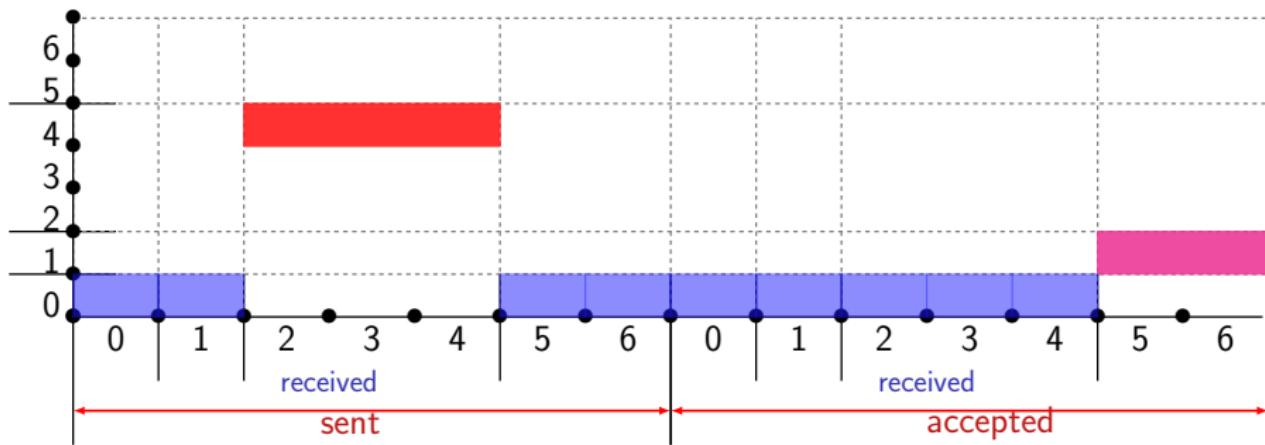
Local state is  $(sv, nrcvd)$ ,  
where  $sv \in \{sent, accepted\}$  and  $0 \leq rcvd \leq n$

# Data + counter abstraction over parametric intervals

$$n = 6, t = 1, f = 1$$

$$t + 1 = 2, n - t = 5$$

nr. processes (counters)



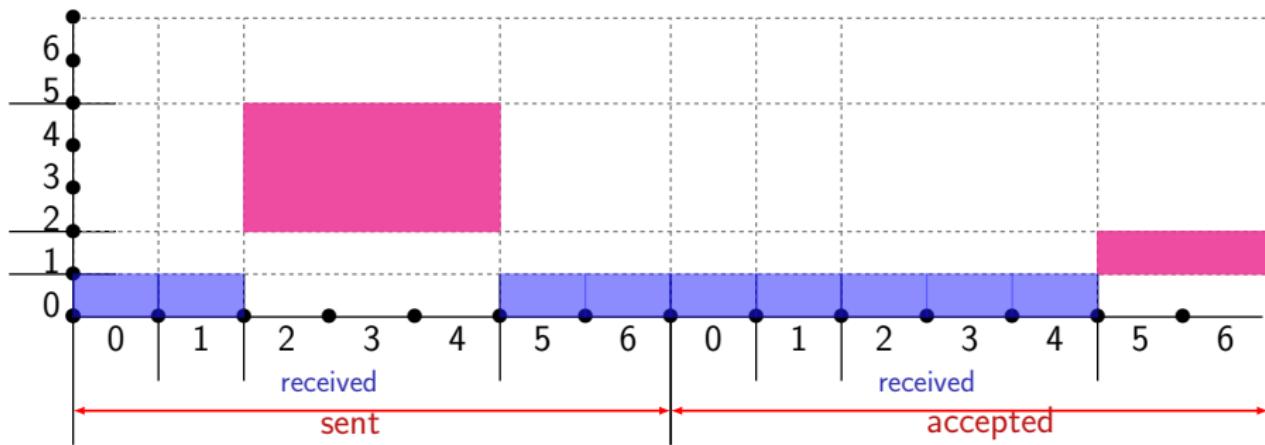
Local state is  $(sv, nrcvd)$ ,  
where  $sv \in \{sent, accepted\}$  and  $0 \leq rcvd \leq n$

# Data + counter abstraction over parametric intervals

$$n = 6, t = 1, f = 1$$

$$t + 1 = 2, n - t = 5$$

nr. processes (counters)



Local state is  $(sv, nrcvd)$ ,  
where  $sv \in \{sent, accepted\}$  and  $0 \leq rcvd \leq n$

# Data + counter abstraction over parametric intervals

$$n=6, t=1, f=1$$

$$n > 3 \cdot t \wedge t \geq f$$

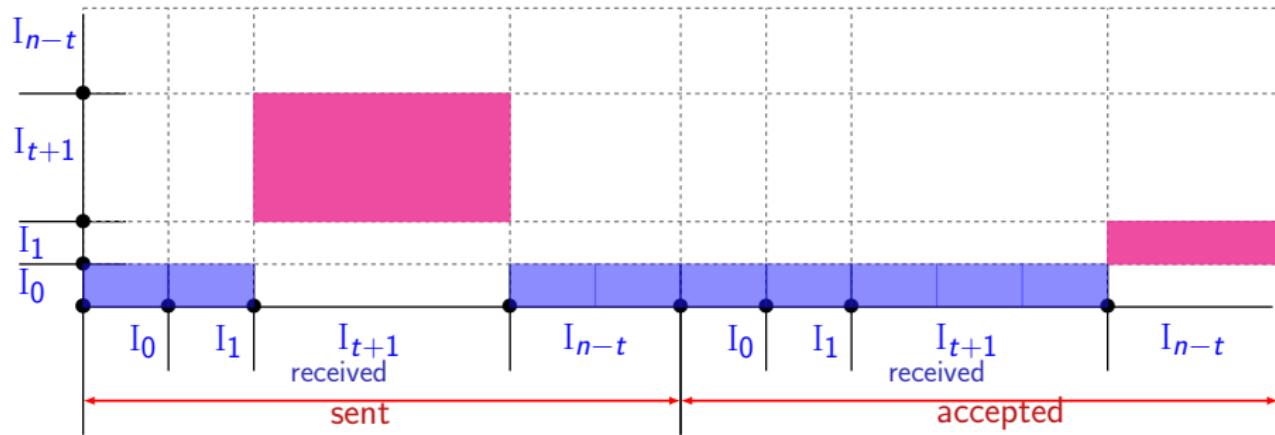
nr. processes (counters)

Parametric intervals:

$$I_0 = [0, 1) \quad I_1 = [1, t+1)$$

$$I_{t+1} = [t+1, n-t)$$

$$I_{n-t} = [n-t, \infty)$$



A local state is  $(sv, nrcvd)$ ,  
where  $sv \in \{sent, accepted\}$  and  $nrcvd \in \{I_0, I_1, I_{t+1}, I_{n-t}\}$

# Data + counter abstraction over parametric intervals

Parametric intervals:

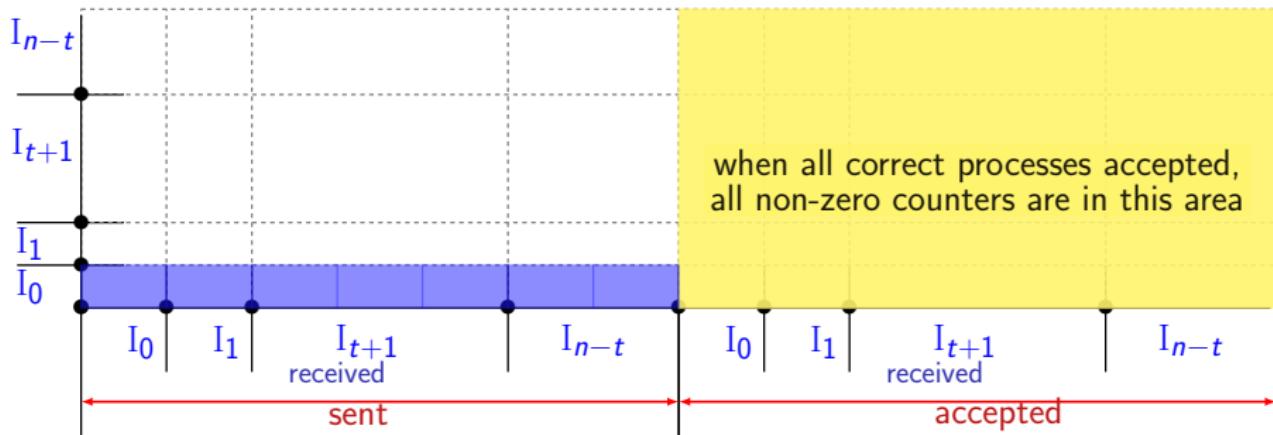
$$n > 3 \cdot t \wedge t \geq f$$

$$I_0 = [0, 1] \quad I_1 = [1, t+1]$$

$$I_{t+1} = [t+1, n-t)$$

nr. processes (counters)

$$I_{n-t} = [n-t, \infty)$$



A local state is  $(sv, nrcvd)$ ,  
where  $sv \in \{sent, accepted\}$  and  $nrcvd \in \{I_0, I_1, I_{t+1}, I_{n-t}\}$

# Abstraction refinement

# Spurious behavior

abstraction adds behaviors (e.g.,  $x' = x + 1$  may lead to  $x'$  being equal to  $x$ )

# Spurious behavior

abstraction adds behaviors (e.g.,  $x' = x + 1$  may lead to  $x'$  being equal to  $x$ )

⇒ specs that hold in concrete system may be violated in abstract system

- spurious counterexamples
- we have to reduce the behaviors of the abstract system  
make it more concrete
- ... based on the counterexamples = CEGAR

# Spurious behavior

abstraction adds behaviors (e.g.,  $x' = x + 1$  may lead to  $x'$  being equal to  $x$ )

⇒ specs that hold in concrete system may be violated in abstract system

- spurious counterexamples
- we have to reduce the behaviors of the abstract system  
make it more concrete
- ... based on the counterexamples = CEGAR

Three sources of spurious behavior

- # processes decreasing or increasing
- # messages sent ≠ # processes which have sent a message
- unfair loops

# Spurious behavior

abstraction adds behaviors (e.g.,  $x' = x + 1$  may lead to  $x'$  being equal to  $x$ )

⇒ specs that hold in concrete system may be violated in abstract system

- spurious counterexamples
- we have to reduce the behaviors of the abstract system  
make it more concrete
- ... based on the counterexamples = CEGAR

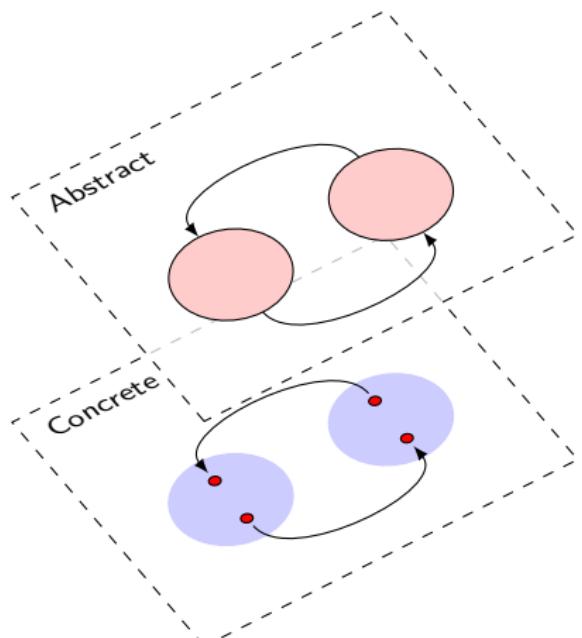
Three sources of spurious behavior

- # processes decreasing or increasing
- # messages sent ≠ # processes which have sent a message
- unfair loops

... and a new abstraction phenomenon

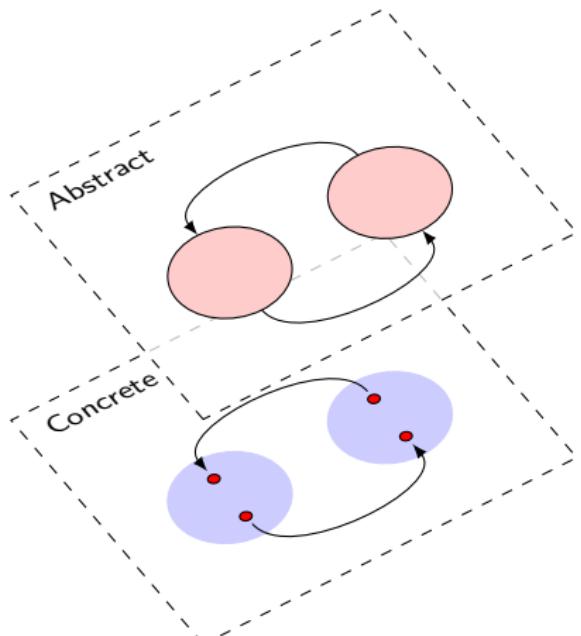
# Parametric abst. refinement — uniformly spurious paths

Classic case:

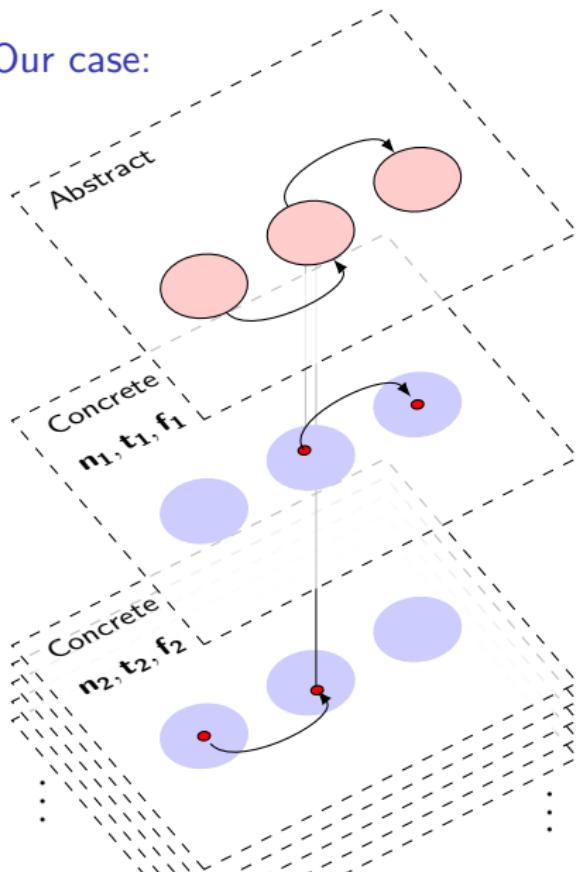


# Parametric abst. refinement—uniformly spurious paths

Classic case:



Our case:



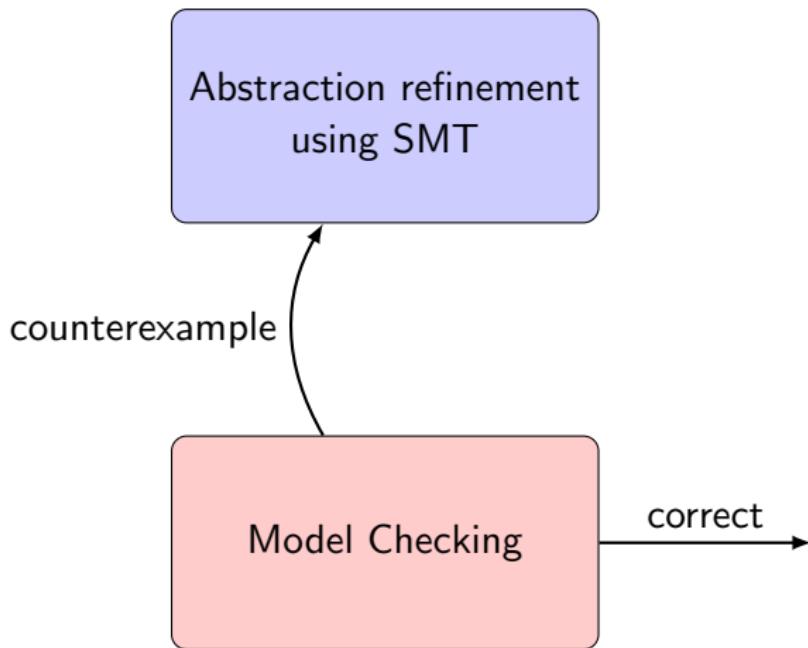
# CEGAR—automated workflow

Model Checking

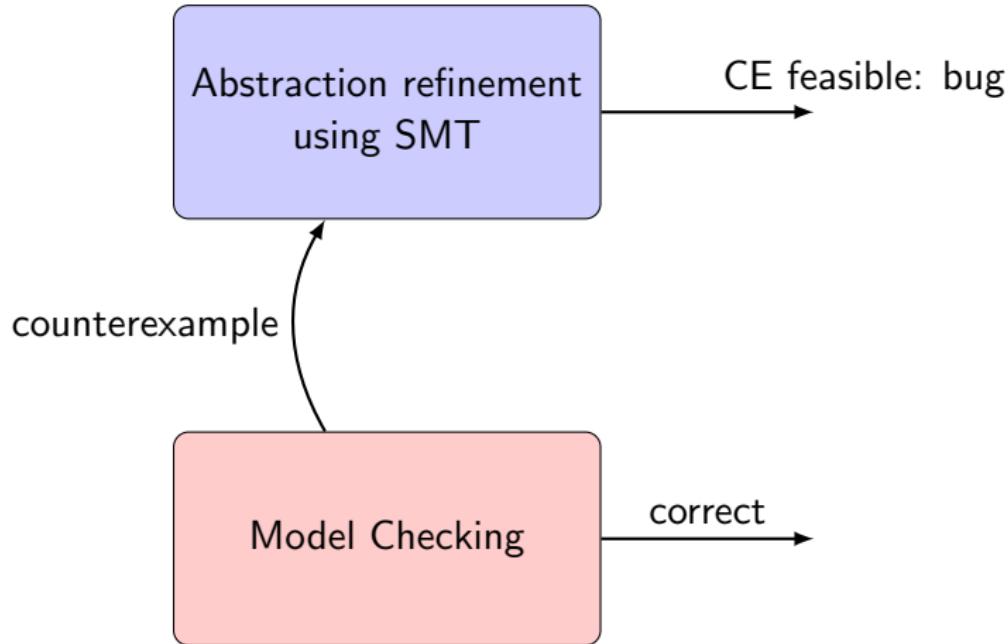
# CEGAR—automated workflow



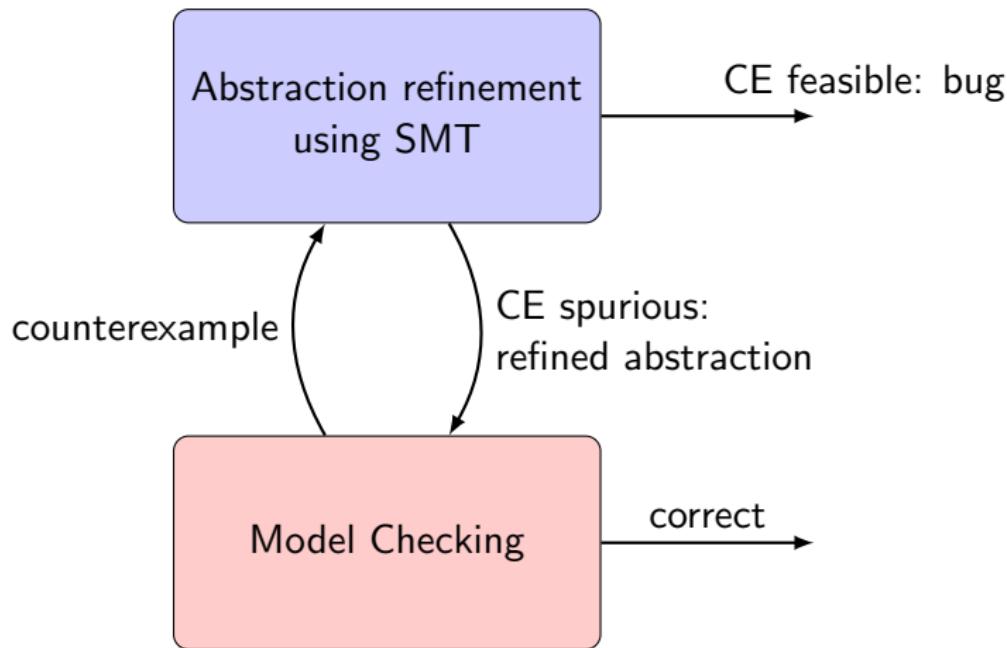
# CEGAR—automated workflow



# CEGAR—automated workflow



# CEGAR—automated workflow



# What is SMT?

recall SAT:

- given a Boolean formula, e.g.,  $(\neg a \vee \neg b \vee c) \wedge (\neg a \vee b \vee d \vee e)$
- is there an assignment of TRUE and FALSE to variables  $a, b, c, d, e$  such that the formula evaluates to TRUE?

# What is SMT?

recall SAT:

- given a Boolean formula, e.g.,  $(\neg a \vee \neg b \vee c) \wedge (\neg a \vee b \vee d \vee e)$
- is there an assignment of TRUE and FALSE to variables  $a, b, c, d, e$  such that the formula evaluates to TRUE?

Satisfiability Modulo Theories (SMT) :

- here just linear arithmetics
- given a formula, e.g.,

$$x = y \wedge y = z \wedge u \neq x \wedge (x + y \leq 1 \wedge 2x + y = 1) \vee 3x + 2y \geq 3$$

- is there an assignment of values to  $u, x, y, z$  such that formula evaluates to TRUE?
- practically efficient tools: YICES, Z3

## Counter example: losing processes

Output of data abstraction: **16 local states**:  $L = \{(sv, \hat{rcvd})\}$   
with  $sv \in \{v0, v1, sent, accepted\}$  and  $\hat{rcvd} \in \{I_0, I_1, I_{t+1}, I_{n-t}\}\}$

An abstract **global** state is  $(\hat{k}, \hat{nsnt})$ ,

where  $\hat{nsnt} \in \{I_0, I_1, I_{t+1}, I_{n-t}\}$  and  $\hat{k} : L \rightarrow \{I_0, I_1, I_{t+1}, I_{n-t}\}$

**Consider an abstract trace:**

$$\hat{nsnt}_1 = I_0$$

$$\hat{k}_1(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

$$\hat{nsnt}_2 = I_1$$

$$\hat{k}_2(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_1, & \text{if } \ell = (sent, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

$$\hat{nsnt}_3 = I_{t+1}$$

$$\hat{k}_3(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_{t+1}, & \text{if } \ell = (sent, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

**Encode the last state in SMT** as a conjunction  $T$  of the constraints:

resilience condition

$$n > 3t \wedge t \geq f \wedge f \geq 0$$

zero counters

$$(i \neq 4 \wedge i \neq 8) \rightarrow 0 \leq k_3[i] < 1 \quad \text{UNSAT}$$

non-zero counters

$$n - t \leq k_3[4] \wedge t + 1 \leq k_3[8] < n - t$$

system size

$$n - f = k_3[0] + k_3[1] + \cdots + k_3[15]$$

## Counter example: losing processes

Output of data abstraction: **16 local states:**  $L = \{(sv, \hat{rcvd})\}$   
with  $sv \in \{v0, v1, sent, accepted\}$  and  $\hat{rcvd} \in \{I_0, I_1, I_{t+1}, I_{n-t}\}\}$

An abstract **global state** is  $(\hat{k}, \hat{nsnt})$ ,

where  $\hat{nsnt} \in \{I_0, I_1, I_{t+1}, I_{n-t}\}$  and  $\hat{k} : L \rightarrow \{I_0, I_1, I_{t+1}, I_{n-t}\}$

**Consider an abstract trace:**

$$\hat{nsnt}_1 = I_0$$

$$\hat{k}_1(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

$$\hat{nsnt}_2 = I_1$$

$$\hat{k}_2(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_1, & \text{if } \ell = (sent, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

$$\hat{nsnt}_3 = I_{t+1}$$

$$\hat{k}_3(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_{t+1}, & \text{if } \ell = (sent, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

**Encode the last state in SMT** as a conjunction  $T$  of the constraints:

resilience condition

$$n > 3t \wedge t \geq f \wedge f \geq 0$$

zero counters

$$(i \neq 4 \wedge i \neq 8) \rightarrow 0 \leq k_3[i] < 1 \quad \text{UNSAT}$$

non-zero counters

$$n - t \leq k_3[4] \wedge t + 1 \leq k_3[8] < n - t$$

system size

$$n - f = k_3[0] + k_3[1] + \cdots + k_3[15]$$

## Counter example: losing processes

Output of data abstraction: **16 local states:**  $L = \{(sv, \hat{rcvd})\}$   
with  $sv \in \{v0, v1, sent, accepted\}$  and  $\hat{rcvd} \in \{I_0, I_1, I_{t+1}, I_{n-t}\}\}$

An abstract **global state** is  $(\hat{k}, \hat{nsnt})$ ,

where  $\hat{nsnt} \in \{I_0, I_1, I_{t+1}, I_{n-t}\}$  and  $\hat{k} : L \rightarrow \{I_0, I_1, I_{t+1}, I_{n-t}\}$

**Consider an abstract trace:**

$$\hat{nsnt}_1 = I_0$$

$$\hat{k}_1(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

$$\hat{nsnt}_2 = I_1$$

$$\hat{k}_2(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_1, & \text{if } \ell = (sent, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

$$\hat{nsnt}_3 = I_{t+1}$$

$$\hat{k}_3(\ell) =$$

$$\begin{cases} I_{n-t}, & \text{if } \ell = (v1, I_0) \\ I_{t+1}, & \text{if } \ell = (sent, I_0) \\ I_0, & \text{otherwise} \end{cases}$$

**Encode the last state in SMT** as a conjunction  $T$  of the constraints:

resilience condition

$$n > 3t \wedge t \geq f \wedge f \geq 0$$

zero counters

$$(i \neq 4 \wedge i \neq 8) \rightarrow 0 \leq k_3[i] < 1 \quad \text{UNSAT}$$

non-zero counters

$$n - t \leq k_3[4] \wedge t + 1 \leq k_3[8] < n - t$$

system size

$$n - f = k_3[0] + k_3[1] + \cdots + k_3[15]$$

## Remove transitions

- We ask the SMT solver:  
is there a satisfiable assignment for  $T$ ?
- if yes,  
then the state is OK, may be part of a real counterexample
- if not, then the state is spurious  
remove transitions to that state in the abstract system

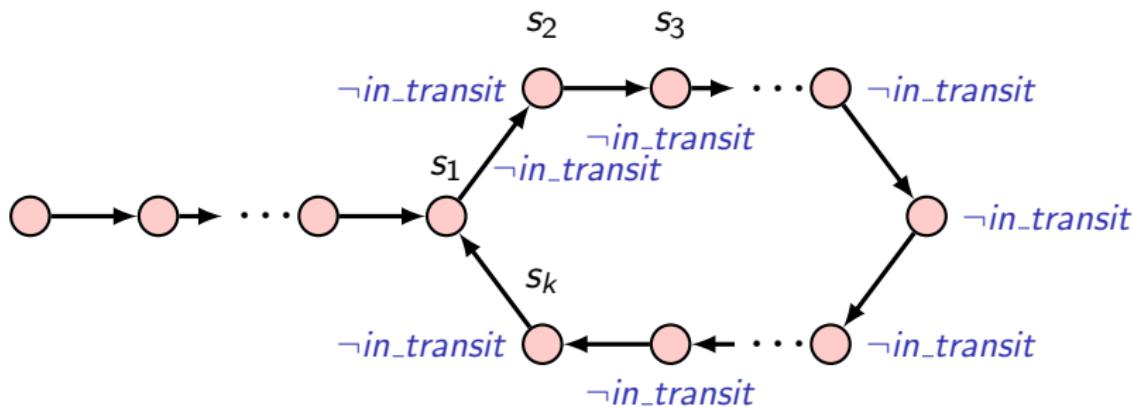
# Liveness

- distributed algorithm requires reliable communication
- every message sent is eventually received
- $\neg \text{in\_transit} \equiv [\forall i. \text{nrcvd}_i \geq \text{nsnt}]$
- fairness  $\mathbf{FG} \neg \text{in\_transit}$  necessary to verify liveness,  
e.g.,  $(\mathbf{FG} \neg \text{in\_transit} \rightarrow (\mathbf{G}([\forall i. sv_i = v1] \rightarrow \mathbf{F}[\forall i. sv_i = accept])))$

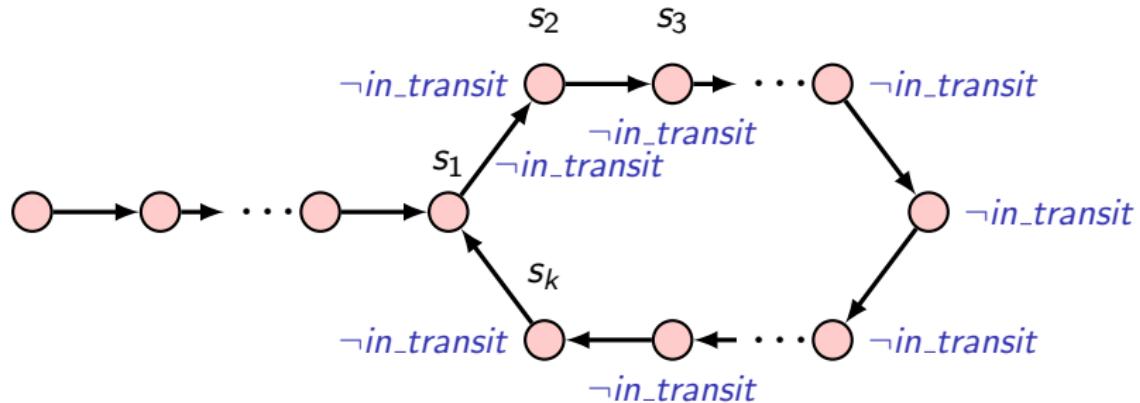
# Liveness

- distributed algorithm requires reliable communication
- every message sent is eventually received
- $\neg \text{in\_transit} \equiv [\forall i. \text{nrcvd}_i \geq \text{nsnt}]$
- fairness  $\mathbf{FG} \neg \text{in\_transit}$  necessary to verify liveness,  
e.g.,  $(\mathbf{FG} \neg \text{in\_transit} \rightarrow (\mathbf{G}([\forall i. sv_i = v1] \rightarrow \mathbf{F}[\forall i. sv_i = \text{accept}])))$

counter example (lasso):

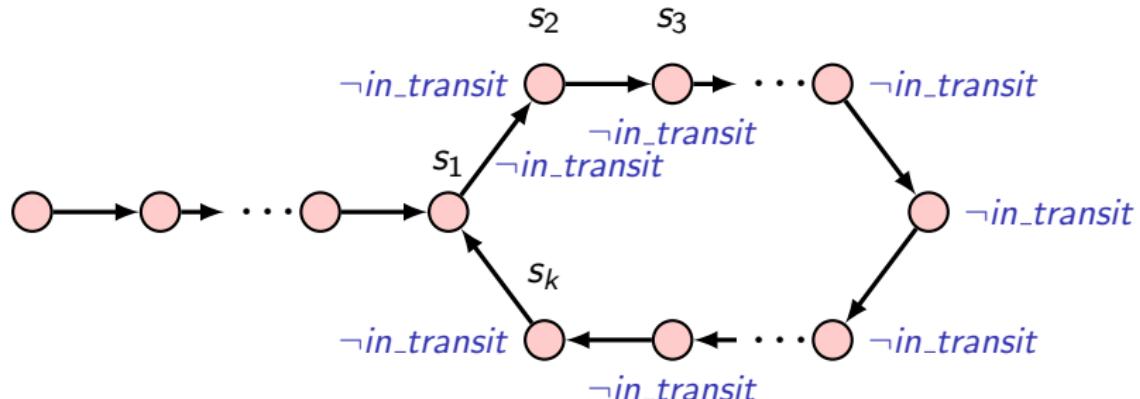


## Liveness—fairness suppression



if there is a spurious  $s_j$  (all its concretizations violate  $\neg \text{in\_transit}$ ), then the loop is spurious.

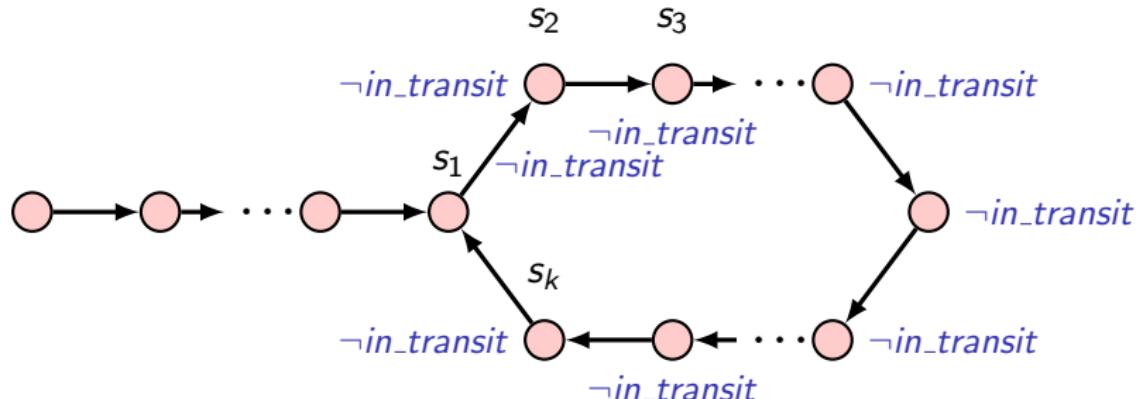
## Liveness—fairness suppression



if there is a spurious  $s_j$  (all its concretizations violate  $\neg \text{in\_transit}$ ),  
then the loop is spurious.

refine fairness to  $\mathbf{F} \mathbf{G} \neg \text{in\_transit} \wedge \mathbf{G} \mathbf{F} \left( \bigwedge_{1 \leq j \leq k} \text{"out of } s_j'' \right)$

## Liveness—fairness suppression

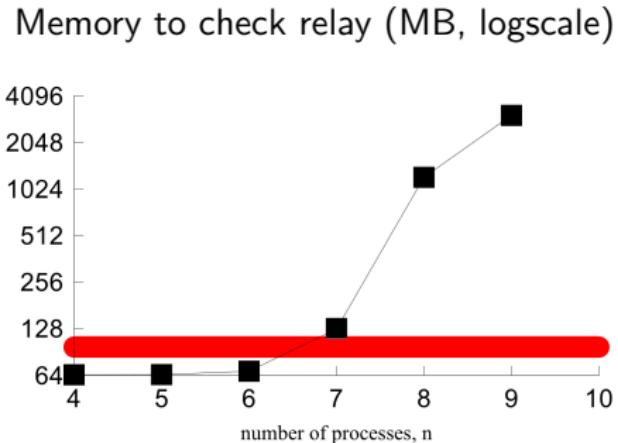
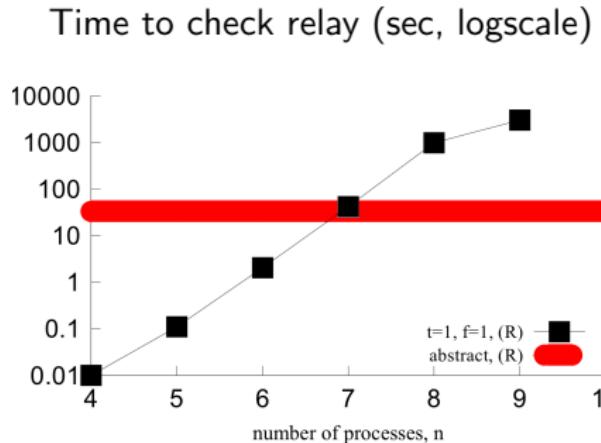


if there is a spurious  $s_j$  (all its concretizations violate  $\neg \text{in\_transit}$ ),  
then the loop is spurious.

refine fairness to  $\mathbf{F} \mathbf{G} \neg \text{in\_transit} \wedge \mathbf{G} \mathbf{F} \left( \bigwedge_{1 \leq j \leq k} \text{"out of } s_j'' \right)$

# experimental evaluation

# Concrete vs. parameterized (Byzantine case)



- Parameterized model checking performs well (the red line).
- Experiments for fixed parameters quickly degrade ( $n = 9$  runs out of memory).
- We found counter-examples for the cases  $n = 3t$  and  $f > t$ , where the resilience condition is violated.

# Experimental results at a glance

Algorithm	Fault	Resilience	Property	Valid?	#Refinements	Time
<b>ST87</b>	BYZ	$n > 3t$	<b>U</b>	✓	0	4 sec.
<b>ST87</b>	BYZ	$n > 3t$	<b>C</b>	✓	10	32 sec.
<b>ST87</b>	BYZ	$n > 3t$	<b>R</b>	✓	10	24 sec.
<b>ST87</b>	SYMM	$n > 2t$	<b>U</b>	✓	0	1 sec.
<b>ST87</b>	SYMM	$n > 2t$	<b>C</b>	✓	2	3 sec.
<b>ST87</b>	SYMM	$n > 2t$	<b>R</b>	✓	12	16 sec.
<b>ST87</b>	OMIT	$n > 2t$	<b>U</b>	✓	0	1 sec.
<b>ST87</b>	OMIT	$n > 2t$	<b>C</b>	✓	5	6 sec.
<b>ST87</b>	OMIT	$n > 2t$	<b>R</b>	✓	5	10 sec.
<b>ST87</b>	CLEAN	$n > t$	<b>U</b>	✓	0	2 sec.
<b>ST87</b>	CLEAN	$n > t$	<b>C</b>	✓	4	8 sec.
<b>ST87</b>	CLEAN	$n > t$	<b>R</b>	✓	13	31 sec.
<b>CT96</b>	CLEAN	$n > t$	<b>U</b>	✓	0	1 sec.
<b>CT96</b>	CLEAN	$n > t$	<b>A</b>	✓	0	1 sec.
<b>CT96</b>	CLEAN	$n > t$	<b>R</b>	✓	0	1 sec.
<b>CT96</b>	CLEAN	$n > t$	<b>C</b>	✗	0	1 sec.

# When resilience condition is wrong...

Algorithm	Fault	Resilience	Property	Valid?	#Refinements	Time
<b>ST87</b>	BYZ	$n > 3t \wedge f \leq t+1$	<b>U</b>	<b>X</b>	9	56 sec.
<b>ST87</b>	BYZ	$n > 3t \wedge f \leq t+1$	<b>C</b>	<b>X</b>	11	52 sec.
<b>ST87</b>	BYZ	$n > 3t \wedge f \leq t+1$	<b>R</b>	<b>X</b>	10	17 sec.
<b>ST87</b>	BYZ	$n \geq 3t \wedge f \leq t$	<b>U</b>	<b>✓</b>	0	5 sec.
<b>ST87</b>	BYZ	$n \geq 3t \wedge f \leq t$	<b>C</b>	<b>✓</b>	9	32 sec.
<b>ST87</b>	BYZ	$n \geq 3t \wedge f \leq t$	<b>R</b>	<b>X</b>	30	78 sec.
<b>ST87</b>	SYMM	$n > 2t \wedge f \leq t+1$	<b>U</b>	<b>X</b>	0	2 sec.
<b>ST87</b>	SYMM	$n > 2t \wedge f \leq t+1$	<b>C</b>	<b>X</b>	2	4 sec.
<b>ST87</b>	SYMM	$n > 2t \wedge f \leq t+1$	<b>R</b>	<b>✓</b>	8	12 sec.
<b>ST87</b>	OMIT	$n \geq 2t \wedge f \leq t$	<b>U</b>	<b>✓</b>	0	1 sec.
<b>ST87</b>	OMIT	$n \geq 2t \wedge f \leq t$	<b>C</b>	<b>X</b>	0	2 sec.
<b>ST87</b>	OMIT	$n \geq 2t \wedge f \leq t$	<b>R</b>	<b>X</b>	0	2 sec.

# Summary of results

- Abstraction tailored for distributed algorithms
  - threshold-based
  - fault-tolerant
  - allows to express different fault assumptions
- Verification of threshold-based fault-tolerant algorithms
  - with threshold guards that are widely used
  - Byzantine faults (and other)
  - for all system sizes

## Related work: non-parameterized

Model checking of the small size instances:

- clock synchronization [Steiner, Rushby, Sorea, Pfeifer 2004]
- consensus [Tsuchiya, Schiper 2011]
- asynchronous agreement, folklore broadcast, condition-based consensus [John, Konnov, Schmid, Veith, Widder 2013]
- and more...

## Related work: parameterized case

Regular model checking of fault-tolerant distributed protocols:

[Fisman, Kupferman, Lustig 2008]

- “First-shot” theoretical framework.
- No guards like  $x \geq t + 1$ , only  $x \geq 1$ .
- No implementation.
- Manual analysis applied to folklore broadcast (**crash faults**).

## Related work: parameterized case

Regular model checking of fault-tolerant distributed protocols:

[Fisman, Kupferman, Lustig 2008]

- “First-shot” theoretical framework.
- No guards like  $x \geq t + 1$ , only  $x \geq 1$ .
- No implementation.
- Manual analysis applied to folklore broadcast (**crash faults**).

Backward reachability using SMT with arrays:

[Alberti, Ghilardi, Pagani, Ranise, Rossi 2010-2012]

- Implementation.
- Experiments on Chandra-Toueg 1990.
- No resilience conditions like  $n > 3t$ .
- Safety only.

# Our current work

Discrete synchronous	Discrete partially synchronous	Discrete asynchronous	Continuous synchronous	Continuous partially synchronous
One instance/ finite payload		<b>one-shot broadcast, c.b.consensus</b>		
		core of {ST87, BT87, CT96}, MA06 (common), MR04 (binary)		
Many inst./ finite payload				
Many inst./ unbounded payload				
Messages with reals				

# Future work: threshold guards + orthogonal features

	Discrete synchronous	Discrete partially synchronous	Discrete asynchronous	Continuous synchronous	Continuous partially synchronous
One instance/ finite payload			<b>one-shot broadcast, c.b.consensus</b> core of {ST87, BT87, CT96}, MA06 (common), MR04 (binary)		
Many inst./ finite payload	DHM12	CT96 (failure detector)	<b>broadcast</b> ST87, BT87, CT96, with failure-detectors	ELPS13	<b>clock sync</b> FSDK06
Many inst./ unbounded payload	ST87	DLS86, ST86, L98 (Paxos)		WS07	WS09
Messages with reals	AK00		<b>approx. agreement</b> DLPSW86	ST87 (JACM)	

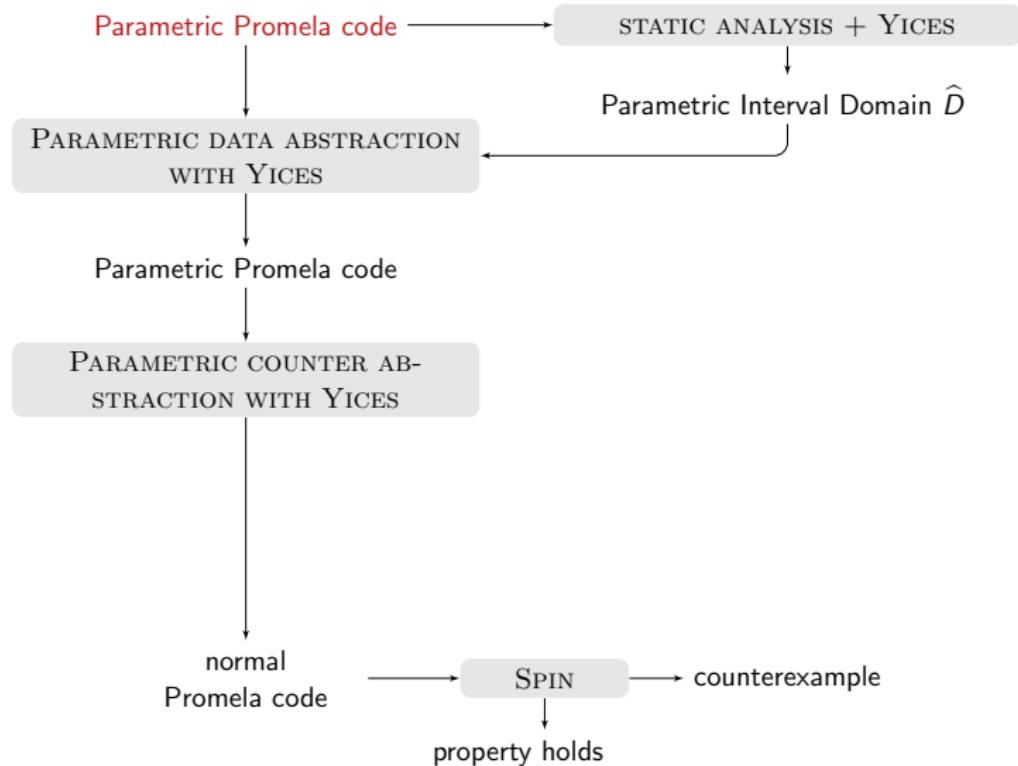
The diagram illustrates the relationship between various consensus and agreement protocols across different categories. Red arrows point from the 'approx. agreement' row to the 'clock sync' column, indicating a connection between them.

# Thank you!

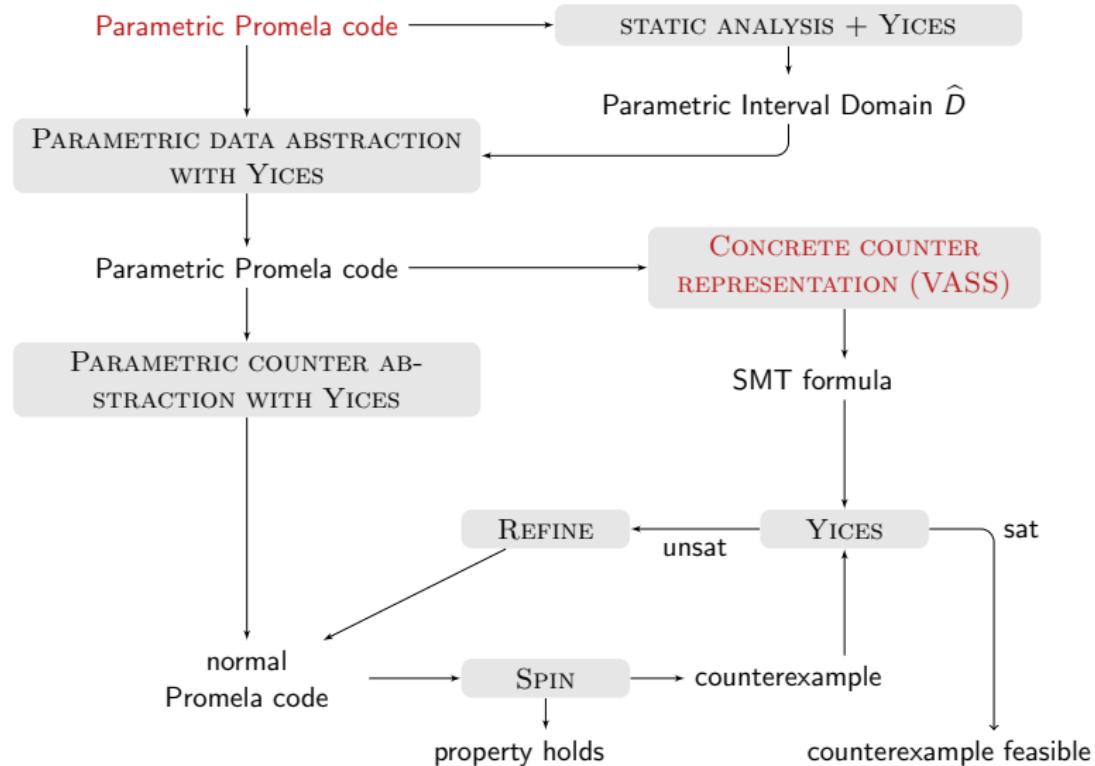
[ <http://forsyte.at/software/bymc> ]

# the implementation

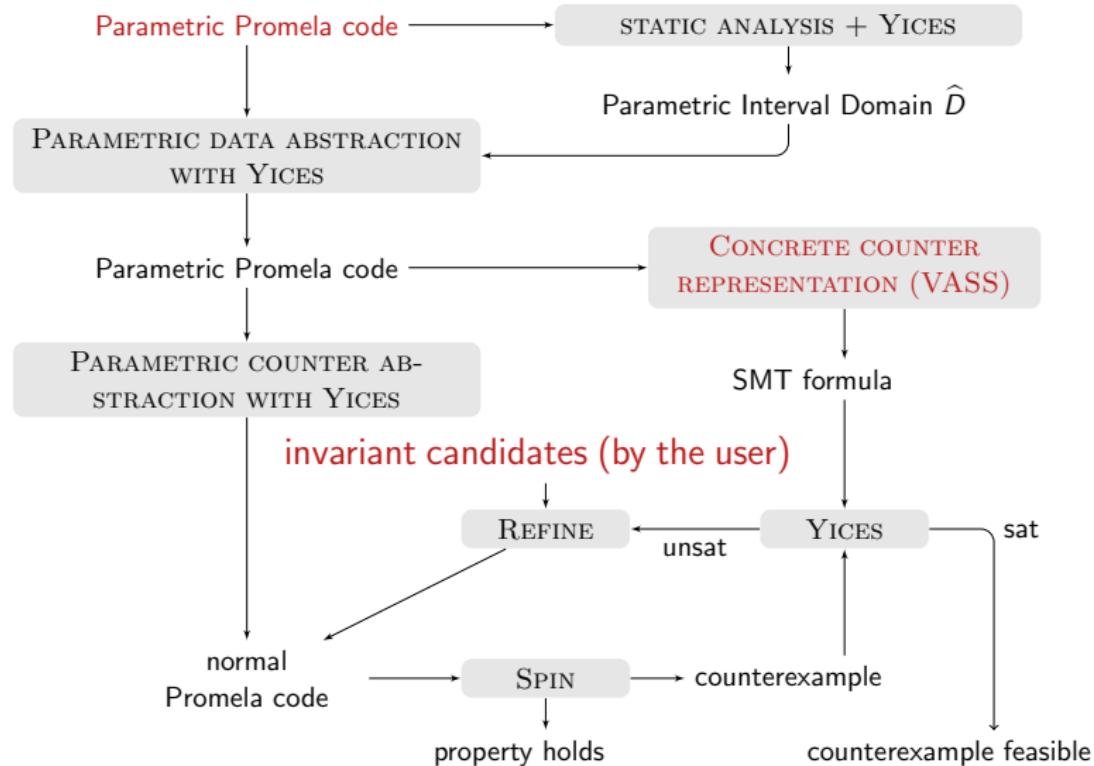
# Tool Chain: BYMC



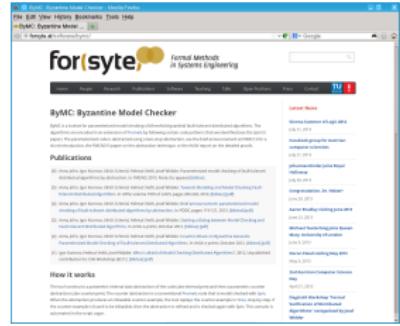
# Tool Chain: BYMC



# Tool Chain: BYMC



# Experimental setup

A screenshot of a Windows desktop environment showing a virtual machine window titled "Ubuntu 12.04 LTS (Precise Pangolin) (forsyte)" running on VMWare. The window displays the ByMC interface, which includes a sidebar with tabs like "File", "Edit", "Run", "Help", and "About". The main area shows a code editor with OCaml code and a terminal window below it showing command-line output.A screenshot of a terminal window titled "Terminal" showing a large amount of OCaml code. The code is annotated with comments explaining the steps of the algorithm, such as "Regenerating the counter abstraction", "Extracting the abstract domain", and "Constructing interval abstraction". The code uses various symbols like Proc, nrcrd, and scratch to represent different components of the algorithm.

The tool (source code in OCaml),  
the code of the distributed algorithms in Parametric Promela,  
and a virtual machine with full setup

are available at: <http://forsyte.at/software/bymc>

## Running the tool—concrete case

- user specifies parameter value
- useful to check whether the code behaves as expected
- `$bymc/verifyco-spin "N=4,T=1,F=1" bcast-byz.pml relay`
  - model checking problem in directory  
`“./x/spin-bcast-byz-relay-N=4,T=1,F=1”`
  - in `concrete.prm`
    - parameters are replaced by numbers
    - process prototype is replaced with  $N - F = 3$  active processes

# Running the tool—parameterized model checking

- PIA data and counter abstraction
- finite-state model checking on abstract model
- \$bymc/verifypa-spin bcast-omit.pml relay
  - model checking problem in directory  
“./x/bcast-byz-relay-yyymmdd-HHMM.\*”
  - directory contains
    - abs-interval.prm: result of the data abstraction;
    - abs-counter.prm: result of the counter abstraction;
    - abs-vass.prm: auxiliary abstraction for abstraction refinement;
    - mc.out: the last output by SPIN;
    - cex.trace: the counterexample (if there is one);
    - yices.log: communication log with YICES.

# Fairness, Refinement, and Invariants

- In the Byzantine case we have  $in\_transit : \forall i. (nrcvd_i \geq nsnt)$  and  $\mathbf{GF} \neg in\_transit$ .
- In this case communication fairness implies computation fairness.
- But in the abstract version  $nsnt$  can deviate from the number of processes who sent the echo message.
- In this case the user formulates a simple state invariant candidate, e.g.,  $nsnt = K([sv = SE \vee sv = AC])$  (on the level of the original concrete system).
- The tool checks automatically, whether the candidate is actually a state invariant.
- After the abstraction the abstract version of the invariant restricts the behavior of the abstract transition system.

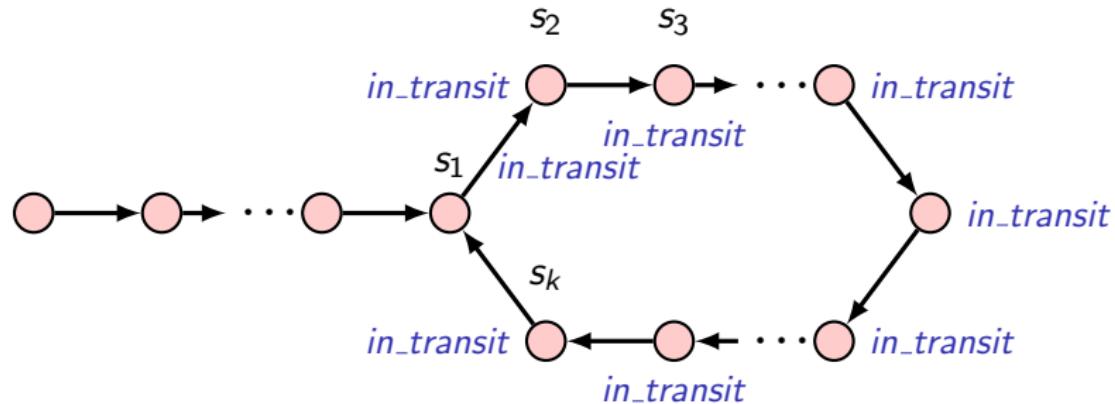
# Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \textit{in\_transit}$  necessary to verify liveness

# Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \text{in\_transit}$  necessary to verify liveness

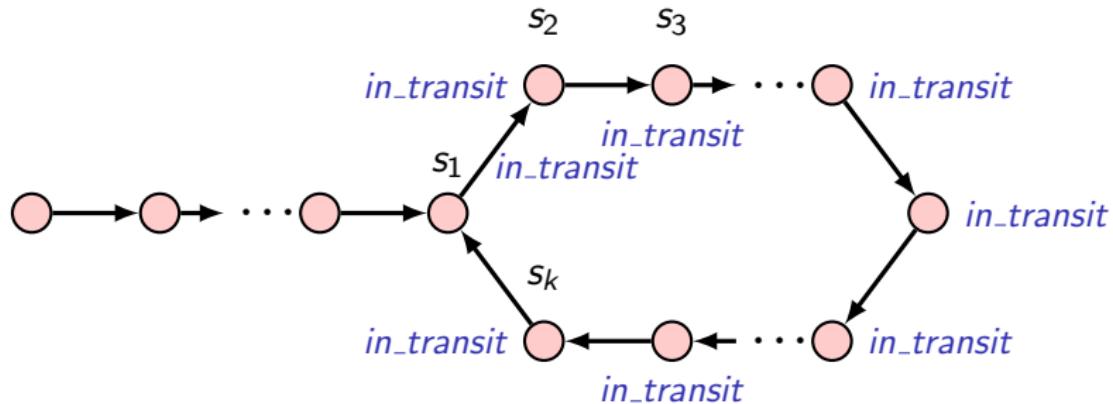
counter example:



if  $\forall j$  all concretizations of  $s_j$  violate  $\neg \text{in\_transit}$ , then CE is spurious.

# Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \text{in\_transit}$  necessary to verify liveness  
counter example:

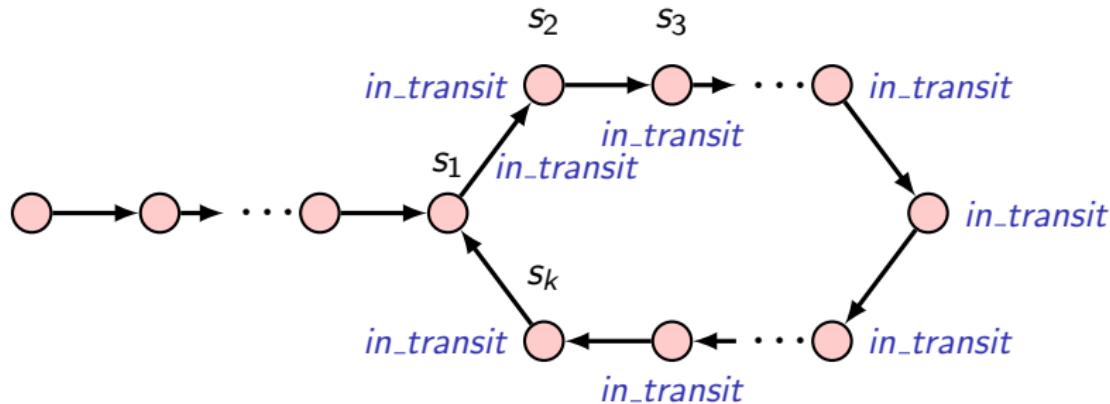


if  $\forall j$  all concretizations of  $s_j$  violate  $\neg \text{in\_transit}$ , then CE is spurious.

refine justice to  $\mathbf{G} \mathbf{F} \neg \text{in\_transit} \wedge \mathbf{G} \mathbf{F} \left( \bigvee_{1 \leq j \leq k} \neg \text{at}(s_j) \right)$

# Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \text{in\_transit}$  necessary to verify liveness  
counter example:



if  $\forall j$  all concretizations of  $s_j$  violate  $\neg \text{in\_transit}$ , then CE is spurious.

refine justice to  $\mathbf{G} \mathbf{F} \neg \text{in\_transit} \wedge \mathbf{G} \mathbf{F} \left( \bigvee_{1 \leq j \leq k} \neg \text{at}(s_j) \right)$

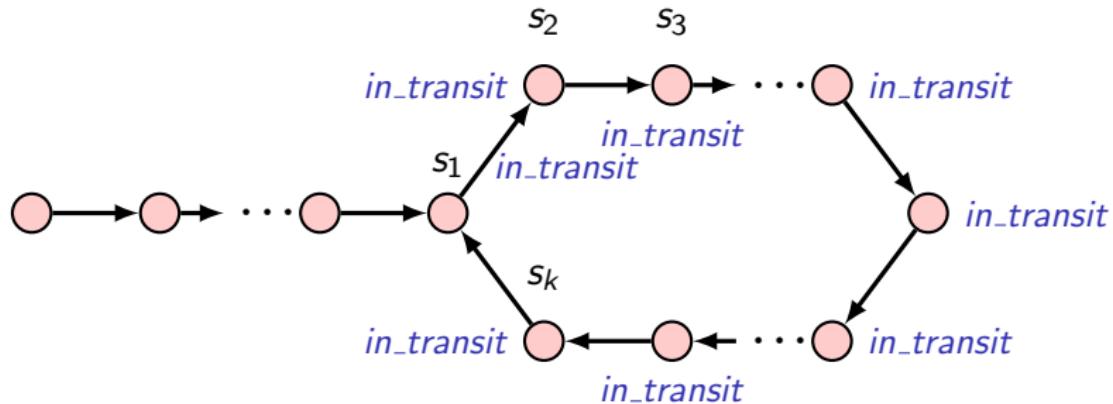
... we use unsat cores to refine several loops at once

## Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \textit{in\_transit}$  necessary to verify liveness

# Parametric abstraction refinement—justice suppression

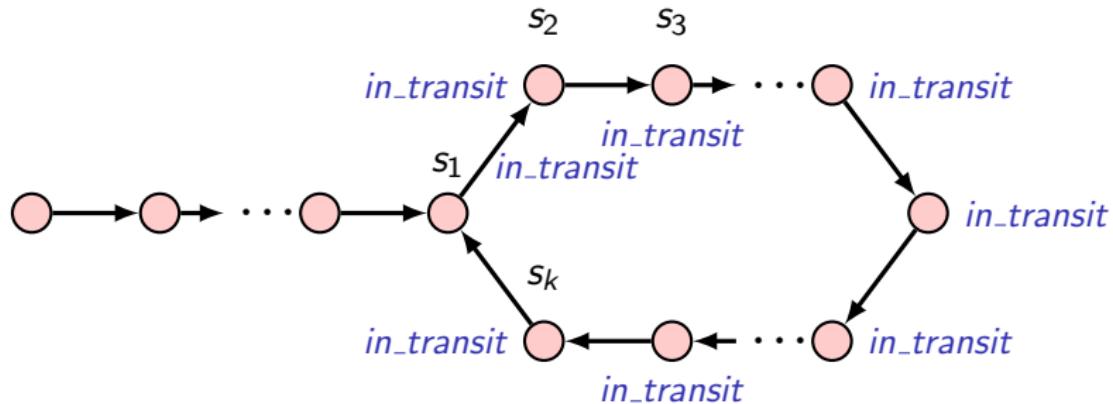
justice  $\mathbf{G} \mathbf{F} \neg \text{in\_transit}$  necessary to verify liveness  
counter example:



if  $\forall j$  all concretizations of  $s_j$  violate  $\neg \text{in\_transit}$ , then CE is spurious.

# Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \text{in\_transit}$  necessary to verify liveness  
counter example:

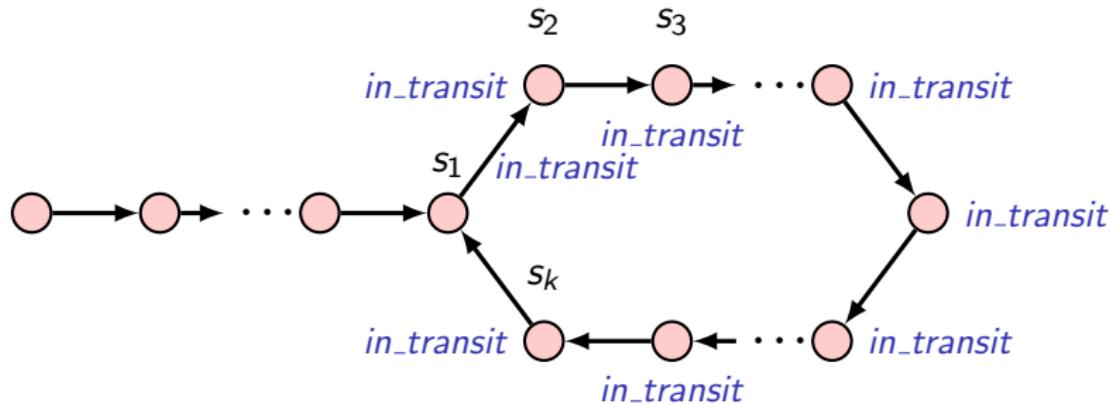


if  $\forall j$  all concretizations of  $s_j$  violate  $\neg \text{in\_transit}$ , then CE is spurious.

refine justice to  $\mathbf{G} \mathbf{F} \neg \text{in\_transit} \wedge \mathbf{G} \mathbf{F} \left( \bigvee_{1 \leq j \leq k} \neg \text{at}(s_j) \right)$

# Parametric abstraction refinement—justice suppression

justice  $\mathbf{G} \mathbf{F} \neg \text{in\_transit}$  necessary to verify liveness  
counter example:



if  $\forall j$  all concretizations of  $s_j$  violate  $\neg \text{in\_transit}$ , then CE is spurious.

refine justice to  $\mathbf{G} \mathbf{F} \neg \text{in\_transit} \wedge \mathbf{G} \mathbf{F} \left( \bigvee_{1 \leq j \leq k} \neg \text{at}(s_j) \right)$

... we use unsat cores to refine several loops at once

# asynchronous reliable broadcast (srikanth & toueg 1987)

the core of the classic broadcast algorithm from the da literature.  
it solves an agreement problem depending on the inputs  $v_i$ .

*Variables of process i*

$v_i$ : {0, 1} init with 0 or 1

$accept_i$ : {0, 1} init with 0

An indivisible step:

**if**  $v_i = 1$

**then** send (echo) to all;

**if** received (echo) from at least

**t + 1** distinct processes

**and not** sent (echo) before

**then** send (echo) to all;

**if** received (echo) from at least

**n - t** distinct processes

**then**  $accept_i := 1$ ;

# asynchronous reliable broadcast (srikanth & toueg 1987)

the core of the classic broadcast algorithm from the da literature.  
it solves an agreement problem depending on the inputs  $v_i$ .

Variables of process  $i$

$v_i$ : {0, 1} init with 0 or 1

$accept_i$ : {0, 1} init with 0

asynchronous

An indivisible step:

**if**  $v_i = 1$   
**then** send (echo) to all;

$t$  byzantine faults

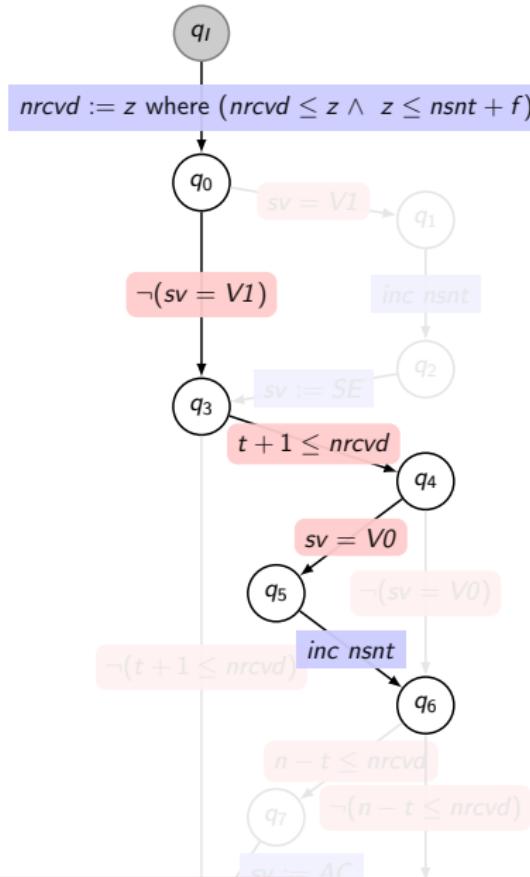
**if** received (echo) from at least  
 $t + 1$  distinct processes  
and not sent (echo) before  
**then** send (echo) to all;

correct if  $n > 3t$   
resilience condition rc

**if** received (echo) from at least  
 $n - t$  distinct processes  
**then**  $accept_i := 1$ ;

parameterized process  
skeleton  $p(n, t)$

# Abstract CFA



# Abstract CFA

