

Synthesis of Distributed Algorithms with Parameterized Threshold Guards

Igor Konnov

Marijana Lazić



Josef Widder



Roderick Bloem




OPODIS'17, December 2017

Verifying fault-tolerant systems

safety critical systems: cars, planes, etc.


- rare but dangerous faults
 - 3 to 7 processes
-

finite-state model checking




datacenters: thousands of computers

- faults happen every day
 - 100–10,000 processes
-



parameterized model checking

Fault-tolerant distributed algorithms



n processes communicate by sending messages

f processes are faulty (unknown)

t is an upper bound on f (known)

resilience condition on n , t , and f ,

e.g., $n > 3t \wedge t \geq f \geq 0$

Reliable broadcast service (informally)

one process broadcasts a message **bcast**

correctness: if all correct processes received **bcast**,
then some correct process **eventually accepts bcast**

111...1

relay: if a correct process accepts **bcast**,
then all correct processes **eventually accept bcast**

011...1

unforgeability: if no correct process received **bcast**,
then no correct process ever **accepts bcast**

000...0

fairness: every sent message is eventually received

Reliable broadcast by Srikanth & Toueg 87

local $myval_i \in \{0, 1\}$ - did process i receive **bcast**?

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
 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
od

resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

Reliable broadcast by Srikanth & Toueg 87

local $myval_i \in \{0, 1\}$ - did process i accept a value

a threshold guard

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
then send ECHO to all

if received ECHO from at least $t + 1$
and not sent ECHO before
then send ECHO to all

distinct processes

if received ECHO from at least $n - t$
then accept
od

a threshold guard

resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

More threshold guards...

Reliable broadcast	$x \geq t + 1$ $x \geq n - t$	[Srikanth, Toueg'86]
Hybrid broadcast	$x \geq t_b + 1$ $x \geq n - t_b - t_c$	[Widder, Schmid'07]
Byzantine agreement	$x \geq \lceil \frac{n}{2} \rceil + 1$	[Bracha, Toueg'85]
Non-blocking atomic commitment	$x \geq n$	[Raynal'97], [Guerraoui'01]
Condition-based consensus	$x \geq n - t$ $x \geq \lceil \frac{n}{2} \rceil + 1$	[Mostéfaoui, Mourgaya, Parvedy, Raynal'03]
Consensus in one communication step	$x \geq n - t$ $x \geq n - 2t$	[Brasileiro, Greve, Mostéfaoui, Raynal'03]
Byzantine one-step consensus	$x \geq \lceil \frac{n+3t}{2} \rceil + 1$	[Song, van Renesse'08]


In general, there is a resilience condition, e.g., $n > 3t$, $n > 7t$

Byzantine model checker

[forsyte.at/software/bymc]

(source code, benchmarks, virtual machines, etc.)

10 **parameterized** fault-tolerant distributed algorithms:



[CAV'15] & [POPL'17]



From verification to synthesis

Different threshold guards for one sketch

local $myval_i \in \{0, 1\}$ - did process i receive **bcast**?

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
 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
od



resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

Different threshold guards for one sketch

local $myval_i \in \{0, 1\}$ - did process i receive **bcast**?

while true **do**

if $myval_i = 1$ **and not** sent ECHO before

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

od

$t + 1$

$2t + 1$



resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

Different threshold guards for one sketch

local $myval_i \in \{0, 1\}$ - did process i receive **bcast**?

while true **do**

if $myval_i = 1$ **and not** sent ECHO before

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

od

$t + 1$ $n - 2t$

$2t + 1$ $n - t$



resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

Different threshold guards for one sketch

local $myval_i \in \{0, 1\}$ - did process i receive **bcast**?

while true **do**

if $myval_i = 1$ **and not** sent ECHO before

then send ECHO to all

if received ECHO from at least $t + 1$

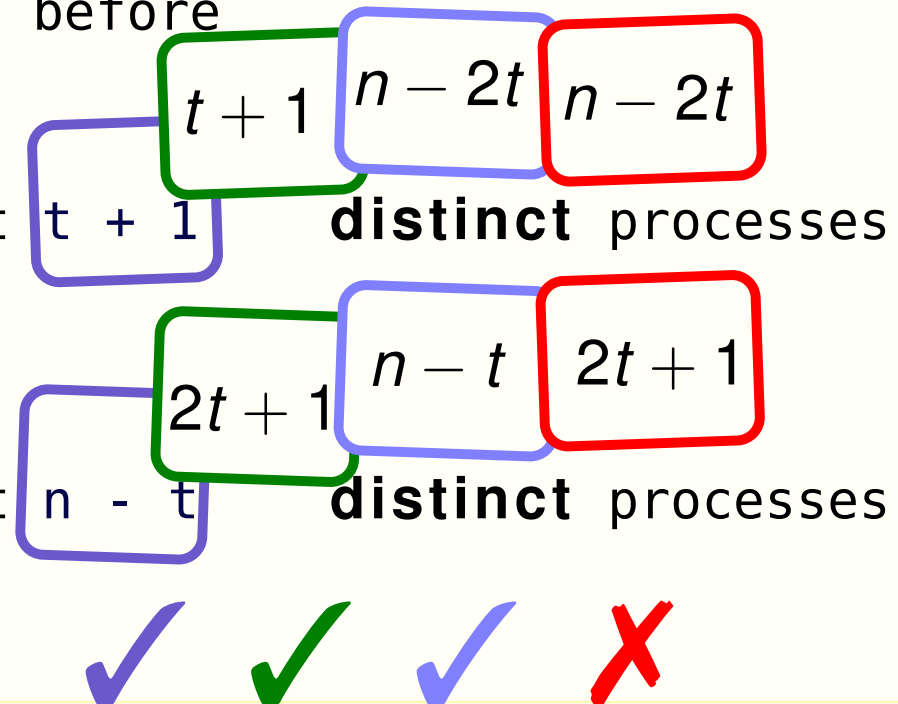
and not sent ECHO before

then send ECHO to all

if received ECHO from at least $n - t$

then accept

od



resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

Find thresholds automatically?

local $myval_i \in \{0, 1\}$ - did process i receive **bcast**?

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
then send ECHO to all

if received ECHO from at least
and not sent ECHO before
then send ECHO to all

if received ECHO from at least
then accept
od

$?$ ^{k}

$?_1 \cdot n + ?_2 \cdot t + ?_3$
distinct processes

$?$ ^{k}

$?_4 \cdot n + ?_5 \cdot t + ?_6$
distinct processes


resilience: of $n > 3t$ processes, $f \leq t$ processes are Byzantine

Synthesis loop


Find $?_1, \dots, ?_k \in \mathbb{Q}$



Generator
infinite
search space



Verifier
Byzantine MC



Synthesis problem

Find a distributed algorithm that satisfies spec φ

for all parameter values n , t , and f that satisfy resilience condition

Input:

sketch algorithm,

if received $?_1 \cdot n + ?_2 \cdot t + ?_3$ echoes
then send echo to all

specification,

e.g., unforgeability, correctness & relay

resilience condition,

e.g., $n > 3t$, $t \geq f \geq 0$

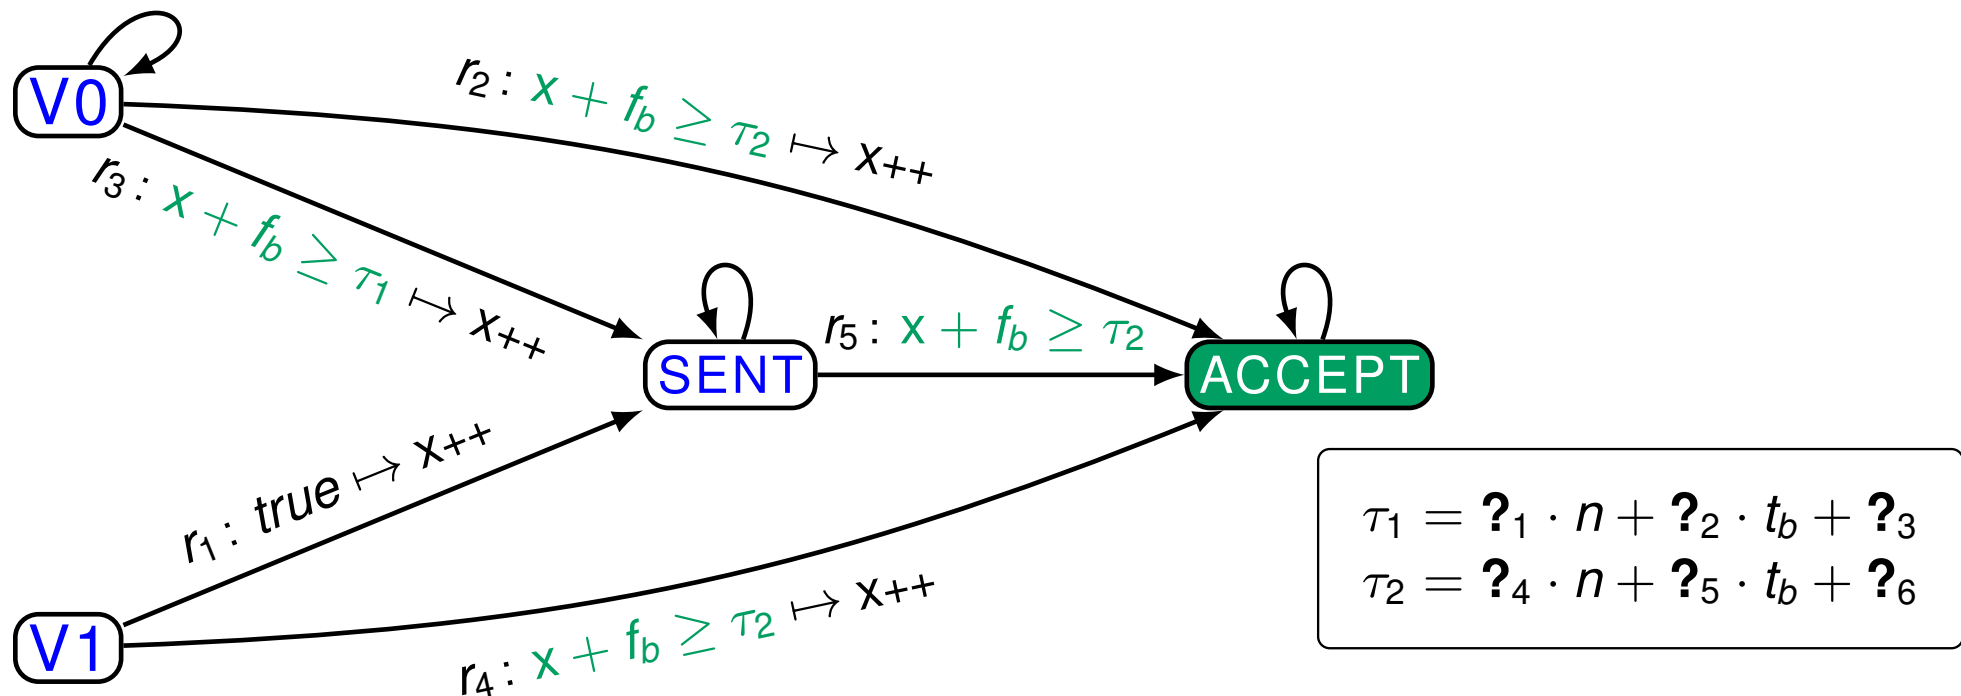
Find (if exist): $a_1, \dots, a_k \in \mathbb{Q}$ for $?_1, \dots, ?_k$

Sketch threshold automata to capture the pseudo-code

Linear temporal logic to formalize the specifications

Linear integer arithmetic to express the resilience condition

Sketch threshold automata




Byzantine faults:

run $n - f_b$ processes,

count messages modulo Byzantine processes, e.g., $x + f_b \geq \tau_2$

Sketch threshold automata



$f_b \leq t_b$ **Byzantine** and $f_c \leq t_c$ **crash faults**:

run $n - f_b$ processes,

resilience condition e.g. $n > 3t_b + 2t_c \wedge t_b \geq 0 \wedge t_c \geq 0$

Relay:

if a correct process accepts **bcast**,
then all correct processes eventually accept **bcast**

Relay:

~~if~~ a correct process accepts **bcast**,
~~then all correct processes eventually accept **bcast**~~
and at least one process never accept **bcast**

Relay:

- if a correct process accepts **bcast**
 - ~~then all correct processes eventually accept **bcast**~~
 - and at least one process never accepts **bcast**
-

$$\mathbf{E} \left(\mathbf{F} \left(\kappa_{\text{ACCEPT}} \neq 0 \right) \wedge \mathbf{G} \left(\kappa_{V1} \neq 0 \vee \kappa_{V0} \neq 0 \vee \kappa_{\text{SENT}} \neq 0 \right) \right) \wedge \mathbf{G} \mathbf{F} \psi_{\text{fair}} \right)$$

Relay:

~~if~~ a correct process accepts **bcast**
~~then all correct processes eventually accept **bcast**~~
and at least one process never accepts **bcast**

$$\mathbf{E} \left(\mathbf{F} \left(\kappa_{\text{ACCEPT}} \neq 0 \wedge \mathbf{G} \left(\kappa_{V1} \neq 0 \vee \kappa_{V0} \neq 0 \vee \kappa_{\text{SENT}} \neq 0 \right) \right) \wedge \mathbf{G} \mathbf{F} \psi_{\text{fair}} \right)$$

Propositional formulas:

- (1) $\bigwedge_{\ell \in S} \kappa_{\ell} = 0$
- (2) $\bigvee_{\ell \in S} \kappa_{\ell} \neq 0$
- (3) $\bigwedge_{S \subseteq T} \bigvee_{\ell \in S} \kappa_{\ell} \neq 0$
- (4) $\text{Bool}(\text{Guards}) \rightarrow (1) \wedge (2) \wedge (3)$

Temporal formulas:

$$\psi ::= \text{prop} \mid \mathbf{G} \psi \mid \mathbf{F} \psi \mid \psi \wedge \psi$$

We call this fragment ELTL_{FT}


Our solution to synthesis

Synthesis loop

Find $?_1, \dots, ?_k \in \mathbb{Q}$



Generator
infinite
search space



Verifier
Byzantine MC



solution

Termination?

sane guards \Rightarrow bounded search space

Efficiency?

Generator learns from counterexamples

Sane guards: thresholds lie in the interval $[0, n]$

Classic threshold guards:

if received $\frac{n}{2}$ messages... ✓

wait for a majority

if received $t + 1$ messages... ✓

wait for a correct process

if received $n - t$ messages... ✓

wait for non-faulty processes

Syntactically correct but meaningless guards:

if received $2n$ messages... ✗

if received -5 messages... ✗

Search space for sane guards

Resilience condition: $n > 3t > 0$

Threshold: $0 \leq ?_a n + ?_b t + ?_c \leq n$

\Rightarrow

$$0 \leq ?_a \leq 1$$

$$-4 \leq ?_b \leq 4$$

$$-8 \leq ?_c \leq 8$$

Theorem

Assume: $n > \sum_{1 \leq i \leq k} \delta_i \cdot t_i$ and $\forall i. t_i \geq 0$ — resilience cond.

$0 \leq ?_a n + \sum_{1 \leq i \leq k} ?_{b_i} \cdot t_i + ?_c \leq n$ — threshold

Then:

$$\left\{ \begin{array}{ll} 0 \leq ?_a \leq 1 \\ -B_i \leq ?_{b_i} \leq B_i & \text{for } B_i = \delta_i + 1 \text{ and } 1 \leq i \leq k \\ -C \leq ?_c \leq C & \text{for } C = k + 1 + 2(\delta_1 + \dots + \delta_k) \end{array} \right.$$

Search space for sane guards

$$\begin{array}{ll}
 \text{Resilience condition: } n > 3t > 0 & \\
 \text{Threshold: } 0 \leq \alpha n + \beta t + \gamma \leq n & \Rightarrow \begin{array}{l} 0 \leq \alpha \leq 1 \\ -4 \leq \beta \leq 4 \\ -8 \leq \gamma \leq 8 \end{array}
 \end{array}$$

Theorem

Assume: $n > \sum_{1 \leq i \leq k} \delta_i \cdot t_i$ and $\forall i. t_i \geq 0$ — resilience cond.

$$0 \leq \alpha n + \sum_{1 \leq i \leq k} \beta_i \cdot t_i + \gamma \leq n \quad \text{— threshold}$$

Then:

$$\left\{ \begin{array}{ll} 0 \leq \alpha \leq 1 \\ -B_i \leq \beta_i \leq B_i & \text{for } B_i = \delta_i + 1 \text{ and } 1 \leq i \leq k \\ -C \leq \gamma \leq C & \text{for } C = k + 1 + 2(\delta_1 + \dots + \delta_k) \end{array} \right.$$

From \mathbb{Q} to a finite search space

Synthesizing thresholds of the form $\frac{n}{2}$ or $\frac{2n}{3}$:

Assume:

Resilience condition: $n > 3t \geq 0$

Threshold: $0 \leq \frac{?_a'}{6}n + \frac{?_b'}{6}t + \frac{?_c'}{6} \leq n$

Then

$$0 \leq ?_a' \leq 6 \cdot 1$$

$$-6 \cdot 4 \leq ?_b' \leq 6 \cdot 4$$

$$-6 \cdot 8 \leq ?_c' \leq 6 \cdot 8$$

$$?_a', ?_b', ?_c' \in \mathbb{Z}$$

Explicit enumeration?

coefficients $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6)$ in reliable broadcast:

$$(2 \cdot 9 \cdot 17)^2 \text{ vectors}$$

coefficients $(0, \alpha_2, 1, 0, \alpha_5, 1)$ in reliable broadcast:

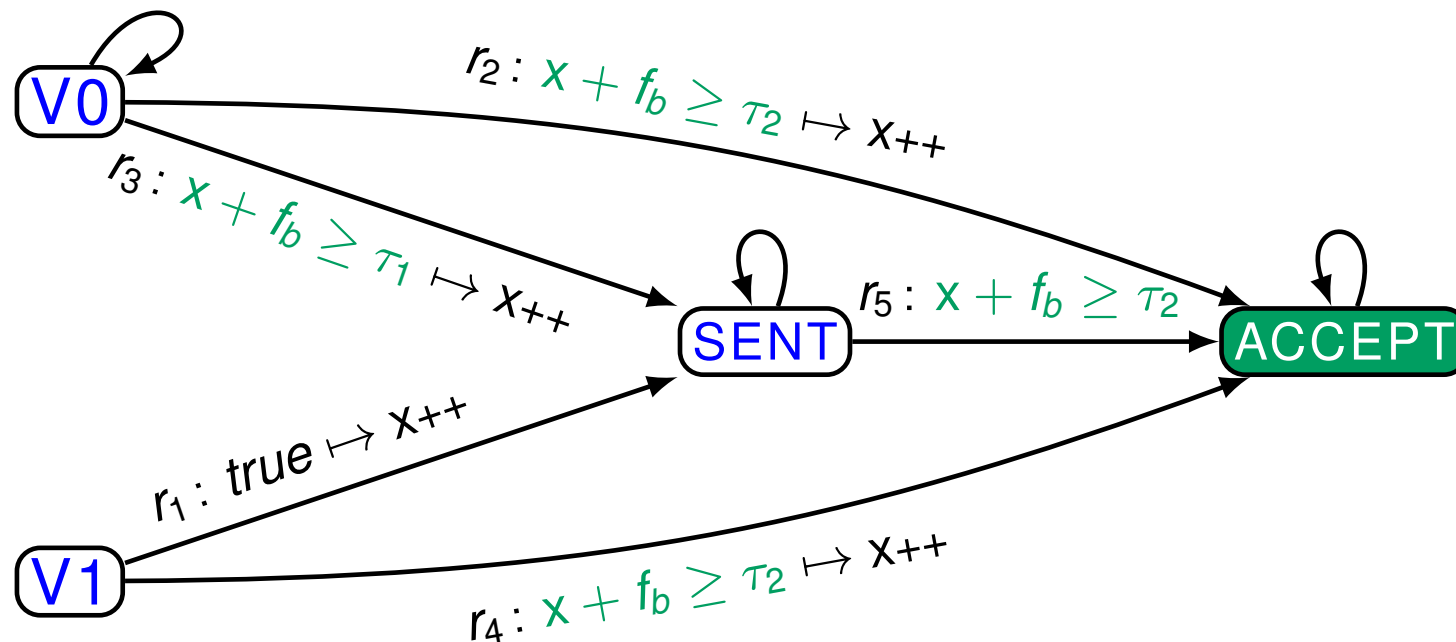
$$9 \cdot 9 \text{ vectors}$$

coefficients $(\alpha_1, \alpha_2, \dots, \alpha_9)$ in one-step consensus (BOSCO):

$$(3 \cdot 17 \cdot 33)^3 \text{ vectors}$$

The generator should learn from the counterexamples!

Sketch of reliable broadcast in 2D




missing coefficients for t :

$$\tau_1 = ?_2 \cdot t + 1$$

$$\tau_2 = ?_5 \cdot t + 1$$

bounded search space



Counterexample to unforgeability

Model checker flags an error:

$$n = 4, t = 1, \text{ and } f = 1$$

State 1. $\kappa_{V0} = 3$, other counters 0




$$\Downarrow r_3 : 0 + 1 \geq 0 \cdot 1 + 1$$

State 2. $\kappa_{\text{SENT}} = 1, x = 1, \kappa_{V0} = 2$

$$\Downarrow r_5 : 1 + 1 \geq 0 \cdot 1 + 1$$

State 3. $\kappa_{\text{ACCEPT}} = 1, \kappa_{\text{SENT}} = 0$

$$?_2 = 0 \text{ and } ?_5 = 0$$



Counterexample to unforgeability

Model checker flags an error:

$$n = 4, t = 1, \text{ and } f = 1$$

State 1. $\kappa_{V0} = 3$, other counters 0




$$\begin{aligned} \Downarrow \text{ } & \cancel{r_3 : 0 + 1 \geq 0 \cdot 1 + 1} \\ \Downarrow \text{ } & r_3 : 0 + 1 \geq ?_2 \cdot 1 + 1 \end{aligned}$$

State 2. $\kappa_{\text{SENT}} = 1$, $x = 1$, $\kappa_{V0} = 2$

$$\begin{aligned} \Downarrow \text{ } & \cancel{r_5 : 1 + 1 \geq 0 \cdot 1 + 1} \\ \Downarrow \text{ } & r_5 : 1 + 1 \geq ?_5 \cdot 1 + 1 \end{aligned}$$

State 3. $\kappa_{\text{ACCEPT}} = 1$, $\kappa_{\text{SENT}} = 0$

$$?_2 = 0 \text{ and } ?_5 = 0$$



Counterexample to unforgeability

Model checker flags an error:

$$n = 4, t = 1, \text{ and } f = 1$$

State 1. $\kappa_{V0} = 3$, other counters 0




$$\begin{aligned} &\Downarrow \text{ ~~$r_3 : 0 + 1 \geq 0 \cdot 1 + 1$~~ } \\ &\Downarrow r_3 : 0 + 1 \geq ?_2 \cdot 1 + 1 \end{aligned}$$

State 2. $\kappa_{\text{SENT}} = 1, x = 1, \kappa_{V0} = 2$


$$\begin{aligned} &\Downarrow \text{ ~~$r_5 : 1 + 1 \geq 0 \cdot 1 + 1$~~ } \\ &\Downarrow r_5 : 1 + 1 \geq ?_5 \cdot 1 + 1 \end{aligned}$$

State 3. $\kappa_{\text{ACCEPT}} = 1, \kappa_{\text{SENT}} = 0$

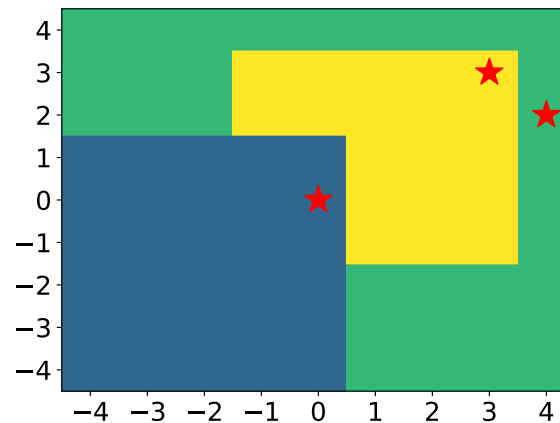
$$?_2 = 0 \text{ and } ?_5 = 0$$




Learning from counterexamples



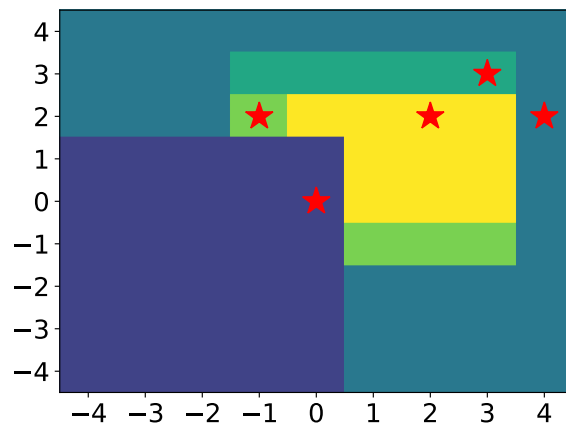
1. unforgability




2. sanity




3. correctness



4. sanity



5. relay



6. relay


found the solution $?_2 = 1$ and $?_5 = 2$

Synthesis loop


Find $?_1, \dots, ?_k \in \mathbb{Q}$



Generator
infinite
search space



Verifier
Byzantine MC




Experiments

We have synthesized

reliable broadcast, hybrid broadcast, and
BOSCO (one-step consensus) using



&



© VSC / Claudia Blaas-Schneider

[forsyte.at/software/bymc]

SYNTHESIZED
Correct-By-Construction

Thresholds for Byzantine reliable broadcast

local $myval_i \in \{0, 1\}$

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
then send ECHO to all

if received ECHO from at least $t + 1$
and not sent ECHO before
then send ECHO to all

if received ECHO from at least $n - t$
then accept
od

3 solutions

16 seconds

31 calls to verifier

$t + 1$

$n - 2t$

distinct processes

$2t + 1$

$n - t$

distinct processes



resilience: $n > 3t, f \leq t$ **Byzantine** faults

Thresholds for Byzantine reliable broadcast

local $myval_i \in \{0, 1\}$

0 solutions

7 seconds

25 calls to verifier

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
then send ECHO to all

if received ECHO from at least
and not sent ECHO before
then send ECHO to all



distinct processes

if received ECHO from at least
then accept



distinct processes

od

$n \geq 3t$

resilience: $n > 3t, t \leq t$ **Byzantine** faults

Thresholds for hybrid reliable broadcast

local $myval_i \in \{0, 1\}$

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
 then send ECHO to all

if received ECHO from at least $t_b + 1$ **distinct** processes
 and not sent ECHO before
 then send ECHO to all

if received ECHO from at least $2t_b + t_c + 1$ **distinct** processes
 then accept
od

3 solutions

50 seconds

34 calls to verifier

$$t_b + 1$$

$$n - 2t_b - 2t_c$$

distinct processes

$$n - t_b + t_c$$

$$n - t_b - t_c$$

$2t_b + t_c + 1$ **distinct** processes



resilience: $n > 3t_b + 2t_c$ $f_b \leq t_b$ **Byzantine** and $f_c \leq t_c$ **crash** faults

Thresholds for hybrid reliable broadcast

local $myval_i \in \{0, 1\}$

0 solutions

24 seconds

29 calls to verifier

while true **do**

if $myval_i = 1$ **and not** sent ECHO before
then send ECHO to all

if received ECHO from at least
and not sent ECHO before
then send ECHO to all



distinct processes

if received ECHO from at least
then accept



distinct processes

od


$$n > 3t_b + t_c$$

resilience: $n > 3t_b + 2t_c$ $f_b \leq t_b$ **Byzantine** and $f_c \leq t_c$ **crash** faults

Reliable broadcast: changing specifications

$$\leq X$$

unforgeability: if ~~no~~ correct process received **bcast**,
then no correct process ever **accepts bcast**

correctness and **relay** like before

$$\underline{X = 2}$$

no solutions for $n > 3t$

3 solutions for $n > 3t + 2$


$$\underline{X = t}$$

no solutions for $n > 3t$

3 solutions for $n > 4t$

each answer found in less than **12 sec.**

BOSCO (one-step consensus)



$$\left. \begin{array}{l} \text{Agreement} \\ \text{Termination} \end{array} \right\} \text{when } n > 3t$$
$$\left. \begin{array}{l} \text{One step} \\ \text{Fast termination} \end{array} \right\} \text{when } n > 7t \quad \text{or} \quad n > 5t, f = 0$$

Found 4 solutions using 4 cluster nodes = 64 cores **24 min.**

No solutions for $n \geq 5t, f = 0$ and $n \geq 7t$ **40 min.**

(the conditions $n > 5t$ and $n > 7$ are tight)

Conclusions

we can **verify** and **synthesize** distributed algorithms that are:

- asynchronous and parameterized,
- subject to faults, sending to all, and
- counting messages and comparing to threshold guards



next steps:

- learning from positive examples,
- geometrical structure of learned regions,
- synthesizing threshold automata, not just thresholds



[forsyte.at/software/bymc]