Parameterized Protocol Analysis and Causal Reasoning

Richard Trefler

D.R. Cheriton School of Comp. Sci.
U. of Waterloo

Motivation I

Model checking/analysis of parameterized protocol designs provides assurance that designs meet their specifications.

Analysis failures/traces may represent program bugs.

Use causal reasoning to explain (complex) bug traces.

Motivation II

Model checkers provide assurance that systems behave as intended.

State explosion in the size of the analyzed model is a significant barrier.

Idea: use symmetry of the system to alleviate state explosion.

Motivation III

Analysis of parameterized protocols is challenging — in general, undecidable.

Goal: use causal reasoning to explain proof failures in analysis of symmetry reduced models.

Example: Dining Philosophers

Process i

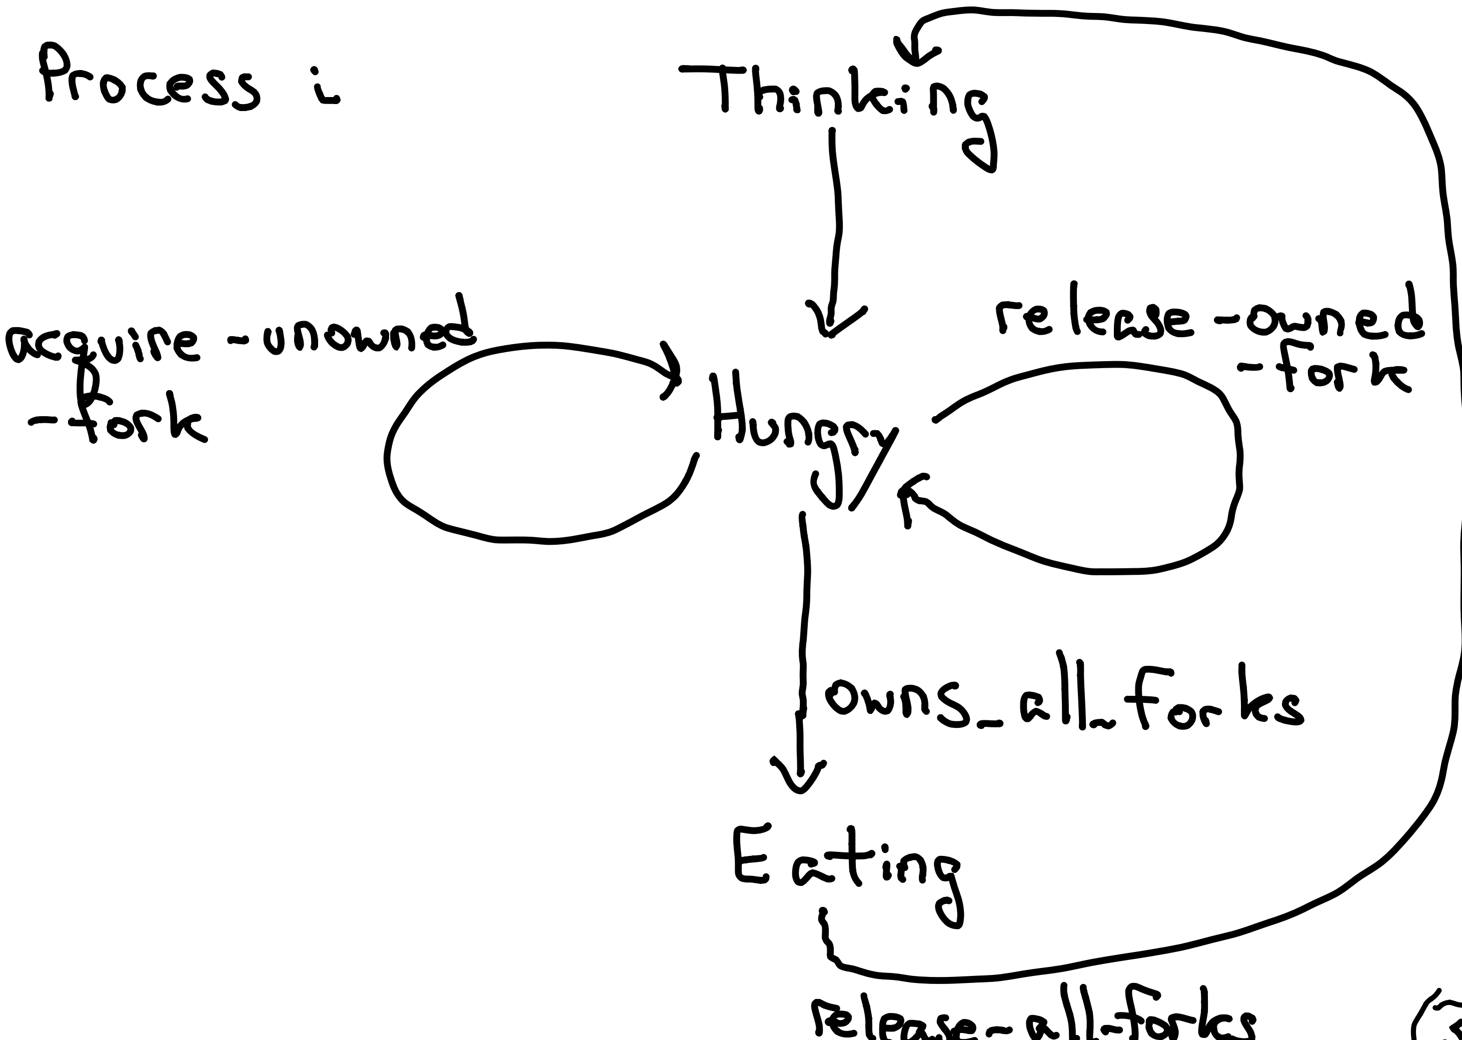

PS

DP

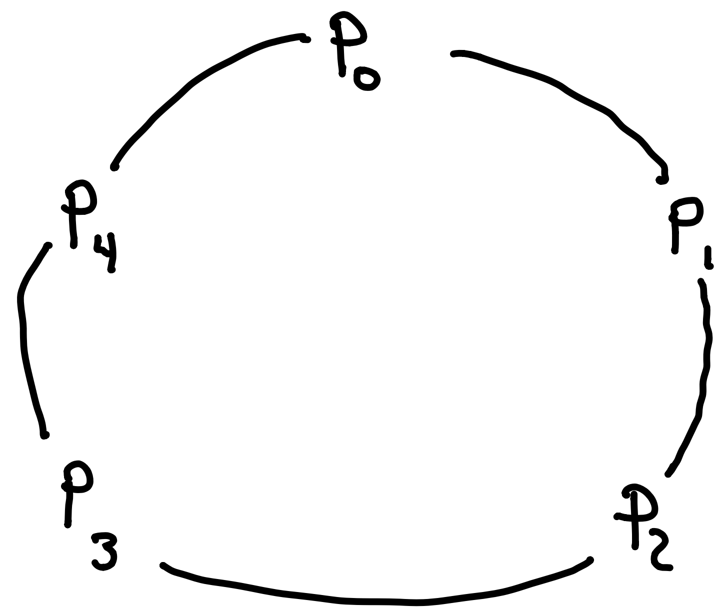

Isomorphic processes

Symmetry [ES96][CEFJ96]:
permutations / automorphisms of PS

Symmetry Reduction

Idea: symmetries define equivalence classes of processes and process state spaces.

Analyze/check only the representatives of the equivalence classes

[ES96][CEFJ96] c.f. [NT12/16/18][ACT19]

Challenges

$$\begin{tikzcd}
& T \arrow[dl] \arrow[dr] & \\
E & & E \\
T \arrow[u] \arrow[rr] & & T \arrow[u]
\end{tikzcd}$$

P_0 is isomorphic to P_1 yet $(0\ 1)$ is not a symmetry of PS.

Challenges

- P5 has only rotational symmetry — global symmetry affords only polynomial reduction.

- Symmetries of P5 are not symmetries of P6.

Local Symmetry [NT12, 15, 16, 18, AGT19]

Processes P_m and P_n have neighborhood symmetry when they have isomorphic programs

isomorphic neighborhood structure

isomorphic local state spaces

— this includes interference from neighbors.

→ 10

Local Symmetry

Symmetry is described by a **balance relation** $\{(m, \beta, n)\}$ — equivalence classes of symmetric processes all of whose neighbors are related by the balance relation.

DP Example

- P5 has a single equivalence class containing all P_i
- P6 has a single equivalence class containing all P_i
- P_1 of P5 is locally similar to P_1 of P6.
- There is a single equivalence class for $\{P_k\}_{k>2}$

Compositional Invariance

$P = \|_{i \in [1..k]} P_i$ process network with k processes

Init: $[I_n(x_n) \to \theta_n(x_n)]$

Step: $[\theta_n(x_n) \wedge T_n(x_n, y_n) \to \theta_n(y_n)]$

Non-Interference: for all $m \in nbr(n)$

$$[\theta_n(x_n) \wedge \theta_m(x_m) \wedge T_m(x, y) \to \theta_n(y_n)]$$

Theorem [NT12] If $\{\theta_n\}$ is a compositional inductive invariant then $\bigwedge_i \theta_i$ is a global inductive invariant.

Local and Global State Spaces

$P = \underset{i \in [1..k]}{\|} P_i$ network program

$G = (S, S_0, R)$ - global state space of P

H_n^{Θ} : restriction of G onto the local states of the neighborhood of P_n that respect Θ — includes interference.

Bisimulation between Local State Spaces

Theorem: Let B be a balance relation respected by both network $P = \| P_i$ and compositional inductive invariant Θ:

For all $(m, \beta, n) \in B$:

H_m^Θ is bisimilar to H_n^Θ

Local mu-calculus

- Atomic propositions local to a process: b
- Propositional variables local to a process: Z
- $\neg \varphi \qquad \varphi \wedge \psi \qquad \varphi \vee \psi$
- $E[\varphi \, U_a \, \psi] \qquad A[\varphi \, W_a \, \psi]$
- $\mu Z. \, \varphi(Z) \qquad \nu Z. \, \varphi(Z)$

Corollary

If B and $P = \parallel P_i$ respect Θ with $(m, \beta, n) \in B$ and $f(i)$ is a parametric mu-calculus formula then

$$H_m^{\Theta}, s \models f(m) \text{ iff } H_n^{\Theta}, \beta(s) \models f(n).$$

Local – Global Simulation

Assume that the processes in $P = \|_i P_i$ have unconditionally fair scheduling.

Theorem: For each P_m in $P = \|_{i \in [1..k]} P_i$ the local space H_m^{θ} simulates the global space G_m up to stuttering.

Corollary

If $f(m)$ is a universal local mu-calculus formula then for all global states, s, and local states t such that $s[m] = t$:

$$H^{\theta}_{m}, t \models f(m) \text{ implies } G_m, s \models f(m).$$

Outward Facing Interactions

Idea: restrict interference conditions to depend only on shared state.

Theorem: With outward facing interactions $H_m^{\otimes}$ is stuttering bisimilar to G_m.

Corollary: $H_m^{\otimes}$ and G_m satisfy the same local mu-calculus formulae.

Model Checking Strategy

1. Find/determine a balance relation, B, of local symmetries for program $P = \|_i P_i$, or for network family $\{P_k\}_{k \geq min}$

2. Model check local specification $f(m)$ on representative P_m for each equivalence class of program or family.

Program Symmetries

P_m is equivalent to P_n if there is a bijective mapping from T_m to T_n.

A permutation, π, of process indices is an automorphism of $P = \|_{i \in [1..k]} P_i$ if P_m is equivalent to $P_{\pi(m)}$ for all $m \in [1..k]$.

Syntactic Symmetry Reduction

CR — communication relation of $P = \|_i P_i$

Each P_i implements template W.

Syntactic global symmetries of CR define global symmetries of G.

Global symmetries $\longrightarrow$ local symmetries.

Can provides exponential savings — e.g. ring/tori CR have single rep. nodes.

Symmetry Patterns

Many 'regular' networks have only little global symmetry — rings, etc.

Enforce local symmetry through patterns of tiles.

Each tile type describes a local neighborhood that enforces neighboring types of interference.

Examples

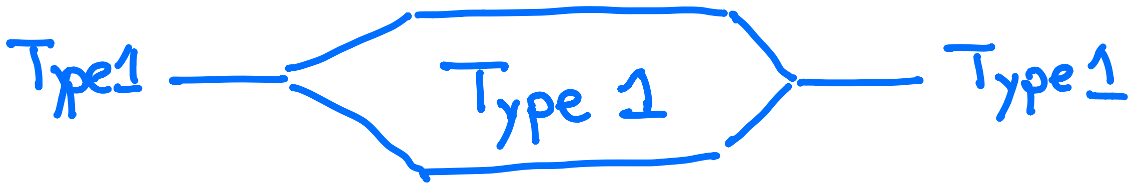

Ring networks

Example

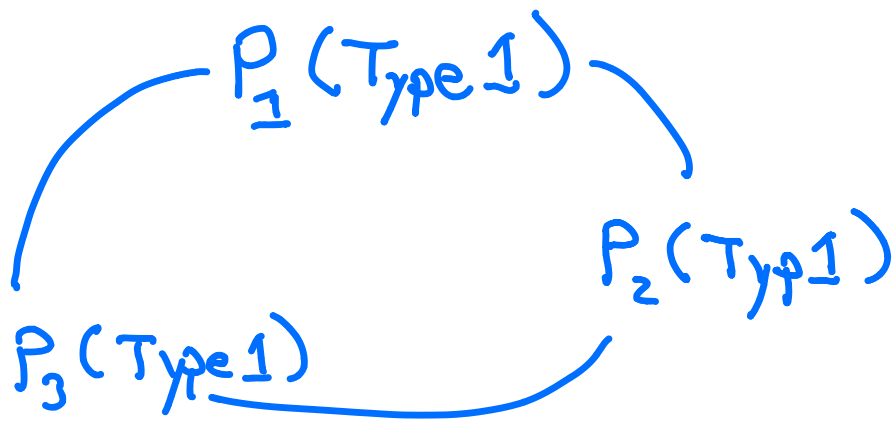

Example:

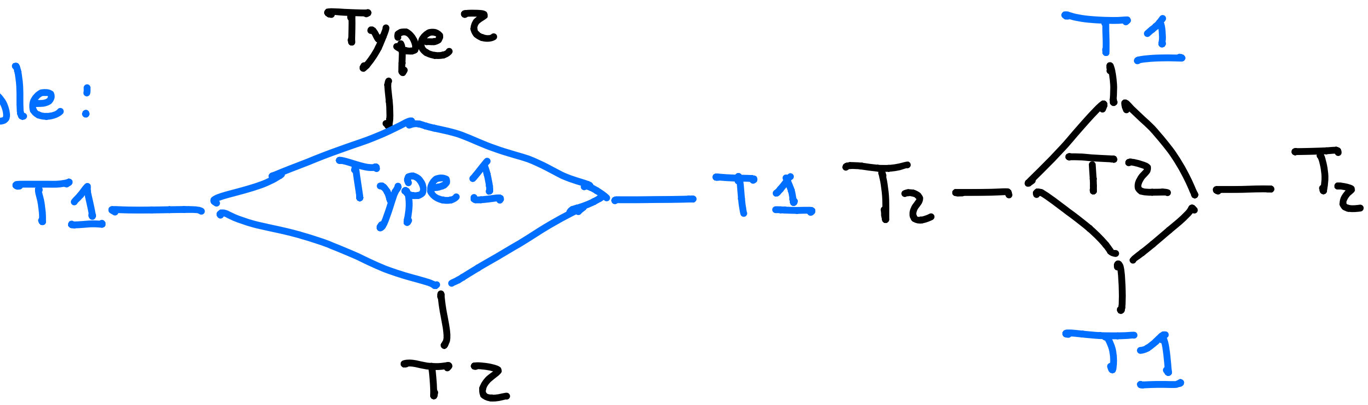

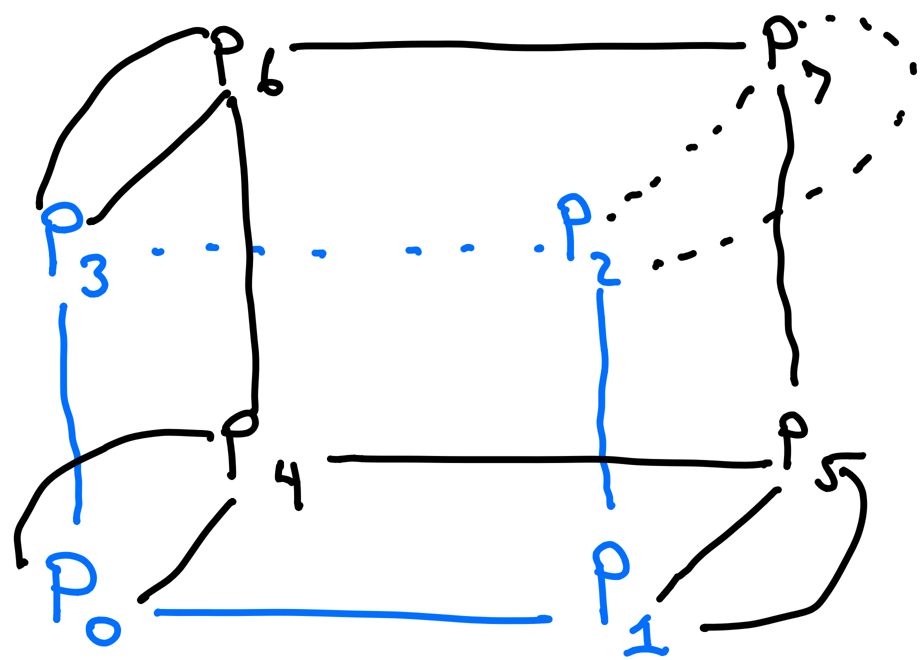

Tiling Network Families

Theorem: Network families built from tiling patterns have a balance relation with finitely many equivalence classes.

Applications

① Token ring networks

for all i : $AG(E_i \rightarrow (x_i = tok))$

for all i : $AG(H_i \rightarrow AF\ E_i)$

(liveness)

② Generalize to 2 tokens, etc.

③ Rings with red/black nodes.

④ Full symmetry counter reductions.

AODV v2

- Ad-Hoc On-Demand Distance Vector Routing
- mobile routing in wireless, multihop networks

Dynamic Networks

- network nodes may come and go
- node neighbors may come and go
- global route table unavailable
- Routes from Origin to Target must react to network change

Dynamic Routing

- No guarantee route discovery from Origin to Target will succeed

- key safety property: at all reachable global states, combined routing tables should not countain a forwarding loop

Route Requests

- Check incoming messages from O to T for freshness and cost
- Prefer latest/cheapest routes

Tricky Part

- What happens to route table entries when links fail?
- Accepting any route can lead to a routing loop (cause of failure).
- Accept routes that are no worse than current broken route

Example Broken Route

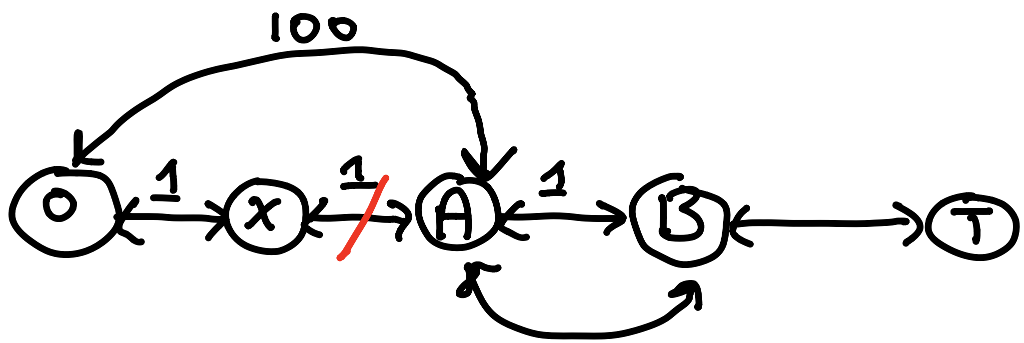

- Initial route: O; X; A; B; T
- failure: X; A — all changes lost
- A accepts long route to O
- A accepts late/cheaper route to O through B — loop established

Local Inductive Proof

- key lemma: Invariant — for any node H and node G, if H has a route entry to O with next hop G, then G has a route entry to O that is better than the entry to O at H.

Explaining A Counterexample

- System trace detailing computation that does not satisfy the local specification

- Important special cases include:
  - $\Box$ good
  - $\Box$(start $\rightarrow$ XXend)

Example

Start, End – boolean valued, local variables

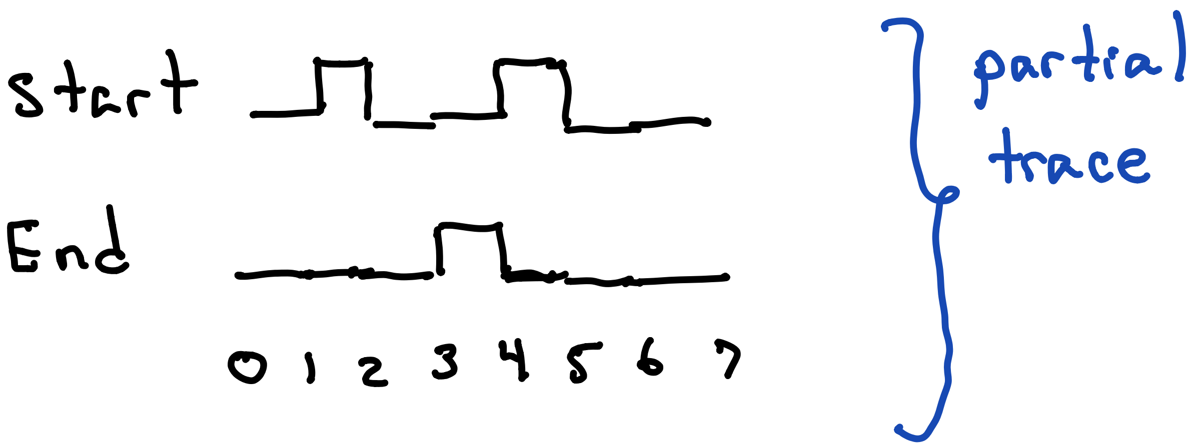

Example

Failure : $\Box(\text{Start} \rightarrow XX\,\text{End})$

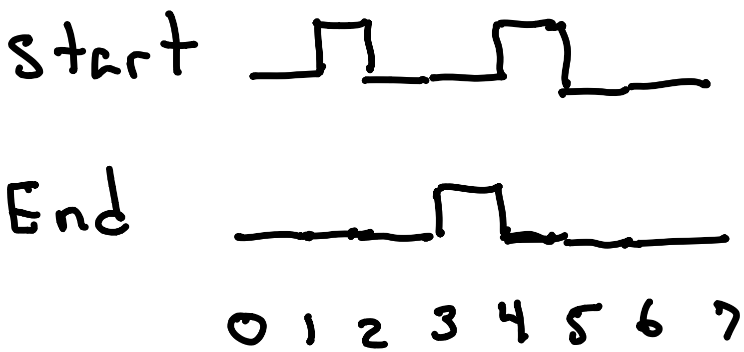

$\langle s, v \rangle$: a state, variable pair

σ a program trace

$\sigma = \sigma_0 \sigma_1 \sigma_2 \ldots$

each σ_i is a state of the program

Given $\langle s, v\rangle$ then $\langle \hat{s}, v\rangle$ is everywhere the same as s except the (boolean) value of v is switched

A pair $\langle s, v\rangle$ is <u>critical</u> for the failure of φ on σ if $\sigma \not\models \varphi$ but $\sigma^{\langle \hat{s}, v\rangle} \models \varphi$

Suppose $\sigma \not\models \varphi$

A pair $\langle s, v \rangle$ of σ is a cause of the failure if there exists a set of pairs, A, such that:

$\langle s, v \rangle \notin A$

and

$\langle s, v \rangle$ is critical for $\sigma^{\hat{A}} \not\models \varphi$

and

for all $D \subseteq A$, $\sigma^{\hat{D}} \not\models \varphi$

Example

failure : $\Box(\text{Start} \rightarrow XX\,\text{End})$
does not hold on the trace

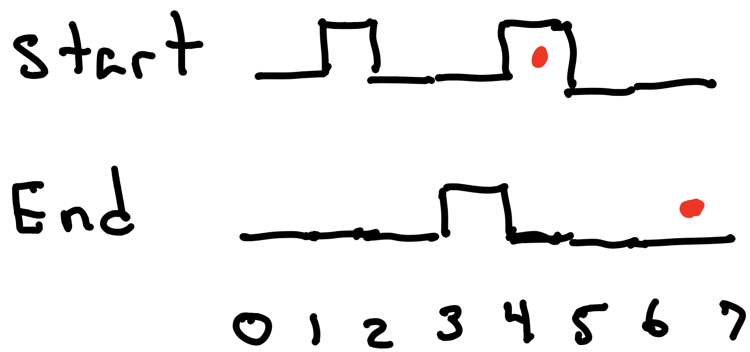

critical causes of the failure

Combing Local Reasoning with Causal Reasoning

Main challenge: local reasoning provides a sound reasoning engine through over-approximation

Challenges: Identify failure cause

① There is a bug in the local component of some k node protocol.

② The protocol is correct but local proof fails — too abstract, too local, the protocol is not symmetric enough...

Applications

0. Dining philosophers — on a ring, in dynamic graphs
1. Red/black rings
2. AODVv2 — ad hoc on-demand distance vector routing + bug/fix
3. Leader election — on a ring, local proof using an interactive prover

Refs

NT – K. Namjoshi, R. Trefler

NT	VMCAI	2012
NT	VMCAI	2013
NT	TACAS	2015
NT	FORTE	2015
NT	TACAS	2016
NT	TACAS	2018
ABT	NFM	2019

BB-DCOT FMSD 2012

BB-DCOT – I. Beer, S. Ben David, H. Chockler, A. Orni, R. Trefler