

More Stubborn Set Methods for Process Algebras

Congratulations to Bill!


Antti Valmari

Tampere University of Technology
Department of Mathematics

- 1 Introduction
- 2 Intuition of Stubborn Sets
- 3 Why More Than One Method?
- 4 Tree Failures and Fair Testing
- 5 System Model
- 6 Definition of Stubborn Sets
- 7 Construction of Stubborn Sets
- 8 Why Disabled Actions in Stubborn?
- 9 The Famous Cycle Condition
- 10 Terminal SC Conditions
- 11 Automata-Theoretic Visibility
- 12 Insight on the Ignoring Problem
- 13 Avoiding Terminal SC Conditions
- 14 Remembering Divergences
- 15 Discussion

1 Introduction

Stubborn set methods construct **reduced** state spaces preserving certain properties



Abbreviated history

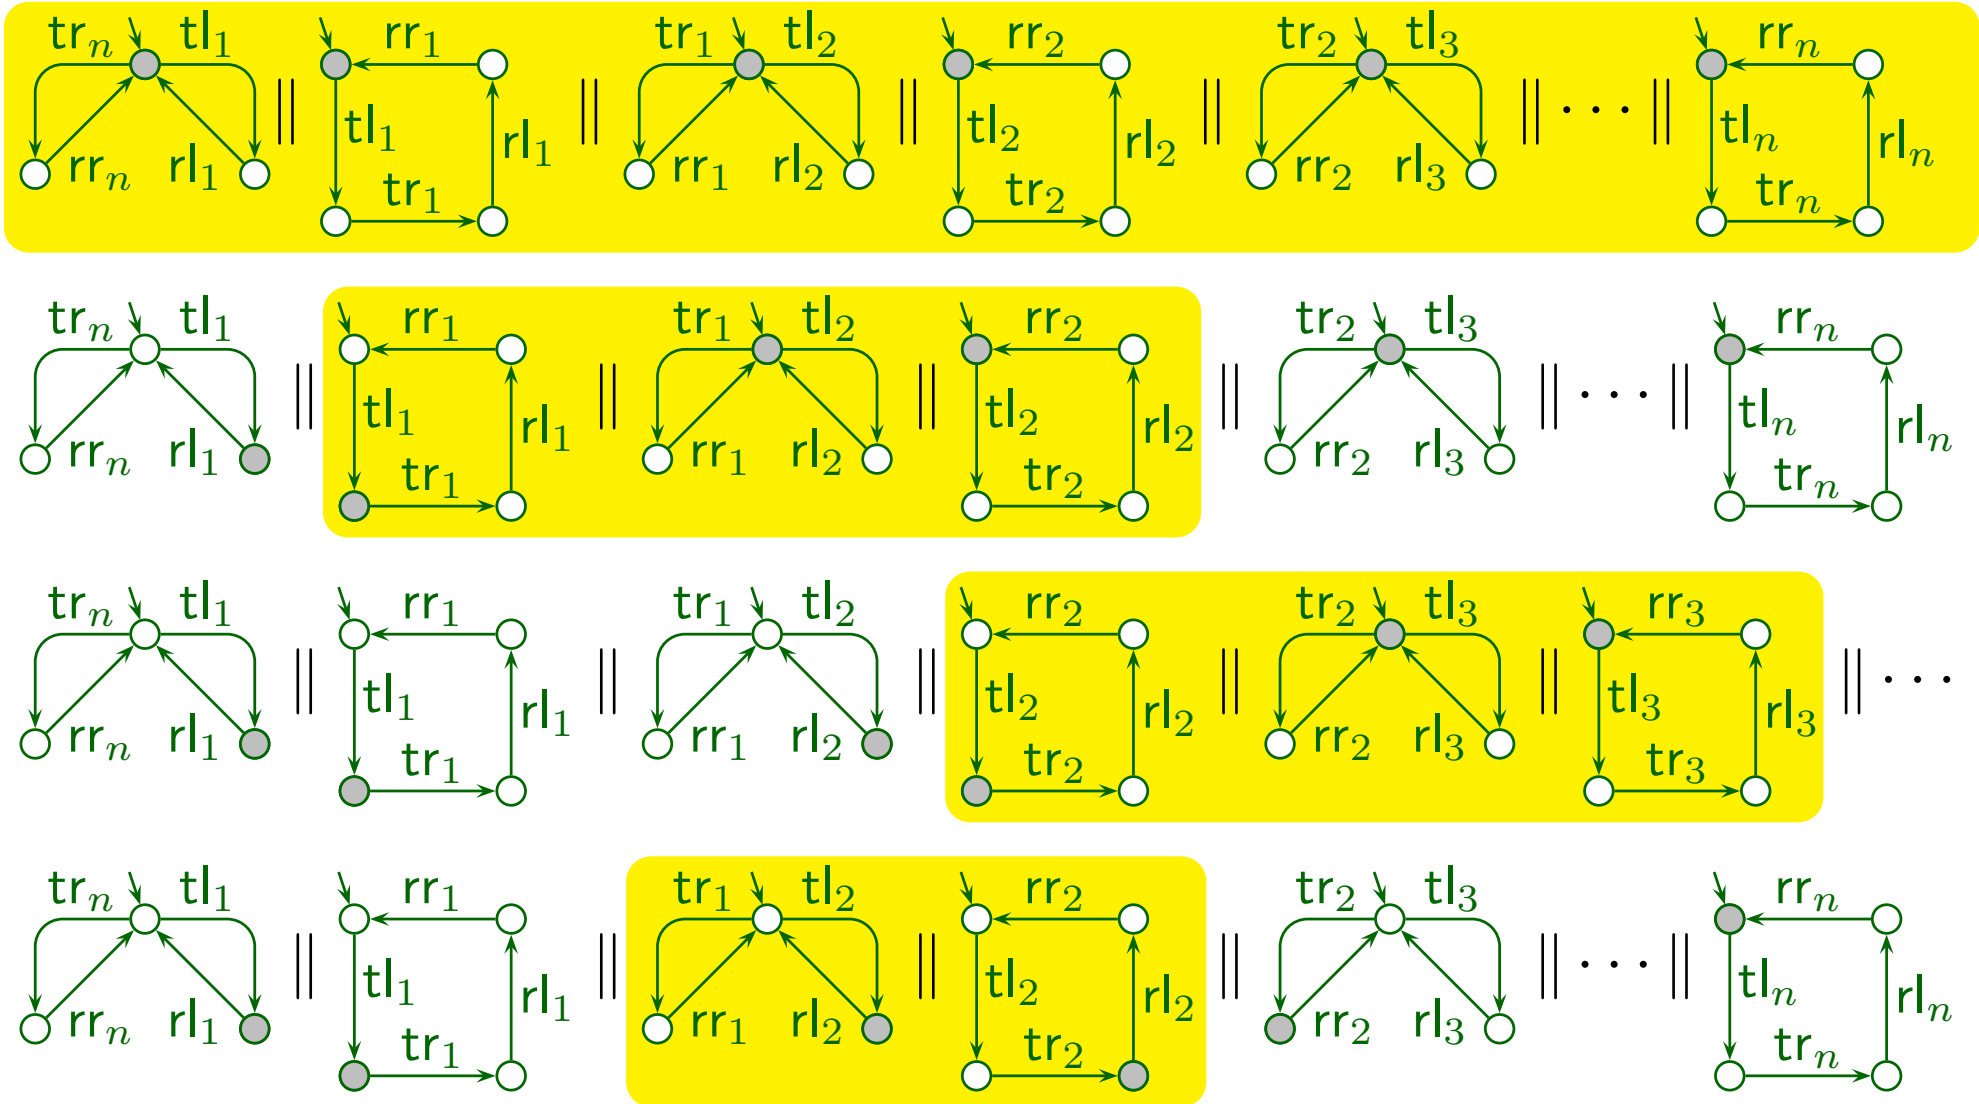
- basic stubborn sets, algorithm 1988
 - safety properties 1989
 - linear temporal logic 1990
 - nondeterministic actions 1992
 - ⇒ **Stubborn Set Methods for Process Algebras 1996**
 - persistent sets 1990, ample sets 1993
- deadlocks
 \exists inf. execution
- stable failures
 - CSP failures divergences
 - traces
 - chaos-free failures diverg.s
 - branching bisimilarity

This talk: new results 2015 – 2017

- fair testing equivalence, tree failures \Rightarrow ordinary failures
- an improvement to the terminal strong component condition for traces and fair t.
 - visibility-driven stubborn sets
- automata-theoretic visibility
- how to avoid the costly cycle / terminal sc conditions in many cases
- “remembering” divergences instead of re-constructing them in later states

2 Intuition of Stubborn Sets

n dining philosophers



$3^n - 1 \rightsquigarrow 3n^2 - 3n + 2$ states $\rightsquigarrow 3n - 1$ with also symmetries


3 Why More Than One Method?

The basic method preserves all deadlocks and yes/no to “are there infinite executions?”

Not more because


1. the ordering of concurrent events may be important to the property

- let e_i and l_i denote entering and leaving a critical section
- this problem is rather mild
- one of the new results addresses this directly



2. there is the

ignoring problem



- the basic method may terminate after investigating just the τ -cycle
- this is a tough and versatile one!
- two of the new results address this directly


4 Tree Failures and Fair Testing Equivalence

Fair Testing Equivalence [Vogler 1992, Rensink & Vogler 2007]

- is the weakest congruence that preserves ordinary (i.e., not necessarily stable) failures
- is the weakest congruence that preserves **AG EF** a
 - a canonical example of a branching-time property
- the only congruences below it are the traces, the alphabet, and the dullest
- facilitates checking a useful notion of progress without making fairness assumptions

About its definition

- tree failure: like ordinary, but a set of strings is refused
- tree failure equivalence compares them and the alphabets
- fair testing equivalence allows one system to “enter” the refusal set of the other system and refuse the suffixes



[Valmari & Vogler 2016]: the traditional method that preserves the traces also preserves fair testing

- a useful notion of progress for no additional cost
- surprising, because branching-time

[2017]: A small variant preserves also tree failures

5 System Model

$$(\bar{L}_1 \parallel \dots \parallel \bar{L}_N) \setminus (H \cup \{\tau_1, \dots, \tau_N\})$$

where

- \bar{L}_i are τ -renamed labelled transitions systems $(S_i, \bar{\Sigma}_i, \bar{\Delta}_i, \hat{s}_i)$
 - τ_i is replaced for τ

\Rightarrow each action belongs to a unique set of components

- $\text{Vis} = (\Sigma_1 \cup \dots \cup \Sigma_N) \setminus H$
- $\text{Inv} = H \cup \{\tau_1, \dots, \tau_N\}$
- $\text{Acts} = \text{Vis} \cup \text{Inv}$

Additional notation

- $s -a_1 \dots a_n \rightarrow s'$ path from s to s' with also invisible actions listed
- $s =a_1 \dots a_n \Rightarrow s'$ path from s to s' with invisible actions not listed
- $s -a_1 \dots a_n \rightarrow$ there is s' such that $s -a_1 \dots a_n \rightarrow s'$
- $s -\tau^\omega \rightarrow$ s diverges
- $\text{en}(s) = \{a \mid s -a \rightarrow \text{ in full system}\} =$ enabled actions
- $\text{en}_i(s) = \{a \mid s -a \rightarrow \text{ in } \bar{L}_i\}$

6 Definition of Stubborn Sets


Reduced LTS: r-states, r-transitions, r-paths, ...

- put \hat{s} in S_r
- for each $s \in S_r$, construct $\mathcal{A}(s) \subseteq \text{Acts}$
 - stubborn set
 - for each a and s' , if $a \in \mathcal{A}(s)$ and $s \xrightarrow{a} s'$, put s' in S_r and (s, a, s') in $\bar{\Delta}_r$


Depending on the method, $\mathcal{A}(s)$ must satisfy **D1**, (variant of) **D2**, and other conditions

D0 if there are enabled actions, $\mathcal{A}(s)$ must contain at least one


D1 $a \in \mathcal{A}(\bullet)$, $a_i \notin \mathcal{A}(\bullet)$ and




\Rightarrow there is \bullet such that



D2 $a \in \mathcal{A}(\bullet)$, $a_i \notin \mathcal{A}(\bullet)$ and



\Rightarrow there is \bullet such that




V if $\mathcal{A}(s)$ contains an enabled visible action, then $\text{Vis} \subseteq \mathcal{A}(s)$

I if there are enabled invisible actions, $\mathcal{A}(s)$ must contain at least one

● some condition for preventing ignoring: **Sen, SV, L, C3**

7 Construction of Stubborn Sets

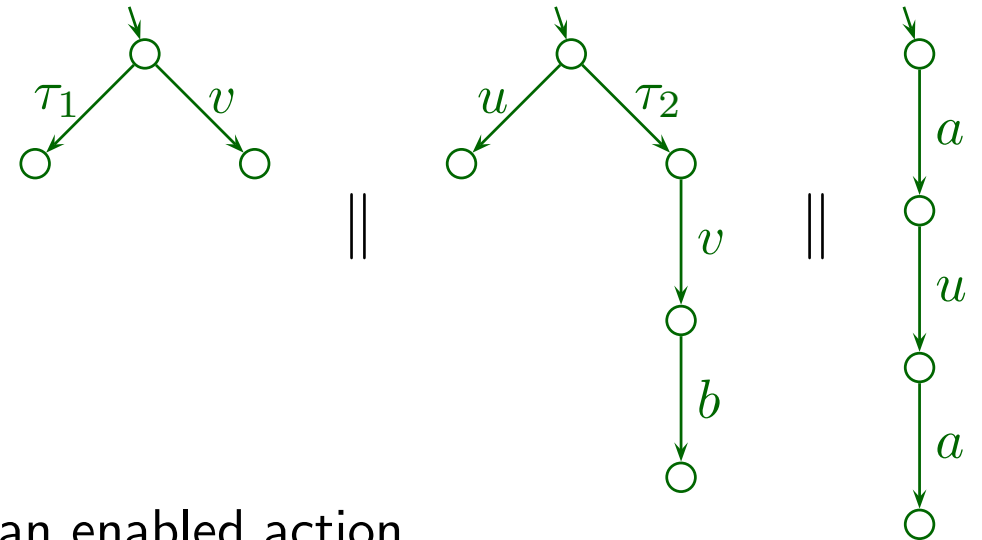


" \rightsquigarrow_s " \subseteq Acts \times Acts

- if $a \notin \text{en}(s)$, choose i such that \bar{L}_i disables a , and make $a \rightsquigarrow_s b$ for every $b \in \text{en}_i(s)$
- if $a \in \text{en}(s)$, then, for every i such that $a \in \bar{\Sigma}_i$ and for every $b \in \text{en}_i(s)$, make $a \rightsquigarrow_s b$
- it does not matter whether $a \rightsquigarrow_s a$

clsr(s, a) = the closure of a w.r.t. " \rightsquigarrow_s "

- satisfies **D1** and **D2**
- satisfies also **V** if $a \rightsquigarrow_s b$ is added for every $a \in \text{Vis} \cap \text{en}(s)$ and $b \in \text{Vis}$




gsc(s, a, \dots) = "good strong component"

- finds a \subseteq -minimal closed set that contains an enabled action or replies that such a set does not exist
- additional parameters tune it for future needs (may be called more than once in the same state)
- $O(|\rightsquigarrow|)$

8 Why Disabled Actions in Stubborn Sets?

Only **D0**, **D1**, and **D2** are obeyed, starting points from left to right



Also **V** is obeyed, where $V = \{a, b\}$

- although initially $a \in V$ is taken, τ_1 is not

\Rightarrow **V** is better than **C2** in ample set theory

- if $\text{ample}(s)$ contains an enabled visible action, then $\text{ample}(s) = \text{en}(s)$

\Rightarrow Allowing disabled actions in stubborn sets facilitates formulation of better conditions

- and better algorithms: $\rightsquigarrow_s, \text{gsc}$

9 The Famous Cycle Condition for Liveness

L For every $a \in \text{Vis}$, every r-cycle must contain a state s such that $a \in \mathcal{A}(s)$



C3 Every r-cycle must contain a state s such that $\text{ample}(s) = \text{en}(s)$

Implementation of C3 [Clarke & al. 1999]

- construct r-states and r-transitions in depth-first order
- if $a \in \text{ample}(s)$, $s \xrightarrow{a} s'$, and s' is in depth-first stack, choose $\text{ample}(s) = \text{en}(s)$

A discouraging example


- try components from left to right
- sticking to a component helps a bit
 - [1999] does not tell to do so
 - fails badly with 3-dimensional case



We can expand s' instead (Theorem: DFS), seems better

This issue has received too little attention!

- observed in [Evangelista & Pajault 2010] (another example)
- nobody knows how serious it really is
- we are not told how to deal with it




10 Terminal Strong Component Conditions for Safety

Sen For each r -terminal strong component C and each $a \in \text{en}(r)$ where r is the root of C , there is $s \in C$ such that $a \in \mathcal{A}(s)$

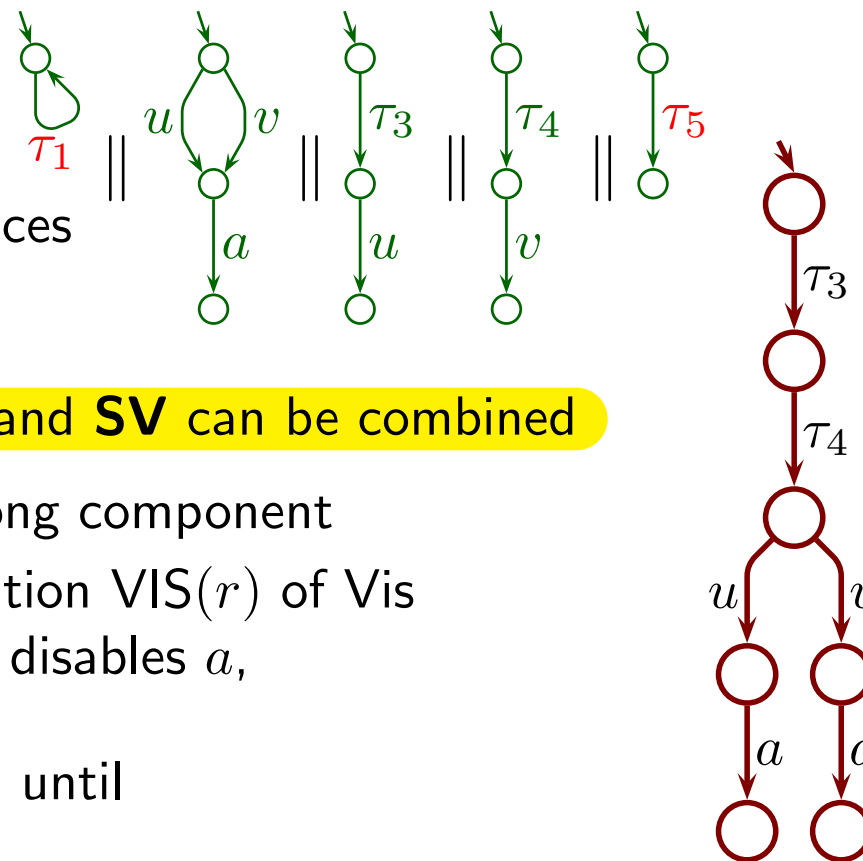
SV For each r -terminal strong component C and each $a \in \text{Vis}$, there is $s \in C$ such that $a \in \mathcal{A}(s)$

- implemented efficiently using depth-first order and Tarjan's algorithm



Each suffers from a problem

- Sen** may force to fire irrelevant actions
 - e.g., when $\text{clr}(\text{Vis}) = \emptyset$ when preserving traces
- SV** may yield big stubborn sets



[Valmari & Hansen 2016]: The advantages of **Sen** and **SV** can be combined

- consider extending the root r of a terminal strong component
- compute an “enabling-closed” upper approximation $\text{VIS}(r)$ of Vis
 - if $a \in \text{VIS}(r) \setminus \text{en}(r)$, choose i such that \bar{L}_i disables a , and add every $b \in \text{en}_i(r)$ to $\text{VIS}(r)$
 - try $\text{gsc}(r, a, \dots)$ for each $a \in \text{en}(r) \cap \text{VIS}(r)$ until r is no longer a root or $\text{VIS}(r)$ is exhausted

11 Automata-Theoretic Visibility

Assume that bad states are recognized immediately

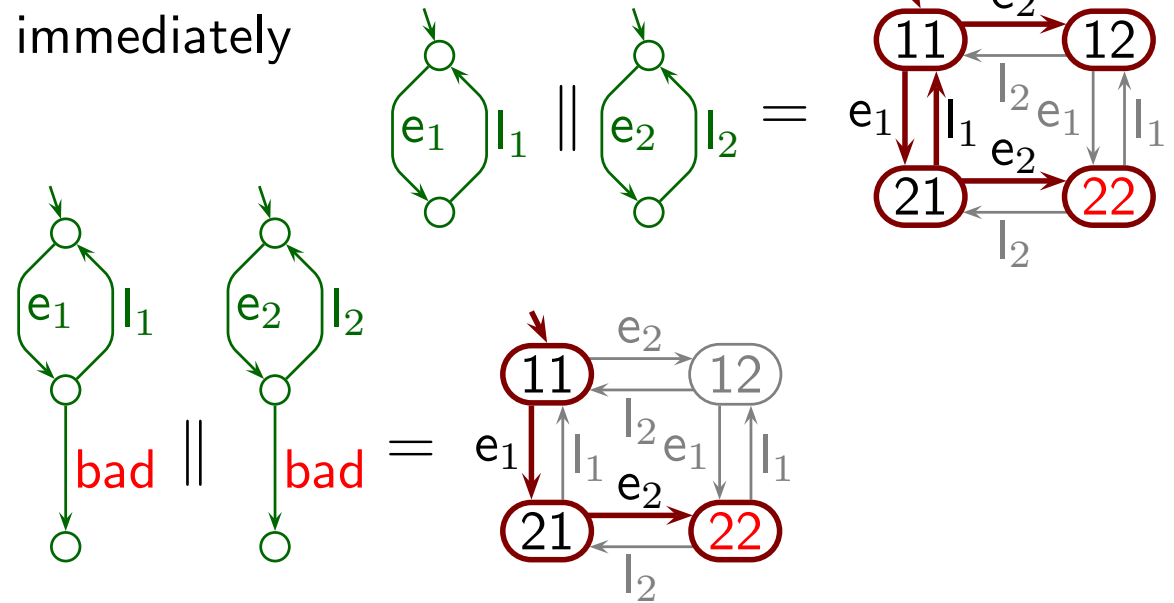
With the trace-preserving method

- $Vis = \{e_1, e_2, l_1, l_2\}$
- \mathbf{V} forces $e_1 \rightsquigarrow_{\hat{s}} e_2 \rightsquigarrow_{\hat{s}} e_1$

Using **bad** to drive analysis

- $e_1 \not\rightsquigarrow_{\hat{s}} e_2 \not\rightsquigarrow_{\hat{s}} e_1$

⇒ **smaller reduced LTS**



⇒ Property- and state-dependent “incomplete” visibility suffices

Finite automata

- sufficiently keep track of the state, catch errors on-the-fly
- block irrelevant branches of the LTS
- problem: “if $a \in en(s)$, ... for every $b \in en_i(s)$, make $a \rightsquigarrow_s b$ ” acts like \mathbf{V}

[Hansen & Valmari 2016]: the automaton need not cause $a \rightsquigarrow_s b$, if every error-detecting path starting with b in it remains error-detecting if a is added or moved to its front

- implemented by pre-processing the automaton

12 Insight on the Ignoring Problem

Infinite executions

- if traces are preserved and the reduced LTS is finite, then infinite traces are preserved
 - König's Lemma type of reasoning
 - **I** (with **D0**, **D1**, **D2b**, and **V**) suffices for preserving minimal divergence traces
- ⇒ the only problematic infinite executions are non-minimal divergence traces
- nothing better is known than cycle conditions
 - nothing is known for fairness assumptions

Finite executions

- **D0** or $\text{gsc}(\text{Vis})$ (with **D1**, **D2r**, and **V**) suffices for traces that lead to stable states
 - $\text{gsc}(s, a, \dots)$ for each $a \in \text{Vis}$ until an enabled action is found or Vis is exhausted
 - the only problematic case is **permanently diverging** states
 - s is permanently diverging iff every state reachable from s is diverging
 - that is, no stable state is reachable from s
- ⇒ without **S** and **L**,
- either all traces (and fair testing and ...) are preserved,
 - or **the method tells that** the system may reach a permanently diverging state
- ⇒ **if we want the system not to have such states, S is not needed**

13 Avoiding Terminal Strong Component Conditions

So the following works for traces, fair testing equivalence, and so on

- construct a reduced LTS using $\text{gsc}(\text{Vis})$
- if the result contains a terminal strong component that is not a deadlock and does not contain occurrences of visible actions, fix the system and try again

It seems that the actions in the component could be permanently frozen and $\text{gsc}(\text{Vis})$ executed again, but this is future research

For the time being: **Stop It, and Be Stubborn!** [Valmari 2015]


- add stop-transitions to terminal states to chosen states of chosen \bar{L}_i
 - they model clients deciding not to request for service
 - often this has to be done anyway, to not lose certain progress errors
- compute the reduced LTS
- if it contains a terminal sc that is not a deadlock, go back to first step
 - or of the above kind
- otherwise remove its stop-transitions

Deadlocks automatically guarantee **Sen** and **SV**

This trick works also for some liveness properties

14 Remembering Divergences

I may force to construct “the same” divergence many times




Solution (new):

- obey **D0**, **D1**, **D2**, **V**, and **L**
 - **D0** and **D2** may be weakened a bit
- store the union of $\mathcal{A}(s) \cap \text{en}(s)$ of the states of the cycle to each state of the cycle
 - a state may have many div-sets
 - (or, to simplify implementation, at most one)
- when constructing a new r-state s' via $s \xrightarrow{a} s'$, copy from s the div-sets that lack a
- apply **I** on states where divergence matters that have no div-sets

By **D2**, each state with div-sets diverges

Example

- assume divergence matters always after a
- it is remembered after b



15 Discussion

The conditions in the present talk have the following roles:

- permutation correctness conditions: **D1**, **V**
- idling correctness conditions: variants of **D2**
- driving force: **D0**, the new **S**, **I**
- useful progress guarantee: variants of **S**, **L**

Most results of the talk are about avoiding progressing in a useless direction

- save from some state sub-explosions

The fair testing (and tree failures) result provides a useful fairness notion with surprisingly little additional cost

- the first realistic ample / persistent / stubborn set method for fairness
- surprising step towards branching time

Thank you for attention!
Questions?