# Verifying Timed, Asynchronous Circuits Using ACL2

## Yan Peng<sup>1</sup> Mark R. Greenstreet<sup>1</sup>

<sup>1</sup>Department of Computer Science University of British Columbia

May 14th 2019

We verify timed circuits with the generality of ACL2 while achieving performance comparable to dedicated tools by using Smtlink, a SMT solver interface





- 3 Modeling asP\* Pipelines
- 4 Verifying the asP\* Pipelines



## Timed Asynchronous Circuits are Great

- Timed asynchronous circuits exploit verifiable timing relationships in a circuit to improve performance and reduce power and area.
- Many timed pipeline designs have been proposed: micropipeline[Sut89], asP\* pipelines[MJ<sup>+</sup>97], GasP[SF01], Mousetrap[SN07].



Let's look at the example of an asP\* pipeline.



full; asserts the stage is full and empty; asserts the stage is empty.

Let's look at the example of an asP\* pipeline.



When the first stage is full,  $\overline{\text{goEmpty}_1}$  and  $\overline{\text{goFull}_2}$  are excited to go low.

Let's look at the example of an asP\* pipeline.



The SR-flops for stages 1 & 2 are both excited to change. The pipeline transfers data from stage 1 to stage 2.

Let's look at the example of an asP\* pipeline.



The first stage is now empty and the second stage full. The data ripples through the pipeline.

What can go wrong when timing constraints are violated?



Let's consider a scenario where the turn-around on the right of the NAND gate is significant longer than the left of it, for example, due to long wires in the layout.

What can go wrong when timing constraints are violated?



empty<sub>1</sub> goes high, enabling  $\overline{\text{goFull}_2}$  to return high before full<sub>2</sub> has a chance to change.

What can go wrong when timing constraints are violated?



Now, full<sub>1</sub> has gone low, but full<sub>2</sub> never went high. This disables stage two from going full. A data value is lost.  $\bigcirc$ 

Timing properties are crucial for correctness of timed asynchronous circuits. But how do we know we've defined enough timing constraints?

As is pointed out in  $[MJ^+97]$ :

For the circuit presented here, determining the delay conditions that must be satisfied for reliable operation was difficult. We have used a mixture of ad hoc analysis and hSpice simulation to achieve our results ...

- We propose a framework for modeling timed asynchronous circuits using *timed traces* and *trace recognizers*. It naturally models **nondeterminism**.
- We verified timing properties of three configurations of *asP\* pipelines*, for each configurations:
  - We defined and proved the inductive timing invariants.
  - Using the invariants, we are able to prove hazard-freedom.
- Our model can handle **loop structures** and allows **parameterized verification**.

We verify timed circuits with the generality of ACL2 while achieving performance comparable to dedicated tools by using Smtlink, a SMT solver interface



2 Timed Traces

- 3 Modeling asP\* Pipelines
- 4 Verifying the asP\* Pipelines
- 5 Conclusions



Every circuit contains a set of signal paths, representing nodes in the circuit.





- A state of the circuit is a table mapping from node paths to their values.
- A state can contain other signals that are outside of the component, allowing composability.



| <i>r</i> <sub>0</sub> :F | a <sub>0</sub> :F    | <i>r</i> <sub>1</sub> :F         | a <sub>1</sub> :F | $\overline{a_1}$ :F         | <i>r</i> <sub>2</sub> :F | a <sub>2</sub> :F | $\overline{a_2}$ :F |  |
|--------------------------|----------------------|----------------------------------|-------------------|-----------------------------|--------------------------|-------------------|---------------------|--|
| <i>r</i> ₀: <b>⊤</b>     | a <sub>0</sub> :F    | <i>r</i> <sub>1</sub> :F         | a <sub>1</sub> :F | $\overline{a_1}$ : <b>T</b> | <i>r</i> <sub>2</sub> :F | a <sub>2</sub> :F | $\overline{a_2}$ :F |  |
| <i>r</i> <sub>0</sub> :T | <i>a</i> ₀: <b>⊺</b> | <i>r</i> <sub>1</sub> : <b>T</b> | a <sub>1</sub> :F | $\overline{a_1}$ :T         | <i>r</i> <sub>2</sub> :F | a <sub>2</sub> :F | $\overline{a_2}$ :F |  |
|                          |                      |                                  |                   |                             |                          |                   |                     |  |

A trace of the circuit is a list of states.

Trace recognizers are functions defined for each component. It specifies that for this component, every two consecutive states in the trace must make valid transitions:

$$\begin{array}{c|c} in_1 & c\_step(prev,next) \triangleq \\ in_2 & out & prev[out] = next[out] \\ & \lor(prev[in1] = prev[in2] \land prev[out] \neq prev[in1]) \end{array}$$

- For example, for the component C-element,
  - Either the output hasn't changed, making no assumption about the inputs.
  - Or the output changes only when the inputs match and the output doesn't.
- The trace recognizers are receptive to all possible inputs[Dil88].
- It defines allowed behaviour on the outputs based on the inputs.

## Introducing Time - Traces



| 0   | <i>r</i> ₀: <mark>F</mark> ,0         | a <sub>0</sub> :F,0          | <i>r</i> <sub>1</sub> :F,0           | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ :F,0          |  |
|-----|---------------------------------------|------------------------------|--------------------------------------|----------------------------|--------------------------------|--|
| 0.3 | r <sub>0</sub> : <b>T,0.3</b>         | <i>a</i> <sub>0</sub> :F,0   | <i>r</i> <sub>1</sub> :F,0           | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ :F,0          |  |
| 4   | <i>r</i> ₀: <b>T</b> ,0.3             | <i>a</i> <sub>0</sub> :F,0   | <i>r</i> <sub>1</sub> :F,0           | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ : <b>T</b> ,4 |  |
| 10  | <i>r</i> <sub>0</sub> : <b>T</b> ,0.3 | a <sub>0</sub> : <b>T,10</b> | <i>r</i> <sub>1</sub> : <b>T</b> ,10 | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ : <b>T</b> ,4 |  |
|     | •••                                   |                              |                                      |                            | •••                            |  |

We extend the traces to include timing information.

- Add a time for each state.
- For each node signal in a state, we associate its value with a time representing the most recent time the signal acquired that value.

Peng & Greenstreet (UBC)

## Introducing Time - Traces



| 0   | <i>r</i> ₀: <mark>F</mark> ,0         | a <sub>0</sub> :F,0          | <i>r</i> <sub>1</sub> :F,0           | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ :F,0          |       |
|-----|---------------------------------------|------------------------------|--------------------------------------|----------------------------|--------------------------------|-------|
| 0.3 | <i>r</i> <sub>0</sub> : <b>T,0.3</b>  | <i>a</i> <sub>0</sub> :F,0   | <i>r</i> <sub>1</sub> :F,0           | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ :F,0          |       |
| 4   | <i>r</i> <sub>0</sub> :T,0.3          | <i>a</i> <sub>0</sub> :F,0   | <i>r</i> <sub>1</sub> :F,0           | a <sub>1</sub> :F,0        | <u>a₁</u> : <b>Т,4</b>         |       |
| 10  | <i>r</i> <sub>0</sub> : <b>T</b> ,0.3 | a <sub>0</sub> : <b>T,10</b> | <i>r</i> <sub>1</sub> : <b>T</b> ,10 | <i>a</i> <sub>1</sub> :F,0 | $\overline{a_1}$ : <b>T</b> ,4 |       |
|     |                                       |                              |                                      |                            |                                | • • • |

A few notes about time:

- Time is non-negative and monotonically increasing.
- Time of each signal must not be larger than the state time.
- Every trace can be extended so that time progresses without bound ("non-Zenoness" [AL94]).

Peng & Greenstreet (UBC)

## Introducing Time - Trace Recognizers

For each component, we use the delay model of a symbolic delay bound  $[\delta_{lo}, \delta_{hi})$  where  $\delta_{lo} > 0$  and  $\delta_{lo} \leq \delta_{hi}$ .

The trace recognizer now defines valid behaviour of a component in terms of both logic values and transition time. For example:

$$\begin{array}{c} \underset{n_{1}}{\underset{n_{2}}{in_{2}}} & \text{out} \end{array} \quad \begin{array}{c} c\_\text{step}(\texttt{prev},\texttt{next}) \triangleq \\ & \ddots & \cdot \\ & \vee(\texttt{prev}[\texttt{in1}].\texttt{v} = \texttt{prev}[\texttt{in2}].\texttt{v} \land \texttt{prev}[\texttt{out}].\texttt{v} \neq \texttt{prev}[\texttt{in1}].\texttt{v} \\ & \wedge\texttt{prev}[\texttt{out}].\texttt{v} \neq \texttt{next}[\texttt{out}].\texttt{v} \\ & \wedge\texttt{max}(\texttt{prev}[\texttt{in1}].\texttt{t},\texttt{prev}[\texttt{in2}].\texttt{t}) + \delta_{lo} \leq \texttt{next}.\texttt{t} \\ & \wedge\texttt{next}[\texttt{out}].\texttt{t} = \texttt{next}.\texttt{t} \\ & \wedge\texttt{next}.\texttt{t} < \texttt{max}(\texttt{prev}[\texttt{in1}].\texttt{t},\texttt{prev}[\texttt{in2}].\texttt{t}) + \delta_{hi}) \end{array}$$

This says, when the output of the C-element changes, it should happen between delay bounds.

## Nondeterminism



- Because of continuous time, a state in a trace can have an infinite number of successors.
- Our timed trace model naturally captures nondeterminism.

Peng & Greenstreet (UBC)

Verifying Timed Asynchronous Circuits



2 Timed Traces

- 3 Modeling asP\* Pipelines
- 4 Verifying the asP\* Pipelines
- 5 Conclusions

## The asP\* Pipelines

We recall the asP\* pipeline is:



We want to find a way to break the pipeline into stages.

## The asP\* Pipelines

Apply a few logical equivalences to get:



The obvious way is to split the AND gate into two halves:



But we don't want to model half of an AND gate.

## The asP\* Pipelines

Instead we duplicate the AND gates:



This transformation includes all possible behaviours of the pipeline and additional behaviours due to the duplicated AND gates.

# A Single Stage

So a single stage becomes:



- *eOut* stands for current stage being empty, and *fln* stands for previous stage being full. In this case, the current stage can set
- *fOut* stands for current stage being full, and *eln* stands for next stage being empty. In this case, the current stage can reset



- To verify a pipeline, we need to consider actions of the environment.
- A 1Env is a "left-environment"
  - It acts like an asP\* stage,
  - But it doesn't have a left-neighbour.
  - An empty 1Env can spontaneously go full any time after a minimum delay after going empty.
  - An internal signal li allows the environment to have a pending output event.
- A rEnv is similar it can spontaneously go empty.

## rEnv + lEnv = Stage

A stage can be built with a right and a left environment with the constraint that ri == li:



• Signal ri == li represents the internal state of a stage when it has acquired the value but hasn't set/reset Q or  $\overline{Q}$  yet.

This gives us the final model for the asP\* pipeline:





2 Timed Traces

- 3 Modeling asP\* Pipelines
- 4 Verifying the asP\* Pipelines
- 5 Conclusions

## Three configurations

• Configuration 1: a lenv with a renv, 0 stage



• Configuration 2: a lenv, an arbitrary number of stages, and a renv



• Configuration 3: a ring with an arbitrary number of stages



We define the timing invariant for an empty pipeline, and the invariant for the other two configurations can be defined recursively. The timing invariant is composed of two parts:

- Timing constraints within each individual component: lenv and renv Obvious, based on the trace recognizers.
- Timing constraints on the interactions between components The interesting part.

We prove the inductive invariant by proving the theorem:

```
invariant_env(s1) \land valid_step(s1,s2) \Rightarrow invariant_env(s2)
```

(1)

## Interaction invariants



## Definition 1 (Hazard)

Between two consecutive states, if some component output, y, is excited to change in the first state, but has neither changed nor is excited to make that change in the second state, we call this a hazard.

For example, the hazard-free step condition for signal li of lenv from state s1 to state s2 is:

(2)

## Definition 1 (Hazard)

Between two consecutive states, if some component output, y, is excited to change in the first state, but has neither changed nor is excited to make that change in the second state, we call this a hazard.

To show that a module's invariant ensures hazard freedom, we prove the theorem:

invariant\_mod(s1) ∧ valid\_step\_mod(s1,s2) ⇒ hazard\_free\_mod(s1, s2)

Inductively, we prove that every step of a valid trace satisfies the hazard-free property.

(2)

# The Proof Tree

We verify timed circuits with the generality of ACL2 while achieving performance comparable to dedicated tools by using Smtlink, a SMT solver interface



- This graph represents all theorems proved.
- The proof is greatly automated by using the Smtlink package.
- Counter-example driven manual learning of invariant.

# Conclusion and Future Work

- ACL2 is a general purpose theorem prover that can be used to verify correctness of timed, asynchronous circuits.
- Using Smtlink, proofs are nearly automatic:
  - Prove each invariant is maintained by a single step (Smtlink).
  - Prove the invariant holds for a valid trace (ACL2 excels at induction).
- The performance is comparable to dedicated tools:
  - Most proofs complete within a second, the longest was less than 30 seconds.
  - The verification problem is formulated in the ACL2 logic rather than writing special-purpose timing analysis algorithms.
- We can verify parameterized designs:
  - An arbitrary number of pipeline or ring stages.
  - Symbolic constraints on timing delays.

- Integrate with existing ACL2 Verilog front-end (SV and VL).
- Use model checking algorithms to automatically generate invariants.
- Explore more sophisticated timing models including metastability.
- Combine with verification of functionality:
  - Prove that timed designs implement untimed abstractions.
  - Reason about functionality in untimed model (e.g. [CCS19]).
- Continue to extend Smtlink and make proofs even more automatic.

# Future Work

We verify timed circuits with the generality of ACL2 while achieving performance comparable to dedicated tools by using Smtlink, a SMT solver interface

- Integrate with existing ACL2 Verilog front-end (SV and VL).
- Use model checking algorithms to automatically generate invariants.
- Explore more sophisticated timing models including metastability.
- Combine with verification of functionality:
  - Prove that timed designs implement untimed abstractions.
  - Reason about functionality in untimed model (e.g. [CCS19]).
- Continue to extend Smtlink and make proofs even more automatic.

## Thank You!

## References I

#### Martín Abadi and Leslie Lamport.

An old-fashioned recipe for real time.

ACM Transactions on Programming Languages and Systems, 16(5):1543–1571, September 1994.

Matt Kaufmann Marly Roncken Cuong Chau, Warren Hunt and Ivan Sutherland.

A hierarchical approach to self-timed circuit verification.

May 2019.



David L. Dill.

*Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits.* 

PhD thesis, School of Computer Science, Carnegie Mellon University, 1988. Published in book form as part of the ACM Doctoral Dissertation Award Series by the MIT Press, Cambridge, MA, 1989.

## References II

#### Charles E. Molnar, Ian W. Jones, et al.

#### A FIFO ring oscillator performance experiment.

In Proceedings of the Third International Symposium on Advanced Research in Asynchronous Circuits and Systems, pages 279–289. IEEE Computer Society Press, April 1997.

### Ivan Sutherland and Scott Fairbanks.

#### GasP: A minimal FIFO control.

In Proceedings of the Seventh International Symposium on Asynchronous Circuits and Systems, pages 46–53, April 2001.

#### M. Singh and S. M. Nowick.

Mousetrap: High-speed transition-signaling asynchronous pipelines.

*IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 15(6):684–698, June 2007.



### Ivan E. Sutherland.

Micropipelines.

Communications of the ACM, 32(6):720–738, June 1989.

Turing Award lecture.

## Definition 2 (Deadlock-freedom)

For all valid states, there always exist a valid next state that's different from the current state.

We prove deadlock-freedom by defining a witness function that provides a valid next state.

To show that a module's invariant ensures deadlock-freedom, we prove the theorem:

 $\label{eq:sinvariant_mod(s1)} \begin{array}{l} \Rightarrow \quad \texttt{witness\_mod(s1)} \neq \texttt{nil} \\ \land \texttt{valid\_mod(s1, witness\_mod(s1))} \\ \land \texttt{changed(s1, witness\_mod(s1))} \end{array} \tag{3}$ 



For example, this is an invalid trace:

| <i>r</i> <sub>0</sub> :F         | <i>а</i> <sub>0</sub> :F | <i>r</i> <sub>1</sub> :F | <i>а</i> 1:Е | $\overline{a_1}$ :F | <i>r</i> <sub>2</sub> :F | a <sub>2</sub> :F | $\overline{a_2}$ :F |  |
|----------------------------------|--------------------------|--------------------------|--------------|---------------------|--------------------------|-------------------|---------------------|--|
| <i>r</i> <sub>0</sub> : <b>T</b> | a <sub>0</sub> :F        | <i>r</i> <sub>1</sub> :F | <i>а</i> 1:Е | $\overline{a_1}$ :F | <i>r</i> <sub>2</sub> :F | a <sub>2</sub> :F | $\overline{a_2}$ :F |  |
| $r_0:T$                          | <i>a</i> ₀: <b>⊤</b>     | <i>r</i> <sub>1</sub> :F | $a_1$ :F     | $\overline{a_1}$ :F | <i>r</i> <sub>2</sub> :F | a <sub>2</sub> :F | $\overline{a_2}$ :F |  |
|                                  |                          |                          |              |                     |                          |                   |                     |  |

 $a_0$  as the output of the first C-element makes an unexplained transition, violating the trace recognizer for the first C-element:

## Datatypes

In ACL2, we use algebraic datatypes for modeling traces:<sup>1</sup>

| Trace       | list of TimedState           |
|-------------|------------------------------|
| TimedState  | product of Time and StateMap |
| StateMap    | map from Path to TimedValue  |
| Path        | list of PathElement          |
| PathElement | product of Name and Index    |
| TimedValue  | product of Time and Value    |

Using this timed trace model, when proving properties of a circuit, usually we will prove a theorem in the shape:

<sup>1</sup>Basic types: Time - rationalp; Value - booleanp; Name - Symbolp; Index - integerp

Peng & Greenstreet (UBC)

Timed automaton based approaches, relative timing, theorem proving, other parameterized methods.

- Comparing to timed automaton based approaches, our method allows symbolic delays on gates and allows parameterized verification.
- Comparing to relative timing, our approach models full timing information that allows verification of a larger set of circuits.
- Comparing to theorem proving based methods, our use of an SMT solver at the backend greatly reduced human effort.
- Existing parameterized methods suffers from loss of precision, we can achieve precision by using induction proofs.

## Cool Animation



## • This is an awesome animation.

## Cool Animation



## • This is an awesome animation.

## **Cool Animation**



• This is an awesome animation.