

A Framework for Asynchronous Circuit Modeling and Verification in ACL2

Cuong Chau

ckcuong@cs.utexas.edu

Department of Computer Science
The University of Texas at Austin

September 22, 2017

- 1 Introduction
- 2 The DE System
- 3 Modeling and Verification Approach
- 4 Case Studies
- 5 Future Work and Conclusions

- 1 Introduction
- 2 The DE System
- 3 Modeling and Verification Approach
- 4 Case Studies
- 5 Future Work and Conclusions

Introduction

Synchronous circuits (or clock-driven circuits): changes in the state of storage elements are synchronized by **a global clock signal**.

Asynchronous circuits (or self-timed circuits): there is **no** global clock signal distributed in asynchronous circuits. The communications between state-holding elements are performed via **local communication protocols**.

Introduction

Synchronous circuits (or clock-driven circuits): changes in the state of storage elements are synchronized by **a global clock signal**.

Asynchronous circuits (or self-timed circuits): there is **no** global clock signal distributed in asynchronous circuits. The communications between state-holding elements are performed via **local communication protocols**.

Why asynchronous?

Synchronous circuits (or clock-driven circuits): changes in the state of storage elements are synchronized by a **global clock signal**.

Asynchronous circuits (or self-timed circuits): there is **no** global clock signal distributed in asynchronous circuits. The communications between state-holding elements are performed via **local communication protocols**.

Why asynchronous?

- Low power consumption,
- High operating speed,
- Elimination of clock skew problems,
- Better composability and modularity for large systems,
- ...

Our goal: developing **scalable** methods for reasoning about the **functional correctness** of self-timed systems using ACL2.

Our goal: developing [scalable](#) methods for reasoning about the [functional correctness](#) of self-timed systems using ACL2.

- We use [the DE system](#) [Hunt:2000], which is built in ACL2, to specify and verify self-timed circuit designs.

Our goal: developing **scalable** methods for reasoning about the **functional correctness** of self-timed systems using ACL2.

- We use **the DE system** [Hunt:2000], which is built in ACL2, to specify and verify self-timed circuit designs.
- Developing a **hierarchical verification** approach to support scalability.

Our goal: developing **scalable** methods for reasoning about the **functional correctness** of self-timed systems using ACL2.

- We use **the DE system** [Hunt:2000], which is built in ACL2, to specify and verify self-timed circuit designs.
- Developing a **hierarchical verification** approach to support scalability.
- Exploring strategies for reasoning with **non-deterministic** circuit behavior.

- 1 Introduction
- 2 The DE System**
- 3 Modeling and Verification Approach
- 4 Case Studies
- 5 Future Work and Conclusions

The DE System

DE is a formal occurrence-oriented [hardware description language](#) developed in ACL2 for describing Mealy machines [Hunt:2000].

The DE System

DE is a formal occurrence-oriented [hardware description language](#) developed in ACL2 for describing Mealy machines [Hunt:2000].

The DE system supports [hierarchical verification](#):

- Each time a module is specified, there are two lemmas need be proven: a [value lemma](#) specifying the module's outputs and a [state lemma](#) specifying the module's next state.

The DE System

DE is a formal occurrence-oriented [hardware description language](#) developed in ACL2 for describing Mealy machines [Hunt:2000].

The DE system supports [hierarchical verification](#):

- Each time a module is specified, there are two lemmas need be proven: a [value lemma](#) specifying the module's outputs and a [state lemma](#) specifying the module's next state.
- If a module doesn't have an internal state (purely combinational), only the [value lemma](#) need be proven.

The DE System

DE is a formal occurrence-oriented [hardware description language](#) developed in ACL2 for describing Mealy machines [Hunt:2000].

The DE system supports [hierarchical verification](#):

- Each time a module is specified, there are two lemmas need be proven: a [value lemma](#) specifying the module's outputs and a [state lemma](#) specifying the module's next state.
- If a module doesn't have an internal state (purely combinational), only the [value lemma](#) need be proven.
- These lemmas are used to prove the correctness of yet larger modules containing these submodules, **without the need to dig into any details about the submodules.**

The DE System

DE is a formal occurrence-oriented [hardware description language](#) developed in ACL2 for describing Mealy machines [Hunt:2000].

The DE system supports [hierarchical verification](#):

- Each time a module is specified, there are two lemmas need be proven: a [value lemma](#) specifying the module's outputs and a [state lemma](#) specifying the module's next state.
- If a module doesn't have an internal state (purely combinational), only the [value lemma](#) need be proven.
- These lemmas are used to prove the correctness of yet larger modules containing these submodules, **without the need to dig into any details about the submodules**.
- This approach has been demonstrated its **scalability** to large systems, as shown on contemporary x86 designs at Centaur Technology [Slobodova et al.:2011].

- 1 Introduction
- 2 The DE System
- 3 Modeling and Verification Approach**
- 4 Case Studies
- 5 Future Work and Conclusions

- No global clock signal
- Local communication protocols
- Non-deterministic behavior due to variable delays in wires and gates

- No global clock signal
 - ⇒ Adding local signaling to state-holding devices
- Local communication protocols

- Non-deterministic behavior due to variable delays in wires and gates

- No global clock signal
⇒ Adding local signaling to state-holding devices
- Local communication protocols
⇒ Modeling the [link-joint model](#) introduced by Roncken et al., a universal communication model for various self-timed circuit families [Roncken et al.:2015]
- Non-deterministic behavior due to variable delays in wires and gates

- No global clock signal
⇒ Adding local signaling to state-holding devices
- Local communication protocols
⇒ Modeling the [link-joint model](#) introduced by Roncken et al., a universal communication model for various self-timed circuit families [Roncken et al.:2015]
- Non-deterministic behavior due to variable delays in wires and gates
⇒ Employing an oracle, which we call a collection of [go](#) signals

The Link-Joint Model

We model self-timed systems as **finite state machines** (FSMs) representing networks of **communication links**.

Links communicate with each other locally via **handshake components**, which are called **joints**, using the **link-joint model**.

The Link-Joint Model

We model self-timed systems as **finite state machines** (FSMs) representing networks of **communication links**.

Links communicate with each other locally via **handshake components**, which are called **joints**, using the **link-joint model**.

- **Links** are communication channels in which **data** and **full/empty states** are stored.
- **Joints** are handshake components that implement **flow control** and **data operations**.

The Link-Joint Model


We model self-timed systems as **finite state machines** (FSMs) representing networks of **communication links**.

Links communicate with each other locally via **handshake components**, which are called **joints**, using the **link-joint model**.

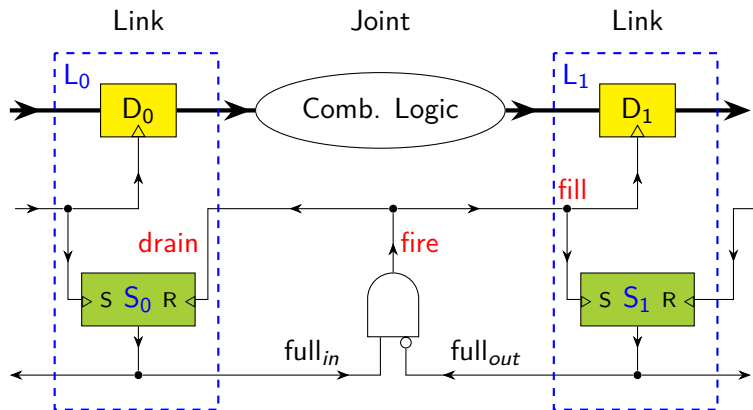
- **Links** are communication channels in which **data** and **full/empty states** are stored.
- **Joints** are handshake components that implement **flow control** and **data operations**.

Joints are the meeting points for links to **coordinate states** and **exchange data**.

The Link-Joint Model



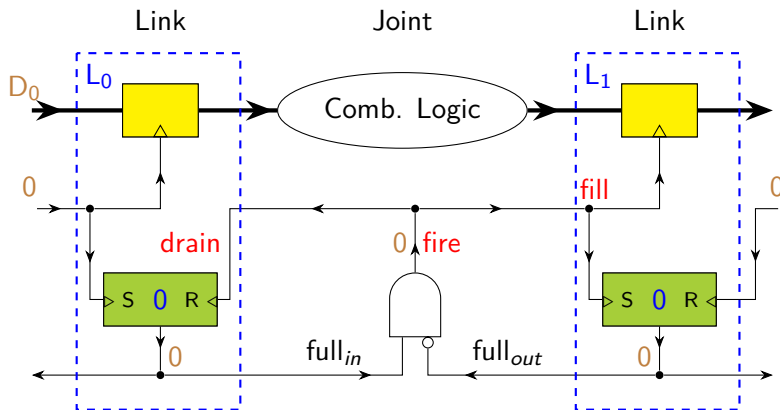
The Link-Joint Model



A joint can have several input and output links connected to it.

Necessary conditions for a joint to fire: all of its input links are **full** and all of its output links are **empty**.

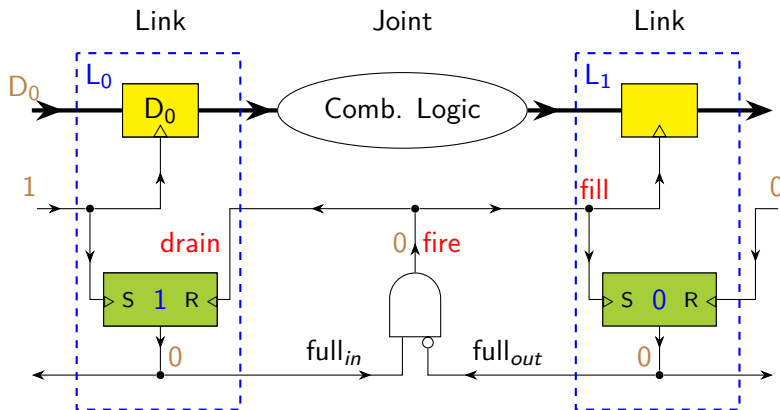
The Link-Joint Model



A joint can have several input and output links connected to it.

Necessary conditions for a joint to fire: all of its input links are **full** and all of its output links are **empty**.

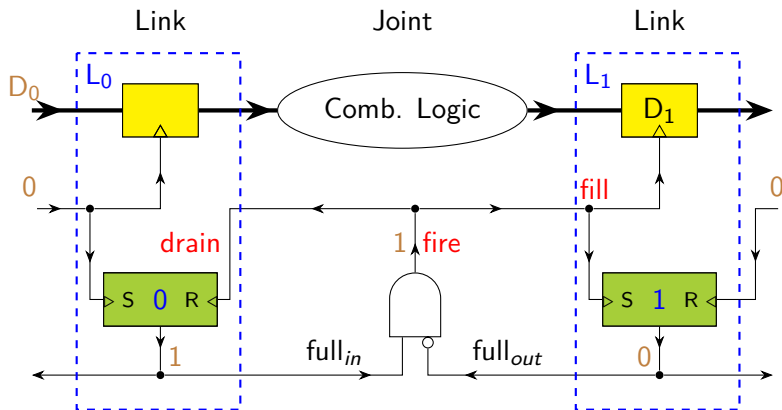
The Link-Joint Model



A joint can have several input and output links connected to it.

Necessary conditions for a joint to fire: all of its input links are **full** and all of its output links are **empty**.

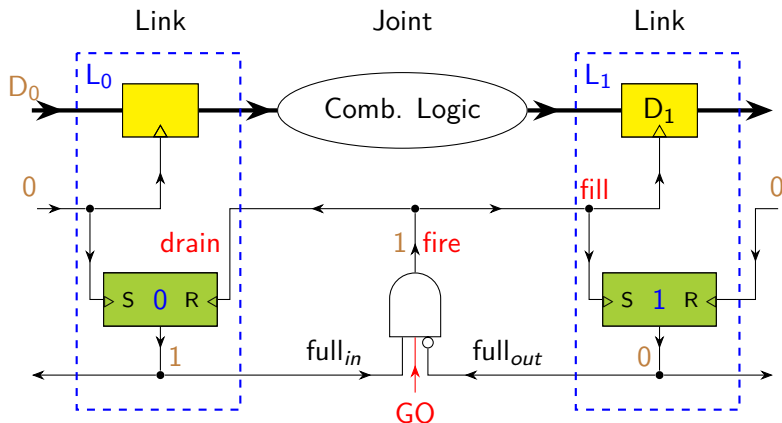
The Link-Joint Model



A joint can have several input and output links connected to it.

Necessary conditions for a joint to fire: all of its input links are **full** and all of its output links are **empty**.


The Link-Joint Model



A joint can have several input and output links connected to it.

Necessary conditions for a joint to fire: all of its input links are **full** and all of its output links are **empty**.

The Link-Joint Model




When a joint fires, the following three actions will be executed in parallel:

- transfer data computed from the input links to the output links,
- **fill** the output links, make them **full**,
- **drain** the input links, make them **empty**.


Our framework applies a [hierarchical verification](#) approach to formalizing single transitions of circuit behavior (simulated by [se](#) and [de](#) functions).

- The output and next state of a module are formalized using the formalized outputs and next states of submodules, without delving into details about the submodules.
- Self-timed modules can be abstracted as “complex” links or “complex” joints.

Self-Timed Modules



A complex link: an adder



A complex joint: a queue Q_2 of two links

Compositional reasoning:

- Functional properties of self-timed systems may involve multi-step executions that are quite burdensome to establish directly.
- Decompose the executions into smaller steps in such a way that sub-properties after executing each of these smaller steps can be carried out much easier.
- The desired properties are then established by simply composing these sub-properties.

Compositional reasoning:

- Functional properties of self-timed systems may involve multi-step executions that are quite burdensome to establish directly.
- Decompose the executions into smaller steps in such a way that sub-properties after executing each of these smaller steps can be carried out much easier.
- The desired properties are then established by simply composing these sub-properties.

Induction:


- We use induction to prove properties of systems over time, for instance, the relationship between input and output sequences.
- We also apply induction to establishing **loop invariants** of iterative circuits, i.e., circuits with feedback loops.

Outline

- 1 Introduction
- 2 The DE System
- 3 Modeling and Verification Approach
- 4 Case Studies**
- 5 Future Work and Conclusions

Circuits with No Feedback Loops


Q3



For self-timed circuits with no feedback loops, we verify their functional correctness in terms of **the relationship between their input and output sequences**.

Circuits with No Feedback Loops

Q3




For self-timed circuits with no feedback loops, we verify their functional correctness in terms of **the relationship between their input and output sequences**.

Let **in-act** denote the **fire** signal from the AND gate in the control logic of joint **in**, **out-act** denote the **fire** signal from the AND gate in the control logic of joint **out**.

Circuits with No Feedback Loops

Q3




For self-timed circuits with no feedback loops, we verify their functional correctness in terms of **the relationship between their input and output sequences**.

Let **in-act** denote the **fire** signal from the AND gate in the control logic of joint **in**, **out-act** denote the **fire** signal from the AND gate in the control logic of joint **out**.

Module Q3 will accept a new input if the **in-act** signal is **high**. In this case, we call this input **valid**. We define a (valid) input sequence as a sequence of valid inputs.

Circuits with No Feedback Loops

Q3



For self-timed circuits with no feedback loops, we verify their functional correctness in terms of **the relationship between their input and output sequences**.


Let **in-act** denote the **fire** signal from the AND gate in the control logic of joint **in**, **out-act** denote the **fire** signal from the AND gate in the control logic of joint **out**.

Module Q3 will accept a new input if the **in-act** signal is **high**. In this case, we call this input **valid**. We define a (valid) input sequence as a sequence of valid inputs.

Similarly, we define a (valid) output sequence as a sequence of valid outputs. Q3 will report a valid output when the **out-act** signal is **high**.

Circuits with No Feedback Loops

Q3



We define the function $q3\$extract-data(st)$ that extracts valid data from state st of Q3, i.e. extracts data from links that are **full** at state st . The following equality states the functional correctness of Q3:

$$q3\$extract-data(q3\$run(input-list, st, n)) ++ out-seq = \\ in-seq ++ q3\$extract-data(st)$$

where

```
q3$run(input-list, st, n) :=
  if (n <= 0)
    st
  else q3$run(rest(input-list),
              q3$step(first(input-list), st),
              n - 1)
```

Circuits with No Feedback Loops

$$q3\$extract-data(q3\$run(input-list, st, n)) ++ out-seq = \\ in-seq ++ q3\$extract-data(st) \quad (1)$$

Our ACL2 proof of (1) uses **induction** and the following **single-step-update** property of $Q3$ as a supporting lemma:


$$q3\$extract-data(q3\$step(input, st)) = q3\$step-spec(input, st) \quad (2)$$

where $q3\$step-spec(input, st) :=$

$$\begin{cases} q3\$extract-data(st), & \text{if } in-act = F \wedge out-act = F \\ [input] ++ q3\$extract-data(st), & \text{if } in-act = T \wedge out-act = F \\ remove-last(q3\$extract-data(st)), & \text{if } in-act = F \wedge out-act = T \\ [input] ++ remove-last(q3\$extract-data(st)), & \text{otherwise} \end{cases}$$

Circuits with No Feedback Loops


C



In terms of input-output relationship, C simply performs the bitwise OR operation on the two input operands. The operation of C over the input sequence is defined as follows:


```
c$op(in-seq) :=  
  if (in-seq is empty) nil  
  else {  
    let {  
      input := first(in-seq)  
      a := first(input)  
      b := second(input)}  
    [v-or(a, b)] ++ c$op(rest(in-seq))}
```

Circuits with No Feedback Loops



```
c$extract-data(st) :=  
  c$op(extract-data(st.a0 ++ st.Q2 ++ st.a1) ⊗  
    extract-data(st.b0 ++ st.Q3 ++ st.b1))  
  ++  
  extract-data(st.s)
```

Circuits with No Feedback Loops



```
c$extract-data(st) :=  
  c$op(extract-data(st.a0 ++ st.Q2 ++ st.a1) ⊗  
    extract-data(st.b0 ++ st.Q3 ++ st.b1))  
  ++  
  extract-data(st.s)
```


If the current state st of C satisfies the following condition:

$$\begin{aligned} \text{size}(\text{extract-data}(st.a0 ++ st.Q2 ++ st.a1)) = \\ \text{size}(\text{extract-data}(st.b0 ++ st.Q3 ++ st.b1)) \end{aligned} \quad (3)$$

then the following functional property of C holds:

$$\begin{aligned} c\$\text{extract-data}(c\$\text{run}(\text{input-list}, st, n)) ++ \text{out-seq} = \\ c\$\text{op}(\text{in-seq}) ++ c\$\text{extract-data}(st) \end{aligned} \quad (4)$$

Circuits with No Feedback Loops




Our ACL2 proof of (4) uses **induction** and the [single-step-update](#) property (described below, given that (3) holds) as a supporting lemma.

$$c\$extract-data(c\$step(input, st)) = c\$step-spec(input, st) \quad (5)$$

where $c\$step-spec(input, st) :=$

$$\begin{cases} c\$extract-data(st), in-act = F \wedge out-act = F \\ [v-or(input.a, input.b)] ++ c\$extract-data(st), in-act = T \wedge out-act = F \\ remove-last(c\$extract-data(st)), in-act = F \wedge out-act = T \\ [v-or(input.a, input.b)] ++ remove-last(c\$extract-data(st)), otherwise \end{cases}$$

Circuits with No Feedback Loops



Our ACL2 proof of (4) uses **induction** and the [single-step-update](#) property (described below, given that (3) holds) as a supporting lemma.

$$c\$extract-data(c\$step(input, st)) = c\$step-spec(input, st) \quad (5)$$

Our proof of (5) does not concern the details of $Q2$ and $Q3$. All we need to know about these two modules in proving (5) is their [single-step-update](#) properties.

In other words, we employ a [hierarchical strategy](#) in proving the [single-step-update](#) property of a self-timed module.

Circuits with No Feedback Loops

In summary, for each self-timed module that has no feedback loops, we prove the following two properties:

- 1 the single-step-update property (proved by using hierarchical reasoning),
- 2 the relationship between the input and output sequences (proved by using induction and the single-step-update property).

Circuits with Feedback Loops

Reasoning with highly non-deterministic behavior in self-timed systems with feedback loops is very challenging.

- Computing **loop invariants** in these systems becomes much more complicated than in synchronous systems.

Circuits with Feedback Loops

Reasoning with highly non-deterministic behavior in self-timed systems with feedback loops is very challenging.

- Computing **loop invariants** in these systems becomes much more complicated than in synchronous systems.

We impose **design restrictions** on iterative circuits to reduce non-determinism, and consequently reduce the complexity of the set of execution paths:

- These restrictions enable our framework to verify **loop invariants** efficiently via **induction** and subsequently verify the **functional correctness** of self-timed circuit designs.

Circuits with Feedback Loops

Reasoning with highly non-deterministic behavior in self-timed systems with feedback loops is very challenging.

- Computing **loop invariants** in these systems becomes much more complicated than in synchronous systems.

We impose **design restrictions** on iterative circuits to reduce non-determinism, and consequently reduce the complexity of the set of execution paths:

- These restrictions enable our framework to verify **loop invariants** efficiently via **induction** and subsequently verify the **functional correctness** of self-timed circuit designs.

Design restrictions: A module is ready to communicate with other modules only when it finishes all of its internal operations and becomes quiescent.

32-Bit Self-Timed Serial Adder Verification

We demonstrate our framework by modeling and verifying the functional correctness of a [32-bit self-timed serial adder](#) [Chau:2017].

We prove that the self-timed serial adder indeed performs the addition under an appropriate initial condition.

- When the adder finishes its execution, the result is proven to be the sum of the two 32-bit input operands and the carry-in.

32-Bit Self-Timed Serial Adder Verification

We demonstrate our framework by modeling and verifying the functional correctness of a [32-bit self-timed serial adder](#) [Chau:2017].

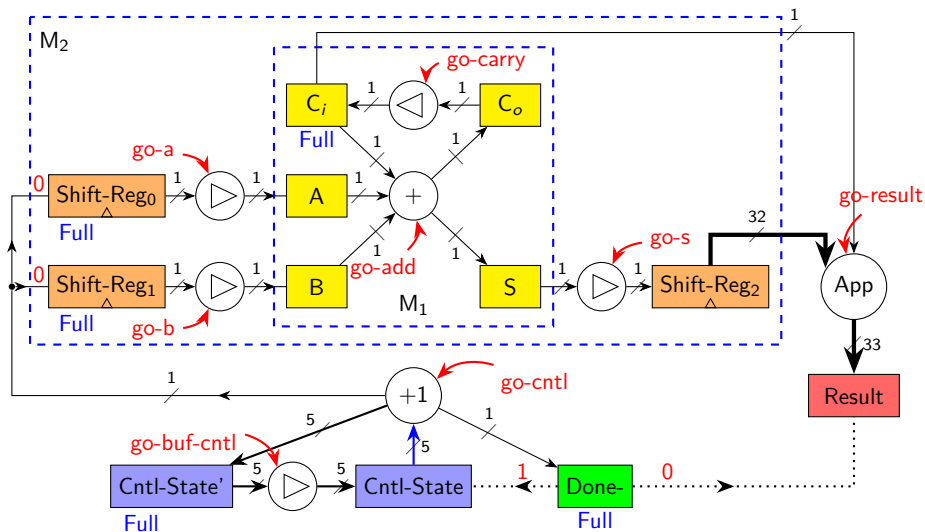
We prove that the self-timed serial adder indeed performs the addition under an appropriate initial condition.

- When the adder finishes its execution, the result is proven to be the sum of the two 32-bit input operands and the carry-in.

Our verification approach applies [compositional reasoning](#).

- Divide the adder's execution into two parts: the loop part and the exit part (the execution after exiting the loop),
- Formalize a loop invariant for the loop part and the adder behavior during the exit part,
- Prove the functional correctness of the adder by glueing these two parts together.

Data Flow of a 32-Bit Self-Timed Serial Adder



Theorem 1 (Partial correctness).

$$\text{async_serial_adder}(\text{netlist}) \wedge \quad (6)$$

$$\text{init_state}(st) \wedge \quad (7)$$

$$(\text{operand_size} = 32) \wedge \quad (8)$$

$$\text{interleavings_spec}(\text{input_seq}, \text{operand_size}) \wedge \quad (9)$$

$$(st' = \text{run}(\text{netlist}, \text{input_seq}, st, n)) \wedge \quad (10)$$

$$\text{full}(\text{result_status}(st')) \quad (11)$$

$$\Rightarrow (\text{result_value}(st') = \text{shift_reg_0_value}(st) + \\ \text{shift_reg_1_value}(st) + \\ \text{ci_value}(st))$$

Theorem 2 (Termination).

$$\text{async_serial_adder}(\text{netlist}) \wedge \quad (6)$$

$$\text{init_state}(st) \wedge \quad (7)$$

$$(\text{operand_size} = 32) \wedge \quad (8)$$

$$\text{interleavings_spec}(\text{input_seq}, \text{operand_size}) \wedge \quad (9)$$

$$(st' = \text{run}(\text{netlist}, \text{input_seq}, st, n)) \wedge \quad (10)$$

$$(n \geq \text{num_steps}(\text{input_seq}, \text{operand_size})) \quad (11')$$

$$\Rightarrow \text{full}(\text{result_status}(st'))$$

- 1 Introduction
- 2 The DE System
- 3 Modeling and Verification Approach
- 4 Case Studies
- 5 Future Work and Conclusions**

Future Work

We are developing a new proof technique for [partial correctness](#) of iterative self-timed circuits that does not have any conditions on the values of **go** signals.

We are developing a new proof technique for **partial correctness** of iterative self-timed circuits that does not have any conditions on the values of **go** signals.

For **termination** proofs, we need a constraint on **go** signals guaranteeing that **delays are bounded**.

We are developing a new proof technique for **partial correctness** of iterative self-timed circuits that does not have any conditions on the values of **go** signals.

For **termination** proofs, we need a constraint on **go** signals guaranteeing that **delays are bounded**.

We intend to follow a **hierarchical approach** to prove module-level properties of iterative circuits of the following form:

- Given an initial state of the module, the module's **final state** meets its specification after that module completes execution.

Conclusions

We have presented a framework for modeling and verifying self-timed circuits using the DE system.

Our goal is to develop a methodology that is capable of verifying the functional correctness of self-timed circuit designs at large scale.

- This work also provides libraries for analyzing self-timed systems in ACL2.

We model self-timed systems as networks of links communicating with each other locally via joints, using the [link-joint model](#) introduced by Roncken et al.

We also model the **non-determinism of event-ordering** in self-timed circuits by associating each joint with an external [go](#) signal.

Our key proof techniques are hierarchical reasoning, compositional reasoning, and induction.



C. Chau, W. Hunt, M. Roncken, and I. Sutherland (2017)
A Framework for Asynchronous Circuit Modeling and Verification in ACL2
HVC 2017, to appear.



W. Hunt (2000)
The DE Language
Computer-Aided Reasoning: ACL2 Case Studies, Kluwer Academic Publishers
Norwell, MA, USA, 151 – 166.



M. Roncken, S. Gilla, H. Park, N. Jamadagni, C. Cowan, I. Sutherland (2015)
Naturalized Communication and Testing
ASYNC 2015, 77 – 84.



A. Slobodova, J. Davis, S. Swords, and W. Hunt (2011)
A Flexible Formal Verification Framework for Industrial Scale Validation
MEMOCODE 2011, 89 – 97.

Questions?