

Temporal Analysis of Time Bounded Digital Systems

Alan R. Martello and Steven P. Levitan

University of Pittsburgh, Department of Electrical Engineering, Pittsburgh, PA 15261

Abstract. To perform verification of digital systems with time bounded delays, it is essential to characterize the space of all possible system behaviors. In this paper, we describe our analysis technique which accepts a behavioral specification of the timing of a digital system and generates the set of all possible behaviors for the system. The ability to represent and reason about time ranges for events is a distinguishing characteristic of our technique and gives our analysis method both its power and complexity.

1 Introduction

Our research focuses on the analysis of digital system descriptions which include time bounded delays. This analysis is part of the problem of verification: given behavioral specifications of a system, check that no possible behavior of the system would violate any requirement constraints imposed on the system. The requirements may come either from the designer or from the needs of the components (e.g., setup and hold constraints) which make up the system. In general, our model of digital systems with time bounded delays allows for the specification of digital systems, the generation of all possible system behaviors given this specification and an initial system state, and the analysis of the resulting system behaviors.

In this paper we focus on the key part of this analysis, the generation of all possible system behaviors from a system specification. We are primarily concerned with the temporal behavior of digital systems, not their function. We are only concerned with “data values” to the extent that one signal could enable or disable the occurrence of an event on another signal in the system.

Attempts to reason about digital systems have been ongoing[1–3]. Various methods for modeling systems with time delay have been proposed and formal verification methods investigated[4–7]. Methods based on simulation have also been used to verify circuits[8,9]. Other recent work has taken a symbolic approach to attempt to reason about circuits[10,11] and our earlier work[12, 13] investigated a more simplistic verification technique upon which this work is based. Recent work in asynchronous verification [14] performs an analysis similar to our search process and also depends on the need to perform bifurcations (or splits) as system behaviors are generated.

2 Overview & Example

To provide an overview of our analysis technique, we present an example of an RS latch constructed from two nor gates, each with gate delay in the range 1 to 3 (Fig. 1(a)), and with the system behavioral specification shown in Fig. 1(b). The RS latch starts with R, S, and Q low and with QBAR high. A rising transition occurs on S at $t = 0$. Figure 1(c) illustrates the timing waveform for this input transition of the latch.




Fig. 1. RS Latch (a), behavioral description (causalities) (b), and timing waveform (c)

If we were doing verification of the latch, we would be checking this waveform against a system requirement like “Q must rise within 5 of S rising.” This requirement may be imposed by other modules in the system or by external constraints on the system. In order to perform verification, all system requirements are compared with all system behaviors generated, to verify proper system operation. The most difficult part of this verification process is the generation of all possible system behaviors.

In order to generate all possible system behaviors for digital systems with time bounded delays, we must enumerate (or search) the space of all possible interactions among signals which often have mutual interdependencies on their transitions. The search space is bounded by identifying system states which have previously occurred, effectively handling cycles and reconvergent behaviors. Conversely, the size of the search space grows based on the interaction among signals and the time ranges associated with their interdependent transitions.

As an example of such a search space explosion, we revisit the RS latch previously shown. The gates which comprise the latch in this example have gate delays

of 1 to 3 and both inputs fall from the “illegal input state” at $t = 0$. The four behaviors shown in Fig. 2 are representative of the 26 behaviors resulting from our analysis of 128 system states (taking 16.90 cpu seconds on a Sparcstation 2). The many behaviors explored are caused by the splitting or bifurcation of nodes in the search space of possible behaviors. From one starting state and initial transitions of R and S the search space grows, caused by alternative interactions based on system state and dissection of the time ranges.




Fig. 2. Sampling of system behaviors if both S and R fall at $t = 0$

Our research shows that a useful model for the analysis and verification of digital systems must include both time ranges and signal interaction based on system state. The interaction of these two elements implies bifurcation and growth of the space of system behaviors which must be explored in this analysis. The next section provides a formal description of our specification and analysis technique followed by an overview of the algorithm developed to perform this analysis.¹

3 Background & Definitions

For our analysis we define a system as a collection of components and the signals which interconnect them. To characterize the behavior of the system we define the system specification in terms of the events which occur on the signals of the system. In this work, the analysis process is dependent on an underlying assumption that the system is completely specified (i.e., all information regarding the system is known). This assumption is necessary and sufficient for the generation of all possible system behaviors. It is necessary since permitting the possibility

¹ The complete algorithm is fully specified in [15].

of unknown information in the system effecting its operation may result in the generation of system behaviors which cannot occur or the absence of system behaviors which can occur.

A partially specified system is one where not all information is known about how all signals in the system interact. Taken to an extreme, this implies that signals may change values in an undetermined manner, at any time, precluding the possibility of generating all possible system behaviors. The assumption that the system is fully specified is sufficient for the generation of all possible system behaviors. If all information is known about the system, then by definition no unanticipated action can occur in the system which affects the generation of possible system behaviors, within the scope of the model used.

We capture the specification as a set of causalities, each describing under what circumstances an event on one signal “causes” an event on another signal in the system. Informally causalities are of the form: “clock rising when D is high causes Q to go high after 10-20ns when reset is low.” More formally we describe the temporal specification of the system using the following definitions.

Definition 1. The *specification of the system* is the set of all causal relationships among the signals of the system: $\mathcal{C} \equiv \{c \mid c = \langle \varepsilon_1, \varepsilon_2, \delta \rangle\}$; i.e. the set of causalities are the rules by which an event on one signal causes an event on another signal.

Definition 2. A *causality*, $c \equiv \langle \varepsilon_1, \varepsilon_2, \delta_c \rangle$ where ε_1 and ε_2 are two guarded events and $\delta_c \equiv (\alpha_c, \beta_c)$ is a time range.

For a given causality, the second event, ε_2 , is specified to occur within δ_c relative to the first event, ε_1 . The time range δ_c must contain two positive values.

A causality reflects a characteristic of the system being verified; it specifies that if ε_1 is enabled, ε_2 will occur provided it is also enabled within the time range δ_c after ε_1 .

Definition 3. A *guarded event*, $\varepsilon \equiv \langle s, e, w \rangle$, is an event, e , on a signal, s , and an enabling expression, w , which identifies when the event is enabled. A guarded event is enabled at some time t if the enabling expression evaluates to *true* at time $t - 1$.

Definition 4. A *time range*, δ , specifies a time window for events where $\delta \equiv (min, max)$ and $min \leq max$.

Time is treated as a discrete quantity measured in an integral number of arbitrary time units. Time ranges are specified relative to some time, t . The fact that time is not continuous is not a strong constraint since the unit by which time is interpreted can be made arbitrarily small, approaching a closer approximation of continuous time, if needed.

Definition 5. A *signal*, s , is an information channel which is used to interconnect modules. Signals may connect multiple modules or a module to itself. The set of all values any signal may have is $\{high, low\} \equiv \{true, false\} \equiv \{1, 0\}$.

A signal is a physical entity (e.g., a wire) or a virtual identifier (e.g., a variable) which can have two values, high and low.

Definition 6. An *event*, e , is a rising or falling transition of a signal which occurs instantaneously.

There is no rise time or fall time associated with events, however, events themselves can occur at any point within a time range. This time range may have a zero width in which case the event occurs at a specific time. The transition of s from high to low is called the event s falling ($s\downarrow$) and the transition of signal s from low to high is the event s rising ($s\uparrow$).

Definition 7. An *enabling expression* (also referred to as a *when clause*), w , is a (possibly null) set of conjunctive (i.e., logical “and”) expressions of signal names.

Definition 8. A *causal graph* is an augmented signal transition graph (STG) where the nodes are events and the directed arcs are the causalities (causal relationships between events). The graph is augmented such that the arcs contain the two enabling expressions, w_1 and w_2 , and the time range, δ specified by the causality. Multiple arcs may exist between any two nodes in the graph.

With the causal graph and a set of initial conditions, it is possible to construct a system history graph which describes all possible behaviors of the system. The algorithm for the construction of the history is given in the next section. Once a history graph is generated, it is possible to traverse the graph to generate a visual representation (timing diagram) of each system history.

Definition 9. A *system history graph*, \mathcal{H} , is a directed graph which describes all possible behaviors from an initial state. Each node in the history graph is a dynamic state, \mathcal{D} , and the arcs specify a temporal ordering of the states.

Definition 10. A *dynamic state* of the system is $\mathcal{D} \equiv \langle \mathcal{S}, \mathcal{P}, t, \mathcal{S}_0, \mathcal{P}_0, \langle ds\text{-}flags \rangle \rangle$ where \mathcal{S} is the static state of the system, \mathcal{P} is the set of pending transitions, \mathcal{S}_0 is the initial static state when \mathcal{D} is created, \mathcal{P}_0 is the initial dynamic state when \mathcal{D} is created, t is the time of the earliest occurring transition in \mathcal{P}_0 , and $\langle ds\text{-}flags \rangle \equiv \langle ProcessedFlag \rangle$ is used in the construction of \mathcal{H} . We say that t is the time at which \mathcal{D} occurs.

The dynamic state holds all the information about the state of the system. Therefore, given any dynamic state (and the causal graph), all possible behaviors evolving from that state can be determined.

For any two dynamic states, \mathcal{D}_i and \mathcal{D}_j , we define $\mathcal{D}_i = \mathcal{D}_j$ if $\mathcal{S}_{0:i} = \mathcal{S}_{0:j}$ and $\mathcal{P}_{0:i} = \mathcal{P}_{0:j}$. That is, these two states are equivalent such that the entire set of derivable system behaviors which can evolve from each of them is identical. Note that for all transitions in a dynamic state, all times are relative to t for that state.

The history graph is complete when $\forall i$, all successors (children) to \mathcal{D}_i already appear in the graph or there are no successors to \mathcal{D}_i . We call such a node a terminal node. The predecessor of \mathcal{D}_i , called \mathcal{D}_i ’s parent, is identified as \mathcal{D}_i^P .

The history graph has the following properties.

- There is a single starting dynamic state, \mathcal{D}_0 ; this is the only node with no parent.
- $\forall i \neq j, \mathcal{D}_i \neq \mathcal{D}_j$.
- $\forall i \neq 0, t_i^P < t_i$.

Definition 11. The *static state* of the system, \mathcal{S} , is the last transition for all signals in the system relative to t , the time of the dynamic state: $\mathcal{S} \equiv \{l(s_i) \mid \forall \text{signals}, s_i\}$. The static system state is a snapshot of the most recent transition for every signal in the system relative to t .

Definition 12. A *signal transition*, τ , is a tuple $\langle \epsilon, \delta, \delta_0, \eta, \eta_c, \langle st\text{-}flags \rangle \rangle$ where ϵ is a guarded event, $\delta \equiv (\alpha, \beta)$ is the current time range for the transition, $\delta_0 \equiv (\alpha_0, \beta_0)$ is the initial time range for the transition, η is a unique serial number for the transition, η_c is the serial number for the transition which caused this transition, and $\langle st\text{-}flags \rangle \equiv \langle TryFalseFlag, DidOccurFlag \rangle$ are used in processing the transition.

Definition 13. A *pending signal transition*, p , is a signal transition with $\alpha > 0$ which may occur in the future.

Pending transitions are created which match the right hand side of a causality when an enabled transition occurs which matches the event specified by the left hand side of a causality. The set of all pending transitions in the system for a particular dynamic state, \mathcal{D} , is \mathcal{P} .

Definition 14. The *last transition* for a signal, $l(s)$, relative to t is a signal transition τ such that $\beta \leq 0$ and β is a maximum relative to t .

Definition 15. A *system history* is defined as any path through the completed system history graph from the starting node, \mathcal{D}_0 to a terminal node.

4 Algorithm

The behavior of the system is completely specified by the history graph, \mathcal{H} . The construction of the history graph from the causal graph begins with the specification of an initial state, \mathcal{D}_0 . Given this initial state, the main processing loop of the algorithm simply chooses any node in the graph which has not yet been processed and processes that node.

Pending transitions in the current dynamic state are processed by first testing if the transition is enabled, moving enabled transitions to the static state, and possibly creating a bifurcation in the history graph. A bifurcation is a splitting of a dynamic state, \mathcal{D} , in the system history graph into two separate dynamic states. The composition (or union) of the two dynamic states is equivalent to the original dynamic state. The split is necessary to disambiguate possible system history alternatives. The procedure specified by *ProcessNode* shown in Fig. 3 handles the processing of a given history node, \mathcal{D} , which has not yet been processed.

```

1  ProcessNode( $\mathcal{D}$ ) {
2      foreach  $p \in \mathcal{P}$  s.t.  $\alpha = 0$  {
3          if (TransitionOccurred?( $p, \mathcal{D}$ )) { ProcessCausalities( $p, \mathcal{D}$ ) }
4           $\alpha_p \leftarrow 1$ 
5          RemoveTransitionIfApplicable( $p, \mathcal{D}$ ) }
6      # check if done processing, i.e.  $\mathcal{D}$  has no pending transitions
7      if ( $\mathcal{P} = \emptyset$ ) { return }
8      # create the new node by copying the parent
9       $\mathcal{D}_{new} \leftarrow \langle \mathcal{S}, \mathcal{P}, t, \mathcal{S}, \mathcal{P}, \langle \text{false} \rangle \rangle$ 
10     # find the time offset between the new node and its parent
11      $t_{Delta} \leftarrow \min(\alpha \in \mathcal{P})$ 
12     # set the new time for the new node
13      $t_{new} \leftarrow t_{Delta} + t$ 
14     # adjust all times in  $\mathcal{D}_{new}$  to be relative to  $t_{new}$  rather than  $t$ 
15     foreach  $\alpha, \beta \in \mathcal{S}_{new}, \mathcal{P}_{new}, \mathcal{S}_{0:new}, \mathcal{P}_{0:new}$ , subtract  $t_{Delta}$ 
16     # make sure a node the same as  $\mathcal{D}_{new}$  doesn't already exist
17     if ( $\exists \mathcal{D}_i \in \mathcal{H}$  s.t.  $\mathcal{D}_i = \mathcal{D}_{new}$ ) {
18         point  $\mathcal{D}$  to  $\mathcal{D}_i$ , destroy  $\mathcal{D}_{new}$ , return }
19     else { point  $\mathcal{D}$  to  $\mathcal{D}_{new}$ , return } }

```

Fig. 3. Procedure to process a dynamic history node

First, all pending transitions which might occur at t (i.e., $\alpha = 0$) are processed (line 2). If the pending transition occurred, causalities which are enabled by the transition are processed, possibly creating additional pending transitions (line 3). If a particular transition occurs, the causality graph is checked for all new transitions which may be caused by this transition and the new transitions are added to \mathcal{P} . Finally, all pending transitions processed are removed from \mathcal{P} if they can no longer occur (i.e., $\beta = 0$) (line 5).

Next, the node is checked for pending transitions. If no pending transitions are present, then no further processing need be performed on \mathcal{D} (line 7). If \mathcal{D} has pending events, then a successor (child node) to \mathcal{D} is created by copying \mathcal{S} into both \mathcal{S}_{new} and $\mathcal{S}_{0:new}$ and copying \mathcal{P} into both \mathcal{P}_{new} and $\mathcal{P}_{0:new}$ (line 9). The initial copies are saved in $\mathcal{S}_{0:new}$ and $\mathcal{P}_{0:new}$ to permit backtracking or rollback of time to a previous state during history bifurcation. Next, the time of \mathcal{D}_{new} is determined by taking the minimum starting time for all pending events in \mathcal{D}_{new} and all times in \mathcal{D}_{new} are adjusted to be relative to t_{new} (lines 11-15).

A check is then made to see if \mathcal{D}_{new} is equivalent to a node which already exists in the graph, \mathcal{D}_i (i.e. if \mathcal{S}_0 and \mathcal{P}_0 of the two history nodes match). If \mathcal{D}_{new} is equivalent to \mathcal{D}_i , then no new histories would be generated by processing \mathcal{D}_{new} , therefore \mathcal{D}_{new} is destroyed and \mathcal{D} is pointed to \mathcal{D}_i (line 18). If \mathcal{D}_{new} does not match some other node, \mathcal{D}_i in the graph, then processing continues, linking \mathcal{D}_{new} into the history graph by pointing \mathcal{D} to \mathcal{D}_{new} (line 19).

The procedure *TransitionOccurred?* seen in Fig. 4 and Fig. 5 handles the processing of a given transition, $p \in \mathcal{P}$. First, a check is made to see if all signals used in w_p are defined. For each signal that is unknown, a bifurcation occurs

```

# process a single transition pending now, p
1 TransitionOccurred?(p, D) {
2   # check if all signals used in enabling expressions are defined
3   if ((SigList ← ReturnListOfUndefinedSignals(wp, D)) ≠ null) {
4     SplitUndefinedEnablingSignals(p, SigList, D) }
5   # only process p if it is enabled now (at t)
6   if (not EnabledNow?(wp, ηp, D)) {
7     # try the false case as long as p did not cause its own
8     # enabling expression to become false
9     if (not TransitionCausedEnablingChange?(p, D)) {
10      (TryFalseFlag) ← (true) }
11    # not enabled ⇒ no additional processing for p
12    return false }
13   if (∃pi ∈ P s.t. αpi = 0 ∧ βp ≠ βpi ∧ pi is enabled ∧ p ≠ pi ∧ sp = spi) {
14     # split since both p and pi can occur
15     SplitOtherChangesPendingNow(p, D) }
16   # check for another enabled transition, pi, pending now
17   # (α = 0) with direction opposite of p (↑ vs. ↓)
18   if (∃pi ∈ P s.t. αpi = 0 ∧ pi is enabled ∧ sp = spi ∧ ep ≠ epi) {
19     # both p and pi can occur, try the false case later
20     (TryFalseFlag) = (true) }

```

Fig. 4. Procedure to process a single transition – part 1

for the signal to take on a high and low value (lines 3-4). Next, the enabling expression in p is checked to see if the event is enabled at t (line 6).

If p is not enabled, then a flag is set indicating that a bifurcation may be necessary when p is removed from \mathcal{P} (i.e., when $\beta = 0$) and the processing of p is complete at t (line 12). If p can occur, then a test is made to see if another transition is pending in \mathcal{P} , on the same signal s , of the same event type as e , which can occur now ($\alpha = 0$) and is enabled (line 13). If all these conditions are true, then a bifurcation in \mathcal{H} occurs (line 15).

In *TransitionOccurred?*, the next step is to check if another transition pending now on p , is enabled, and is opposite to e (i.e., rising vs. falling) (line 18). If such a transition is pending, then a flag is set in p indicating that a bifurcation is necessary when p is removed from \mathcal{P} (line 20). After this processing (after line 20), the dynamic state is updated to reflect the occurrence of p .

In the next part of *TransitionOccurred?* shown in Fig. 5, if the transition on s_p in the static state, τ_{last} , has the same serial number, η as p in the dynamic state, then the current transition is the same as the past transition (line 24). In this case, the ending time of the transition is simply updated in τ_{last} if s_p is not used in any enabling expressions in \mathcal{P}_0 (line 33). If s_p is used in an enabling expression in \mathcal{P}_0 , then a bifurcation is necessary to guarantee that the value of s_p is stable at the time used, $t - 1$ (line 29).

If p does not match the past transition in \mathcal{S} , there are two cases. The first case is that p was already processed and some other intervening transition occurred on p (line 37). If this first case is true and the the past transition on s and the


```

21  # find the previous transition in the static state  $\mathcal{S}$ 
22   $\tau_{last} \leftarrow l(s_p)$ 
23  # check if  $p$  is a continuation (elongates)  $\tau_{last}$ 
24  if ( $\varepsilon_p = \varepsilon_{\tau_{last}} \wedge \eta_p = \eta_{\tau_{last}} \wedge \beta_{\tau_{last}} = -1$ ) {
25      # the current transition is the same as the last transition,
26      # a split is needed if  $s_p$  is used in any enabling expression
27      # in a pending transition in  $\mathcal{P}_0$ 
28      if (SignalUsedInEnablingExpression?( $p, \mathcal{D}$ )) {
29          SplitSignalUsedInEnablingExpression( $p, \mathcal{D}$ )
30          return false }
31      else {
32          # elongate  $\tau_{last}$  by one time unit
33           $\beta_{\tau_{last}} \leftarrow \beta_{\tau_{last}} + 1$ 
34          return true } }
35  else {
36      #  $p$  is not a simple elongation of  $\tau_{last}$ 
37      if (( $\langle DidOccurFlag \rangle = \langle true \rangle$ )) {
38          # try the false case if  $p$  and  $\tau_{last}$  aren't directly related,
39          # i.e. if  $\tau_{last}$  didn't cause  $p \wedge p$  didn't cause  $\tau_{last}$ 
40          if (not MutualCause?( $\tau_{last}, p, \mathcal{D}$ )) {
41               $\langle TryFalseFlag \rangle \leftarrow \langle true \rangle$  }
42          # part of  $p$  already occurred, ignore  $p$  for now
43          return false }
44      else {
45          #  $p$  did not previously occur, add the new transition
46           $\langle DidOccurFlag \rangle \leftarrow \langle true \rangle$ 
47           $p_{copy} \leftarrow \text{copy of } p$ 
48           $\alpha_{p_{copy}} \leftarrow \beta_{p_{copy}} \leftarrow 0$ 
49           $\tau \leftarrow l(s_p)$  and replace  $\tau \in \mathcal{S}$  with  $p_{copy}$ 
50          return true } } }

```

Fig. 5. Procedure to process a single transition – part 2

pending transition on s, p , do not have a direct causal relationship, then set the flag indicating a bifurcation is necessary when p is removed from \mathcal{P} (line 41).

The second case handled is if p did not previously occur, then the leading edge of the pending transition is peeled from p and the last transition in \mathcal{S} is replaced by a new transition specified by the leading edge of p (lines 46-50).

The routine *ProcessCausalities* shown in Fig. 6 is executed only if p occurred at t . In this case, all causalities which have p as their left hand side are evaluated. For each such causality, the enabling expression is tested and each causality with a true enabling expression on its left hand side has the transition specified by its right hand side placed in a new pending transition, p_{new} (lines 7-10). A check is made to see if p_{new} is an extension of an existing pending transition in \mathcal{P} ; if so, the time is updated in the existing transition (lines 13-16). Otherwise if a transition identical to p_{new} does not already exist, then p_{new} is added to \mathcal{P} (line 21).

```

1  ProcessCausalities( $p, \mathcal{D}$ ) {
2      foreach  $c \in \mathcal{C}$  {
3          # check if all signals used in causality lhs are defined
4          if ( $(\text{SigList} \leftarrow \text{ReturnListOfUndefinedSignals}(w_{1,c}, \mathcal{D})) \neq \text{null}$ ) {
5              SplitUndefinedEnablingSignals( $p, \text{SigList}, \mathcal{D}$ ) }
6          # check if the lhs transition of the causality occurred
7          if ( $s_{1,c} = s_p \wedge e_{1,c} = e_p \wedge \varepsilon_1$  is enabled) {
8              # the causality can occur, create a new pending transition
9              # specified by the rhs of the causality
10              $p_{\text{new}} \leftarrow \langle \langle s_{2,c}, e_{2,c}, w_{2,c} \rangle, \delta_c, \delta_c, \eta_{\text{unique}}, \eta_p, \langle \text{false}, \text{false} \rangle \rangle$ 
11             # check if the new transition caused is an extension
12             # of an existing transition
13             if ( $\exists p_i$  s.t.  $p_i \in \mathcal{P} \wedge \varepsilon_{p_i} = \varepsilon_{p_{\text{new}}} \wedge \eta_{c:p_i} = \eta_{c:p_{\text{new}}} \wedge$ 
14                  $(\alpha_{p_i} + 1) \leq \alpha_{p_{\text{new}}} \wedge \beta_{p_i} \geq \alpha_{p_{\text{new}}} \wedge \beta_{p_i} \leq \beta_{p_{\text{new}}}$ ) {
15                 # update the ending time of the transition
16                  $\beta_{0:p_i} \leftarrow \beta_{p_i} \leftarrow \beta_{p_{\text{new}}}$  }
17             else {
18                 # make sure  $p_{\text{new}}$  is not redundant in  $\mathcal{P}$ 
19                 if ( $\text{not}(\exists p_i \in \mathcal{P}$  s.t.  $\varepsilon_{p_i} = \varepsilon_{p_{\text{new}}} \wedge \eta_{c:p_i} = \eta_{c:p_{\text{new}}})$ ) {
20                     # add new pending transition
21                     add  $p_{\text{new}}$  to  $\mathcal{P}$ 
22                     # check if this pending transition already occurred
23                     # and has been deleted from  $\mathcal{P}$ 
24                     if ( $\exists \tau \in \mathcal{S}$  s.t.  $\varepsilon_{\tau} = \varepsilon_{p_{\text{new}}} \wedge \eta_{c:\tau} = \eta_{c:p_{\text{new}}}$ ) {
25                          $\langle \text{DidOccurFlag} \rangle_{p_{\text{new}}} \leftarrow \langle \text{true} \rangle$ 
26                          $\eta_{p_{\text{new}}} \leftarrow \eta_{\tau}$  } } } } } }
```

Fig. 6. Procedure to create new transitions from causalities

5 Conclusion & Summary

The verification of digital circuits is a challenging problem. Our method of verifying correct operation is to generate all possible signal interactions and then to check that specified relationships hold between the signals (e.g., setup or hold times). However, to accurately model digital systems, both bounded time ranges and state (i.e., enabling expressions) are a necessity. The ability to represent and reason about time ranges for events is a distinguishing characteristic of our technique and gives our analysis both its power and complexity. Our approach has been shown to significantly reduce the size of the search space over a “brute force” approach (simulation using all possible combinations of gate delay) as shown by the analysis of the C17 benchmark circuit in Fig. 7(a). The table shown in (b) indicates a significant reduction in number of paths explored in generating all possible system behaviors assuming each gate has the specified propagation delay.

In this paper, we have presented a method for generating all possible system behaviors of a bounded delay digital circuit model given a specification and initial conditions. Although our technique explicitly enumerates the entire space

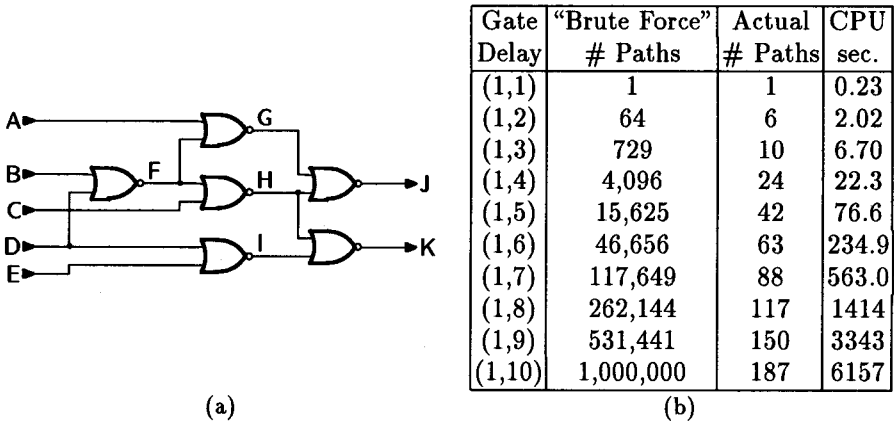


Fig. 7. C17 schematic (a); "brute force" vs. actual number of paths through C17 (b)

of possible behaviors, we have shown that it is reasonable to analyze non-trivial systems in a short period of time. Further, describing systems with causalities is feasible for non-trivial systems. Causalities can often be derived directly from vendor supplied component specifications. Although a large number of dynamic states may be enumerated during the analysis, the automatic generation of timing diagrams from system histories makes it possible for a human to understand the behavior of complex systems.

We have emphasized the generation of system behaviors from the specification of digital circuits since this is key to any analysis, not only verification. Our future work consists of investigating ways of automatically generating appropriate initial states and transitions to enumerate all states which could result in requirement violations, thus providing a total verification system for time bounded digital circuits.

We gratefully acknowledge the support of this research by the National Science Foundation grant number MIP-9102721.

References

1. Thomas M. McWilliams, "Verification of Timing Constraints on Large Digital Systems," *Proc. 17th Design Automation Conference* (June, 1980).
2. Robert B. Hitchcock, Sr., "Timing Verification and the Timing Analysis Program," *Proc. 19th Design Automation Conference* (June, 1982).
3. John J. Wallace, "On Automatic Verification of SLIDE Descriptions," Design Research Center, Carnegie-Mellon University, DRC-01-2-80, Pittsburgh, PA, Aug., 1979.

4. David E. Wallace, "Abstract Timing Verification for Synchronous Digital Systems," Computer Science Division, Univ. of Calif. at Berkeley, UCB/CSD 88/425, Berkeley, CA, June 27, 1988.
5. David L. Dill, "Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits," Dept. of Computer Science, Carnegie-Mellon University, CMU-CS-88-119, Pittsburgh, PA, Feb., 1988.
6. M. Browne, E. Clarke, D. Dill & B. Mishra, "Automatic Verification of Sequential Circuits Using Temporal Logic," Dept. of Computer Science, Carnegie-Mellon University, CMU-CS-85-100, Pittsburgh, PA, Dec., 1984.
7. Patrick C. McGeer & Robert K. Brayton, "Timing Analysis in Precharge / Unate Networks," *Proc. 27th Design Automation Conference* (June, 1990).
8. Tod Amon & Gaetano Borriello, "On the Specification of Timing Behavior," *Proceedings of the 1990 ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU '90)*, New York, NY, SIGDA (Aug., 1990).
9. Dimitrie Doukas & Andrea S. LaPaugh, "CLOVER: A Timing Constraints Verification System," *Proceedings of the 1990 ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU '90)*, New York, NY, SIGDA (Aug., 1990).
10. F. Mavaddat & T. Gahlinger, "On Deducing Tight Bounds from Partial Timing Specifications," *Proceedings of the 1990 ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU '90)*, New York, NY, SIGDA (Aug., 1990).
11. Michael C. McFarland, "CPA: Giving an Account of Timed System Behavior," *Proceedings of the 1990 ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU '90)*, New York, NY, SIGDA (Aug., 1990).
12. Alan R. Martello & Steven P. Levitan, "Causal Timing Verification," *Proceedings of the 1990 ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU '90)*, New York, NY, SIGDA (Aug., 1990).
13. Alan R. Martello, Steven P. Levitan & Donald M. Chiarulli, "Timing Verification Using HDTV," *Proc. 27th Design Automation Conference* (June, 1990).
14. Srinivas Devadas, Kurt Keutzer, Sharad Malik & Albert Wang, "Verification of Asynchronous Interface Circuits with Bounded Wire Delays," *Proc. 1992 Inter. Conf. for Computer-Aided Design* (Nov., 1992).
15. Alan R. Martello, "Temporal Specification Verification," Dept. of Elect. Eng., University of Pittsburgh, PhD. Dissertation (in preparation).