

Testing Two-phase Transition Signaling based Self-timed Circuits in a Synthesis Environment

Prabhakar Kudva

Department of Computer Science
University of Utah
Salt Lake City, UT 84112

Venkatesh Akella

Department of ECE
University of California, Davis
Davis, CA 95616

Abstract

The problem of testing self-timed circuits generated by an automatic synthesis system is studied. Two-phase transition signaling is assumed and the circuits are targeted for an asynchronous macromodule based implementation as in [4, 11, 1, 6]. The partitioning of the circuits into control blocks, function blocks, and predicate (conditional) blocks, originally conceived for synthesis purpose, is found to be very elegant and appropriate for test generation. The problem of data dependent control flow is solved by introducing a new macromodule called SCANSELECT (SELECT with scan). Algorithms for test generation are based on the Petri-net like representation of the physical circuit. The techniques are illustrated on the high-level synthesis system called SHILPA being developed by the authors. An important contribution of the paper is the technique to test both datapath and control parts (unlike testing only control parts as in most of the existing works in this area) of self-timed circuits in a unified framework by adapting existing algorithms for testing synchronous circuits.

1 Introduction

Asynchronous/Self-timed designs are beginning to attract renewed attention as promising means of dealing with the complexity of modern VLSI designs. The advantages of self-timed systems include (i) absence of global clocking and associated problems of reliable clock generation and distribution and loss of valuable power in clock drivers. (ii) ability to lend themselves better for incremental modifications, as there are no global control schedules. (iii) capability to exhibit better average case performance, as they do not have to wait for the next clock "tick" to arrive before an operation can be triggered [9, 12]. Absence of global clock-

ing and adherence to local communication protocols based on *handshaking* also renders self-timed circuits as *ideal* candidates for language-based (or behavioral) synthesis.

Testing of self-timed circuits has received little attention in recent times as most researchers have been focusing on specification, synthesis, and verification aspects self-timed design. In this paper we investigate a scheme to incorporate testing in a language-based synthesis scheme for self-timed circuits. More specifically, we find that a partitioning technique originally devised for synthesis of self-timed circuits from behavioral descriptions actually lends itself very naturally to an elegant testing strategy. Also, we find that one of the basic macro modules, namely, the SELECT module, which is used to implement data-dependent control flow, can be modified slightly to yield a SCANS-ELECT module which makes the circuit more observable and controllable.

The techniques presented in this paper are based on the high-level synthesis system for self-timed circuits called SHILPA [2]. The input to SHILPA are high-level descriptions in a concurrent HDL called hopCP. hopCP is a notation to describe asynchronous hardware behavior as a collection of concurrent processes communicating through synchronous channels via *handshake* a la CSP and/or through restricted shared variables. The hopCP objects can be conveniently expressed as an annotated Petri net, where the annotations capture the data manipulation aspects of the hardware behavior.

The paper is organized as follows: In Section 2 we compare our approach with the existing work in the area. In Section 3 we introduce the high-level synthesis system SHILPA where this work fits in and provide the necessary background to understand the rest of the paper. In Section 4 we present the details of the test generation algorithms. In Section 5 we provide some

preliminary results to evaluate the performance and feasibility of our technique and finally we conclude in Section 6 by reviewing the major contributions of the work and outlining the directions of future work.

2 Comparison with Related Work

Hazewindus [7] provides a scheme to test delay-insensitive circuits in the Martin synthesis framework. They use a four-phase level-based signaling and gate-level circuits unlike our approach which uses two-phase transition signaling and macromodule based implementation. Keutzer et. al. [8] and Berel et. al. [3] explore testability of circuits generated from STGs. Keutzer et. al. require a full-scan circuit for all state elements while Berel et. al. discuss the relationship between semi-modularity and testing. Both these works address testing of the *control* part only. In contrast, our approach handles both control and data parts. Ginosar and David [5] discuss the advantage of dual-rail encoded circuits in making the circuits self-diagnostic. This involves an overhead for implementing the complete data path in dual rail logic while our approach using single-rail datapaths. Finally, in a recent paper Pagey et. al. [10] outline an approach to test Sutherland’s micropipelines by a partitioning approach similar to that proposed in our paper but they do not provide any detailed algorithms and unlike ours, it is not based on a high-level synthesis system.

3 Background

In this section, we shall provide the necessary background to develop the algorithms and main concepts of the paper. First we define transition signaling, then we provide a brief overview of the SHILPA system which includes the input language hopCP, the intermediate format HFG (hopCP Flow Graph), the compilation scheme called *action refinement* and the final circuit representation called NHFG (normal-form HFG). We present an example of synthesis in SHILPA and define predicate blocks, function blocks, and control blocks. Finally, we discuss fault-modeling in SHILPA.

3.1 Transition Signaling

In transition signaling also known as two-phase or *event-based* signaling, both the up-going (rising) and down-going (falling) transition of a signal have *meaning* in the sense that the circuits respond to both the

```

MODULE Example
TYPE byte: vector 8 of bit
SYNCPORT b?,c! : byte
VARIABLE x,y,z : byte
FUNCTION
BEHAVIOR
P [x,y] <= (x > y) → c!(x + y) → P [x,y]
      | (not(x > y)) → b?x → c!x → P [x,y]
END

```

Figure 1: hopCP description of a simple example

transitions. Natural advantages of transition signaling include higher performance and lower power consumption. Asynchronous circuits generated by SHILPA are organized around *two-phase transition signaling* with the data-bundling assumption.[11, 4]. Data-bundling assumption implies that the data arrives at all the receivers before the corresponding control signal.

3.2 Overview of SHILPA

Consider the simple hopCP specification shown in Figure 1. In hopCP, hardware is specified as a set of sequential processes communicating over a set of synchronous channels (a la CSP) or via shared variables. In the module example in Figure 1, there is a single process P, which communicates with the external environment using synchronous channels b?,c! (where ? denotes *input* channel and ! denotes an *output* channel). x and y denote the internal variable (datapath registers for example). <= denotes definition of a process, -> denotes sequencing, and | denotes a conditional. The behavior of the module is specified in the **BEHAVIOR** section and is read as follows: if the value of x is greater than y then, the value denoted by (x+y) is output on channel c!, otherwise, a new value is read in from channel b? and stored in the register x. The same value is output on channel c!. After that, the module repeats its behavior.

This behavior is conveniently represented in the form of an annotated Petri net called HFG, where the places denote the *state* of the system, the actions such as c!(x+y), b?x, denote the transitions of the Petri net. The variables and the functions which get modified by the execution of the system are captured by the annotations.

The HFG is converted into a self-timed circuit using a syntax-directed translation procedure called *action refinement* [1]. Action refinement consists of a set of Petri-net based transformations to convert the HFG




Figure 2: NHFG for the example

into an RTL description denoted by a NHFG (normal-form HFG) and a set of resources. A typical NHFG is shown in Figure 2

In our synthesis environment SHILPA, every block is modeled as an *action block* which implements a hopCP action, and presents a *self-timed interface* consisting of an *initiate* and a *completion* signals. If a transition is received on the *initiate* terminal, a transition is produced on the *completion* terminal *after* an unspecified amount of time. If the action block being considered produces *data* output, the completion signal is produced *only after* valid data is produced at its output. There are three types of action blocks in SHILPA. *Control Action Blocks* (CABs) model the control flow, *Function Action Blocks* (FABs) implement functions and *Predicate Action Blocks* (PABs) implement Boolean predicates. The models for FAB and PAB are illustrated in Figure 3 respectively.

Figure 3(a) illustrates the architecture of a FAB which is based on the following assumption. The control *req* signal will be asserted only after the arguments to the functions (denoted by *d_val*) are available (data-bundling assumption). The control signal *ack* will be asserted after the result (denoted by *out_val*) is produced. The *delay* denotes the processing delay of the block. Figure 3(b) illustrates the architecture of a PAB which implements a predicate (Boolean function) which has the following protocol.




Figure 3: Models for FAB and PAB

After the arguments (*d_val*) are available, the control signal *req* is asserted and depending on whether the predicate evaluates to true or false, the corresponding control signal (*T_ack* or *F_ack*) is asserted. As shown in the figure a PAB can be thought of as a combination of a logic block and a SELECT element. The logic block produces a transition on the *sel* input of SELECT module, which in turn produces a control signal *T_ack* or *F_ack* at the output of the PAB.

It is important to note that this style of partitioning the whole circuit into PAB, FAB or CAB offers the following advantage. The control and data only interact inside a PAB. The interaction is localized to the *sel* signal of the SELECT modules. By making this signal controllable and observable one can solve the problem of handling data dependent control flow during testing. That is exactly what is done in our approach by modifying a SELECT module to a SCANSELECT (which is described later in this section).

It is important to note that, the NHFG captures the netlist (or connectivity) of the various circuit elements *implicitly*. For example a transition in the NHFG corresponds to a transition in the circuit. This gives us the advantage that we can use the NHFG to directly model the control behavior and the netlist of the circuit

3.3 Fault-Modeling in SHILPA

In order to test the circuits generated by SHILPA completely, two types of faults have to be considered: (a) stuck at faults on the various nets and (b) delay faults which result in the violation of the data-bundling assumption. In this paper, we only consider stuck-at faults. In case of transition signaling, if a node is stuck-at 1 or 0, it *blocks* any events flowing through that *path*. It is useful to visualize a transition-signaling based self-timed circuit as a set of computing elements connected by *hollow pipes* with *marbles* (denoting events) flowing through them. A stuck-at fault blocks or obstructs the marble. This property can be used to test the circuits more efficiently when




Figure 4: Interface of SELECT and SCANSELECT modules

compared to standard level-based circuits.


The approach taken in this paper is the following. We claim that if one could *successfully* propagate a $0 \rightarrow 1$ and a $1 \rightarrow 0$ transition through a node, then the node is devoid of both stuck-at-0 and stuck-at-1 faults. A test for a node would need two transitions through a path from primary input to primary output passing through this node. Furthermore, all the nodes in that particular path are also *fault-free*. That means we need not model faults on each and every node in the circuit. This results in a considerable saving in the total number of tests required to test the circuit fully.

However, propagating a $1 \rightarrow 0$ transition or $0 \rightarrow 1$ transition through an arbitrary node, in presence of data dependent control flow, could be tricky, i.e., one may not be able to always find a suitable justification sequence. We solve this problem by modifying the design of the SELECT module in the library of asynchronous modules suggested by Sutherland [11]. This is possible in SHILPA because the only place control flow and data flow interact is within a SELECT module.

Finally, we make the following assumption in our testing strategy. The basic macromodules are considered to be *atomic* in the sense that we do not try to model faults inside the macromodules.

3.4 SCANSELECT

Figure 4(a) shows the normal SELECT module described in [11] and used in SHILPA. **sel** denotes the output of the datapath (logic) block which gets connected to the input of the SELECT module. The behavior of the SELECT module is as follows. A transition on the request input (denoted by R) after the value on the **sel** input is stable, yields a transition on



Mode signal: Used to select operation

Operation:

- 1) Load shift registers Reg n, Reg n-1
 $i_{n-1} = i_n$
 - 2) Run Test
 $i_{n-1} = \text{Sel}$
- For data path testing:
- 3) Latch Propagated Reg_sel to shift out
 $i_n = \text{Reg_sel}$
 - 4) Shift out latched values
 $i_{n-1} = i_n$

Figure 5: Internal Organization of SCANSELECT

the **T_ack** or **F_ack** output, depending on whether the value on the **sel** input is a logic "1" or "0". In general, **sel** signal is not fully observable or controllable and lies buried in the PABs. To make the **sel** signal *fully* observable and controllable we make the following modification to evolve a SCANSELECT module which is shown in Figure 4(b). The SCANSELECT module has additional control and data signals to set up scan-paths across the **sel** signals of all the SELECTs in the circuit. The **shift_req**, **shift_ack**, and **shift_data** pins are used to string the SCANSELECT modules in a circuit into a FIFO. This helps us to observe a given signal or set any desired **sel** signal. The SCANSELECT module has *four* modes of operation which are controlled by the **mode** signal. The different modes of operation are illustrated in the text accompanying Figure 5. A SCANSELECT consists of a regular SELECT module, a latch (REG in the figure), and some control logic (shown as CL Block in the figure).

The operation of a SCANSELECT module is typically as follows: First the values that are required during testing in registers $\text{Reg}_1 \dots \text{Reg}_n$, are shifted in by connecting terminals i_{n-1} and i_n . This is achieved by setting the **mode** signal appropriately. The values are shifted into the FIFO serially, using a sequence of request acknowledge pairs. After the registers are loaded, the **mode** signal is changed such that the output of Reg_{n-1} is connected to the **sel** signal of the

internal SELECT module. This renders the **sel** signal fully controllable and effectively decouples the datapath from the control part. This is sufficient for testing the control part (CAB). While testing the datapath of a circuit one finds that the **Reg_sel** (an output of the data path) value needs to be observed. For this, the **mode** signal is set such that **Reg_sel** is connected to terminal i_n which allows the latching of the value on **Reg_sel** in Reg_n . This can then be shifted out by setting the **mode** signal such that i_{n-1} to i_n are connected.

The SCANSELECT module as shown above needs seven additional pins. But, one can see that the various modes of operation are mutually exclusive. So, in practice, the pins can be multiplexed (at the expense of additional control logic) and one needs only four additional pins. Since, the number of SELECT modules (which are a consequence of conditionals in the high-level description) is typically not very large in any given circuit, the overhead is reasonable. As noted in the previous subsection, we will assume that a SCANSELECT module is *atomic* i.e., we will not attempt to test for faults inside this module.

4 Algorithm for Test Generation

In this section, we present the details of the algorithms to test for stuck-at faults in a self-timed circuit based on its NHFG representation. First, the circuit is partitioned into control part consisting of only the CABs and the data part consisting of FABs and PABs. Then the nodes in each part are modeled separately in the NHFG and the test vectors are generated by traversing the corresponding paths in the NHFG. Initially, we present the top level algorithm which invokes procedures **testCab** and **testDatapath** to test the control and data parts separately. The pseudo-code of the algorithms is presented.

4.1 Top-Level Algorithm

```

GLOBAL:
    tested-hfg-nodes, choices-made
INPUT:
    NHFG, netlist, reslist
OUTPUT:
    P /* set of physical test vectors for datapath */
    T /* set of traces */
    C /* choices-made for the scan-path */
METHOD:

```

```

    ( $N_c, R_c$ ) = extract-control(NHFG, netlist, reslist);
    ( $N_d, R_d$ ) = extract-datapath(NHFG, netlist, reslist);
     $T_1$  = testCab(NHFG,  $N_c, R_c$ )
    ( $P, T_2$ ) = testDatapath(NHFG,  $N_d, R_d$ )
    return ( $T_1, (P, T_2), C$ )

```

The input for the top level algorithm is the output of the SHILPA system, namely, the NHFG, the set of resources (hardware modules), and the physical netlist. **tested-hfg-nodes** and **choices-made** are two global variables used by all the routines. The output of the algorithm are (i) physical test-vectors for all the stuck-at faults in the datapath (denoted by P), (ii) sequence of traces or control sequences (denoted by T) to test the control part of the circuit and to set up conditions to test the datapath with P , and (iii) the choices made at points of data-dependent control (denoted by C) to be used in setting up the scan chain appropriately during testing.

The first step of the algorithm is to partition the circuit into control part and data part (accomplished by appealing to functions **extract-control** and **extract-datapath**). The control parts consists of all the CABs and forms a meaningful circuit if the **request** and **acknowledge** pins are connected together. The data part consists of pure logic blocks (PABs and FABs) with hooks to the control part for the corresponding **request** and **acknowledge** signals.

4.2 Testing Control Circuitry

The procedure for testing the control part of the circuitry, called **testCab** is shown below.

```

procedure testCab (nhfg, netlist, reslist)
    traces = nil
begin
    For each node in netlist
        hnode = Extract-Node (nhfg, node)
        if hnode ∈ tested-hfg-nodes skip
        else
            new-trace = Derive-Trace(nhfg, netlist, hnode)
            tested-hfg-nodes = insert(tested-hfg-nodes,
            new-trace)
            traces = cons(new-trace, traces)
        end
    return(traces)

```

The function **extract-node** finds the transition in the NHFG corresponding to a given name from the netlist. If the node has already been tested (by virtue of lying in the path of a different tested node) then we do nothing, otherwise, the function **Derive-Trace** is

invoked to generate a test sequence for the node. The rest of the functions are for book-keeping.

```

procedure Derive-Trace (nhfg,netlist,hnode)
  initial-node = get-initial-node(nhfg)
  trace = Visit(initial-node,hnode,nhfg)
  return(trace)

procedure Visit(init-node,hnode,nhfg)
begin
  path1 = find-shortest-path(init-node,hnode)
  path2 = find-shortest-path(hnode, hnode)
  /* here we are computing the sequence of
  transitions from hnode to the next occurrence of hn-
  ode */
  if (path2 = null) then
    begin
      path3 = find-shortest-path-to-primary-
      output(hnode)
      test-path = concatenate(path1,path3)
      result = concatenate(test-path,test-path)
    end
  else
    begin
      path3 = find-shortest-path-to-primary-
      output(hnode)
      result = concatenate(path1,path2,path3)
      choice = getChoicenodes (result)
      /* extracts all choice nodes in the result */
      choices-made = cons(choices-made,choice)
      /* choices-made records all data dependent choices
      for later use to set up the scan paths */
    end
    /* list-of-preconditions is generated during the
    find-shortest-path
    routines when transitions with more than one input
    PLACE are encountered. */
    while (list-of-preconditions-to-be-satisfied !=
    nil)
      begin
        new-trace = Visit(init-node, head(list-
        of-preconditions-to-be-satisfied),nhfg)
        result = append(result, new-trace)
      end
    return(result)
  end

```

The algorithm **testCab** works as follows: First a physical node is selected from the netlist and the corresponding transition is marked on the NHFG, if the node has not been tested previously as a consequence of some other test. Let n be the node being tested and t be the corresponding transition in the NHFG. Then a path (sequence of “enabled” transitions in the

NHFG) is selected from the initial state (state of the system on global reset) to the state in which t is *enabled* (in the Petri net sense). Call this $path_1$. Another path called $path_2$ is selected such that we begin in a state in which t is *enabled* and then t is revisited. These are analogous to the justification sequences in synchronous testing. $path_3$ is yet another path which corresponds to the sequence of transitions from t to a primary output. The significance of $path_2$ is that it takes care of the requirement that the node n being tested be subjected to an $1 \rightarrow 0$ transition and a $0 \rightarrow 1$ transition. If $path_2$ is empty, then $path_1 :: path_3$ is traversed *twice* to get the same effect (where $::$ denotes concatenation of sequences).

The final test sequence then is $path_1 :: path_2 :: path_3$. In generating the paths (or sequences) one should remember that NHFGs are not simple graphs but Petri nets, which means that some transitions could have more than one input place (for example, nodes involving C elements). This results in the additional requirement to find all possible paths to enable the given transition. Note that this is accomplished by recursively invoking procedure **visit**.

Finally, in the computation of the paths, one can encounter *choice* nodes (which are implemented as part of a PAB in SHILPA). The nodes are recorded in the global variable to derive the possible scan vectors.

4.3 Testing Data Part

The data path can be viewed to consist of PABs and FABs with control signals acting as the interface to the CABs (in the control part). The only point at which, the data has effect on control is at the SELECT element inside PABs. The top level algorithm for testing the data part is given below.

```

procedure testDatapath(nhfg,netlist,reslist)
begin
  datapath = GetDatapathFrom-
  Netlist(netlist,reslist)
  for each FAB/PAB F in datapath
    begin
      (I, O) = AdaptedDalgorithm(F,Fault)
      J = set of FAB/PAB affected in justifying
      I
      P = set of FAB/PAB affected in
      propagating O
      control-points = FindAffectedControl-
      Points(F)
      markedHfg = MarkHfg(nhfg, control-point)
      Primary-input = justify-data(J,I,markedHfg)
    end
  end

```

```

Primary-output = propagate-
data(P,O,markedHfg)
end
listOfTraces = findCover(markedHfg)
/* Find a path(s) covering all the marked nodes
in the nhfg starting from the initial node */
choices-made = getChoicenodes (listOf-
Traces)
/* extracts the choice nodes from listOfTraces */
test = (list Fault, ListofTraces, choices-
made,Primary-input, Primary-output)
test-list = (cons test test-list)
end
return test-list;
end

```

We *unlike* [10] modify the standard combinational test generation procedure to *process* only those FABs and PABs that are *active* for a given stimulus. This is done by the procedure *AdaptedDalgorithm*. This takes one PAB/FAB at a time in order to keep track of the sequence of control signals that need to be applied.

The output of the algorithm is a list of *Fault*, *Choices-made*, *ListofTraces*, *Primary-input*, *Primary-output*. Given these outputs the testing is accomplished in the following manner. For each *Fault*, we first set up the scan chain as identified by *Choices-made* thus decoupling the data and the control. We then apply the input vector on the input data pins *Primary-input*. The control sequences in the *ListofTraces* are applied. This moves the data from the primary inputs, through the fault to the primary outputs. The Primary-outputs are then observed. In addition to the primary outputs, the outputs of the scan chain are observed to look for particular *sel* signals of SELECT modules.

The algorithm works as follows. We begin by first identifying the fault in a given FAB/PAB, and then applying the D algorithm on the datapath of the FAB. The result represented, by (I, O) denotes the input needed and the outputs produced respectively to sensitize and test the fault. The inputs I are either primary inputs or the outputs of some other FABs (set J). In which case, the wires values in I need to be justified through this set of FABs. This needs to be recursively performed until primary inputs are encountered and the values at these primary inputs are noted. This recursive part is implemented by the function *justify-data* (given later in the paper). It recursively justifies backwards through the FAB's generating new sets of *affected* FABs (J') and justifying the outputs at these FAB's (I') to their respective inputs. Propagation is similarly done by the function *propagate-data*

till primary outputs are encountered. It is important to note here that a *sel* input to a SELECT element encountered during propagation is considered a *primary output* due to the existence of the scan chain. During both justification and propagation, we have to note the control points that need to be activated to move the data from input to the output in each FAB. These control points are marked on the NHFG. The markings on the NHFG indicate a sequence of control actions that need to be performed in order to move the data from the primary inputs, through the fault to the primary outputs. At this stage, we use graph based algorithms to obtain traces that cover all the marked points on the NHFG. While obtaining the paths, we also find that we traverse choice nodes, that need to be decided independent of the data values (since we are testing the data path).

```

procedure justify-data(J,I,markedHfg)
for each FAB/PAB f in J
begin
I' = AdaptedDalgorithm(f,I)
control-points = FindAffectedControlPoints(f)
markedHfg = markHfg(markedHfg,control-points)
J' = set of FAB/PAB affected in justifying I'
if null J' return I'
else justify-data(J',I',markedHfg)
end
procedure propagate-data(P,O,markedHfg)
for each FAB/PAB f in P
begin
O' = AdaptedDalgorithm(f,O)
control-points = FindAffectedControlPoints(f)
markedHfg = markHfg(markedHfg,control-points)
P' = set of FAB/PAB affected in propagating
O'
if null P' return
else propagate-data(P',O',markedHfg)
end

```

5 Experimental Results

In this section we will report the results of our algorithms on a variety of examples. Since, we do not have a reliable set of benchmarks in the area of high-level synthesis of asynchronous circuits, we restrict ourselves to examples from relevant literature. Although we do find that for each SELECT module in the circuit, we need two additional multiplexers, a C element and a transition latch in order to use our technique, the number of SELECT modules in a circuit are

Table 1: Table of experimental results

Circuit Name	Num gates	Control testpaths	data path test vectors/paths	Percent testable
GTAD	233	3	338/3	100
PMLT	377	4	602/6	100
QR42	64	2	0/0	100
FACT	379	4	626/6	100
SALU	348	6	640/6	100

very small compared to the size of the circuit. Therefore the overall circuit overhead as a percentage of the size of the whole circuit is very small. The pin overhead if we were to assign a different pin for each extra signal to the chip would be 7 pins. Since some of the signals will not be used at the same time, we feel that the number of pins can be reduced to about 4 pins. The size of each circuit in terms of the equivalent two input AND/OR gates is shown in Table 1. In the same table, we find that the control sequences to test the control part is very small for even large circuits. This is an important saving in our method and we believe it is due to our fault-modeling strategy (described in section 3).

6 Conclusions

There are two major contributions to this work. The first is a set of algorithms to test self-timed circuits based on two-phase transition signaling and macro-modules and the second is the idea of a scannable SELECT module which makes the circuits *fully* controllable and observable. Our algorithms are based on a Petri-net like model of the physical circuit and are general in the sense that they are compatible with other asynchronous behavioral synthesis approaches like Brunvand [4], Ebergen [6] and any circuit which can be modeled by a NHFG.

References

- [1] AKELLA, V. *An Integrated Framework for the Automatic Synthesis of Efficient Self-timed Circuits from Behavioral Specifications*. PhD thesis, Department of Computer Science, University of Utah, 1992.
- [2] AKELLA, V., AND GOPALAKRISHNAN, G. SHILPA: A High-Level Synthesis System for Self-Timed Circuits. In *International Conference on Computer-aided Design, ICCAD 92* (Nov. 1992), pp. 587–591.
- [3] BEEREL, P. A., AND MENG, T. H.-Y. Semi-modularity and self-diagnostic asynchronous control circuits. In *Advanced Research in VLSI, Proceedings of the 1991 University of California/Santa Cruz Conference* (1991), The MIT Press.
- [4] BRUNVAND, E. *Translating Concurrent Communicating Programs into Asynchronous Circuits*. PhD thesis, Carnegie Mellon University, Nov. 1991.
- [5] DAVID, I., GINOSAR, R., AND YOELI, M. Self-timed is self-diagnostic, 1990.
- [6] EBERGEN, J. C. *Translating Programs into Delay Insensitive Circuits*. Center for Mathematics and Computer Science, Amsterdam, 1989.
- [7] HAZEWDINDUS, P. J. *Testing Delay-Insensitive Circuits*. PhD thesis, California Institute of Technology, 1992.
- [8] KEUTZER, K., LAVAGNO, L., AND SANGIOVANNI-VINCENTELLI, A. Synthesis for testability techniques for asynchronous circuits. In *International Conference on Computer-Aided Design* (Nov. 1991), pp. 326–329.
- [9] MEAD, C. A., AND CONWAY, L. *An Introduction to VLSI Systems*. Addison Wesley, 1980. Chapter 7, entitled “System Timing”.
- [10] PAGEY, S., VENKATESH, G., AND SHERLEKAR, G. Issues in fault modeling and testing of micropipelines. In *First Asian Test Symposium, Hiroshima, Japan* (Nov. 1992), pp. 107–111.
- [11] SUTHERLAND, I. E. Micropipelines. *Communications of the ACM* (June 1989). *The 1988 ACM Turing Award Lecture*.
- [12] WILLIAMS, T. E., AND HOROWITZ, M. A zero-overhead self-timed 160ns 54bit cmos divider. *IEEE Journal of Solid State Circuits* 26, 11 (Nov. 1991), 1651–1661.