IMPROVED SIGNAL INTEGRITY IN EMBEDDED IEEE 1149.1 BOUNDARY-SCAN DESIGNS

by

Efren J. Taboada

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

School of Technology

Brigham Young University

December 2004

Copyright © 2004 Efren J. Taboada

All Rights Reserved

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Efren J. Taboada

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory

______________________________	______________________________
Date	C. Richard Helps, Chair
______________________________	______________________________
Date	Barry M. Lunt
______________________________	______________________________
Date	Michael G. Bailey

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Efren J. Taboada in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

Date

C. Richard Helps
Chair, Graduate Committee

Accepted for the Department

Thomas L. Erekson
Director, School of Technology

Accepted for the College

Douglas M. Chabries
Dean, Ira A. Fulton - College of Engineering and Technology

ABSTRACT

IMPROVED SIGNAL INTEGRITY IN EMBEDDED

IEEE 1149.1 BOUNDARY-SCAN

DESIGNS

Efren J. Taboada

School of Technology

Master of Science

This work is an analysis of solutions to problems derived from inherent timing and signal integrity issues in the use and application of the IEEE 1149.1 Standard at the board level in conjunction with its test system. Setup or hold times violations may occur in a boundary scan chain using IEEE 1149.1 compliant devices. A practical study of the TDI-TDO scan data path has been conducted to show where problems may arise in relationship to a particular board topology and test system. This work points to differences between passing and failing scan path tests for problem characterization. Serial data flow is then analyzed and suitability is discussed. Within certain conditions, a solution is proposed. This work has been shown to work on the test system. Recommendations are made based on this experimental approach.

ACKNOWLEDGEMENTS

This work is certainly not complete without this page. I have now finished what I started in Fall of 1996. There is absolutely no question that without my wife's support throughout all these years I would not have finished. She showed, and continues to show, an unsurpassed willingness to sacrifice what other moms wouldn't and at the same time to hold fast to our marriage and keep it alive. I pay tribute and will devote my eternities to her. Ivelisse, I dedicate this work to you.

To my advisors Richard Helps and Barry Lunt without whom this work would not have been possible as well. For their encouragement and willingness to remotely support me, and for all these years of learning and strict teaching that have now enabled me to learn on my own on almost any subject for any application in my field of studies.

To Robert Mejia, former member of the Fibre Channel Standard committee, for his outstanding insight on how to capture signal integrity results in Chapter 4. Robert currently works in the Multi-Function Printer Lab at Hewlett-Packard as Hardware Designer and signal integrity expert resource.

To the Hewlett-Packard Co. since this project required an extensive investment of time and resources converging to support the thesis statement. The software to drive the scan chain, the oscilloscope with special probes and the logic analyzer were some of the tools used and were all provided by the company.

TABLE OF CONTENTS

TABLE OF CONTENTS viii
LIST OF TABLES x
LIST OF FIGURES xii
CHAPTER 1 1
Introduction 1
Acronyms 2
Problem Context 2
Thesis Statement 7
Limitations 8
CHAPTER 2 11
Research Field Overview 11
CHAPTER 3 19
The Unit Under Test 19
Designing the Buffer Solution 23
Test System Block Diagram 27
The Glue Logic 30
TAP Connecting Board 30
TAP Sequencing Board 32
Thesis Support Statement 34
TAP Controller State Diagram 36
Test Setup Methodology 39
Test System Frequency Settings 41
Timing Measurements Methodology 42
Signal Integrity Measurements Methodology 43
CHAPTER 4 47
Timing Measurements Analysis 47

Signal Integrity Measurements Analysis 51
TCK Signal Integrity Analysis 51
Failure Mode Description (Table 3) 52
Testing with the SDTAP Buffer (Table 4) 55
Testing without the SDTAP Buffer (Table 5) 58
Testing the SDTAP Buffer without ASIC 2 (Table 6) 62
Testing another UUT (Table7) 64
Further Validation of the SDTAP Buffer (Table 8) 67
Other Signal Integrity Measurements 70
TCK Increased Performance 72
Analysis of Results 76
CHAPTER 5 81
Summary and Conclusions 81
Future Research 83
BIBLIOGRAPHY 85

LIST OF TABLES

Table 1 - Boundary scan terminology 2
Table 2 - TCK frequency settings available 41
Table 3 - UUT 1 - Chain configuration 1 without solution 44
Table 4 - UUT 1 - Chain configuration 1 with solution 44
Table 5 - UUT 1 - Chain configuration 2 without solution 44
Table 6 - UUT 1 - Chain configuration 2 with solution 44
Table 7 - UUT 2 - Chain configuration 1 without solution 44
Table 8 - UUT 2 - Chain configuration 1 with solution 44
Table 9 - Improved system hold time margin 49
Table 10 - Setup 2 - 22pF only on TCK 73

LIST OF FIGURES

Figure 1 - Simple boundary scan chain and scan path verification 3
Figure 2 - A PC-based boundary scan test system.. 11
Figure 3 - TAP sequencing & sampling system 15
Figure 4 – UUT 1 – Convoluted TCK path 20
Figure 5 – UUT 1 - Illustration showing trace density 21
Figure 6 - UUT 2 - TCK path & multiplexed TMS signal on second chain.................... 22
Figure 7 - UUT 2 - Trace density... 22
Figure 8 – SDTAP Buffer circuit schematics 23
Figure 9 - The actual SDTAP Buffer - Top assembly 24
Figure 10 - A generic boundary scan test system and SDTAP Buffer operation 26
Figure 11 - A particular boundary scan test system setup 28
Figure 12 - TAP interface at the UUT 29
Figure 13 - TAP concatenation system.. 31
Figure 14 - TAP sequencing system 34
Figure 15 - The TAP controller 37
Figure 16 – ScanWorks ® fail screen (courtesy of ASSET Inter-Tech, Inc.) 40
Figure 17 - ScanWorks ® pass screen (courtesy of ASSET Inter-Tech, Inc.) 40
Figure 18 – Test setup 1... 45
Figure 19 - Test setup 2 46
Figure 20 – Test setup 2: working system hold time margin = 56ns (TCK = 8 MHz).... 48
Figure 21 - Setup 2 - TCK = 1 MHz... 50
Figure 22 - Setup 2 - TCK = 4.16 MHz.. 50
Figure 23 - ASIC 1 - no filtering 53
Figure 24 - ASIC 1 falling edge... 53
Figure 25 - ASIC 2 falling edge... 54
Figure 26 - ASIC 2 rising edge 54

Figure 27 - ASIC 3 falling edge.. 55
Figure 28 - ASIC 1 with solution.. 56
Figure 29 - ASIC 1 falling edge.. 56
Figure 30 - ASIC 2 falling edge.. 57
Figure 31 - ASIC 2 rising edge.. 57
Figure 32 - ASIC 3 falling edge.. 58
Figure 33 - ASIC 1 no filtering... 59
Figure 34 - ASIC 1 falling edge.. 60
Figure 35 - ASIC 2 - falling edge 60
Figure 36 - ASIC 2 rising edge.. 61
Figure 37 - ASIC 3 falling edge.. 61
Figure 38 - ASIC 1 with filtering.. 62
Figure 39 - ASIC 2 (not present) falling edge - with filtering.. 63
Figure 40 - ASIC 2 (not present) rising edge - with filtering 63
Figure 41 - ASIC 3 falling edge with filtering.. 64
Figure 42 - ASIC 1 no filtering... 65
Figure 43 - ASIC 1 rising edge - no filtering... 66
Figure 44 - ASIC 3 - no filtering - both edges.. 66
Figure 45 - ASIC 4 - no filtering 67
Figure 46 - ASIC 1 TCK with filtering... 68
Figure 47 - ASIC 1 rising edge.. 68
Figure 48 - ASIC 3 both edges 69
Figure 49 - ASIC 4 both edges 69
Figure 50 - ASIC 1 - Test Setup 1 - TCK falling edge with no termination 71
Figure 51 - ASIC 1 - Test Setup 1 - TCK falling edges 72
Figure 52 - ASIC 1 – Test setup 2 - Degraded TCK partially works 73
Figure 53 – ScanWorks® scan path verification tests. (Courtesy of ASSET Inter-Tech, Inc.).. 75
Figure 54 – System model for TCK distribution + Test Setup 1 (passing model) 79
Figure 55 – System model for TCK distribution + Test Setup 1 (non-passing model).... 80

CHAPTER 1

Introduction

The IEEE 1149.1 Boundary Scan Standard was first introduced in 1990. The current revision is dated 2001. It addresses the need for node access in dense, small complex boards where traditional "bed of nails" test probes can no longer provide physical test access. It was conceived by the Joint Test Action Group (JTAG), which was made up of companies primarily in Europe and North America. For more than a decade (1990-2004), ASIC implementation of the IEEE 1149.1 Standard has grown enough to lay out a solid foundation that now supports its widely accepted use in the board test industry (Dellecker). This standard is now well rooted in this industry and has served as the basis for other standards such the IEEE 1149.4 that enhances testability of analog signals among others.

The 1149.1 technology embraces very complex digital testing problems at the board level and turns them into easy and manageable solutions that would otherwise require very expensive methods to test the structural integrity of a printed circuit assembly (Kajitani et al.). Some of these methods are three dimensional x-rays and flying probe, which require investments in excess of $500,000.00 and are very slow.

Acronyms

It is important to familiarize the reader with boundary scan terminology before proceeding with any discussion topics. It is important to understand these terms to be able to make inroads into this discussion. These terms are summarized in Table 1.

Table 1 - Boundary scan terminology

ACRONYM	READ OUT	DESCRIPTION
TDO	Test Data Out	Bit-stream coming out of each device in the chain and eventually out of the last device to the test system
TDI	Test Data In	Bit-stream coming into the chain being driven by the test system
TCK	Test ClocK	Clock signal driven by test system
TMS	Test Mode Select	Control signal to the boundary scan controller (Moore machine) to change state
nTRST	Test ReSeT	Asynchronous reset signal driven by test system
TAP	Test Access Port	The grouping of all five boundary scan signals described above
UUT	Unit Under Test	The printed circuit board under scrutiny
PCA	Printed Circuit Assembly	Printed circuit board with components loaded

Problem Context

This study examines a real-life problem that was found to be double clocking caused by reflections. This problem can occur in test systems with many circuit discontinuities along the TDI-TDO path. The delays and distortions associated with control and data signals in a small dense board are important considerations to improve the hold time margin of the TDO signal with respect to the sampling point. Particularly, distortions caused by impedance mismatches can cause signal reflections that affect the scan chain performance in terms of speed and reliability. For reference, a simple scan chain setup is shown in Figure 1.

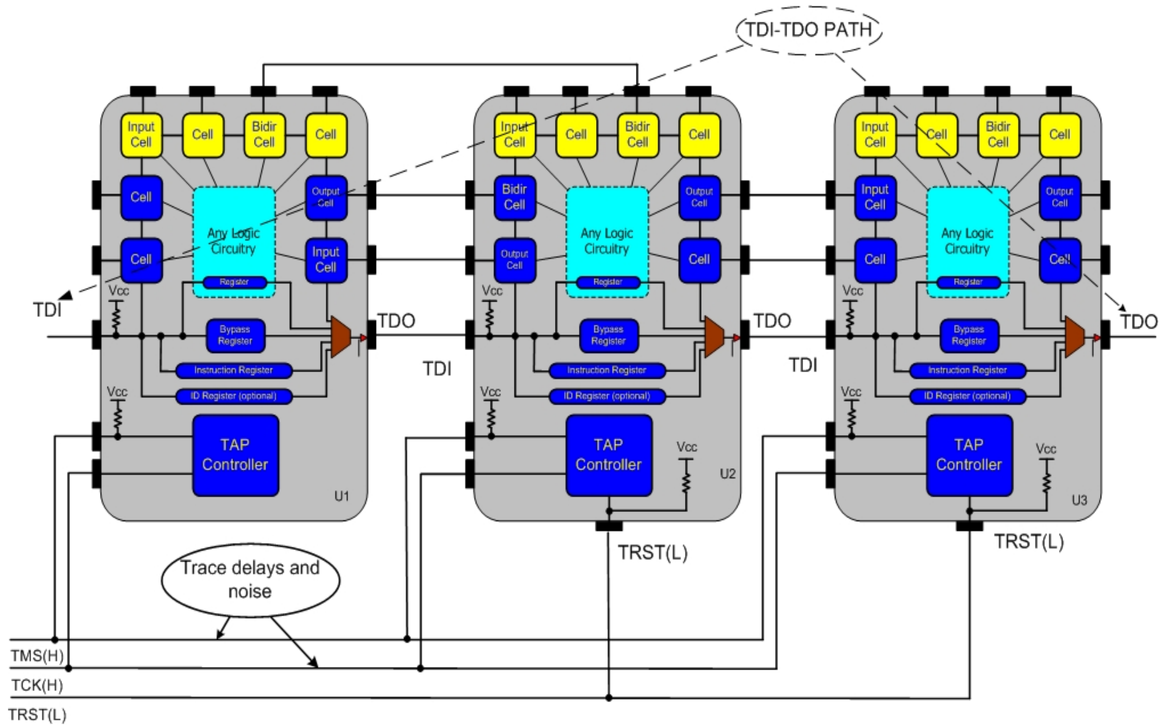

Figure 1 - Simple boundary scan chain and scan path verification

Per standard, the TDI-TDO path samples its inputs at the rising edge of TCK and generates valid TDO data at the falling edge of TCK. This means that data going into any TDI pin should be valid sometime before it is sampled, and data going out of each TDO output pin along the chain is valid sometime after the falling edge of TCK. However, this timing is not specified in the Standard and it is left to the ASIC designer discretion. Of particular interest to this work is the "last" TDO output pin timing relative to TCK coming from the test system connector. The solution presented here slightly increases this delay until there is enough hold time for the signal to travel across the test system, scan chain and back. Then, this signal is received at the sampling point and read back into the PC-based test system for processing. The electrical signal referred to here is the "last" or previous data bit coming out of the chain. In other words, at the sampling moment, the D-F/F (see Figure 3 – Chapter 2) samples the data bit that was driven out by the previous

TCK falling edge. This principle of operation is the basis for our discussion. Based on this premise, the necessary setup and hold times are described within the Glue Logic section of Chapter 3. This description focuses at the system interactions between the PC-based tester and its UUT taking into consideration timing and electrical characteristics for both systems to operate correctly.

In any JTAG compliant board it can be easily inferred, but is often overlooked, that the test data-bit stream has two constraints: path length and data length. A round-trip byte is somewhat affected by the path chosen and the test vector length. It is variable depending on the internal path selected by each individual TAP controller as shown in Figure 1. To identify the different paths, notice that there are several inputs to the TDO multiplexer, which in turns drives the TDO signal out from each path. The length of this test-data bit stream also varies depending on the instruction that is loaded into the Instruction Register and the over shift that sometimes is used to initialize or “flush” the cells and registers in the TAP Controller. Moreover, this is applicable for each TAP controller since they can be operating in different modes at the same time.

Normally a test engineer will find that boundary scan trace layout implementations are unique for each board topology. While the software may be able to adjust to these changes, the test hardware used (which is normally fixed) does not tolerate these changes in all cases. Any changes to the scan path from the sampling point down to the UUT affect the system hold time margin. It is important to understand the effects of buffers and other circuits in this path. This is to make sure that any solution proposed does not negatively impact system operation.

All these combined events may affect hold times and input trigger points designed into an ASIC compliant to the Standard. In turn, this may prevent data capturing at the appropriate time causing erratic operation of the chain as “seen” by the boundary scan master controller. Thus the tester detects a failure in a good system. This failure mode appears as an intermittent failure that is very difficult to correct and diagnose. Improving this system level margin inherent to the 1149.1 technology and its use is the heart of this work.

The TDO bit-stream needs to be adjusted to make sure data at the “last” TDO output can be always stable at a time where the master boundary scan controller is ready to sample. Also, signal reflections need to be filtered. Normally, if data is available after a TCK falling edge, the next stable chance to sample is some time before the upcoming rising edge. Therefore, the author’s goal is to provide a solution that is mostly independent of the scan data path, trace-layout, test system, and software used. This should effectively both shift the TDO data stream with respect to the TCK sampling rate at the PC-based tester connector to “hold” it longer, and filter reflections in the four control and data lines relatively independently of their trace layout or discontinuities.

Although ideal, this sampling point is not necessarily half way through ½TCK cycle. In Figure 3 (Chapter 2), the “previous” bit of data is sampled on the falling edge of TCK just before the “next” one arrives. The solution proposed in this work holds the TDO signal a little bit longer, thereby improving the system hold time margin. This margin is dictated in part by the Flip-Flop in Figure 3. The longer it takes for a TCK falling edge to reach its targets, the better the hold time margin of the test system.

In any boundary scan test system data is supposed to be clocked in at the TCK rate when in reality it is being clocked at (1/TCK) + t, where t is the inherent t_p delay of the circuits carrying the TCK signal to its destination. In turn, this delay helps compensate for the inherent functional variable delays of the TDI-TDO path by ensuring hold times are met with better margin and the TDO bit-stream is stable at the interface between the two systems. Normally, setup times are seldom violated. This is because the dictated Rule 6.2.1.b in the Standard, where data is clocked out at the falling edge of TCK, which renders the setup times for the next flip-flop always significantly less than ½ TCK cycle. From a practical standpoint, the scan path data stream is slowed down to increase the hold time required by the master boundary scan controller to sample valid data. This data comes from the last TDO pin of the UUT. In other words, this is accomplished by shifting the TDO delay of the last chip in the chain relative to the falling edge of TCK at the tester connector to a point in between ½ TCK cycle (assuming 50% duty cycle for TCK). Part of the thesis objective is to improve the system hold time margin, which in turn, improves overall system reliability.

Reasonably, a test engineer will experience difficulties making different board designs work under a wide range of conditions that apply to the test system and the target UUT. In other words, while the PC-based tester remains the same in any boundary-scan application, the UUT topology changes every time.

All these assumptions and conditions make the proposed solution very practical, simple, and powerful to accommodate a wide range of problems for the use and application of embedded boundary scan in almost any digital design.

The test scenario at the board level is similar to those found at the ASIC level where hold time violations may happen when shifting data between scan cells. Very often this issue occurs because of the serial nature of the IEEE 1149.1 test technology. In the ASIC case, appropriate hold times are needed to ensure data is held long enough after the scan cell flip-flops are clocked. In the UUT case, a similar idea applies to all devices in the chain and their relationship to its associated master boundary scan controller. Chapter 2 will explain this concept in more detail.

Thesis Statement

This work will analyze the TDI-TDO path where delays are normally fixed depending on the test hardware of choice, signal traces, and cables. This approach should help focus on improvements to the system hold time margin and to the signal integrity at the board level when applying IEEE 1149.1 technology to testing electronic circuit boards.

To improve the system hold time, delays will be deliberately introduced to hold the TDO signal longer relative to a TCK falling edge. To improve signal integrity a filter will be added to all TAP signals going into the UUT. This delay will be removed from the TDO signal coming back out of the last device in the chain. Notice that this approach proposes filtering of all signals rather than the traditional “TCK only” filtering with a 15pF to 100pF capacitor to ensure balanced operation of the chain. Additionally, the removal of filtering and delaying in the TDO signal serves the purpose of holding TDO longer.

The general problem refers to the inability to run a boundary scan compliant chain consistently when a series of basic scan path tests are applied in a loop fashion. Under

these conditions, signal integrity and hold time violations are most likely the root causes of this type of problem (Stang and Dandapani).

A generic solution must then be successful in correcting faults in bad scan circuit implementations as well as improving signal integrity in good passing circuits. The solution will then take the form of hardware that sits in between a boundary-scan master controller and a UUT. This buffer solution will act like a soft-cushion between the two systems.

Per implementation of the Standard, all signals sourced by the master boundary scan controller going in the outbound-to-UUT direction incur delays, but the TDO signal incurs additional delays in the inbound-from-UUT direction. Based on this idea, this work focuses on a solution that allows controlled delays for any of these two directions. Using this methodology, the goal is that setup and hold timing requirements are always satisfied and signal integrity is guaranteed to be good. These assumptions hold true only within the maximum speed of the slowest device in the chain.

Limitations

As stated before, this thesis is limited to a particular type of hardware. This hardware is a PCA that uses three ASICs compliant to the IEEE 1149.1 Standard and it will be further detailed in Chapter 3. This work does not attempt to prove that the solution is completely universal, but it rather justifies the use of this new technology[1] applied to a variety of boundary-scan test systems driving a target scan chain. In other words, the solution may be particular, but its application is intended to be universal.

---

[1] The new SDTAP Buffer

The framework for this solution is further limited to compliant IEEE 1149.1 systems and it assumes that the ASICs used are compliant to the Standard.

The test setup was designed to accept a maximum TCK frequency of 8.335 MHz. This speed is considered acceptable to satisfy technology requirements and usage. It means that to test PCA structural integrity there is no need to use high speed TCK signals or to design the TAP interface to operate beyond this limit. The test system used is limited to a maximum of 16 MHz.

As a relative measure, the solution was tested on a maximum board size of about 48 square inches. Assuming the shortest path is the goal when laying out TCK and TMS traces, the length of these two traces may be as long as ~7-8 in. Also, the maximum number of nets in a board available for testing was 2200. This should serve as a general measure of the UUT's size limitation used in this research.

CHAPTER 2

Research Field Overview

A thorough understanding of the IEEE 1149.1 Standard is the basis for this work. The functionality of systems A & B from Figure 2 will be described to help illustrate the solution proposed. It is not the objective of this research to describe or analyze in detail the inner workings of the entire system in this figure. However, the interactions between system A and system B are of interest.

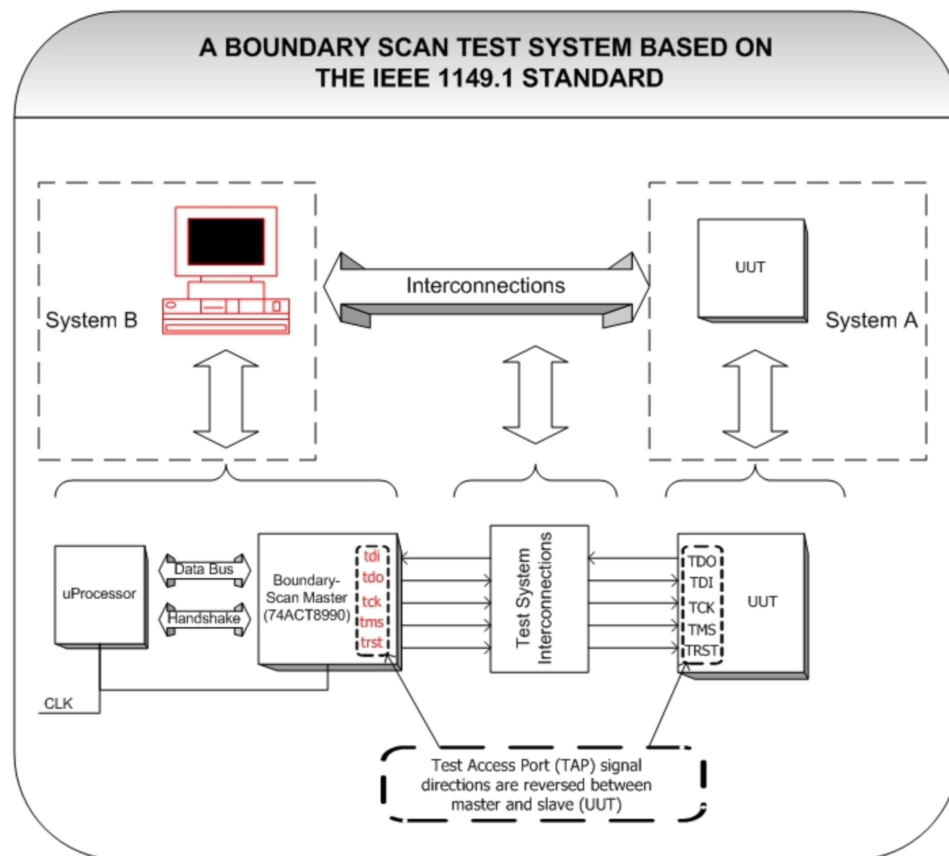

Figure 2 - A PC-based boundary scan test system

An important constituency for this work is the Test Technology Technical Council. This driving force is a team of volunteering professionals from the electronic test industry and it is sponsored by the IEEE Computer Society. The purpose of the TTTC is to share and facilitate the advancement of test technologies

It is widely known in the board test industry that the economics of the 1149.1 technology are very favorable. Caldwell and Langford relates to an early case study where a 30% board yield problem was detected with boundary scan. The root cause of the problem was internal package moisture that would destroy interconnects inside the ASIC during reflow. This was a board with a single chip in the chain; Caldwell and Langford further suggests that the benefits could be even more generous as complexity and component density of the board increases (109).

This work originated with a real world problem. Since the actual symptoms originally appeared to be signal integrity related issues, some basic board level DFT rules for boundary scan testing were followed without good results (Parker). Attention then shifted to the entire test system to look for possible root causes. As stated before, the problem was simply the inability to verify the TDI-TDO scan path functional integrity when looping a series of software tests available to "test" the tester[2]. This inability is a major road block for this technology since it prevents further testing of any Printed Circuit Assembly (PCA) with embedded boundary scan.

The main focus of this work and its solution comes from the analysis of the interactions between systems A and B in Figure 2. These interactions become critical at the system level when you introduce too much delay in the interconnect path between the

[2] The tester in this case refers to the embedded boundary scan logic architecture that ultimately enables structural testing of a complex microprocessor based system.

two systems in relationship to the sampling point. These delays are not only due to long cables, but by extra circuits referred to in this work as the “glue logic.” Given the nature of boundary scan, the commonly recommended approach to solve this problem is to simply lower the TCK frequency to achieve synchronization and allow enough hold time for the TDO signal to travel back to the tester’s sampling flop. Other solutions to this problem include re-timing circuits (TI), where extra bits then have to be accounted for in software to achieve synchronization. The test setup used in this work has a re-timing circuit, but it is not the main interest since different sampling techniques can be used. It is the author’s purpose to keep the proposed solution independent and not detrimental to any particular test system. To make sure it is independent and not detrimental the design needs to consider the TCK frequency and how long it takes a TDO edge to reach the sampling point after a TCK falling edge. This should help achieve a general approach to solve the problem at hand.

Based on these premises, the following questions went unanswered:

1) How does the glue logic, with its delays in place, affect a manufacturing test setup?
2) What could be done if you lower the frequency (130 KHz in this case) and you cannot obtain the hoped for synchronization?

The solution required a system-wide plan of attack. The plan included timing and signal integrity analysis from a system level standpoint. In other words, the questions above actually changed to:

1) How does the TDO sampling point work in relation to the test system glue logic?

2) How does the frequency affect the hold time margin of the sampling system?
3) How does a boundary scan test system deal with a constantly changing board topology?
4) What happens when all subsystems are connected together to perform the finally intended test function? How do they interact?

The questions can actually be extended but are presented here to initiate thought around conditions that affect a complete boundary scan test system.

The idea of improving timing performance in the 1149.1 scan chain is not new. For example, Bhavsar proposed a method for synchronizing the timing of the boundary scan chain logic to the native's clock speed of the chip that contains it. The objective of this effort was to improve speed in the chain functionality (Bhavsar). This author and other application notes from Texas Instruments (TI) show that these timing issues have been widely studied and documented at the ASIC level, but not much at the board level. The author feels that the issues addressed in this work have not been sufficiently documented for system level applications. There is an existing proven layout-based solution to the hold timing problem at the ASIC level. It essentially says that a hold time violation (within a scan chain in an ASIC) can be adequately solved if the internal TDI-TDO data path is laid out opposite to the direction of mainly TCK and TMS signals to ensure data is valid and stable when the TAP controller clocks data between scan cells (Stang and Dandapani). Solving a hold time problem by adding buffering (which adds delay) to the clock and control signals is part of the solution proposed here to solve the same type of problem but at the board level. This principle is then applied to the

relationship between a boundary scan target and its master boundary scan controller to achieve the desired outcome.

There is a difference between the Stang and Dandapani solution and the approach presented in this thesis. The principle is the same, but its application is very different. Stang and Dandapani apply the delay at the ASIC level and focuses attention at cell-to-cell transfers out to TDO. This work uses the same idea at the board level and focuses its attention to the transfers between systems, in this case, the test system and the UUT.

To describe the hold time margin concept, refer to Figures 2 above and Figure 3 below.

Figure 3 - TAP sequencing & sampling system

Data becomes valid some time after the falling edge of TCK during a scan operation. However, the retiming circuit (D-F/F) is sampling back (reading the TDO pin) with the same, but inverted, falling edge. Clearly a race condition may occur here because two events, different in operational nature and separated by necessary interconnect systems, are synchronized by the same falling edge trigger. With the need of capturing the TDO bit-stream that occurs sometime after the falling edge of TCK, but at the same time using synchronous de-skew techniques to compensate for timing delays on TDO off the chain, this method has a potential to create a hold time violation if there isn't enough round-trip delay.

As stated before, most master boundary scan controllers are designed without a system level perspective like the one proposed as a RISC-based microprocessor chip master boundary scan controller (Baang and Hai). Thus, of particular interest is the relationship between the boundary scan master controller and its target boundary scan chain. This relationship has been largely overlooked in most industry applications today where a small delay on TCK with a 100pF capacitor is all that is recommended to delay, filter and attempt to achieve the same results. It is assumed that trace delays at the board level and this small capacitor in the TCK line should be enough to present valid data to the boundary scan master controller. This assumption is mostly based on the claim that low speed signals are not of much concern and therefore layout should not be a concern either. In this work, we will show this is not necessarily true in all cases and that a system level approach and analysis is needed to ensure correct operation of any boundary scan test system.

To ensure 1149.1 compliance, a master boundary scan controller must sample its TDI pin (TDO from the UUT chain) on the rising edge of its clock source, and output its data on the falling edge of TCK. This makes it look like another device in the chain, but it actually drives the control lines as opposed to the target boundary scan devices on the UUT. However, since boundary scan devices on a UUT may be distant from the sampling point (the master boundary scan controller) in a test system, it is almost always necessary to use re-timing D-F/F's as a memory element to capture the naturally delayed data that comes back from TDO. These natural delays are the inherent true minimal delays of the circuits used to drive the chain to accomplish other boundary scan related tasks.

The work by Stang and Dandapani is certainly a starting point for the solution proposed here, but its application at the board and system level is different in nature. The author reasonably infers that the problems inspiring Stang and Dandapani's work should also inspire solutions applicable at the board level in a very similar manner in which they occur at the ASIC level.

CHAPTER 3

The Unit Under Test

The original problem needs to be restated in more general terms to help understand the proposed solution and how it works. The design of a PCA with intelligent processing most likely has a microcontroller, an FPGA, bus controllers, etc. The combination of parts is clearly unique for every UUT. Limiting the number of devices per chain simply to six to use a reasonable number, the question then becomes:

- Can a boundary scan test system be designed such that, within this limit, it will be guaranteed to perform well regardless of board topology, TAP signals layout, or physical discontinuities?

Before presenting an answer let's take a close look at the real life situation that triggered this research. Time-to-market was absolutely critical and there was no time (and probably there will never be) to re-layout signal traces, match impedances or change anything on the board to improve signal integrity of TAP signals. Figures 4 & 5 represent the board in question. These are shown to illustrate the density and real estate problem of today's electronic circuit boards.

A superimposed image of both top and bottom layers is presented in Figure 4. The reader can see the very little room a hardware designer has to route signals. The TAP interface is intended to be used at low speeds and they do not affect PCA functionality. The result of this real situation is that a solid and consistent solution was demanded from

the boundary scan test engineer under very difficult circumstances. These circumstances are line mismatches, skew, noise, timing issues, trace layout, and physical discontinuities in the test fixture.

The TAP interface to UUT 1 in this case must correct the situation and ensure stable functionality of the entire system. The solution was then conceived to be a combination of delays to improve the system hold time margin and of filtering to improve signal integrity, all in one stand-alone board.

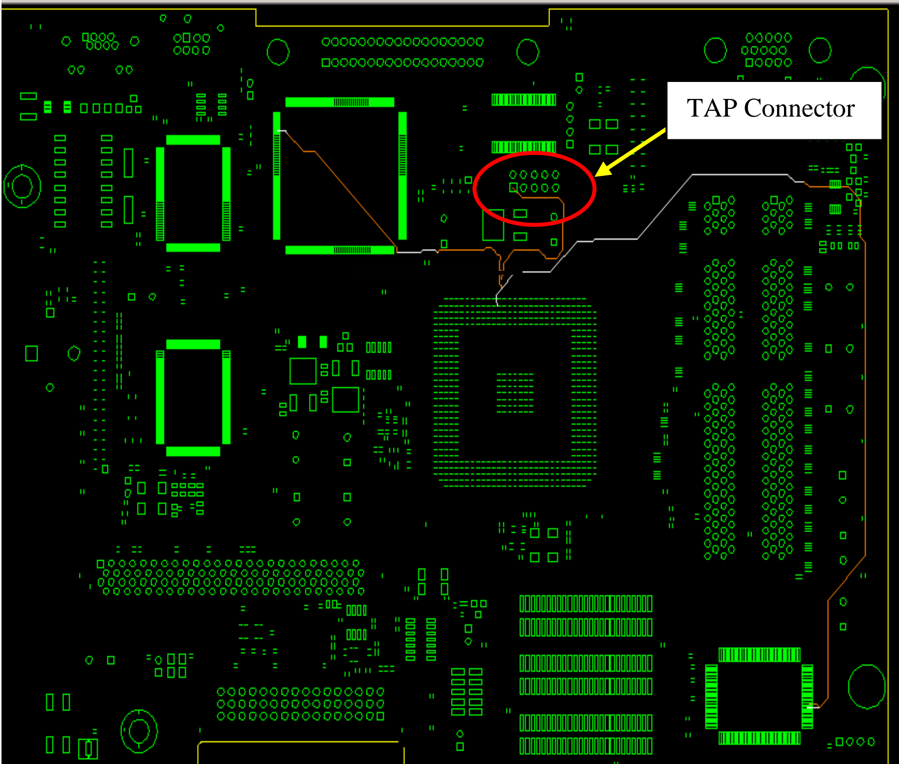

Figure 4 – UUT 1 – Convoluted TCK path

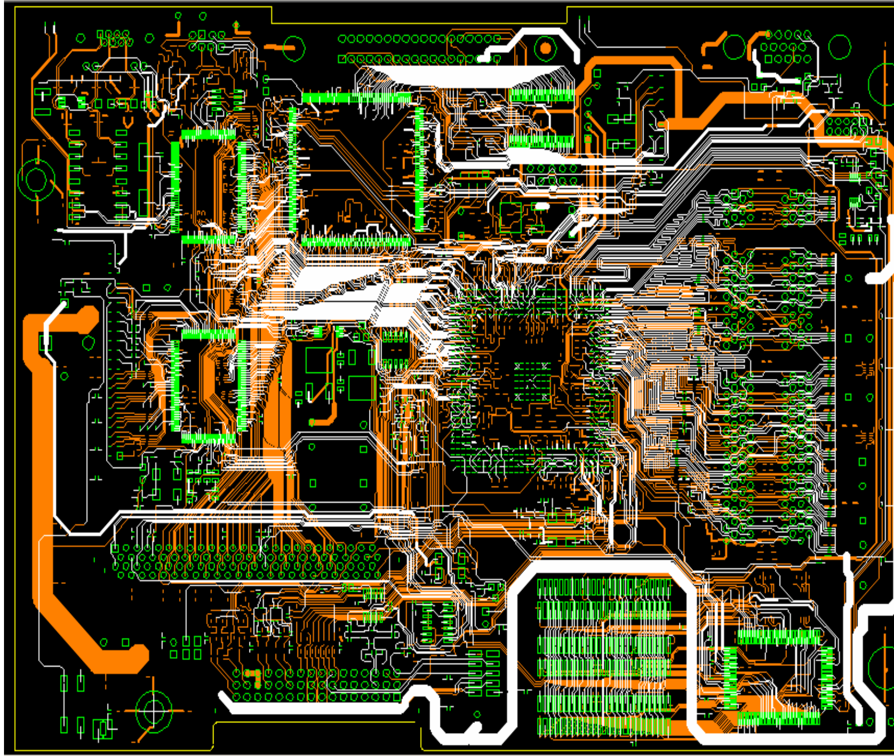

Figure 5 – UUT 1 - Illustration showing trace density

The challenge that these pictures represent is to route TCK, TMS, and nTRST to all three components as shown in Figure 4 and to find the shortest path for TDI-to-TDO around and back to the TAP connector. This proved to be an almost impossible task when consulting with the board designer. The resulting layout was very poor with large stubs that look like antennas when parts are not loaded. These can also cause signal distortions and timing issues. The reason is real and unavoidable; the hardware engineer has no other option but to layout his/her signals first. Then, the TAP signals get laid out last. This is a key motivating factor for the present research.

To further validate the solution, tests were performed in UUT 2 (Test Setup 2). The next two figures are shown to, once again, illustrate the trace density issue.

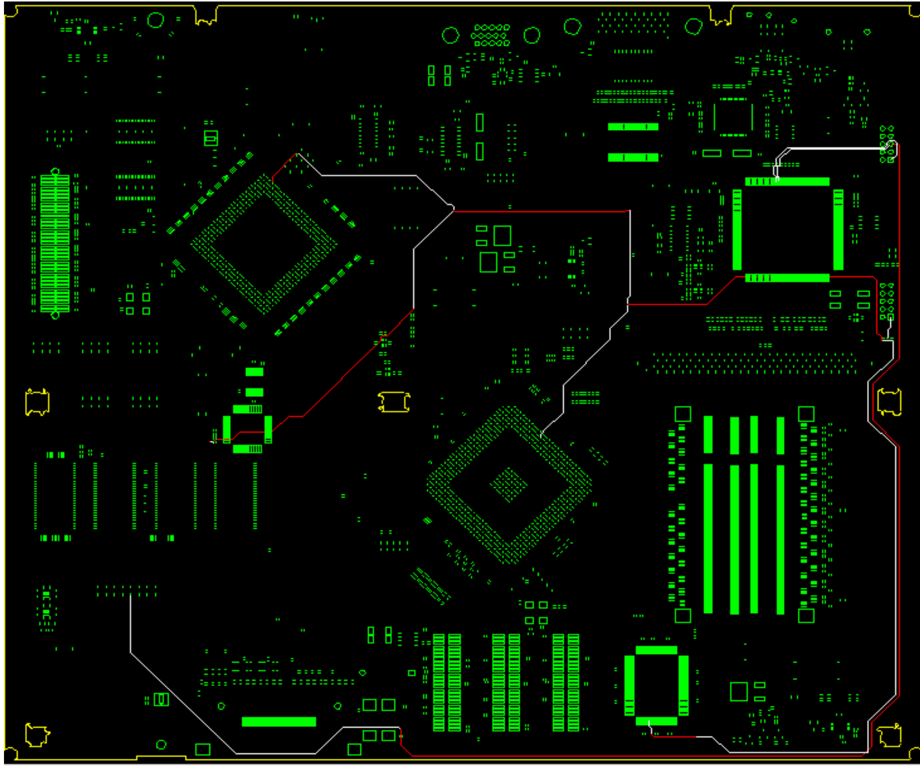

Figure 6 - UUT 2 - TCK path & multiplexed TMS signal on second chain

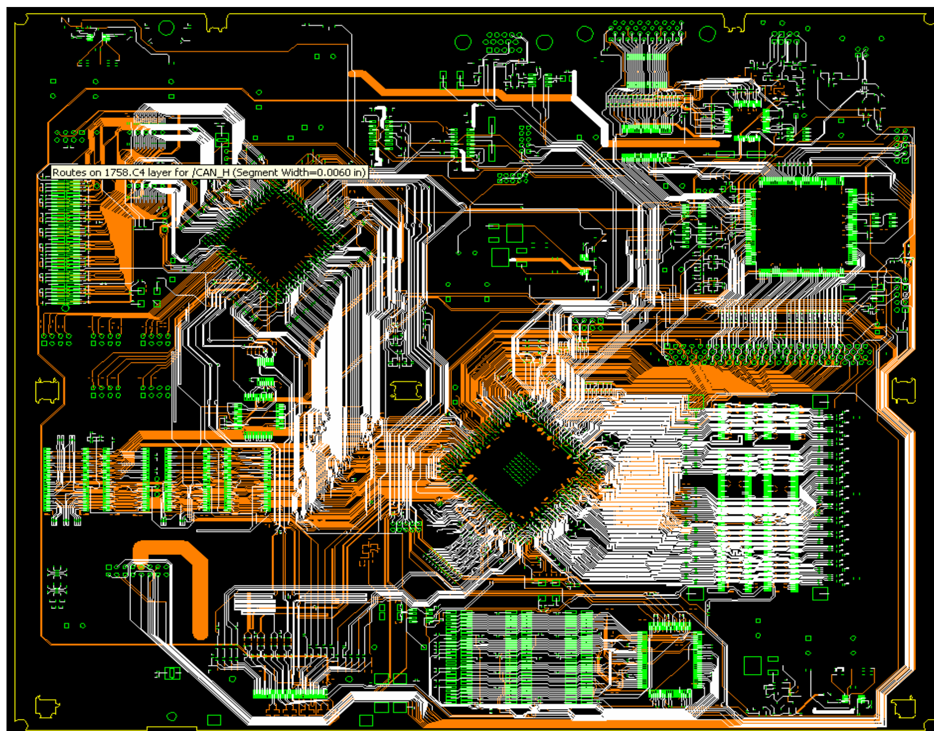

Figure 7 - UUT 2 - Trace density

Designing the Buffer Solution

In order to setup and test the proposed solution the schematic shown in Figure 8 was developed. A circuit board that introduces ~12ns propagation delay plus signal filtering in the outbound-to-UUT direction and ~4ns in the inbound-from-UUT direction as seen by the controller is proposed as the solution. These are measured values because the actual edge propagation time is slightly different than the numbers shown in Figure 8, which are based on the RC constants rather the actual signal propagation time across the SDTAP Buffer. Therefore, the time it takes an edge to reach a gate's trigger level is slightly less than the numbers shown. This delay is then applied to all four control and data signals going into the UUT to ensure balanced chain operation. All RC delays are removed for the TDO signal coming out from the UUT (except the buffer delay), but resistors are left in place to help attenuate any possible ringing.

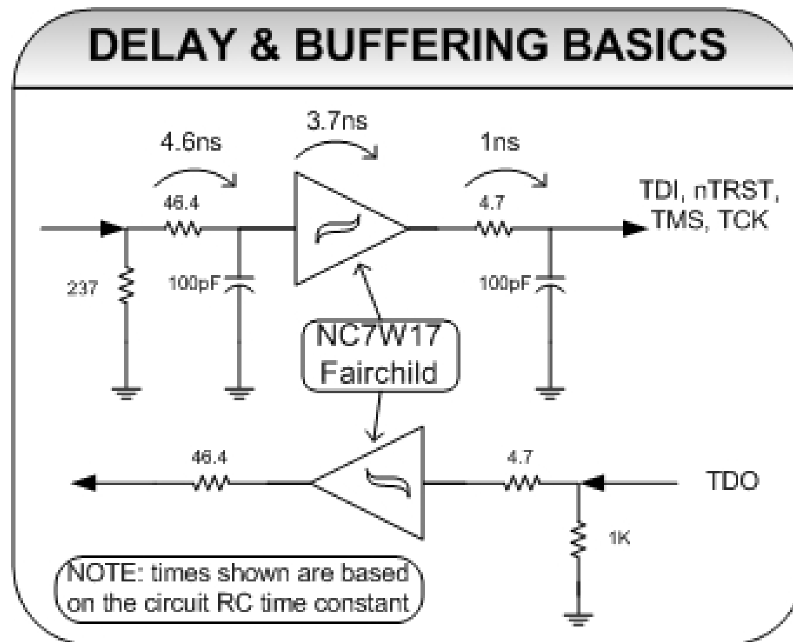

Figure 8 – SDTAP Buffer circuit schematics

Figure 9 shows the actual layout of the final solution proposed in this work. It shows five buffers (U1 to U5) indicating how the solution was implemented in each of the TAP signals.

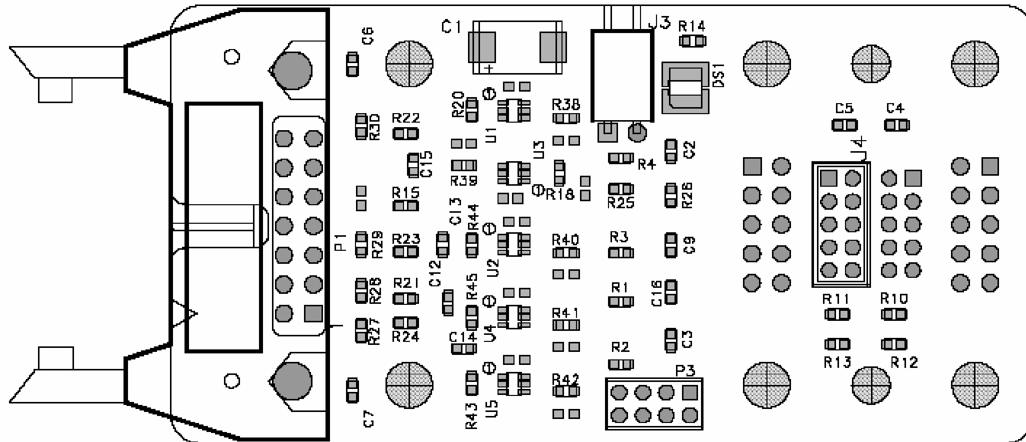

Figure 9 - The actual SDTAP Buffer - Top assembly

The operation of the circuit is very simple. There is a delay introduced by the first low pass filter, the buffer is specified by the manufacturer to 3.7ns, and the last filter provides another small delay. The last filter reduces the slew rate by introducing an approximate 1ns delay, but at the same time filters high order harmonic reflections coming back to the TAP interface. These filters actually are not chosen to any particular frequency, but rather they use small resistor values to prevent line ringing. They also keep voltage drops negligible to ensure that logic voltage levels remain unaffected. The capacitors for the filters were chosen to be within the limits of the buffer’s high output drive capability.

For the TDO signal coming back out of the UUT, resistors were left in place to once again prevent line ringing, but capacitors were removed to minimize line delays.

Normally the trace coming back from the last device in the chain is short and easy to terminate as shown in Figure 8. This differs from the circuit design used in the other four TAP signals because there are no capacitors added. Figure 10 shows a block diagram of the solution proposed and where it fits in the test system. Finally, it is this combination of circuits that makes the heart of the solution proposed in this research.

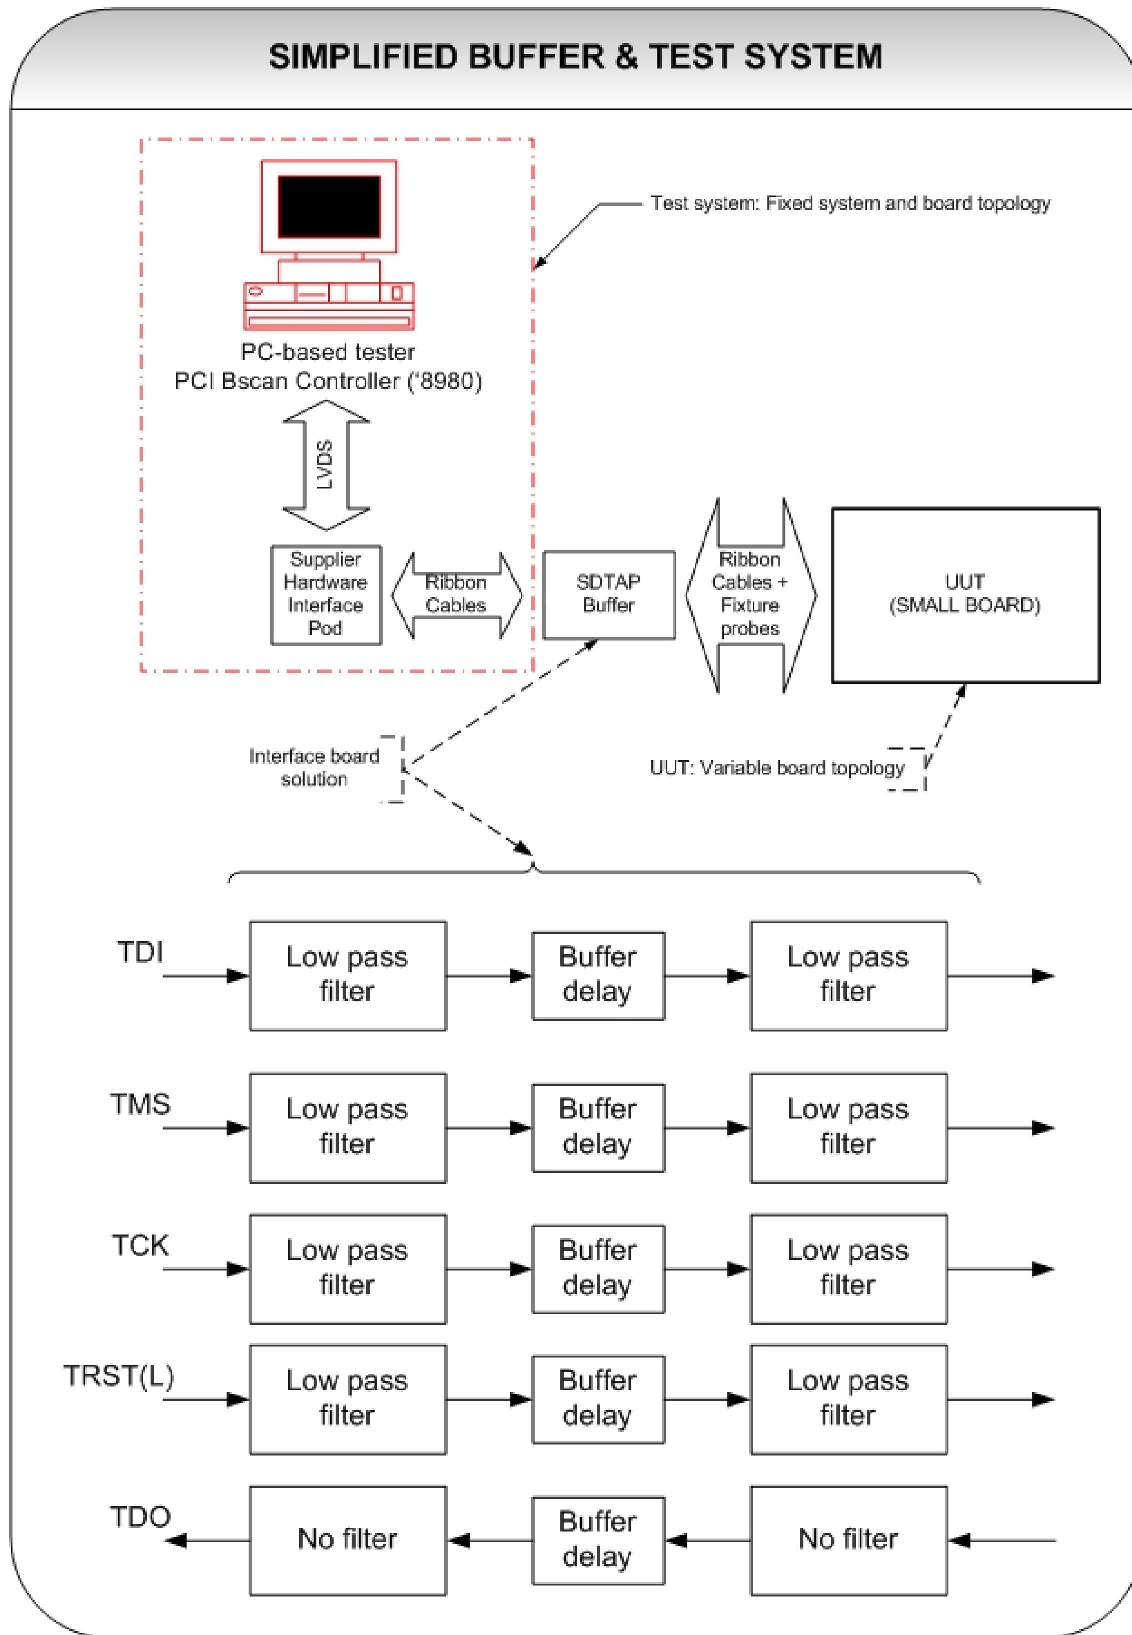

Figure 10 - A generic boundary scan test system and SDTAP Buffer operation

Test System Block Diagram

The SDTAP Buffer board will be plugged into the system as shown in Figure 11. The red-dashed line indicates where all the buffers fit in relationship to the entire system. It also indicates how system A and system B are connected together through the SDTAP Buffer. All measurements will be taken from this setup.

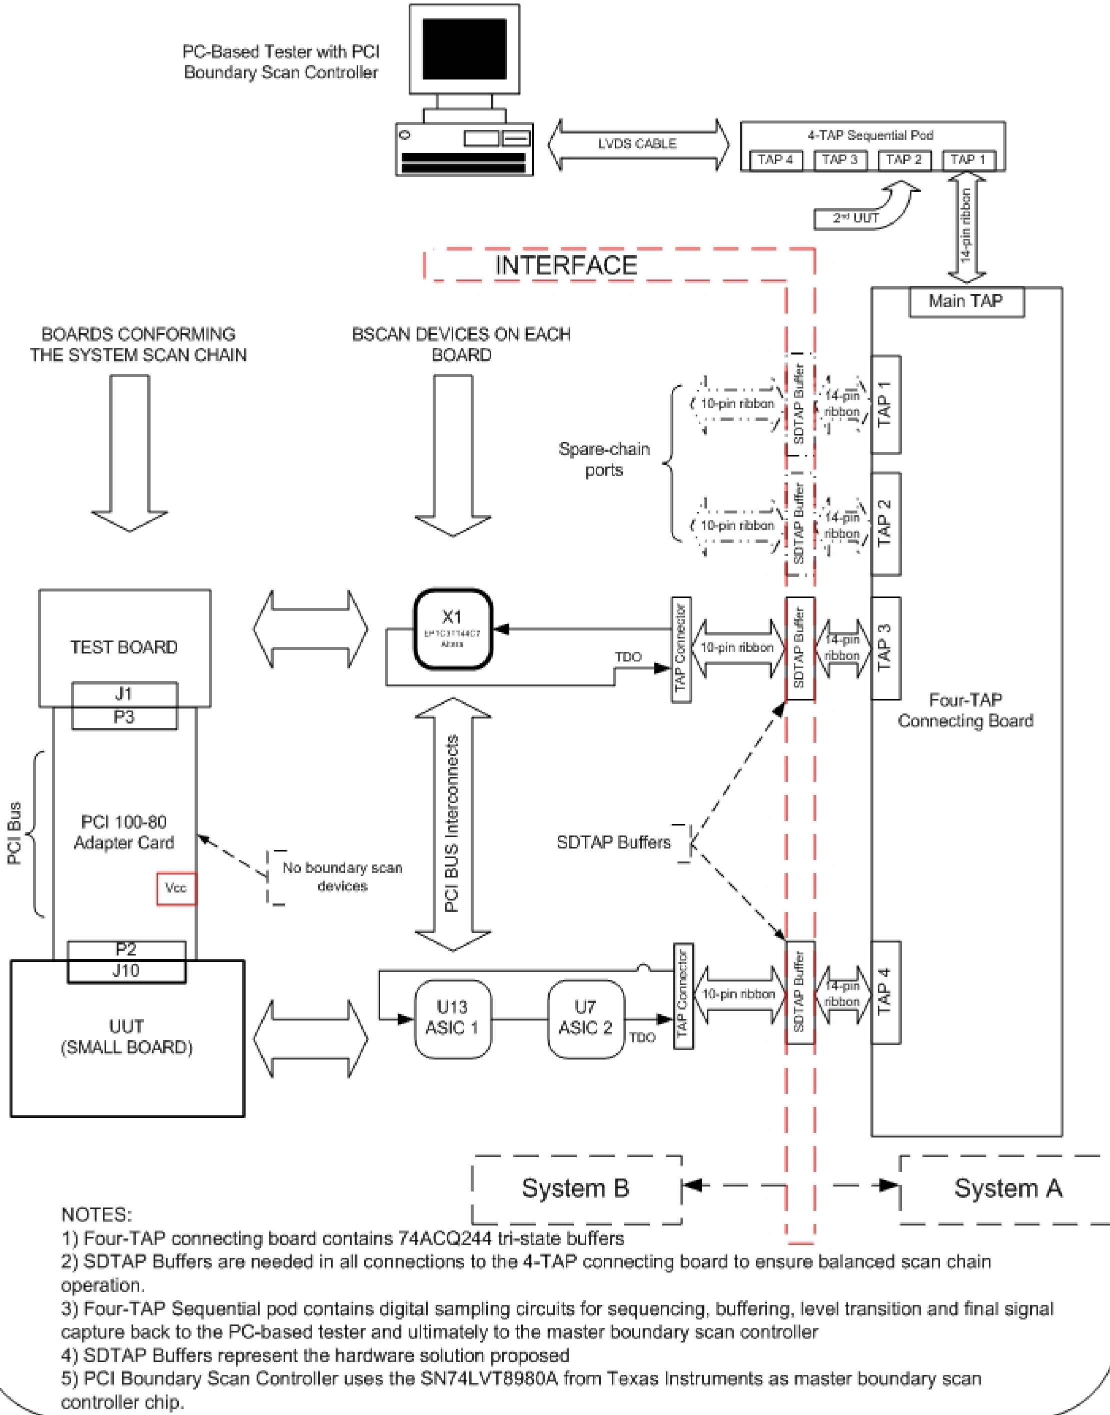

Figure 11 - A particular boundary scan test system setup

The UUT interface to the buffer board was designed as shown in Figure 12 and it resides within the unit under test. This circuit becomes an integral part of the filtering and delay scheme once the SDTAP Buffer board is plugged into the test system.

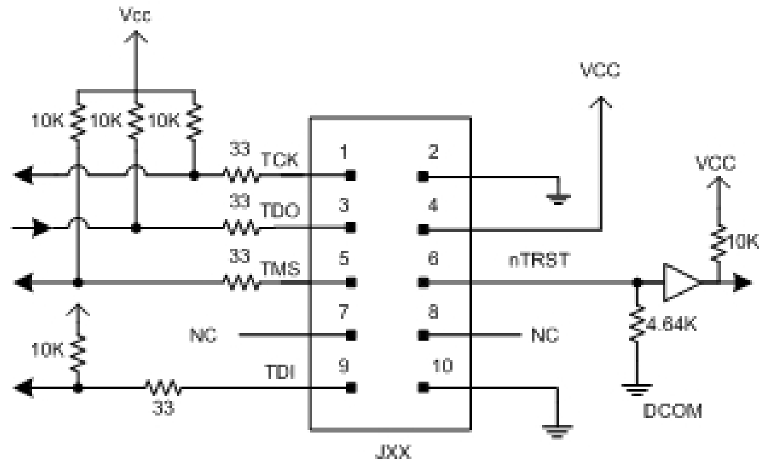

Figure 12 - TAP interface at the UUT

A simple buffering scheme with no filtering (all capacitors removed) will be used to show that the scan chain does not work consistently without filters and that reflections causing serious signal distortions arise even at very low TCK frequencies. This is done to characterize the problem at hand; then, a working solution will be used to show its effectiveness to support the thesis statement.

The idea of a simple buffer is not enough to explain this approach since they are commonly used in electronics. It is a thorough understanding of how an 1149.1 compliant chain functions in relationship to this buffer and the entire test system that makes this an important feature. Notice that the filters are facing both the UUT and the test system for smooth coupling of the two systems. This ensures clean signals are presented to the input side of the 74ACQ244 buffers in the 4-TAP connecting board (see Figure 11 above). This buffer will be called a Special Delay TAP Buffer, the new SDTAP© Buffer.

The Glue Logic

Before further analysis, it is important to describe the individual board functions and their impact on system timing. These are the TAP connecting board and the TAP sequencing board on Figures 13 & 14 respectively. These are referred to as the system "glue logic" boards. They hold the system together and are necessary in most applications. Since a boundary scan master controller is rarely connected directly to its target, real life test scenarios require target interfacing to increase test flexibility and functionality.

TAP Connecting Board

These types of boards are commercially available. The function of any TAP interconnecting board is to concatenate boundary scan chains either in a single PCB or in a system of boards. Normally, this approach is taken when several chains are present in a single PCB, but for some reason they cannot be connected together inside the PCB itself. Typically the individual chains need to be connected into one single chain to achieve maximum fault detection coverage of interconnects inside the UUT. If a test of such nature is developed to deal with only one of the chains, usually the test coverage is low because not all nets on the board will be associated with that particular boundary scan device. For any particular PCB, the goal for maximum structural fault coverage with boundary scan is to have one single chain that touches the maximum number of nets. Figure 13 shows a 4-TAP connecting board and its connections to a potential target.

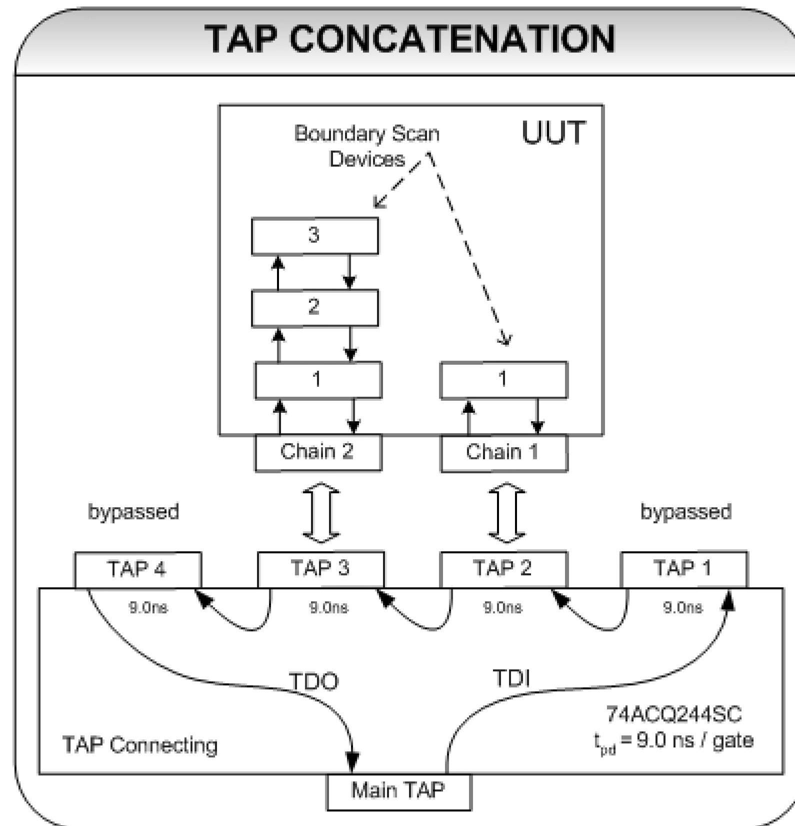

Figure 13 - TAP concatenation system

The concatenating board in this experiment setup uses SN74LVCACQ244SC buffers. These buffers insert propagation delays of up to 9ns that add to those of cables and traces. For convenience in working at 8.335 MHz for TCK, this work ignores ribbon cable and trace delays. As stated in the limitation section, with actual relevant traces of up to 8 in. and ribbon cables of 14 in. maximum length these delays are relatively small when compared to gates and RC delays. According to Johnson and Graham, a typical ribbon cable delay is about 100ps/in and FR4 trace delay about 180ps/in. (599). These delays add once to TCK, TMS, and nTRST; however, they are cumulative for the TDI-TDO path. This is a key idea to improve the system hold time margin since this

cumulative delay is what the solution proposed here partially compensates for. This in turn allows several chains to run at full speed instead of reduced speed.

This glue logic is necessary to perform other functions like the remote selection of internal UUT chains. This function is not used here, but it is available on this board.

The main concept relative to the present research is to consider that these buffers are needed to connect chains. However, at the same time they add a 9.0ns delay in all outbound-to-UUT directions. The TDI-TDO path then suffers incremental 9ns delays as each chain is added. The buffer delay almost removes the delays added by each subsequent concatenation. This is precisely where the buffer delay is beneficial because it subtracts ~6ns from the ~9ns that each concatenation adds. Finally, this reduces by 66% the total TDI-TDO round trip delay seen by the main TAP.

Any chain can be bypassed by means of simple switches provided on board. For this experiment two switches were on; therefore, two chains were connected together and fed back into the main TAP connector. This adds a total of 18ns to the TDO round trip path. The net effect of adding chains is that the TCK frequency may need to be reduced to compensate for those extra delays. However, the SDTAP Buffer compensates this effect by further delaying all TAP signals to try to keep the relative difference between a TCK falling edge and TDO edges about the same as before the insertion of the connecting board. This difference is what the author calls the system hold time margin.

TAP Sequencing Board

The sequencing of the TAP is necessary to provide the ability to test multiple copies of the same board. This is a real situation in PCA manufacturing where one fixture can be used to test multiple electronic assemblies to increase manufacturing capacity by

increasing throughput. Therefore, this commercially available solution is used to select the next image when the test of the previous image has finished, this particular version has room for up to four images.

This board is also used to buffer the signals and introduces some delays as well that are shown in Figure 14 for reference. This particular design uses a retiming method that depends on the falling edge of TCK. Since the Standard requires TDO transfers to occur at the falling edge of TCK the D-F/F is trying to latch the previous bit value, which is later on sampled by the '8980 master boundary scan controller ½ TCK cycle after on the rising edge of TCK.

The test system then uses the concatenating and sequencing boards as key sub-components that need to be in place to enable boundary scan testing of compliant targets. It is important to keep in mind how these systems work together and the effect they have on the overall system operation.

Figure 14 - TAP sequencing system

It is also important to understand the timing structure of Figure 14 to understand how the entire system works; particularly, how it samples TDO back into the PC-based tester for processing. These measurements will be presented and discussed in Chapter 4.

Thesis Support Statement

This solution will support the thesis statement by showing that, independently of tester configurations and/or board topologies, the TDO bit-stream is shifted by some delay "t", to ensure its validity to a certain point between the rising and falling edges of TCK. This is "seen" by the test system at the tester connector to the UUT. In other words, it introduces a small change in the operation of a boundary scan chain at the board level.

It accomplishes this by providing good filtering and enough TDO hold time to facilitate the sampling operation. This is particularly useful for a boundary scan system that samples the TDO pin of the last device in the chain at the falling edge of TCK because it helps the memory element to capture clean data well ahead before TDO changes its current state.

The SDTAP Buffer solution also compensates, by virtue of this known delay, for what the author considers a subtle deficiency in the IEEE 1149.1 Standard. This is where TDO data is assumed to be valid sometime after the falling edge of TCK, but there is no specification of what this time should be. This lack of timing requirements appears in section 4.5.2 of the 1149.1 Standard where TDO operation is described ("IEEE Standard Test Access Port and Boundary-Scan Architecture"). This has probably contributed to implementations of the Standard that are fully compliant but may cause timing problems at the system level. This issue should be a consideration for a master boundary scan chip designer as well. Such designer cannot determine exactly when to sample its TDI input pin because of the relaxation in this rule. Inevitably, chip designers have to limit themselves to the chip technology that they are using and treat the part independently.

Understandably, a boundary scan chip designer cannot know exactly how much delay is needed for TDO to propagate through any particular test system back to its TDI pin. The task at hand is further complicated because UUT topologies are not created equal. It is suggested that it is the test engineer's responsibility to make sure timing and signal integrity requirements are met in a boundary scan test system.

Using general engineering principles and simple analog/digital design, the solution proposed in this work simply makes the TDO data a stable stream ready to be

sampled by the test system. To accomplish this objective the SDTAP Buffer help ensure that the system hold time has sufficient margin and signal integrity is good.

TAP Controller State Diagram

This section is included here for reference. It is important since it is the heart of any boundary scan chain system. Figure 15 shows the TAP controller state diagram as per the 1149.1 Standard (redrawn by the author), which makes possible the operation of a boundary scan chain. It is a finite state machine with mandated and optional built-in instructions to execute certain functions that give access to presence, orientation and bonding of ASICs in any printed circuit board using the “silicon nails” concept (Kajitani, et. al). It is the result of the JTAG work transferred into the IEEE 1149.1-2001 Standard.

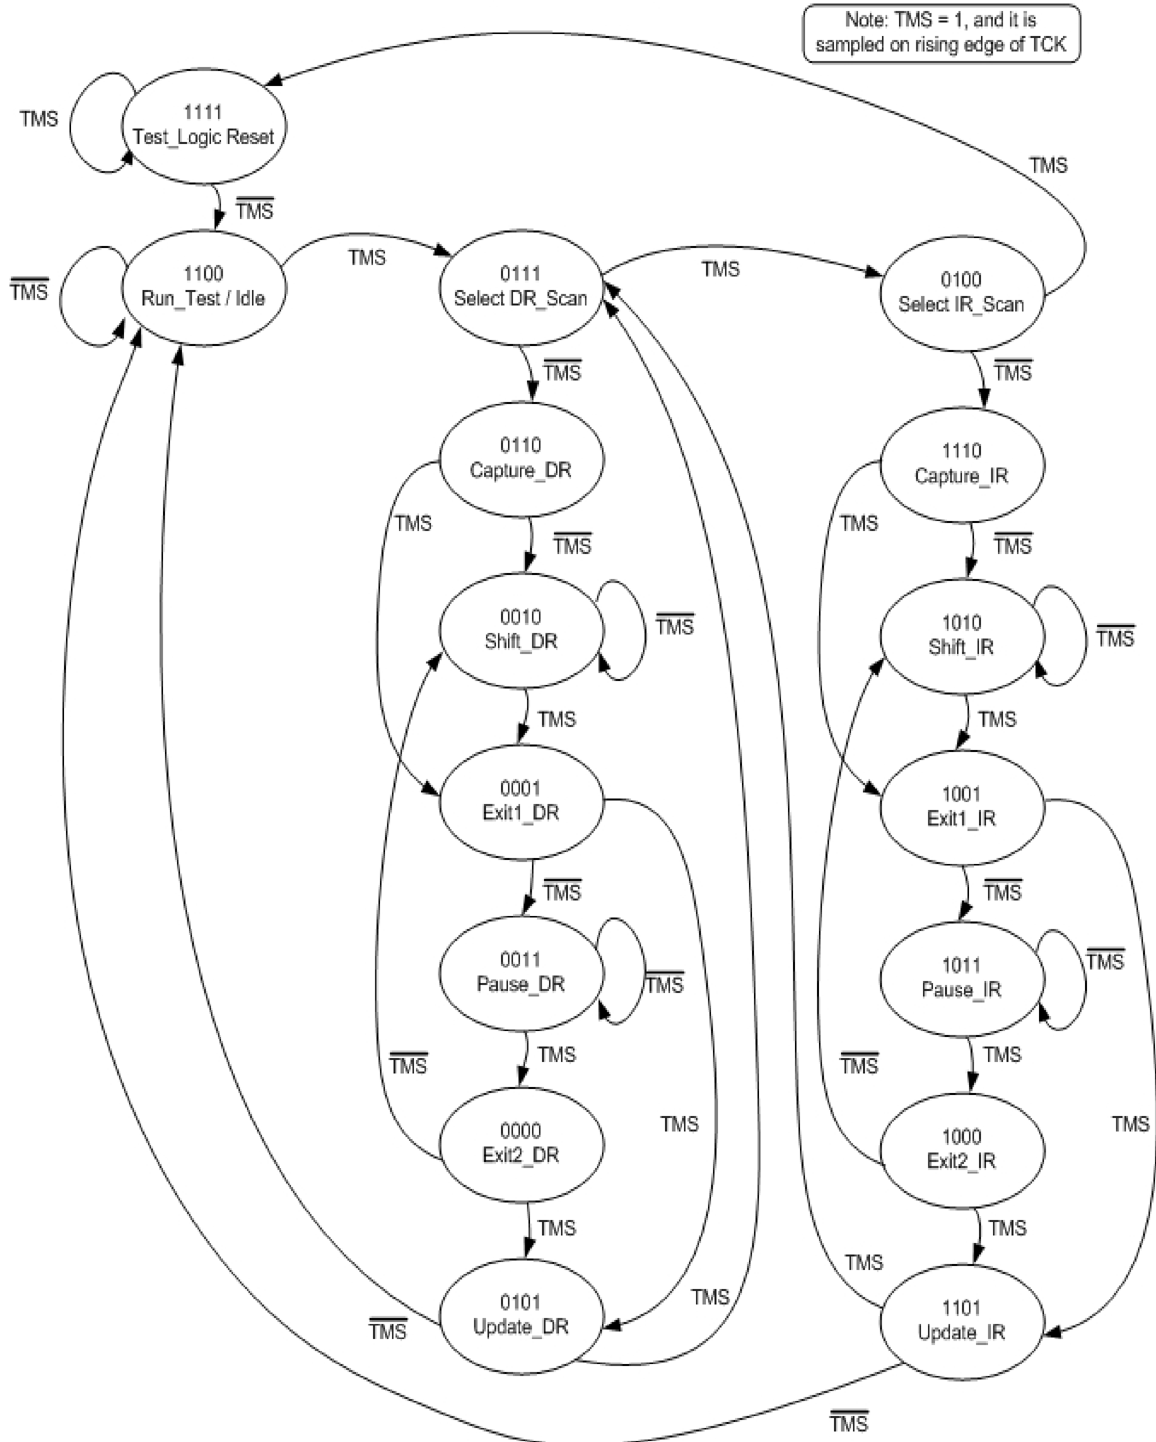

Figure 15 - The TAP controller

This work assumes that this controller is implemented to be fully compliant to the IEEE 1149.1 Standard, but it is not the main focus of this research. Flawless operation of this controller is required to properly exercise the boundary scan testing infrastructure. According to the Standard, actions or events should occur either at the rising or falling edges of TCK. Implementations should use the falling edge of TCK to clock out TDO data and cell-to-cell data transfers (shift operation). With minor exceptions, the ASIC's in this research are all compliant to the 1149.1 Standard.

Incorrect clocking of this state machine is a significant failure mode in this research. It can only be identified with a very sensitive oscilloscope and high-speed low capacitance probes. The following equipment was used to take the appropriate measurements and capture this failure mode.

A 54855A Infiniium Oscilloscope Specifications:

- 7 GHz bandwidth (typical characteristic) with option-008 enhanced bandwidth software
- 20 GSa/s sample rate on all four channels simultaneously
- Up to 1 Mpts MegaZoom deep memory at all sample rates and 32 Mpts MegaZoom deep memory at 2 GSa/s and slower sample rates
- Trigger jitter as low as 1.0 ps rms
- Each InfiniiMax probe amplifier supports both differential and single-ended measurements for a more cost-effective solution
- The 7 GHz InfiniiMax 1134A probe and E2668A/E2669A connectivity kits

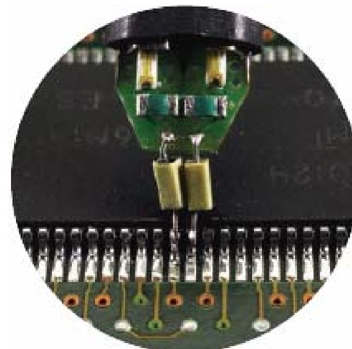

<u>1134A InfiniiMax Probe Specifications:</u>

(See image for details) ⟶

- Bandwidth: 7 GHz
- Dynamic Range: +/- 2.5V
- DC Offset Range: +/- 12V
- Maximum Voltage: +/- 30V

- Differential Input R: 50 KΩ. Differential Input C: 0.27-0.34 pF.
- Single Ended Input R: 25 KΩ. Single Ended Input C: 0.44-0.67 pF.

Figure 18, at the end of this chapter, shows the UUT physical points where the differential probes were connected. This enabled monitoring of the TCK signals as seen by the ASICs very close to the input pads.

Test Setup Methodology

Logic data gathering was based off Figure 18 (at the end of this chapter), which highlights relevant test points. A logic analyzer was used to take time differences at these points when the Scan Path Verification software test was looped. A Scan Path Verification test loads instructions into the Instruction Register and executes tests in a loop to validate the integrity of the boundary scan logic infrastructure.

The same test setup was used for measurements of both working and non-working systems. However, Figures 18 and 19 will serve to illustrate differences between two board topologies and two loading options for each UUT. The loading here refers to the absence or presence a boundary scan device, which is noted with a dash-dotted line pattern within those figures. The need for a Four-TAP connecting board and a TAP-Sequencing board has already been explained in Chapter 3 under the Glue Logic section. The difference between working and non-working systems will be established by a software package indication of failure. The failures will occur when the filtering introduced by the SDTAP Buffer is removed from the test system. A screen output from the software will be our indication of failure. It will show the results of a looped scan path verification test for both failing and passing systems (see Figures 16 and 17). This test is what actually "tests" the tester as mentioned before. The purpose of looping the

test is to make sure the boundary scan logic responds consistently to a repeated stream of control and data signals.

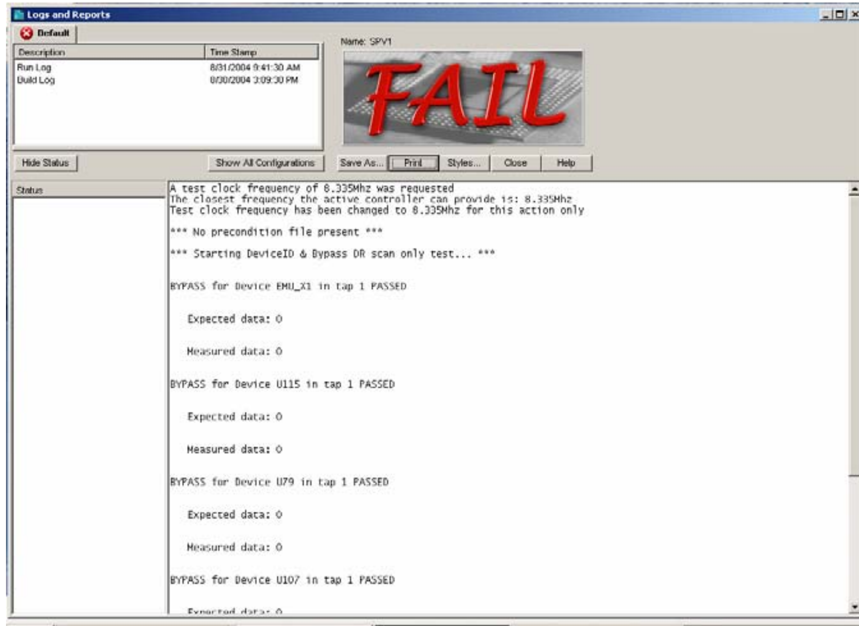

Figure 16 – ScanWorks ® fail screen (courtesy of ASSET Inter-Tech, Inc.)

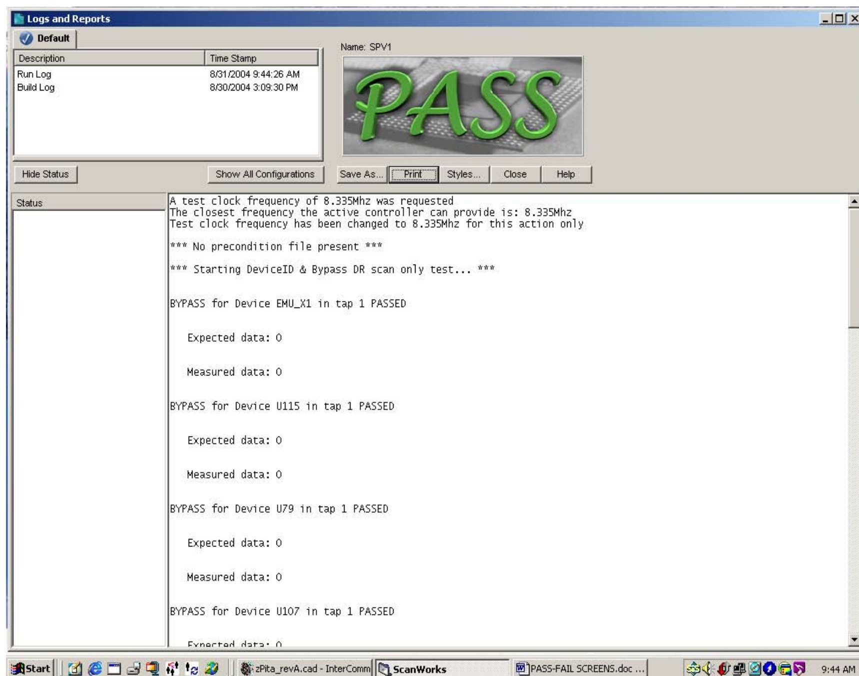

Figure 17 - ScanWorks ® pass screen (courtesy of ASSET Inter-Tech, Inc.)

As described earlier in the Limitations section (Chapter 1), this work is focusing on signal integrity and timing issues and it ignores propagation delay times inherent to traces, cables, or other interconnect discontinuities. The main testing points for logic analysis are clearly defined in Figure 18 and Figure 19 below. These were "A" for after the buffer, "B" for before the buffer and "C" just physically before data is sampled at the sequencing and sampling board. This allows for monitoring of all five TAP signals at each of these points. Timing measurements were then taken to illustrate the improved hold time margin solution proposed as part of this work.

Other more critical points for signal integrity are shown by arrows pointing to the boundary scan devices. These arrows mean that TCK was measured at the closest physical location of the TCK pin for each part. A brief discussion of each measurement will be given and analysis of data in relationship to the entire test system will also be provided.

Test System Frequency Settings

The test system has the following frequency settings available.

Table 2 - TCK frequency settings available

Freq (Hz)	Period (ns)	1/2 Period (ns)	1.5 Period
130000	7692.00	3846.00	11538.00
260000	3846.00	1923.00	5769.00
520000	1923.00	961.50	2884.50
1042000	961.50	480.75	1442.25
2084000	480.75	240.38	721.13
4167000	240.38	120.19	360.56
8335000	120.19	60.09	180.28
16670000	60.09	30.05	90.14

However, this range was not fully exercised because the ASICs used in this experiment could not be operated beyond 8.335 MHz. Therefore, 8.335 MHz is the maximum TCK speed for all experiments. This also means that the UUTs in this research do not allow TCK signals above 8.335 MHz; consequently, the tester could not be run at the next available frequency setting of 16.670 MHz.

Timing Measurements Methodology

The timing measurements are mostly logic measurements. They are included to show benefits derived from adding delays in the TAP path. They will also show how the TDI-TDO path delays introduced by the 4-TAP connecting board are almost compensated by the SDTAP Buffer delay. This is an improvement that impacts the system level performance. These logic measurements did not make sense without a passing system. In other words, all timing measurements refer to a fully working system. This allows for good understanding of how the logic system should work.

This section illustrates how long it takes for a TDO edge to travel to the sampling point and how to vary this delay to enhance the system hold time margin. It confirms that within 130 KHz to 8.335 MHz TCK range the solution’s delays are actually beneficial to the system and not detrimental at all. Contrary to what might be assumed, relative delay of the TCK signal with respect to the sampling device is actually part of a good working 1149.1 system. This is because, per Standard, the TDO signal should change state after the falling edge of TCK. In other words, it is important to understand the time interval between the sampling event and a state change in TDO. Once this was understood, the circuit solution was designed to increase this time. The SDTAP Buffer was then

connected and timing measurements were taken to support the hold time improvement affirmation.

Signal Integrity Measurements Methodology

Because of the extension and complexity of the measurements, a limited set was chosen. This set is explained in the following tables to describe measurements of the TCK signal at different points in both test setups. Test Setup 1 refers to a board that will be called UUT1. This unit under test has two scan chain configurations, one with ASIC 2 present and one without it. Test Setup 2 refers to a board that will be called UUT2. By the same token two scan chain configurations are available but only one was exercised to show improvements in the chain speed. Test Setup 2 is included to illustrate the affirmation that a simple 100pF on TCK is not always enough to clean the TAP signals and allow running of TCK at the maximum UUT chain speed of 8.335 MHz.

Only measurements shown as “X” are discussed below to validate the solution. Measurements shown as “x” were also taken but not used in this research. One reason is that it did not make sense to take an exhaustive set of measurements at different TCK frequencies below 8.335 MHz. These results would have shown a scale change but the signals would actually look and behave the same way. Since embedded boundary scan is a serial chain, it is important to prove that the system can run at its maximum TCK speed. A boundary scan system will always work at any speed below the maximum working frequency. This approach should make this work succinct.

Table 3 - UUT 1 - Chain configuration 1 without solution

	Full cycle	*Rising edge*	*Falling edge*	*Both edges*
ASIC 1	X	x	X	x
ASIC 2	x	X	X	x
ASIC 3	x	x	X	x

Table 4 - UUT 1 - Chain configuration 1 with solution

	Full cycle	*Rising edge*	*Falling edge*	*Both edges*
ASIC 1	X	x	X	x
ASIC 2	x	X	X	x
ASIC 3	x	x	X	x

Table 5 - UUT 1 - Chain configuration 2 without solution

	Full cycle	*Rising edge*	*Falling edge*	*Both edges*
ASIC 1	X	X	X	X
~~***ASIC 2***~~	~~x~~	~~X~~	~~X~~	~~x~~
ASIC 3	x	x	X	x

Table 6 - UUT 1 - Chain configuration 2 with solution

	Full cycle	*Rising edge*	*Falling edge*	*Both edges*
ASIC 1	X	X	X	X
~~***ASIC 2***~~	~~x~~	~~X~~	~~X~~	~~x~~
ASIC 3	x	x	X	x

Table 7 - UUT 2 - Chain configuration 1 without solution

	Full cycle	*Rising edge*	*Falling edge*	*Both edges*
ASIC 1	X	X	x	x
ASIC 3	x	x	x	X
ASIC 4	x	x	x	X

Table 8 - UUT 2 - Chain configuration 1 with solution

	Full cycle	*Rising edge*	*Falling edge*	*Both edges*
ASIC 1	X	X	x	x
ASIC 3	x	x	x	X
ASIC 4	x	x	x	X

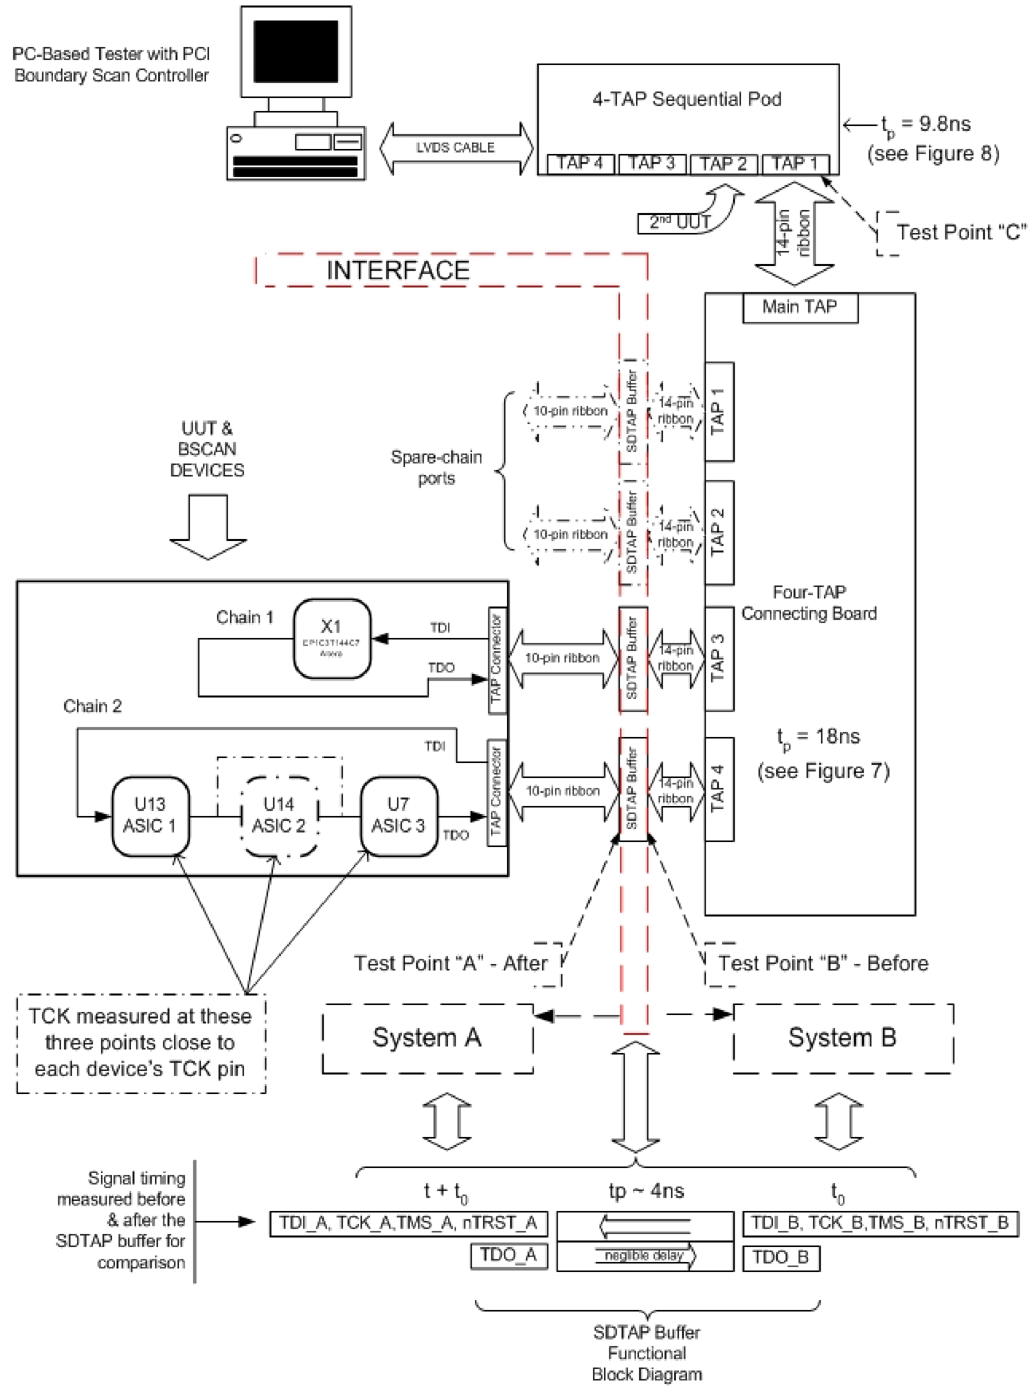

Figure 18 – Test setup 1

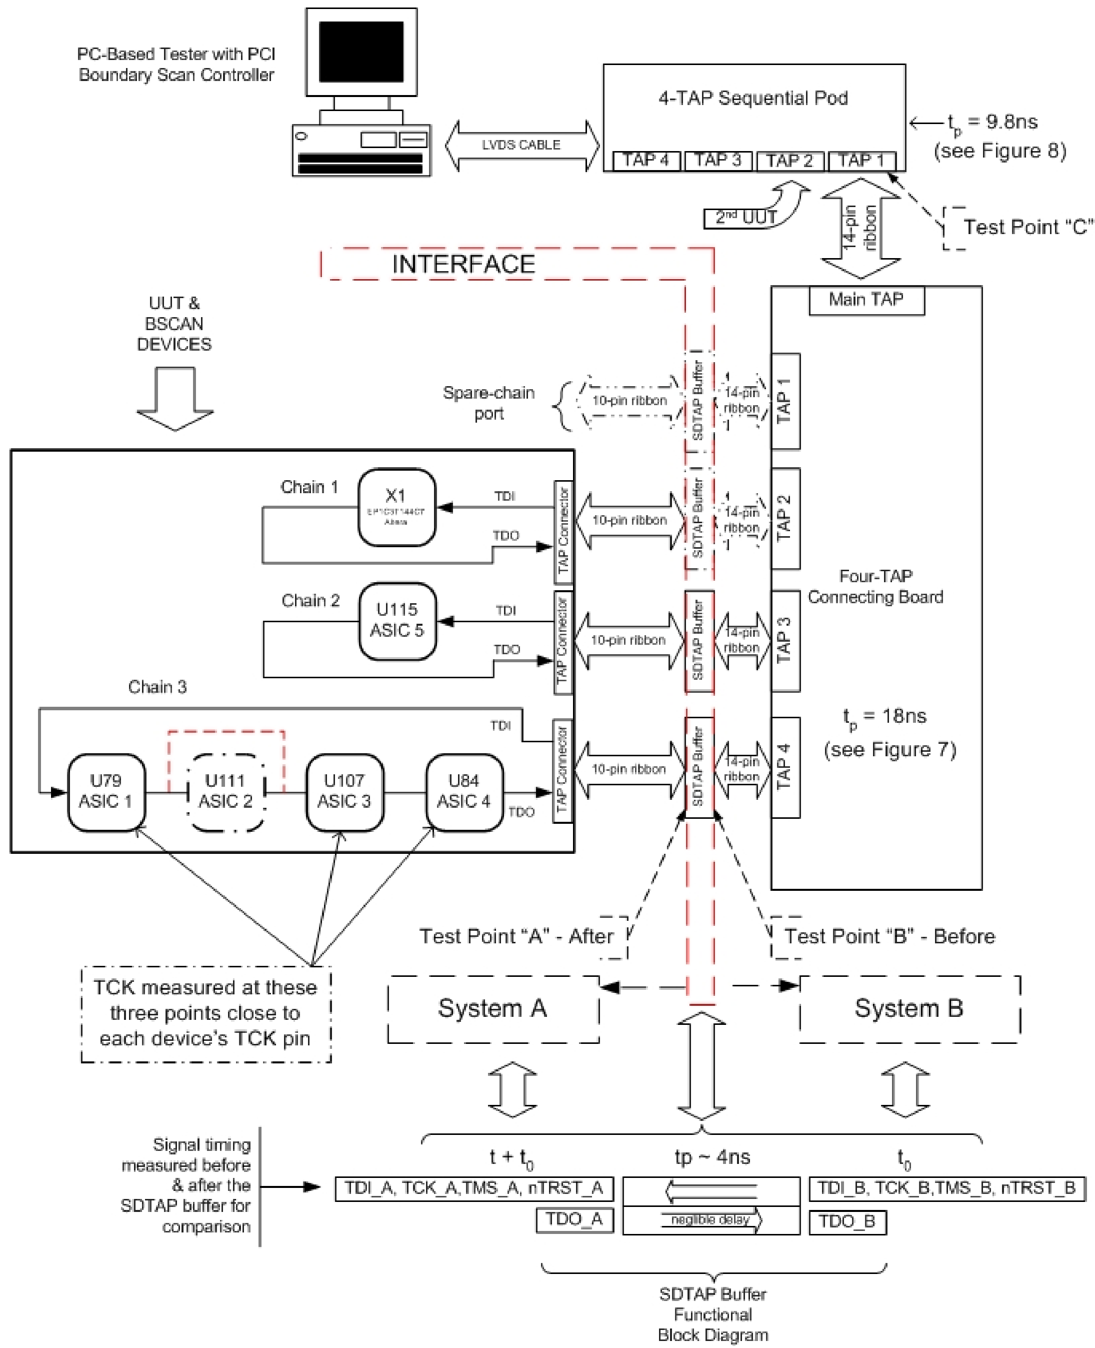

Figure 19 - Test setup 2

CHAPTER 4

Timing Measurements Analysis

Testing the system timing relies on the basic operation of a compliant scan chain. As discussed in previous sections, all TDI pins for each device in the chain are sampled on the rising edge of TCK. Conversely, all TDO pins for each device in the chain drive data out after the falling edge of TCK. These concepts must be kept in mind to understand the measurements in Figure 20, Figure 21 and Figure 22 where TCK was set to 8.335 MHz unless otherwise noted. This section will focus on Test Setup 2 to support and explain the improved timing portion of the thesis statement.

To this end, the hold time margin (time between markers G1 and G2) for Test Setup 2 was measured to be 56ns. This is for a fully working system and includes the ~6ns SDTAP Buffer delay plus glue logic delays. It means that the previous bit at the TDO_F/F = D_F/F input pad is held for 56ns after it is sampled on the rising edge of CLK_F/F. This is the already improved margin in question, which is mostly independent of TCK's frequency. Referring to the D flip-flop in Figure 14, G1 (TCK) is marking the clock input and G2 (TDO) is marking the D input. The longer this time, the better the hold time margin of a boundary scan system. This time is the sum of three time intervals. The time (T1) it takes for a TCK falling edge to reach the UUT connector from the source, the time (T2) it takes a TDO edge to appear at the UUT connector after a falling edge of TCK at the connector, and the time (T3) it takes for that TDO edge to reach the sampling point (D-F/F)

Roughly speaking about Figure 20, notice how the time difference between TDO edges at points A and B is much less than the difference between TCK edges at the same points. This result confirms the buffer design objective of adding delays in the outbound-to-UUT direction (TCK path), but minimizing them on the opposite direction (TDO path). The SDTAP Buffer solution strives to increase T1 for a TCK edge to reach the UUT, but to minimize impact to T3 for a TDO edge to reach the sampling point at the appropriate time. Since by design it is not possible to avoid insertion of delays, it is the relative difference between T1 and T3 that needs to be considered when designing such a buffer. Once again, the SDTAP Buffer increases T1 on purpose to increase the system hold time margin and attempts to minimize additions to T3. Time T2 should remain about the same because it is the UUT response time to present valid TDO data after a TCK falling edge.

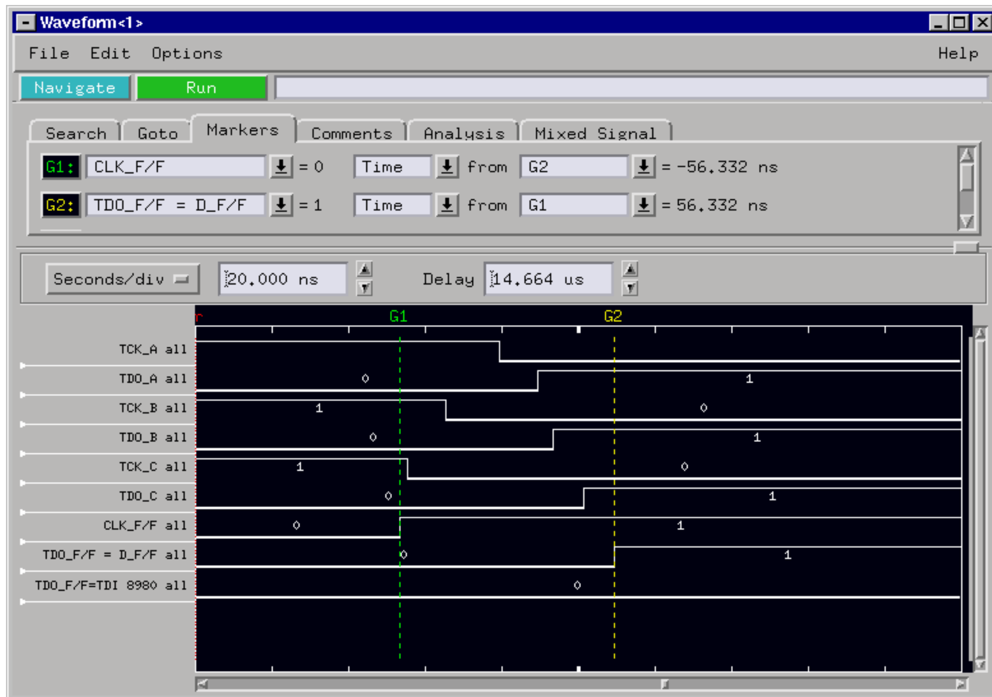

Figure 20 – Test setup 2: working system hold time margin = 56ns (TCK = 8 MHz)

The next two figures show a system with this margin diminished. Figure 21 and Figure 22 confirm that the hold time is hardware fixed and mostly independent of TCK since there is no difference between measurements at two different TCK speeds. The hold time margin is 22ns. This means that TDO will change state (if there is any) 22ns after the previous TDO bit was sampled. These two measurements were taken without the buffers and without the 4-TAP connecting board (one chain active); only one 100pF capacitor on TCK was used. This is why there is no signal showing at point B and it is here to show how the system behaves without glue logic and buffers.

The time difference of Figures 21 & 22 with Figure 20 is 56 ns - 22 ns = 34ns. The following table shows test conditions, system hold time and TCK frequencies used.

Table 9 - Improved system hold time margin

	T1(ns)	*T2(ns)*	*T3(ns)*	*System Hold Time(ns)*	*4-TAP*	*SDTAP*	*TCK (MHz)*
Fig.20	28	8	20	56	Installed	Installed	8
Fig.21	7	8	7	22	Not Installed	Not Installed	1
Fig.22	7	8	7	22	Not Installed	Not Installed	4

The difference between the two systems is explained by the difference in test conditions. The 4-TAP connecting board adds ~9ns to both directions and the SDTAP Buffer adds ~12ns to the TCK, TMS and nTRST path and ~4ns to the return path. According to these propagation times, the system behaves as expected. The timing can be roughly measured from Figures 20, 21, & 22. All logic measurements are about 10% accurate and are reasonably acceptable.

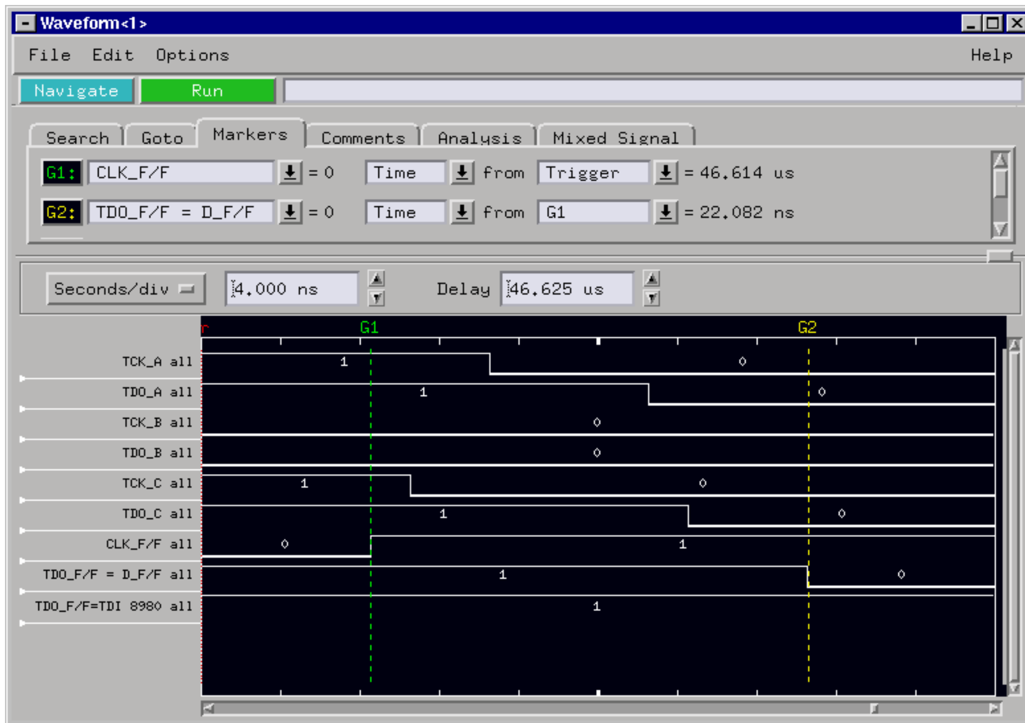

Figure 21 - Setup 2 - TCK = 1 MHz

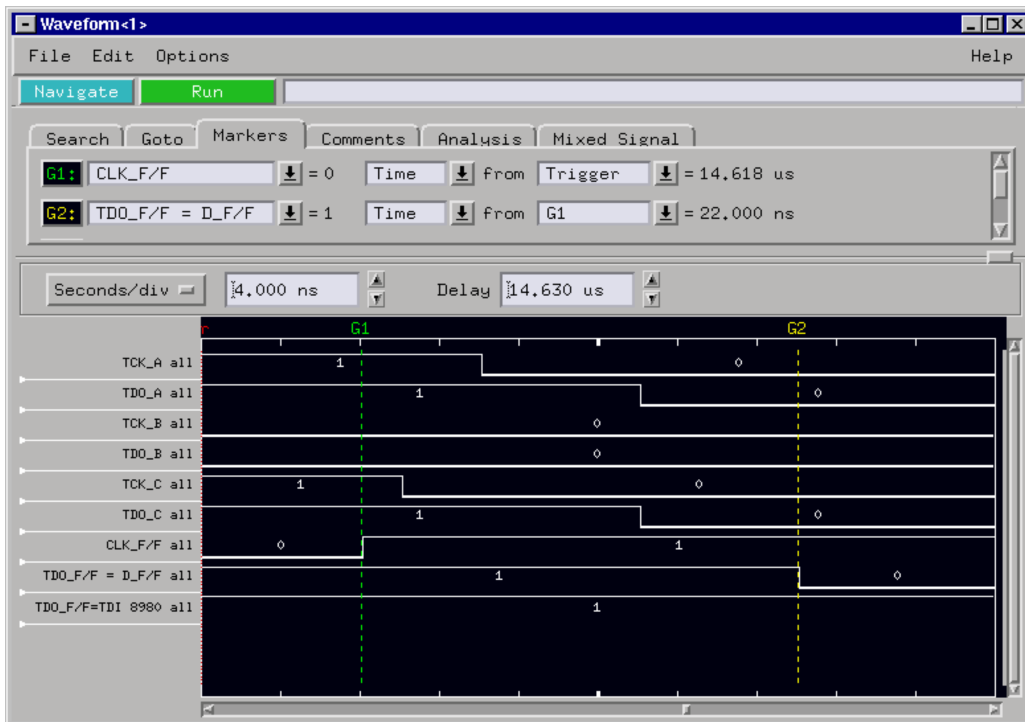

Figure 22 - Setup 2 - TCK = 4.16 MHz

A boundary scan system is more reliable if the system hold time margin is the longest possible with the maximum TCK available. The previous analysis of timing measurements confirms that the system hold time margin is determined by the hardware of choice. If the hardware is carefully chosen to increase delays to TCK, TMS and nTRST, but not to TDO, it is possible to achieve the desired buffer behavior within the system. This sufficiently proves that the SDTAP Buffer is beneficial and not detrimental to timing in a boundary scan system.

Signal Integrity Measurements Analysis

This section of results analysis is very relevant to the solution proposed. It shows that the SDTAP Buffer helps increase chain performance by means of correcting signal integrity. It also shows that the layout of TCK (and TMS) signals has minimal impact on signal integrity when the source termination is a complex impedance (meaning real + imaginary).

TCK Signal Integrity Analysis

Test Setup 1 was used as the test bed for these measurements. These were taken following the tables' order and methodology of Chapter 3. For analysis each figure will be shown and discussed individually. All measurements in this section relate to the test setup in Figures 18 & 19. Keep in mind, as each picture is presented, where the measurement was taken in relation to Figures 3 and 6 as presented before.

It is also very important to recognize that these measurements were taken with special high speed probes with 0.44pF input capacitance soldered with short leads to the physical point of interest. Without the use and availability of these probes, this work

could not have been possible. As we will show later, regular 10-15pF probes are such a high load that they can mask the real signals in the circuits under test.

Failure Mode Description (Table 3)

This table shows measurements for the TCK path in Figure 3. As shown there, three stubs with very different lengths are branching out of a single point to reach boundary scan device. The 100pF cap was removed from all signals and TCK was fed into UUT 1 with SDTAP and 4-TAP connecting board left in place.

It is suspected that there is a signal integrity problem somewhere along the TCK path. Figure 23 shows that TCK is mostly clean with a little bit of ringing. This result led to more measurements to try to locate the source of the problem because this signal looked clean. The scale is shown expanded in Figure 24 to make sure there were no double edges on ASIC 1's TCK pad.

Figure 25 for ASIC 2 shows an irregular edge right at the threshold level of 1.65V and lots of ringing with amplitude of 758mV. This clearly generates extra clocks that incorrectly advance the TAP controller to an unknown state causing erratic chain operation. This can happen even in the presence of good signal integrity on TDO. The rising edge on ASIC 2 looked very similar to the way it is shown in Figure 26.

Although far from the threshold level, Figure 27 shows ASIC 3's falling edge with unacceptable ringing of 771mV. Therefore, with this configuration and no filtering, the system presents a hard fail caused by reflections generated within the two longest stubs.

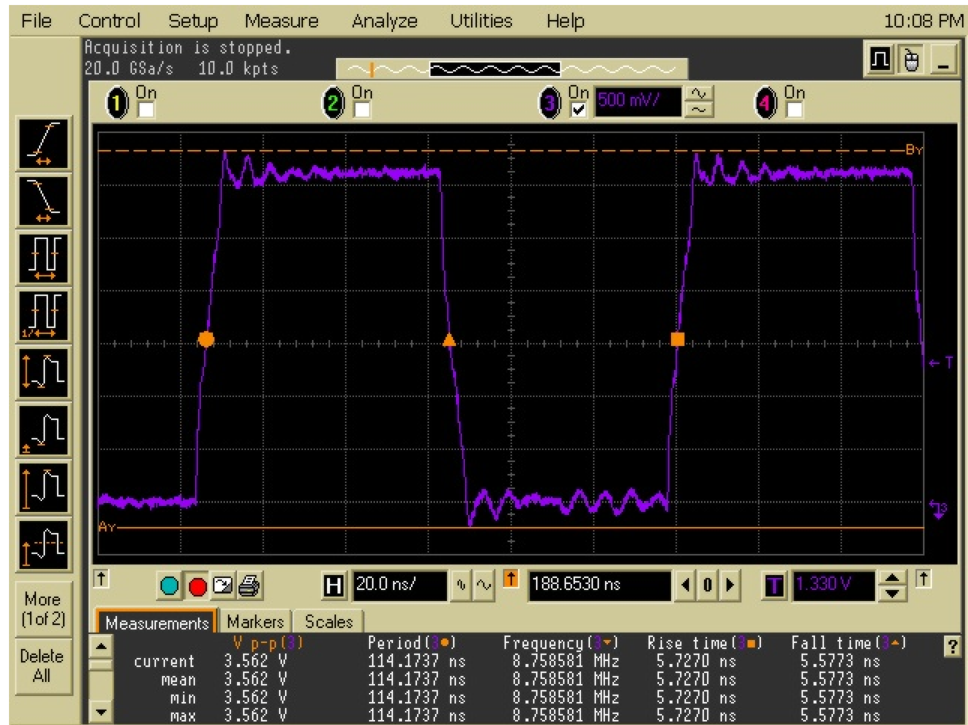

Figure 23 - ASIC 1 - no filtering

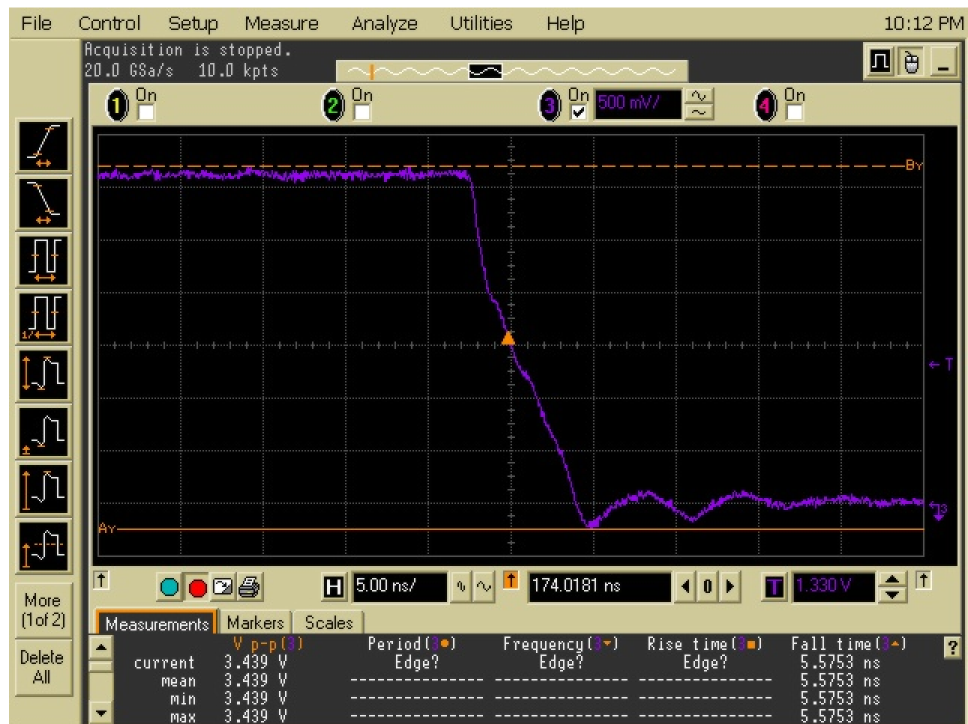

Figure 24 - ASIC 1 falling edge

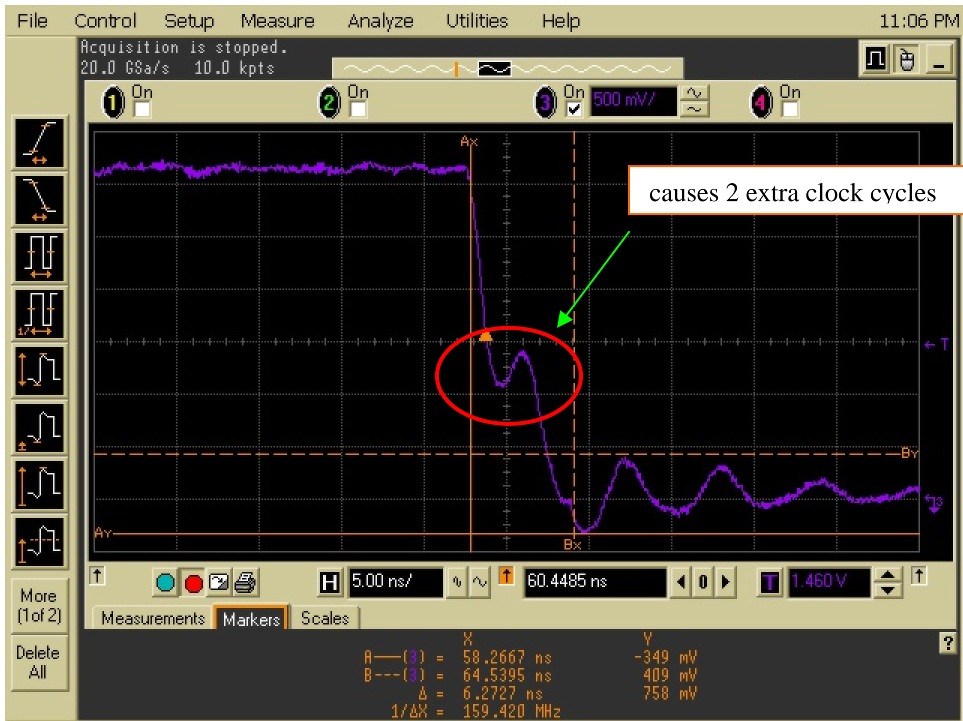

Figure 25 - ASIC 2 falling edge

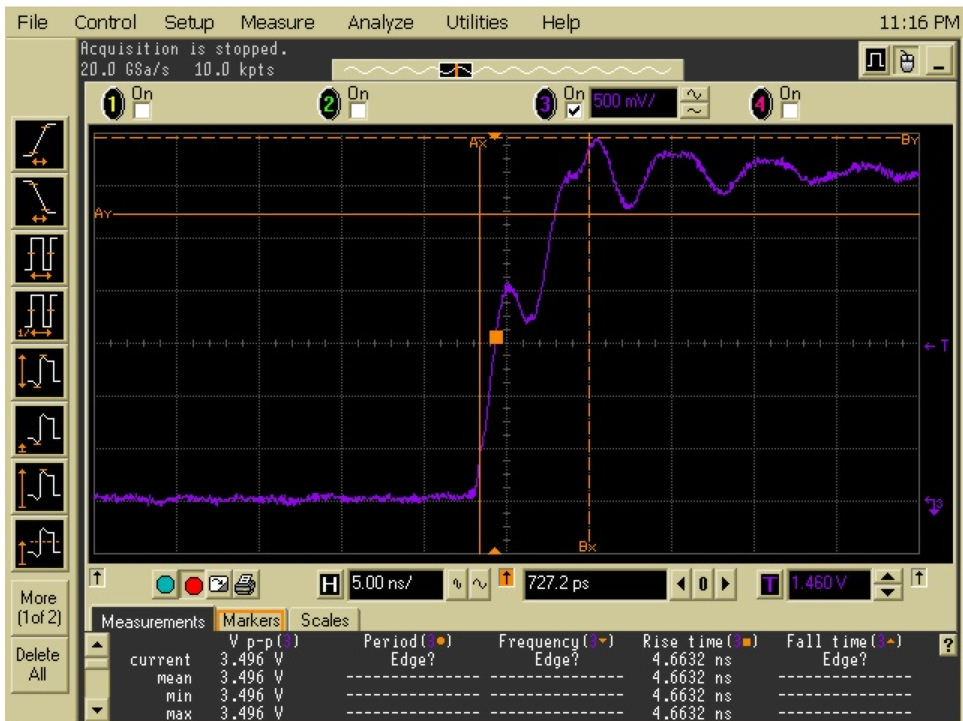

Figure 26 - ASIC 2 rising edge

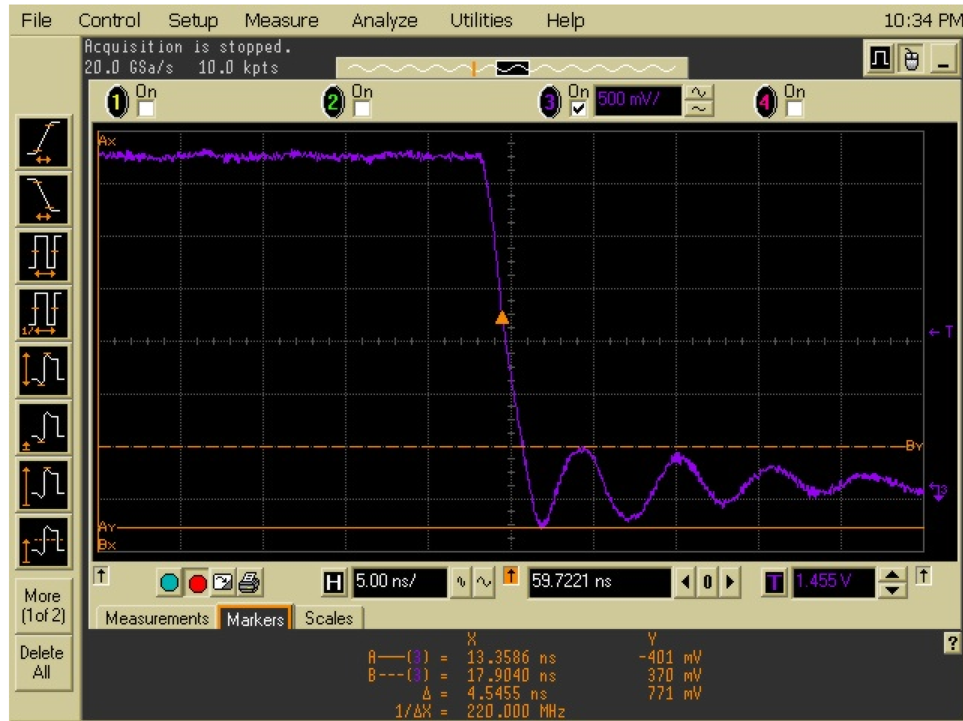

Figure 27 - ASIC 3 falling edge

Testing with the SDTAP Buffer (Table 4)

This table shows all the previous measurements with the SDTAP Buffer solution applied and it refers to Figures 28 through 32. The system was very stable with all six scan path tests enabled or with individual ones as well. The reader is invited to compare measurement sets 3 & 4 to see the effects of the buffer solution.

The measurements speak for themselves. It is worth noticing that reflections in this configuration are filtered and that the edges are slower than the buffer specs of 3.7ns. It is also noticeable that the amplitude of the ringing was greatly reduced.

By reducing the slew rate, this approach reduces the high frequency content of the step function making it easier to have clean edges that allow reliable operation at 8.335 MHz. It is easily understood that operation below this frequency is guaranteed; therefore, slower frequencies were not used.

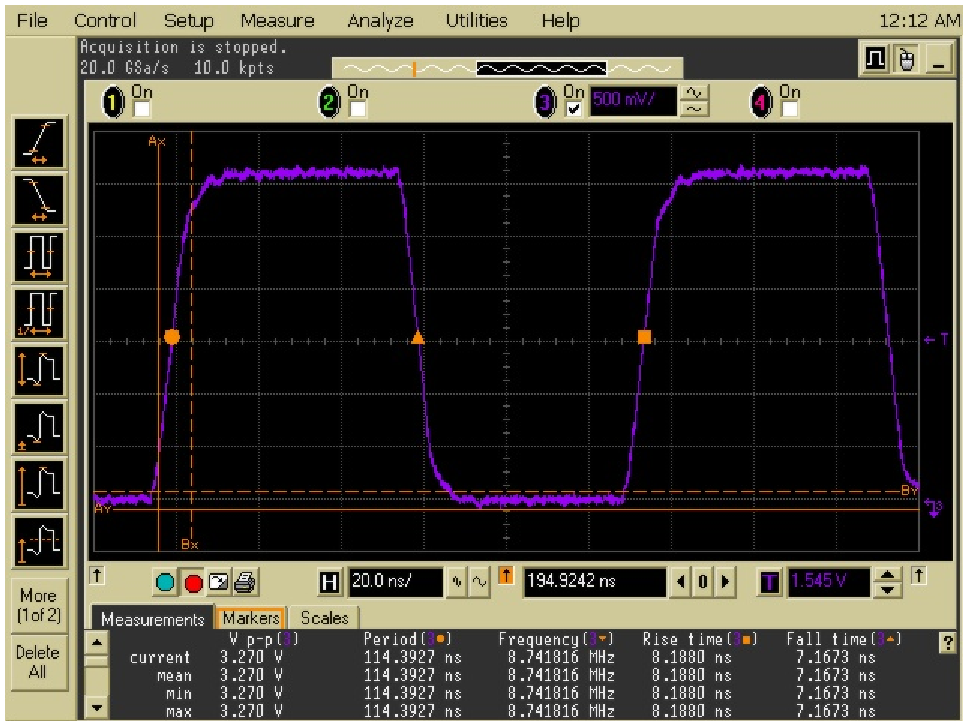

Figure 28 - ASIC 1 with solution

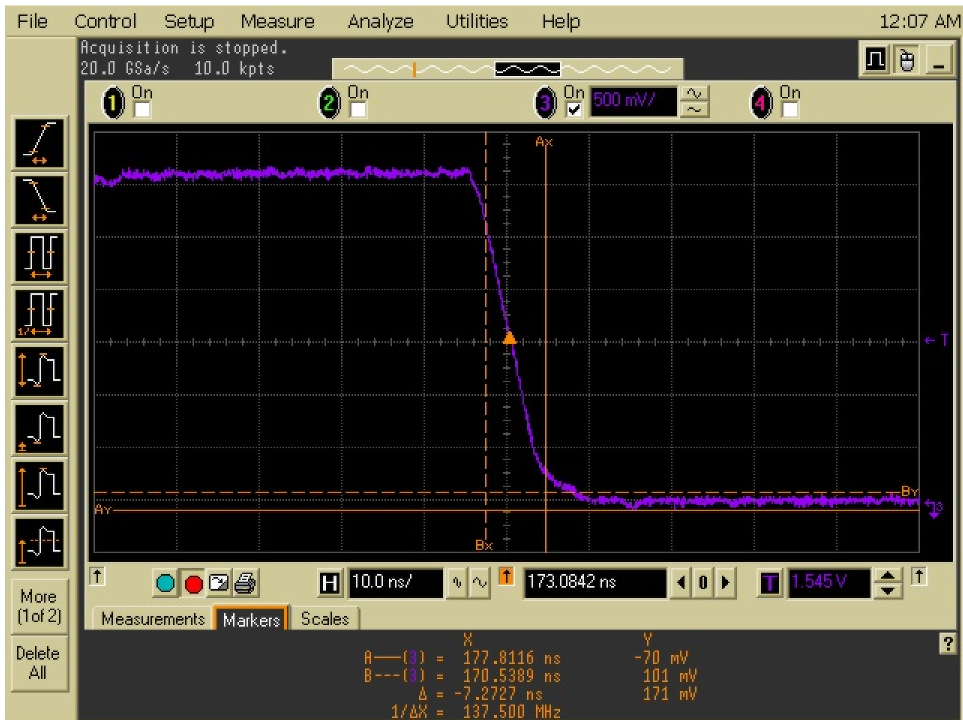

Figure 29 - ASIC 1 falling edge

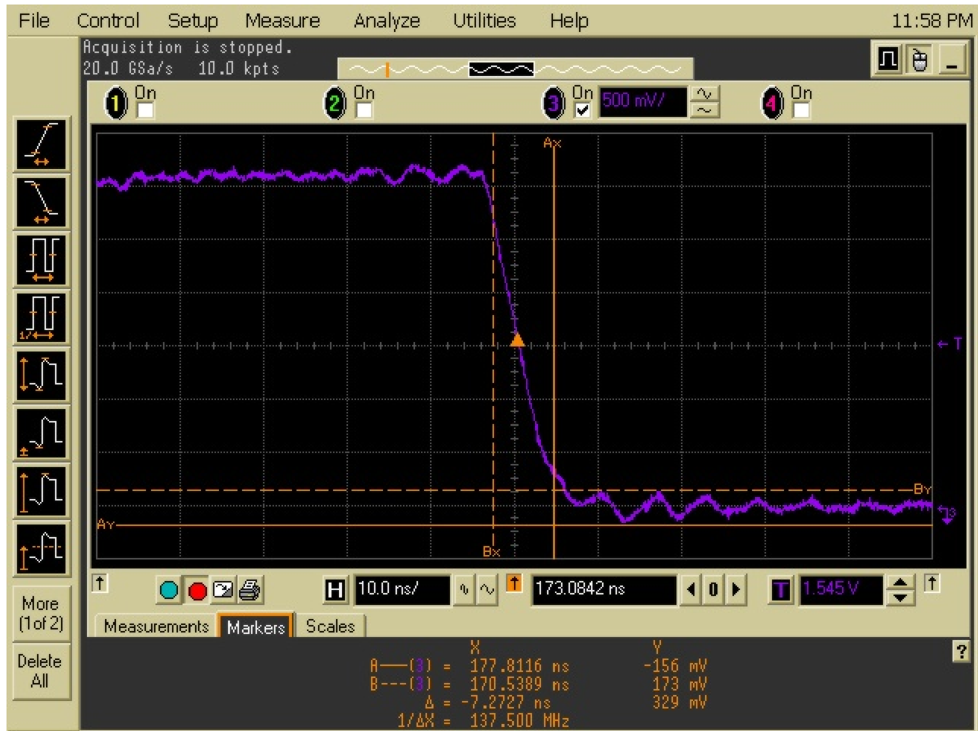

Figure 30 - ASIC 2 falling edge

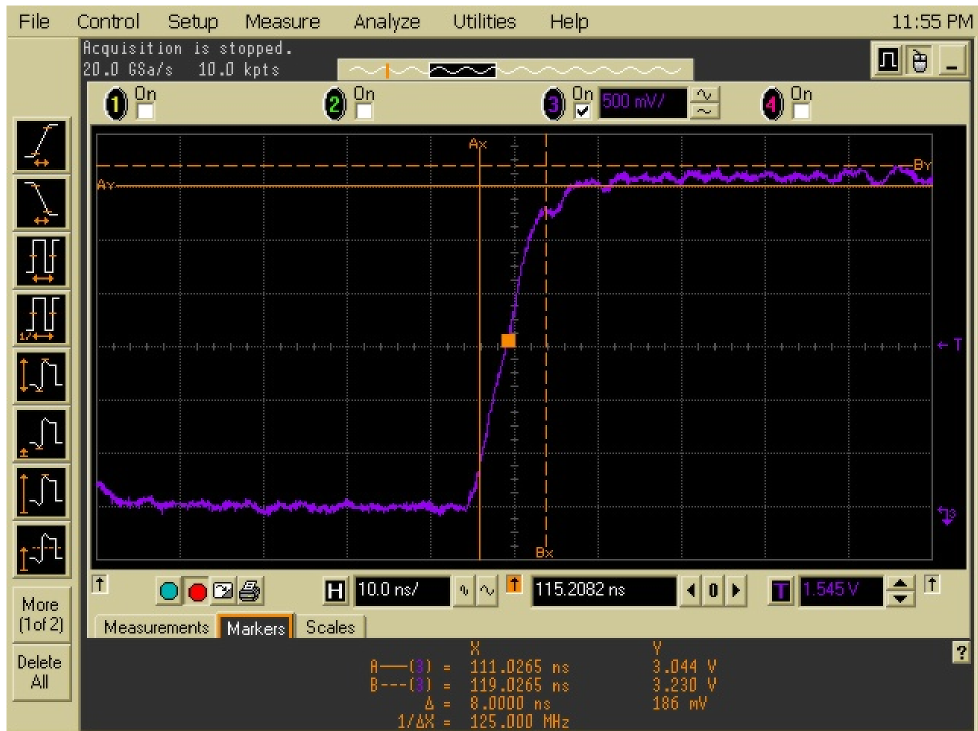

Figure 31 - ASIC 2 rising edge

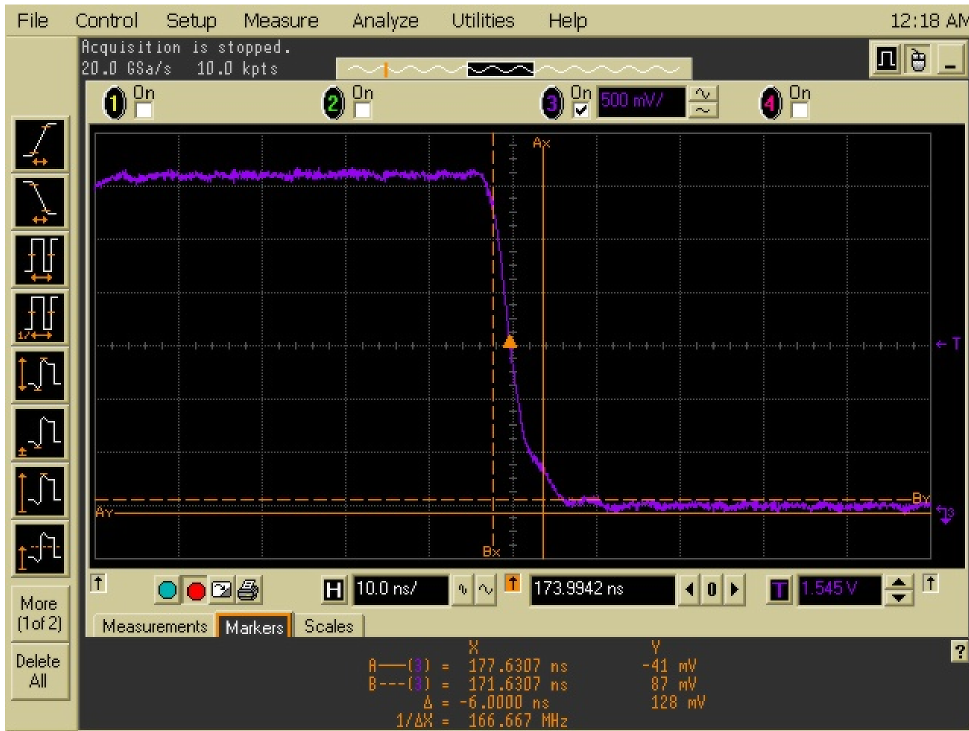

Figure 32 - ASIC 3 falling edge

Testing without the SDTAP Buffer (Table 5)

This table uses Test Setup 1 but ASIC 2 is not loaded. The series termination and capacitor on the TCK, TMS, TDI, and nTRST lines going into this UUT device was removed. Consequently, all energy was fully reflected to the main TAP connector and the other two TCK branches. None of the TAP signals was filtered in this test set.

As shown in Figure 33, it is clear that reflections from the open stub (ASIC 2) are now stronger since they are fully reflected. Input pin capacitance or input impedances for receivers are not normally considered formal terminations. However, this result indicates that the presence of ASIC 2 provides some termination to the line that otherwise becomes infinite if the chip is not loaded. Reflections that were not there with ASIC 2 now appear at ASIC 1 TCK input pin. This shifts the double clocking problem from ASIC 2 to ASIC 1 when no buffering solution is applied.

Although it is not relevant for scan chain operation, a measurement at the end of the TCK stub is shown in Figure 35. It shows that reflections are still present, which now come back to the TAP connector and impact ASIC 1 input as the closest input.

Clearly the presence or absence of ASIC 2 effectively bounces reflections all around the TCK path in UUT 1. As mentioned before, this shows that the termination provided by the TCK inputs cannot be totally disregarded in boundary scan applications.

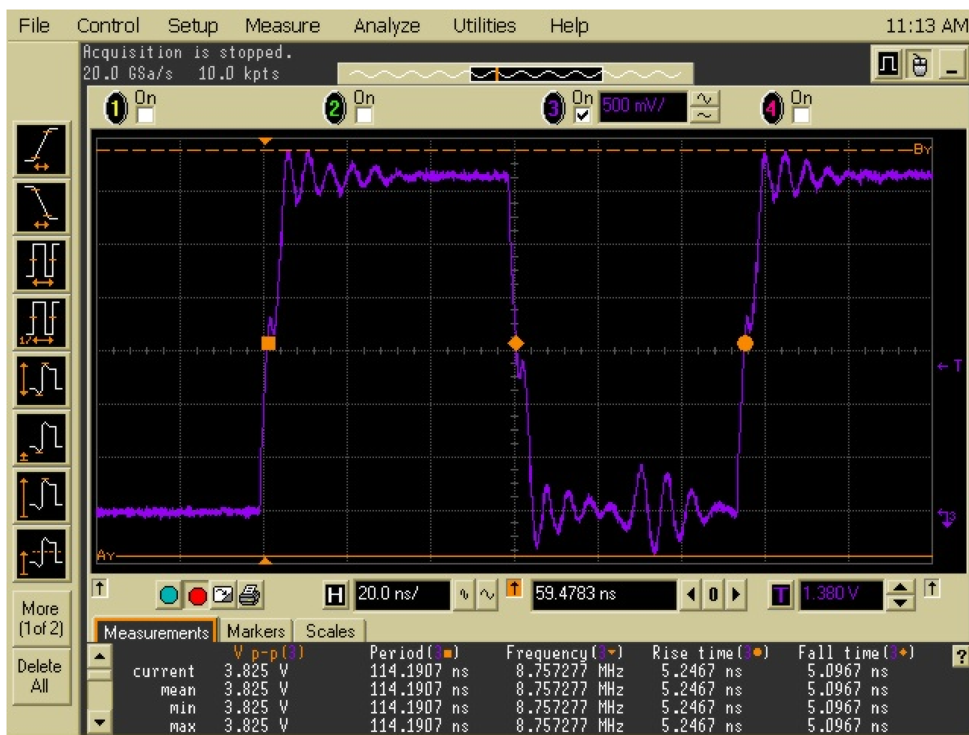

Figure 33 - ASIC 1 no filtering

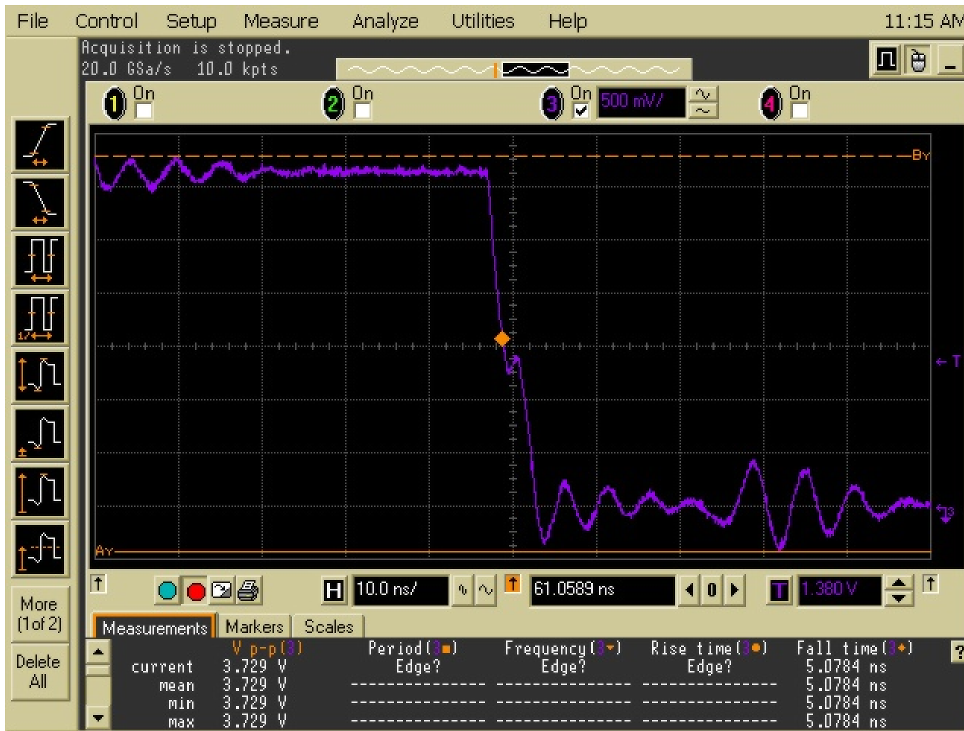

Figure 34 - ASIC 1 falling edge

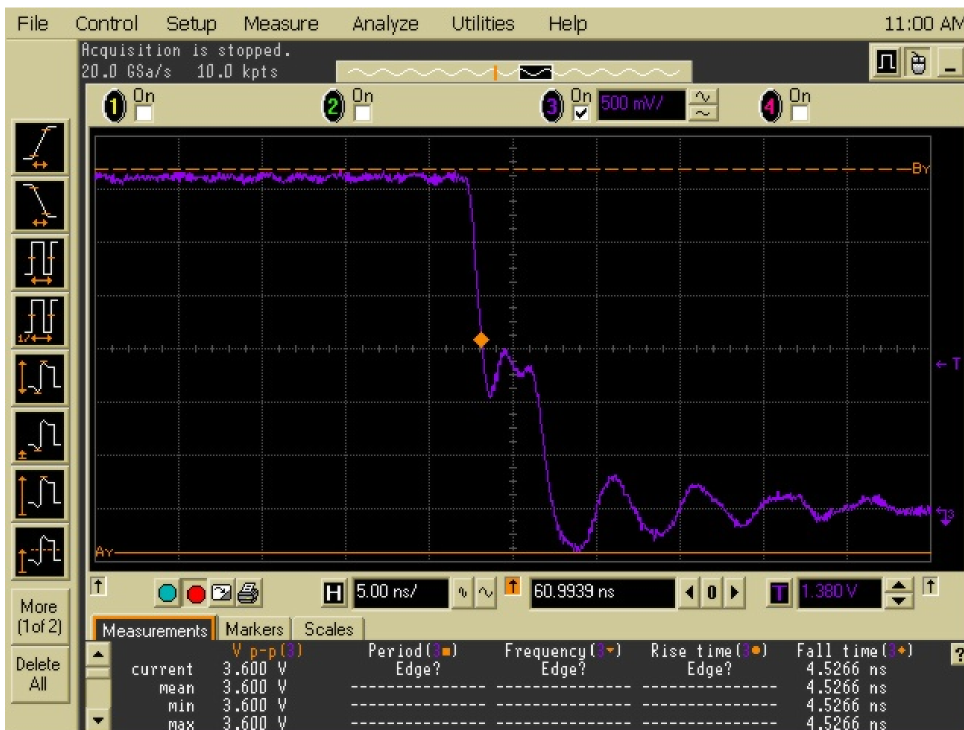

Figure 35 - ASIC 2 - falling edge

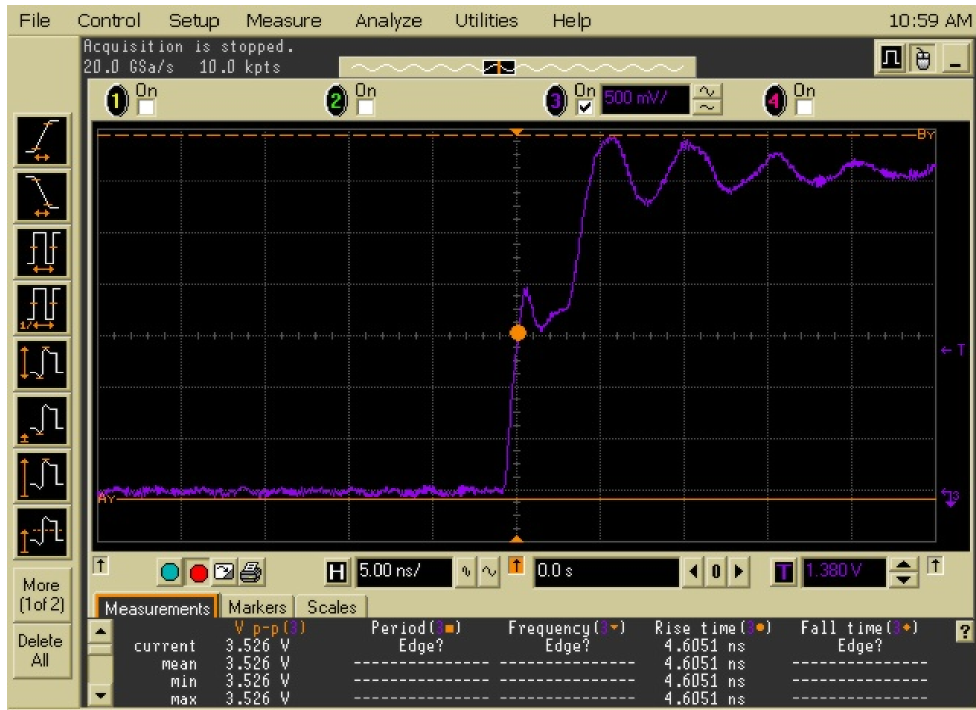

Figure 36 - ASIC 2 rising edge

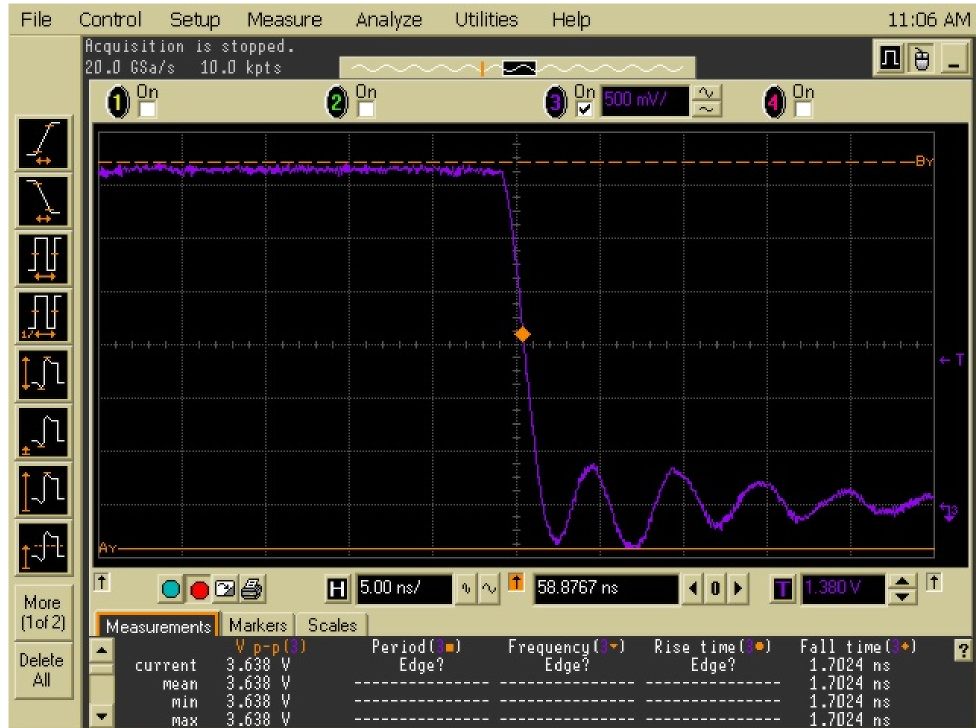

Figure 37 - ASIC 3 falling edge

Testing the SDTAP Buffer without ASIC 2 (Table 6)

This table shows UUT 1 with no ASIC 2 present and the SDTAP Buffer solution applied. The operation of the chain was very stable with this test setup. Comparison of Figures 33 and 38 renders a clear explanation. It shows that TCK now operates with monotonic edges between six to nine nanoseconds and that there is virtually no oscillation and a very flat response in time to the TCK input pin.

The same reasoning can be applied to Figures 39 and 40. Interestingly enough, the absence of ASIC 2 does not preclude the stub signal to be acceptable. This indicates that the solution proposed is fairly independent of the pin's input impedance and trace. When this result is compared to Figures 30 or 31 it looks almost identical further suggesting little impact of the presence or absence of ASIC 2 on circuit performance when the SDTAP Buffer is used.

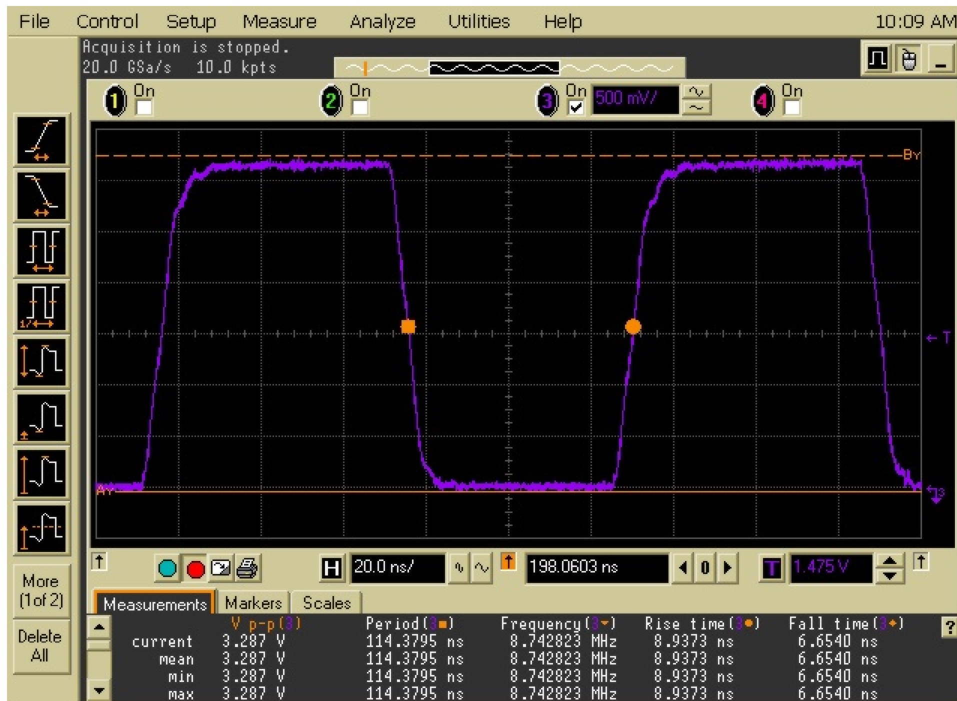

Figure 38 - ASIC 1 with filtering

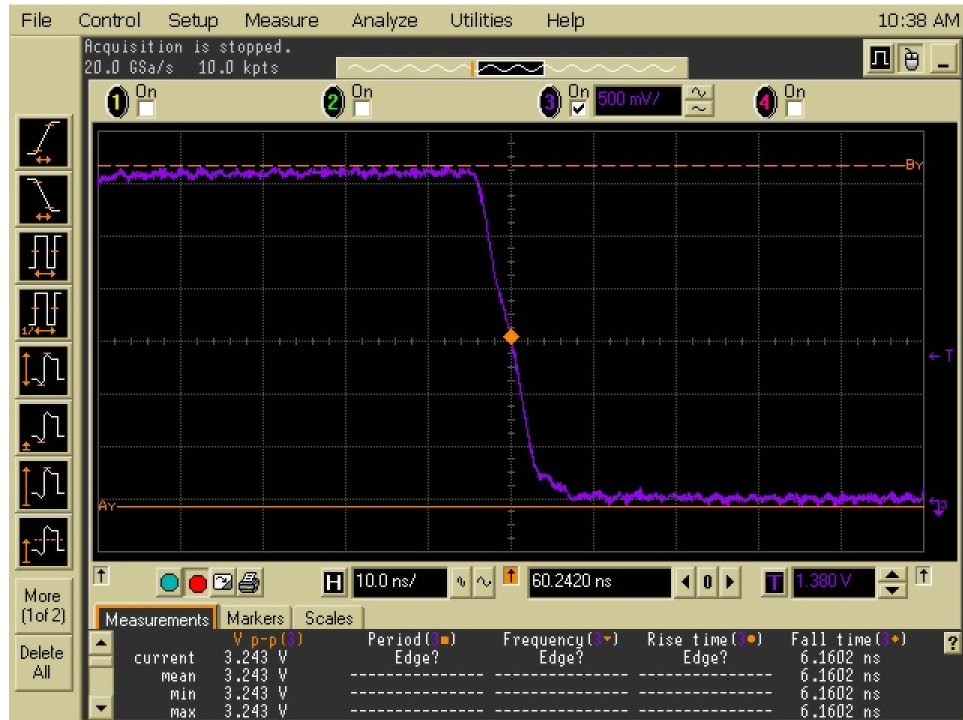

Figure 39 - ASIC 2 (not present) falling edge - with filtering

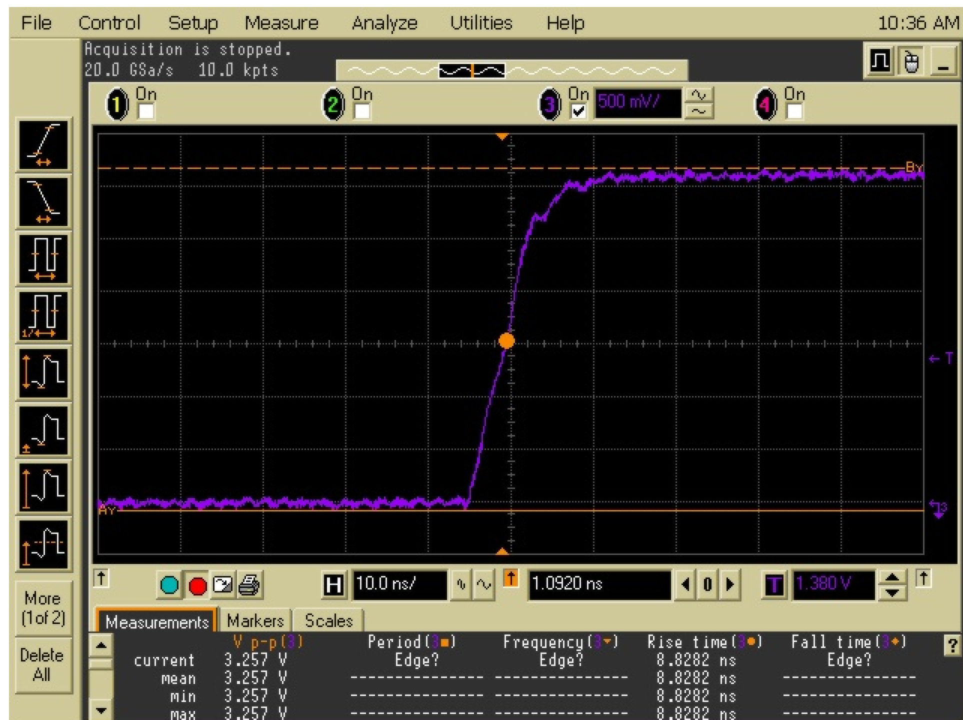

Figure 40 - ASIC 2 (not present) rising edge - with filtering

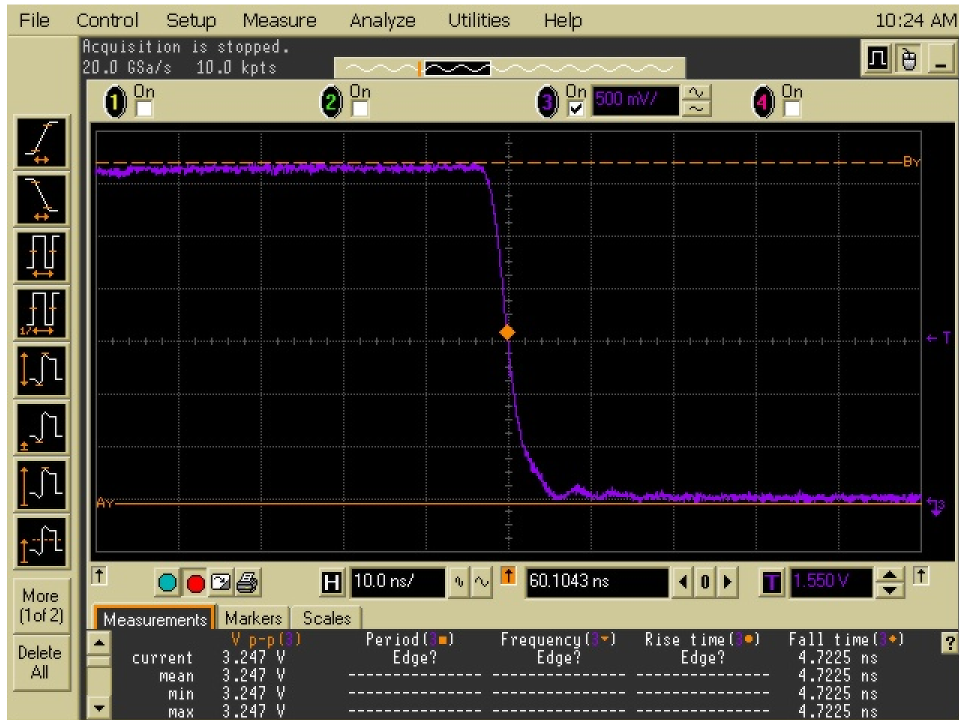

Figure 41 - ASIC 3 falling edge with filtering

Testing another UUT (Table7)

UUT 1 was the trigger board for this research. However, in an effort to further test the validity of the thesis statement, other UUTs were used with and without the SDTAP Buffer solution to show its effectiveness. UUT 2 was chosen and Figures in Appendix A show the TCK path and density of UUT 2. The test bed is shown in Figure 19 as Test Setup 2; refer to this setup since the ASIC numbers refer to ASICs different than those the previous four tables.

This set of measurements represents and emphasizes the challenge that a boundary scan test engineer may face when dealing with dense printed circuit assemblies. An engineer's first reaction is to try to fix the layout. But the reality is that both boundary scan and mission mode layouts do not take the same precedence in the minds of designers

and managers. However, these two electronic circuits must co-exist and function independently.

Before the SDTAP was available UUT 2 could not be run any higher than 4.167 MHz. This was with only a 100pF on TCK at the connector. After testing with the SDTAP Buffer, it could now be run at 8.335 MHz. This result suggests that signal integrity issues may be masked as timing issues effectively deceiving a boundary scan test engineer.

Figures 42 to 45 all show unacceptable TCK signals at 8.335 MHz with lots of reflections and irregular edges. The test condition of no filtering but buffering has been shown to be not reliable and non-working. Results further suggest that to operate a boundary scan chain properly, in the presence of a pseudo-star TCK layout, it must have some form of filtering.

Figure 42 - ASIC 1 no filtering

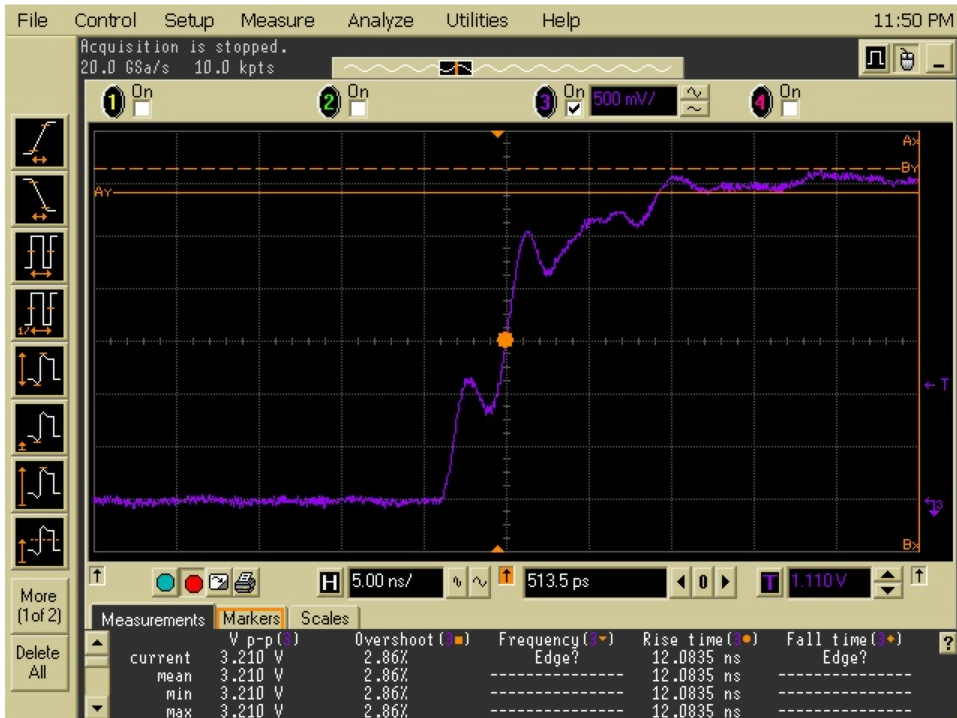

Figure 43 - ASIC 1 rising edge - no filtering

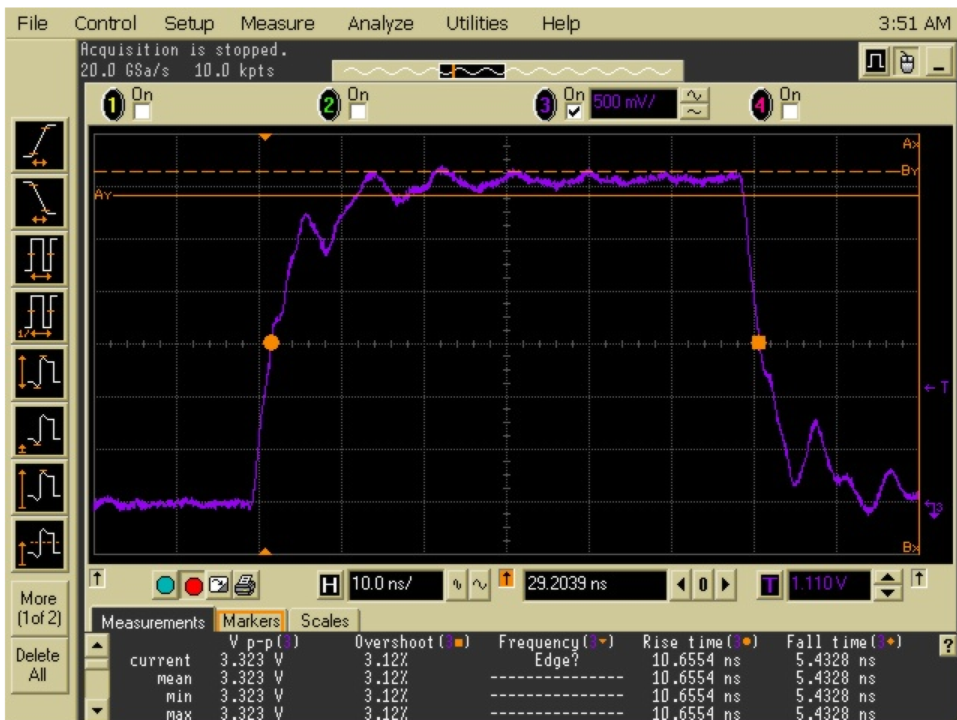

Figure 44 - ASIC 3 - no filtering - both edges

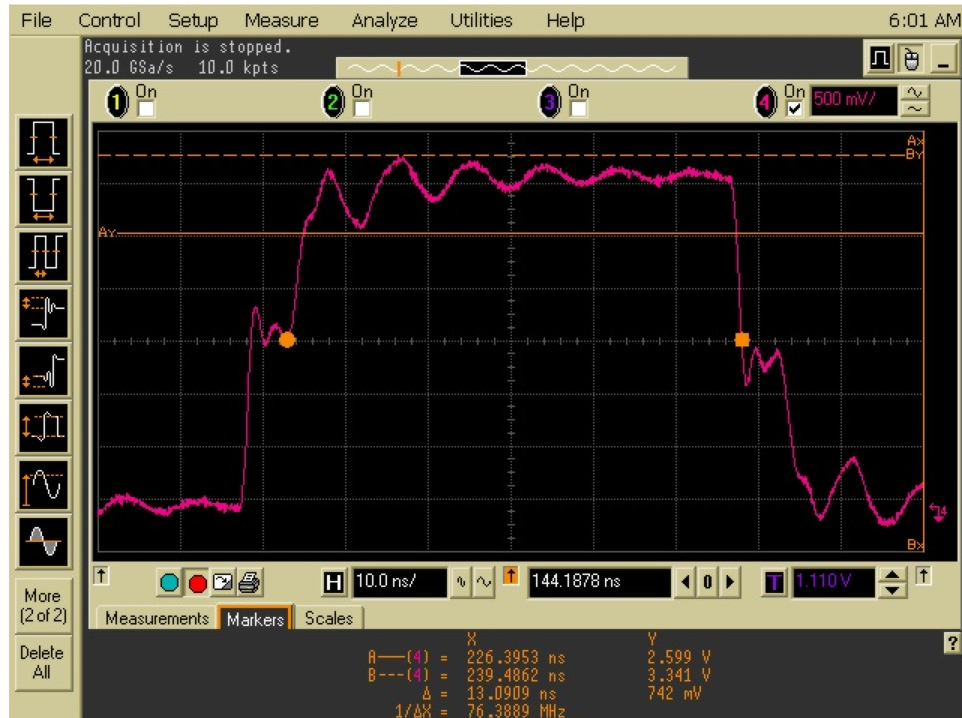

Figure 45 - ASIC 4 - no filtering

Further Validation of the SDTAP Buffer (Table 8)

This set of measurements stands in contrast to those in Table 7. The chain on UUT 2 can now be operated at 8.335 MHz in a very stable manner. Figures 46 to 49 show that even in the presence of minor signal deformities the chain can work properly because of monotonic edges. Thinking about these minor deformities, it can be inferred from the measurements that a slow smooth edge greatly increases the overall system reliability in boundary scan applications. Clearly, reflections at the threshold point will cause incorrect clocking of TCK.

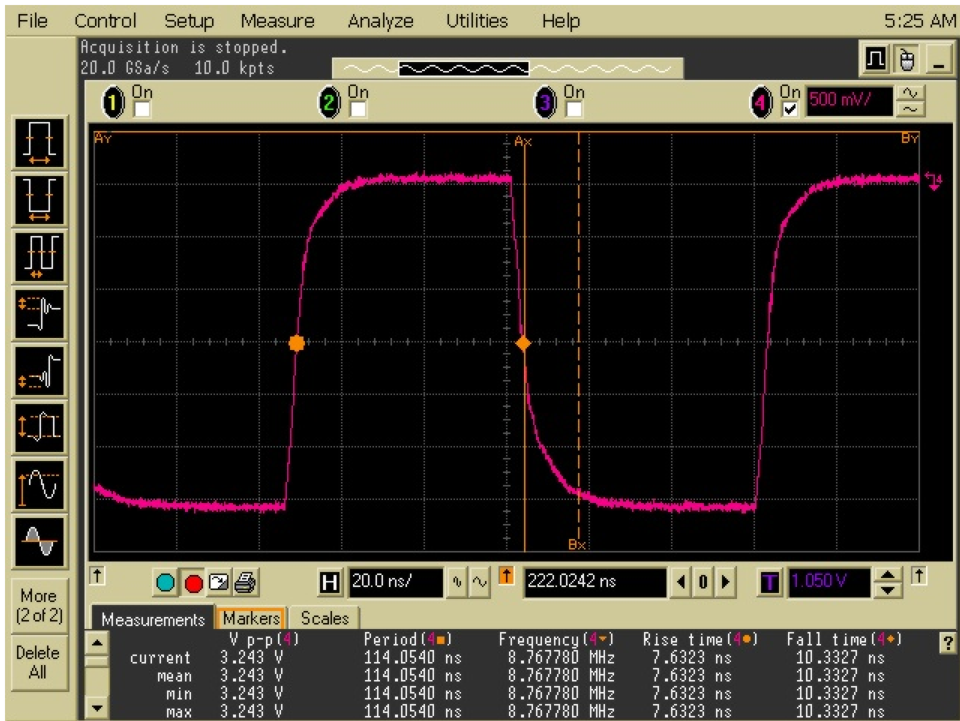

Figure 46 - ASIC 1 TCK with filtering

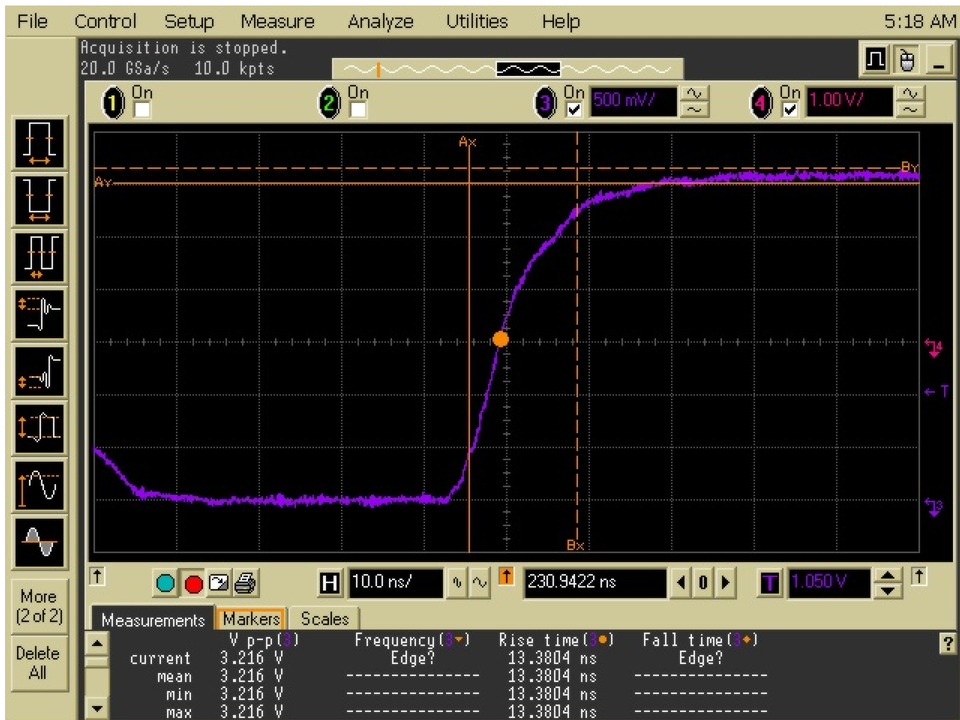

Figure 47 - ASIC 1 rising edge

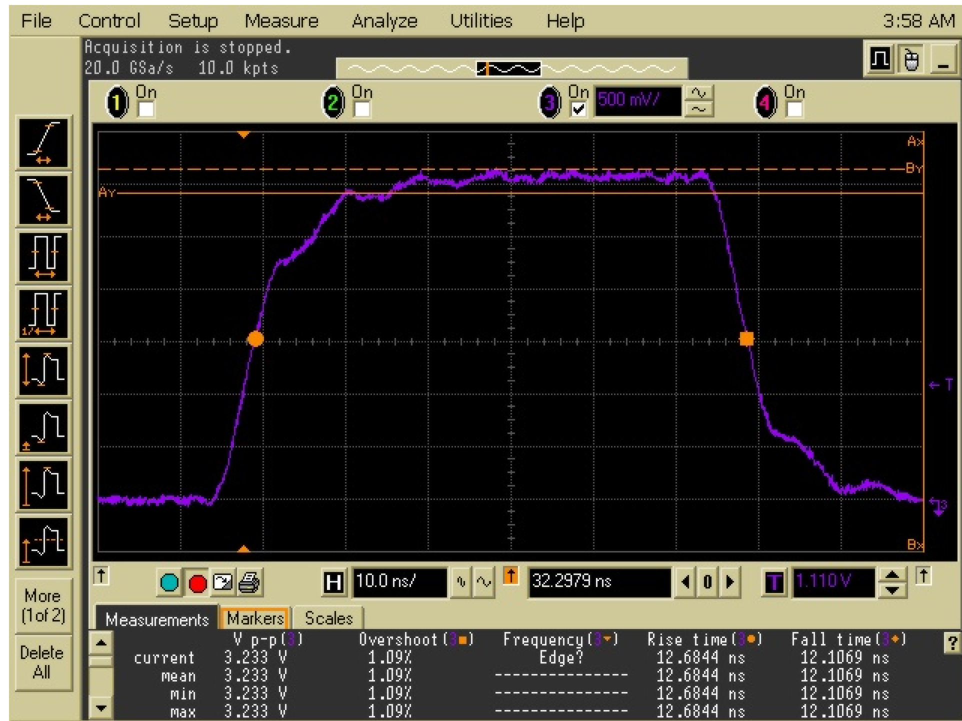

Figure 48 - ASIC 3 both edges

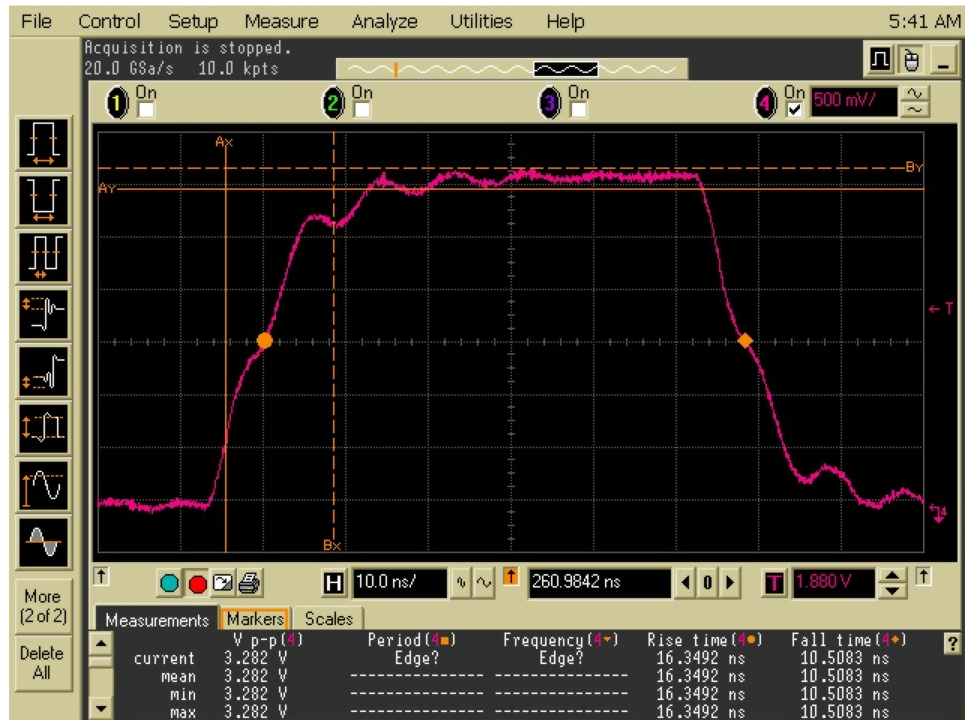

Figure 49 - ASIC 4 both edges

Other Signal Integrity Measurements

To further support these findings other measurements were considered important. An experiment was run to show the impact of a 12pF load on Test Setup 1 with scan chain configuration 2. In other words, since ASIC 2 was not loaded, this presents infinite impedance at the end of its TCK trace stub causing total signal energy reflection. The measurements in this section show the effects of this load on the TCK clock distribution path.

Figures 50 and 51 show the beneficial effects of a 12pF load placed on the TCK path for ASIC 2 and ASIC 3 on Test Setup 1. The load was placed right close the TCK input pin at the end of the stubs. When this load was placed on ASIC 2, the inflection almost disappeared and the chain rendered a passing test. Although the signal edge of this new TCK signal is not very good, it does not oscillate. This suggests that the amplitude of the inflection (118mVp-p Figure 50) is the final root cause of the original problem. This measurement is about 59mVp suggesting that this difference is what caused the double clocking of TCK. It also important to notice the rise time of the reflection: 655 ps. This reveals a frequency component of 1.528 GHz! This measurement shows the necessity for a very fast scope. Notice the 2ns/div scale while TCK was running at only 8.335 MHz.

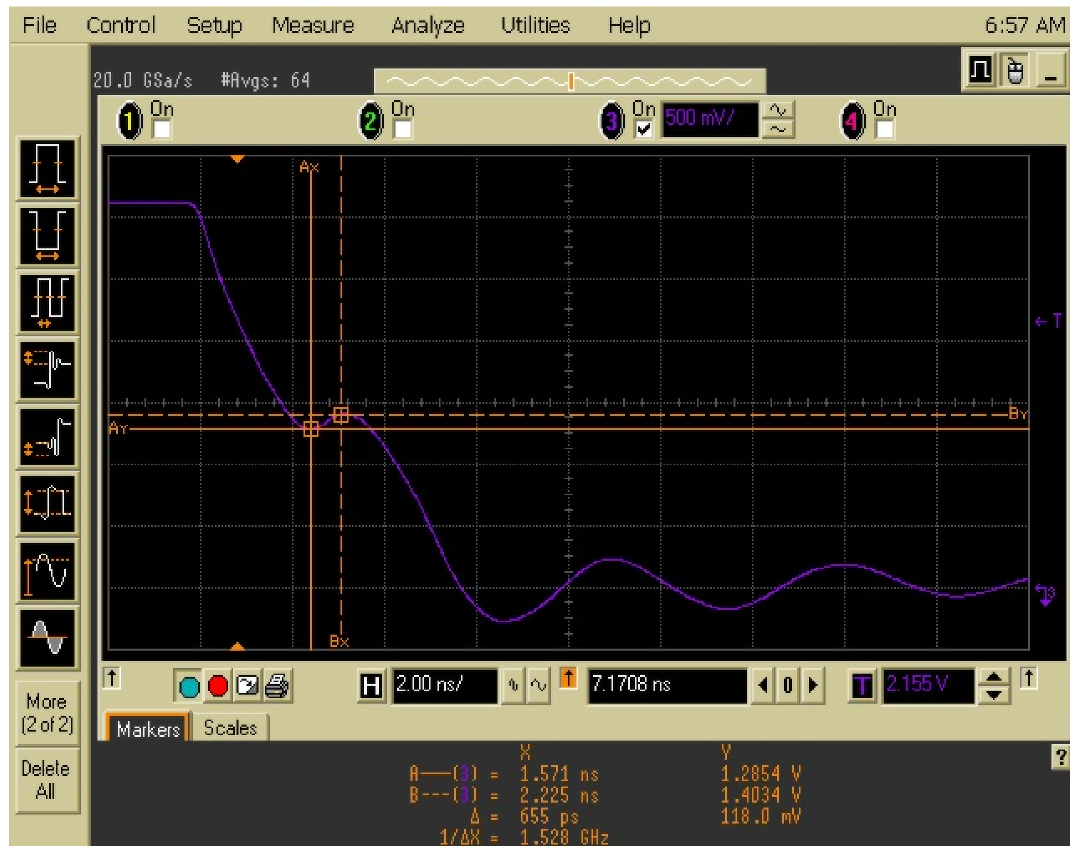

Figure 50 - ASIC 1 - Test Setup 1 - TCK falling edge with no termination

To further verify the validity of the measurement the load was removed and a regular oscilloscope probe with 10-12pF load was placed on the same two stub locations. This procedure rendered similar results when compared to measurements with the actual capacitor loaded. These results show that a simple probe can potentially mislead an inadvertent engineer to believe that there is no problem in the chain because the signals look good. Although this loading effect is widely known, it is often neglected because 8 MHz is considered a low speed signal (Johnson and Graham, High-Speed Digital Design: A Handbook of Black Magic).

Finally, the SDTAP Buffer solution was put in place showing a very clean and smooth edge achieved by placing the filtering at the TAP connector (see filtered tck – figure below).

Figure 51 - ASIC 1 - Test Setup 1 - TCK falling edges

TCK Increased Performance

Table 10 describes practical results obtained when the scan path verification test was looped at two TCK frequencies. The results in this table and those in Figure 52 suggest that a signal integrity problem may initially appear as a timing issue. The scan chain cannot respond to a <u>full</u> Scan Path Verification (SPV) test at once any longer, but it can work at 4.167 MHz with the same degraded TCK.

Table 10 - Setup 2 - 22pF only on TCK

Test Setup 2	*4.167 MHz*	*8.335 MHz*
Device ID + Bypass DR	Pass	Pass
IR Capture	Pass	Pass
Bypass	Pass	Pass
IDCODE	Pass	Fail
Boundary Register	Pass	Fail
USERCODE	Pass	Fail
nTRST	Pass	Fail

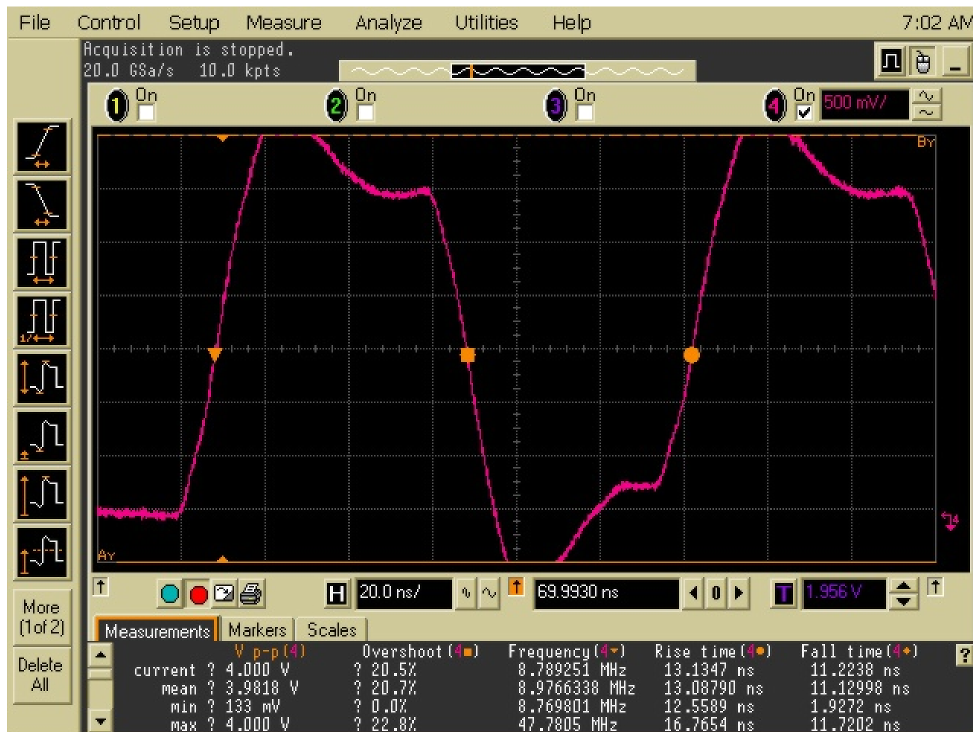

Figure 52 - ASIC 1 – Test setup 2 - Degraded TCK partially works

The SDTAP Buffer was designed to drive a large capacitor, but at the same time to obtain controlled edges close to 10ns. Figure 52 shows that the edges are too slow (~13ns) without the buffer, which is one reason to fail the scan path test at a higher TCK speed. This is because at 8.335 MHz these deformed edges are not allowing for 50% duty cycle, which is critical for any clock driven system. In this case, the TDO data stream is

not well synchronized because the falling TCK appears too soon. This can cause both data missed at the sampling point or not enough time for TDO to travel through longest path like the boundary scan register length test. If the scan path verification test is looped within these conditions the test passes with few tests, but fails as more tests are included.

Keep in mind that the hold time margin described here is relatively independent of the hold time applicable to the D-F/F itself. The reason is that hold times within new chip technologies are now around 0.5 ns and the system designer wants to sample TDO when it is more stable.

The total time for a full scan path verification test is variable depending on the number of scan chain tests applied. In other words, the chain can be exercised in different ways depending on the number of tests applied (see Figure 53). Normally, more tests simply mean more time; however, on a marginally working system (with TCK like in Figure 52) more tests may fail as the number of tests increases.

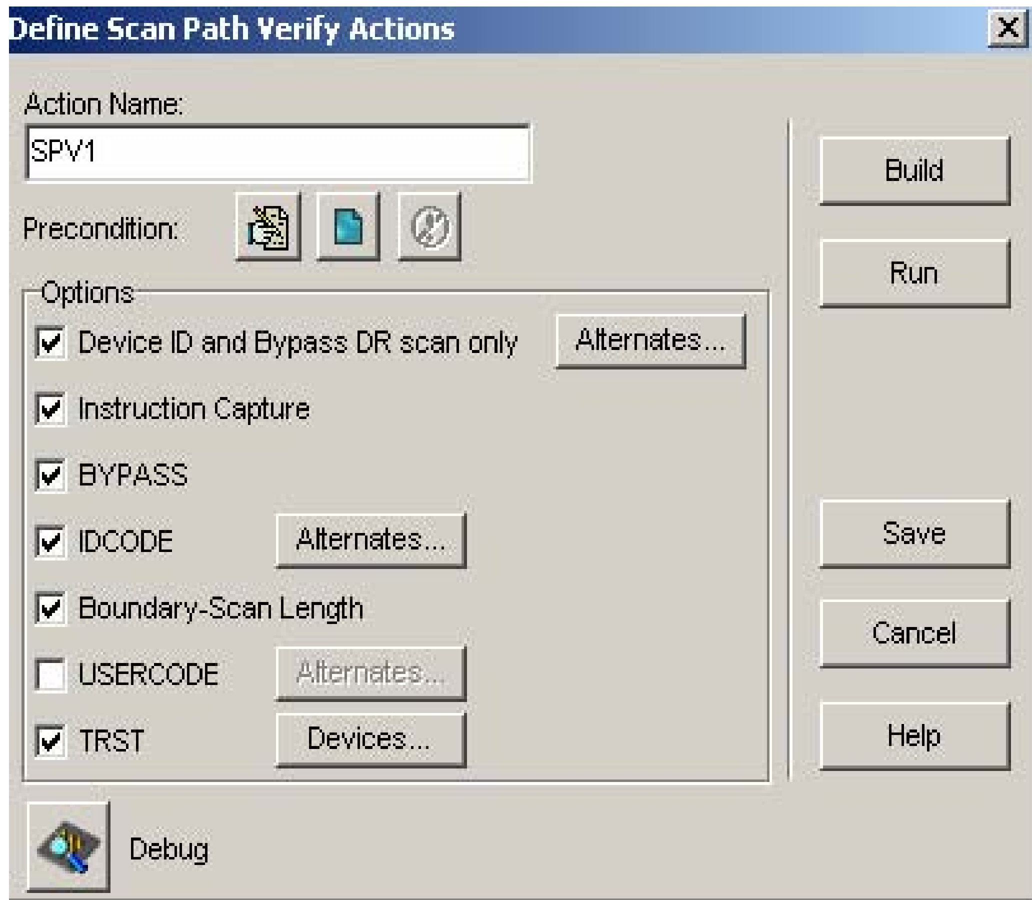

Figure 53 – ScanWorks® scan path verification tests. (Courtesy of ASSET Inter-Tech, Inc.)[3]

The maximum number of tests available is seven for the tool used in these experiments. If the number of tests applied increases and there is not enough system hold time margin because little or no filtering is used, then the scan chain test may fail some tests and pass others (see table 9). These tests apply to Figure 1 and ideally the test engineer would loop the scan path verification test with the maximum number of applicable tests. This is usually six tests as shown in Figure 53.

Please refer the Problem Context section in Chapter 1 for further details. The author believes that this is sufficient proof that the skew introduced by the buffers and

[3] ScanWorks® is a Registered Trademark of ASSET Inter-Tech, Inc.

filters it is actually beneficial to a boundary scan test system. These delays allow for increased TCK to enable reduction in test cycle time in the manufacturing floor. This conclusion comes from the practical observation that, without the SDTAP Buffer, Test Setup 2 at 4.167 MHz would pass all tests, but it will fail some tests at 8.335 MHz.

Finally, the SPV test cycle time for Test Setup 1 is about 5 seconds, which is comprised of 3 sec. setup time and 2 sec. for actual test time without the SDTAP Buffer. With the solution applied the test can now run with an increased TCK of 8.335 MHz (from 4.167 MHz) effectively reducing test time to about 1 sec. The cycle time improvement for this boundary scan test is about 20% since the test can now run twice as fast.

Analysis of Results

A theoretical basis will help to explain the results. According to Johnson and Graham, physically small circuits are considered lumped circuits whereas large circuit structures are considered distributed systems (7). Applying this definition to a boundary scan system we can conclude that a boundary scan system should be treated as a distributed system. Let's look at edge speed and physical trace lengths. The SDTAP Buffer uses a driver specified for 3ns edges. The propagation delay used on FR4 is about 180 ps/in. Using formula [1.3] (Johnson and Graham, High-Speed Digital Design: A Handbook of Black Magic) shows that a rising edge propagating along an FR4 trace is 2.77 in., which is six or more times smaller than the traces shown in Figures 4 and 6. Even if the edge is slower, the relationship still applies because of the 12 in. trace going to ASIC 3 in UUT 1.

$$l = \frac{T_r}{D} \qquad [1.3]$$

where l = length of rising edge, in.
T_r = rise time, ps
D = delay, ps/in.

Distributed systems then have several options for line terminations; those are: source termination, end termination, and both ends termination. Before choosing any of these solutions, it is very important to understand the constraints surrounding the use of the IEEE 1149.1 Standard within any embedded system. Based on the current results, an engineer may attempt to solve the problem using any of these known techniques. However, the required solution needs to render a better system level performance, higher reliability, shorter time-to-market, and meet financial objectives. The realities and complexities of today's electronic assemblies demand a solution that is the more cost-effective, fastest, and less disruptive solution possible. It requires a solution that is re-usable and robust in the manufacturing floor. The author believes Figures 4 and 6 become very important for this part of the discussion.

Below are some real estate constraints found when laying out a boundary scan chain in today's complex electronic circuits (in no particular order):

- No physical spaces open to daisy-chain the TCK, TMS, and nTRST paths or to easily layout any other possible trace corrections.
- A star configuration most likely will not render legs of the same length as has been shown in Figures 4 and 6.
- More careful trace layout is usually not possible, and even so, it will render a path with perhaps an excessive number of vias.

- The board designer *must* connect his/her design first and then try to find a path for the TAP signals.
- When TAP signals are multiplexed with functional signals, there is no way to control the layout of the TAP interface.
- The program manager does not want to spend money, time, or resources on TAP buffering on-board to improve the clock distribution of a signal that will be used for only seconds during the life of the board.

These are real challenges for the test engineer in charge. The solution to this problem must be given within the frequency range in table 2 (max. of 8 MHz in this case) and meet all of the above requirements. The author proposes that the source termination provided by the SDTAP Buffer is adequate for the conditions described in this work and in other similar designs. It shows that the SDTAP Buffer meets and perhaps exceeds the above requirements in all aspects. For example, the buffer circuit (NC7WZ17) from Fairchild was chosen because of its capability to drive large capacitive loads without oscillations, source enough current, Schmitt trigger features. It was also chosen for its use of fast 3.7ns edges. It is certainly easier and more practical to have fast edges and slow them down than to do the opposite.

When connected to the UUT, the SDTAP Buffer completes a source terminated line where the termination is very close to the driver (Johnson and Graham, High-Speed Signal Propagation: Advanced Black Magic). Also, the load impedance (5pF) is much larger than the transmission line impedance for each separate branch in UUT 1 (Figure 3). All these characteristics are shown in Figure 54 and through practical measurements it has been shown that the presence or absence of these end loads certainly has an effect on

the quality of the TCK signal. It is clear that any high frequency reflections coming back towards the connector see a low pass filter (33ohm-100pF). This is what reduces the reflections. The 4.7ohm-100pF filter actually renders slower edges.

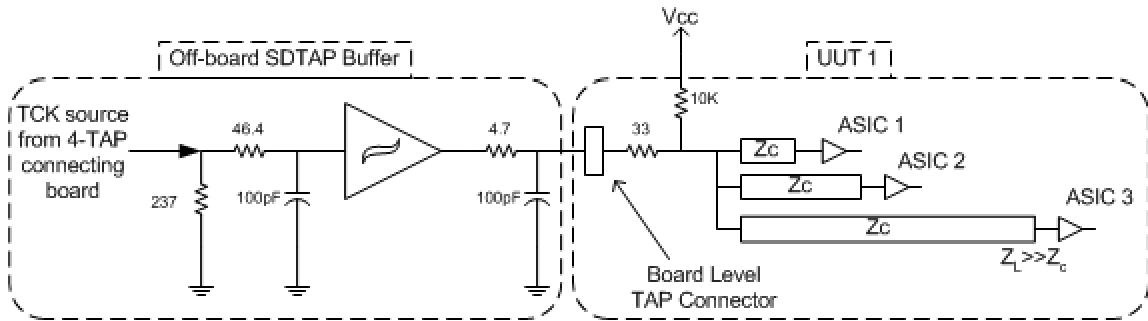

Figure 54 – System model for TCK distribution + Test Setup 1 (passing model)

The SDTAP Buffer solution is based on a system approach supported with simple DFT guidelines as shown in Figure 8. This renders reasonable assurance that the system will work when put together for final testing on the manufacturing floor. The author believes that all previous measurements and explanations clearly support the thesis statement.

The model for the non-passing system is depicted in Figure 55. This model was not able to contain the reflections coming back from the TCK lines. The result indicates that the source impedance, in this case, needs to have an imaginary component as opposed to resistive only. The reader can easily see that any high frequency reflections coming back towards the TAP connector will not be filtered because there is no low pass filter.

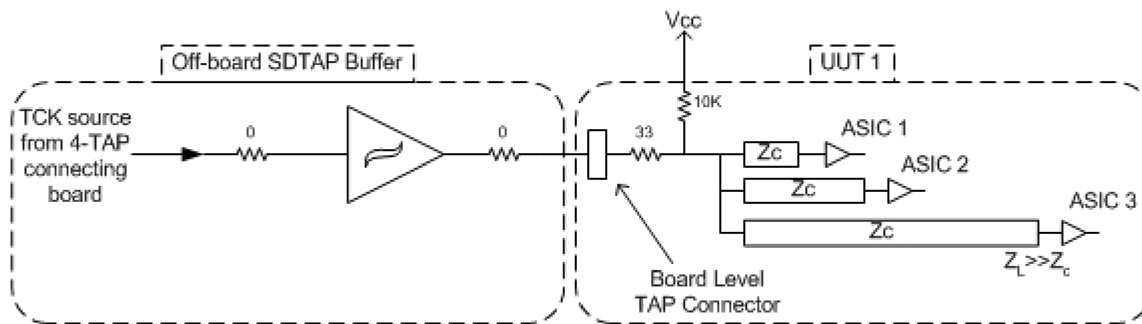

Figure 55 – System model for TCK distribution + Test Setup 1 (non-passing model)

Finally, the timing measurements were provided to clear any questions about the possible negative effects of introducing these delays in TAP interface. It has been shown that these delays are actually beneficial (always within the frequency range) because they improve the system hold time margin. Consequently, this ensures that the TDO signal is sampled after it has been there for a long time and most oscillations have disappeared.

CHAPTER 5

Summary and Conclusions

This research has shown the combined use of general engineering principles to enable a simple practical solution. It is an example of applied technology to provide a solid foundation for both experts and new engineers in the use and application of the IEEE 1149.1 Standard across the board test industry. It explains how to ensure that almost any PCA featuring embedded boundary scan technology can be shielded from the effects of the test system. At the same time, it cleans up the TAP signals almost independently of poor trace layout due to board density or chain configurations. Finally, it presents a clean consistent signal back to the tester for diagnostics and analysis even in the presence of less than ideal TCK signals.

For a boundary scan chain, at the PCA level, this work presents a solution to a signal integrity problem and an improvement to the system hold time margin. It also offers improvements to a scan chain speed performance in marginally working systems. This in turn can improve manufacturing throughput due to higher TCK speeds. The difference with the commonly accepted solution of a single 15pF to 100pF capacitor on TCK is that termination of all four TAP signals going into the UUT is provided to ensure balanced operation of the chain. By design philosophy, the filtering was not included in the return data path (or TDO signal) coming back to the test system. This was done to minimize the impact of time delays to the overall test system.

The solution proposed is in the form of a PCA that also solves mechanical problems at the fixture level. These mechanical problems are beyond the scope of this work, but were also important factors during design and layout. The SDTAP Buffer went through four revision cycles before it was considered finalized. The released revision was installed in a high-volume production system successfully and reliably.

All questions posed in Chapter 1 have been answered and all requirements met with the use of the SDTAP Buffer. Within the 8.335 MHz range, the buffer performs at acceptable levels meeting reasonable expectations for chain reliability. This work represents the sum of about two years of observations and consecutive experiments to thoroughly understand the problem at hand.

If applied correctly, this solution removes the hardware concern from the test engineer and allows for full exploitation of the powerful structural test system: the embedded boundary scan engine. Moreover, the detached approach of the SDTAP Buffer reduces the possibility of chain malfunction due to defective filters or missing parts during board manufacturing. By keeping the scan path with the minimum number of components (i.e., resistors, buffers, etc.), the possibility of incorrect loading or defective buffers is minimized.

The concepts used during this research are not new; however, the author feels that the combination of the different concepts used, as they are presented in this research, is an addition to the board test industry body of knowledge. It represents the use and careful combination of many details and years of experience to produce a solution that would be suitable under many constraints. It is the author's hope that this work helps in the

development of a promising technology that is already the basis for many more standards like the IEEE 1149.4, 1149.6, 1532, P1500 and others.

Future Research

Future research should include testing of Stang and Dandapani's conclusions at the board level with ASICs that follow his DFT advice (749). It would be interesting to see how a chain with a potentially slower 1149.1 ASIC implementation behaves in a star vs. daisy-chain configuration at the board level. The author certainly had a pseudo-star configuration, but it didn't have ASICs that followed Stang and Dandapani's advice. In theory, a star configuration is faster and more reliable than a daisy-chain layout for all control signals.

Another approach to layout is the trace length matching of each TCK branch from a resistor placed close to the board level TAP connector. During this research, this was deemed almost an impossible task because of the board trace density problem, but theoretically it should ensure a totally synchronized UUT chain operation. Such an approach still requires signal integrity analysis and testing of the solution proposed here.

Due to the TCK speed limit of 8.335 MHz, there is some room to study the effects of this solution at higher speeds. A proposed new limit could be 16.670 MHz, which will require an entire new set of measurements and analysis. This is because the 10ns edges used in the solution proposed in this work may become a concern on a 30ns half-cycle TCK.

BIBLIOGRAPHY

Baang, S. X., and Liang S. H. "Design and Implementation of a Jtag Boundary-Scan Interface Controller." Test Symposium, 1993., Proceedings of the Second Asian, 1993. 215-18.

Bhavsar, D. "A Method for Synchronizing IEEE 1149.1 Test Access Port for Chip Level Testability Access." VLSI Design, 1998. Proceedings., 1998 Eleventh International Conference on, 1998. 289-92.

Caldwell, B., and T. Langford. "Is IEEE 1149.1 Boundary Scan Cost Effective: A Simple Case Study." Test Conference, 1992. Proceedings., International, 1992. 106.

Dellecker, R. "Boundary-Scan Bursts into the Modern Production Facility." Aerospace and Electronic Systems Magazine, IEEE 16.6 (2001): 21-24.

"IEEE Standard Test Access Port and Boundary-Scan Architecture." IEEE Std 1149.1-2001, 2001. I-200.

Johnson, H., and Graham, M. High-Speed Digital Design: A Handbook of Black Magic. Upper Saddle River, NJ 07548: Prentice-Hall, Inc., 1993.

Johnson, H., and Graham, M. High-Speed Signal Propagation: Advanced Black Magic. Upper Saddle River, NJ 07548: Prentice Hall, Inc, 2003.

Kajitani, H., et al. "Practical Test Generation with IEEE 1149.1 Boundary-Scan." ATE & Instrumentation Conference. Anaheim, CA, 1992. 10-17.

Parker, K. The Boundary-Scan Handbook Second Edition Analog and Digital. Second ed: Kluwer Academic Publishers, 1998. Fifth.

Stang, D., and Dandapani R. "An Implementation of IEEE 1149.1 to Avoid Timing Violations and Other Practical in-Compliance Improvements." Test Conference, 2002. Proceedings. International, 2002. 746-54.

Texas, I. TexasInstruments. Embedded Test Bus Controllers - IEEE Std 1149.1 (Jtag) Tap Masters with 8-Bit Generic Host Interfaces. April 2002. March 2004 2002. PDF Datasheet. Texas Instruments, Inc. Available: http://focus.ti.com/lit/ds/symlink/sn74lvt8980a.pdf. 05-28-2004 2004.

Filename: Bscan Thesis BEFORE PDF CREATION 1
Directory: C:\Documents and Settings\eft\My Documents\Working Folder
Template: C:\Documents and Settings\eft\Application Data\Microsoft\Templates\Normal.dot
Title: BOARD LEVEL PROPAGATION DELAY TIMES FOR SMALL EMBEDDED BOUNDARY-SCAN DESIGNS
Subject: BOUNDARY SCAN TECHNOLOGY USE
Author: Efren J. Taboada
Keywords:
Comments:
Creation Date: 11/3/2004 6:56 PM
Change Number: 17
Last Saved On: 11/3/2004 8:02 PM
Last Saved By: Efren Taboada
Total Editing Time: 42 Minutes
Last Printed On: 11/3/2004 8:03 PM
As of Last Complete Printing
Number of Pages: 104
Number of Words: 18,752 (approx.)
Number of Characters: 90,950 (approx.)