AFRL-RV-PS-
TR-2011-0152

AFRL-RV-PS-
TR-2011-0152

AC COUPLED INTERCONNECT FOR LOW POWER SPACEBORNE ELECTRONICS

Paul D. Franzon and Steve Lipa

North Carolina State University
Campus Box 7911
Raleigh, NC 27695-7911

18 January 2011

Final Report

AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776

DTIC COPY
NOTICE AND SIGNATURE PAGE

This report was cleared for public release by the Air Force Research Laboratory RV Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RV-PS-TR-2011-0152 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT

//signed//
CHETH OUCH
Program Manager

//signed//
B. SINGARAJU, Ph.D
Deputy Chief, Spacecraft Technology Division
Space Vehicles Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
18-01-2012	Final Report	25-Sep-2006-30-Sep-2011

4. TITLE AND SUBTITLE
AC Coupled Interconnect for Low Power SpaceBorne Electronics

5a. CONTRACT NUMBER
FA9453-06-2-0350

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
63401F

6. AUTHOR(S)
Paul D. Franzon and Steve Lipa

5d. PROJECT NUMBER
5108

5e. TASK NUMBER

5f. WORK UNIT NUMBER
299104

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
North Carolina State University
Campus Box 7911
Raleigh, NC 27695-7911

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory
Space Vehicles Directorate
3550 Aberdeen Ave., SE
Kirtland AFB, NM 87117-5776

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RVSE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-RV-PS-TR-2011-0152

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited. (377ABW-2012-0033, dtd 13 Jan 2012.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The goal of this project was to provide a demonstration of issues related to laminate packaging, and demonstration of a socket system and a connector system using inductively and capacitively coupled AC interconnect. In addition, we were tasked to deliver a test module for a test in near earth orbit on the TacSat-3 spacecraft. The vast majority of the goals of the project were achieved.

15. SUBJECT TERMS
Space Vehicles; High Bandwidth I/O; Capacitively Coupled; Interconnects; AC Coupled; transformers; inductor

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Cheth Ouch
Unclassified	Unclassified	Unclassified	Unlimited	70	**19b. TELEPHONE NUMBER** *(include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. 239.18

(This page intentionally left blank)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

FOREWORD

Government Reviewer's Note: This report contains several shortcomings that have been identified during AFRL review. The principle problem is that there are several unsubstantiated statements made within the report. Beyond that, the document is not sufficiently complete to allow an interested party to repeat the activities executed by North Carolina State University (NCSU).

The work presented here was performed under a U.S. Air Force Grant to NCSU. Consequently, AFRL is unable to correct the shortcomings for the following reasons. First, under the terms of the Grant the authors were only required to submit a Technical Report but were not required to revise the report to correct flaws identified by AFRL. While the authors did make some changes at the government's request, it became apparent that the changes required to meet the AFRL standards would not be forthcoming. Second, in the areas where insufficient details were provided (i.e., claims were unsubstantiated), the authors had not provided sufficient information to the government, nor could they be compelled to do so, to allow AFRL to fill in the missing details. Therefore, there is no way the government can adequately complete the report.

While we would prefer to produce a report that reflects the high technical standards of AFRL, it is apparent that without the active support of the authors, it is not possible to correct the shortcomings identified above.

(This page intentionally left blank)

1.0 INTRODUCTION

This report describes the work done under project FA9453-06-2-0350, described in our 2006 proposal "System Technologies for AC Coupled Interconnect for Low Power SpaceBorne Electronics," submitted in response to Broad Agency Announcement VS-06-01 from the Air Force Research Laboratory Space Vehicles Directorate.

The overall goals of the project, as stated in the proposal were :

1. Demonstrate, to the level of commercial acceptance, a capacitive coupled chip-package integrated structure, operating at 6 Gbps per pin and with a pitch of 65 µm.
    - 6 Gbps/65 µm in a 0.18 µm die on a silicon package.
    - On a laminate package, our objective is to demonstrate 3 Gbps with a 100 µm pad pitch.
2. Demonstrate, to the level of commercial acceptance, an inductively coupled socket, operating at over 1 Gbps and with a pitch of 1.3 mm.
3. Build a space-capable demonstration of a system using AC Coupled Interconnect (ACCI), capable of being flown in a United States Air Force (USAF) test rocket.

The proposal set out seven tasks which were intended to achieve these goals. Each task was comprised of several elements. The vast majority of these elements were achieved. A very small minority of the elements were not completed for reasons that were outside of our control (specifically with sub-contractors). A summary of the tasks and what was achieved is as follows:

- **Task 1. Mechanical and Electrical Demonstrator for a Laminate Package**. The intent of this task was to build a demonstrator for a large capacitively coupled die. Due to cost reasons, the die is a demonstrator, not functioning silicon. This task was only partially achieved.
- **Task 2**. **Circuit Demonstrator with a Laminate Package.** The intent of this task was to build a complete circuit demonstrator showing capacitively coupled signaling over long channels built using laminate packages. This task was achieved.
- **Task 3**. **Capacitive Connections using embedded board capacitors.** The intent of this task was to demonstrate how the circuit principles established in capacativie connections could be applied to structures with capacitors embedded in the Printed Circuit Board. This task was achieved.
- **Task 4. High Density Capacitively Coupled Socket.** The intent of this task was to demonstrate a representative socket structure using capacitively coupled elements. The circuit structures to support this task were demonstated in combination with Task 3.
- **Task 5. High Density Inductively Coupled Socket.** The intent of this task was to demonstrate a representative socket structure using inductively coupled elements. This was demonstrated.
- **Task 6. Inductively Coupled Connector.** The intent of this task was to demonstrate a representative connector structure using inductively coupled elements. This was demonstrated.

- **Task 7. Space Experiment.** The intent of this task was to build a module for deployment in a satellite experiment run by Microsat systems. The module was delivered but we never obtained any results from Microsat systems.

In the following sections we will summarize the work done toward completing these tasks and the results achieved.

2.0 TASKS

2.1 Task 1. Mechanical and Electrical Demonstrator for Laminate Package.

The objective of this task was to show that a laminate capacitively coupled assembly can be built with large die (die sizes up to approximately 28 x 28 mm). This represents die sizes larger than can currently be built using flip-chip solder-bump technology. By showing this, we would enable companies such as IBM (International Business Machines) to integrate larger die into their systems than they are currently able to do. (IBM is a customer of Endicott Interconnect Technologies [EIT]). This structure also permits very high pin counts, both for the direct current (DC) pins in the center and the ACCI pins around the edge. This capability is also of interest to processor companies such as Intel, as it permits the construction of cost effective high pin count packages.

The plan was for North Carolina State University (NCSU) to cooperate with Research Triangle Institute (RTI) and EIT to develop test structures featuring capacitively coupled chips installed onto HyperBGA packages which would be bump attached to RTI substrates. Figure 1 shows the general idea:

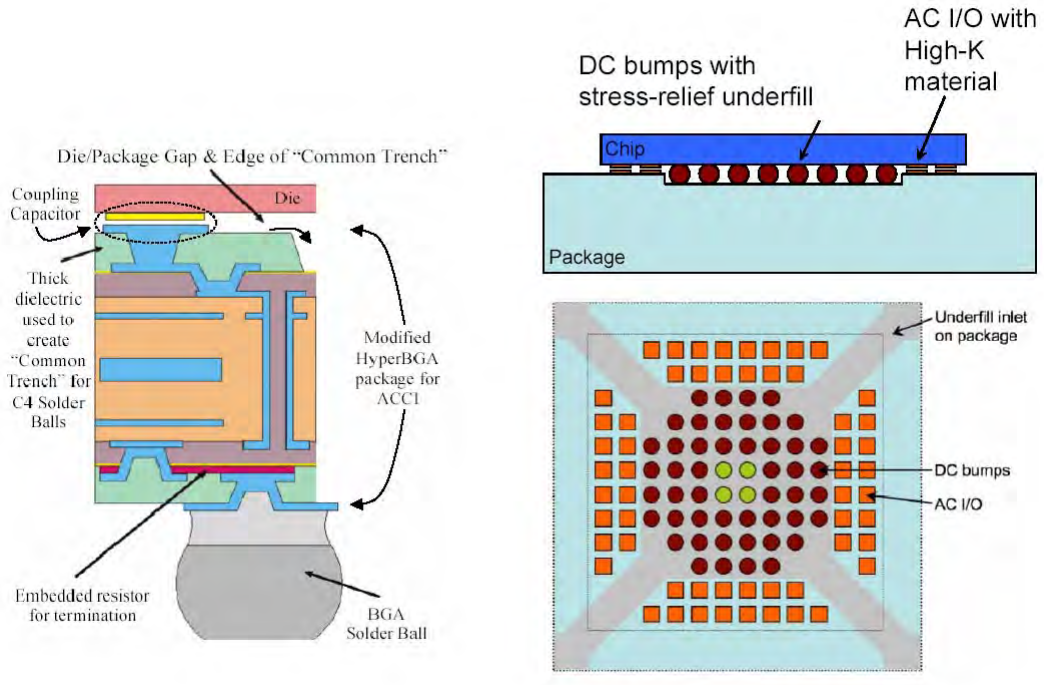

Figure 1 Proposed Test Structures

Our goal was to demonstrate a structure like that in Figure 1 with capacitor pads at a pitch approaching 65 μm. The overall approach to this task was to use a high-K material developed by Angus Kingon and Taeyun Kim in Materials Science at NCSU as an "underfill" that would replace the air gap that would otherwise appear between the chip and the package in the "coupling capacitor" shown in Figure 1. The package is a laminate. Since large chips are very expensive , the chip shown in the top of the figure was a simple structure fabricated on glass, so as to emulate a large chip. Key elements in achieving a high capacitance are (1) that this underfill leaves no airgaps, and (2) the laminate structure have a smooth surface so that the dielectric gap is thin. It is for the second reason that EIT was included as a sub-contractor, as they had a suitable package.

We also worked with EIT to develop some test boards to determine how suitable their process was for the project. Figure 2 shows a typical cross section:

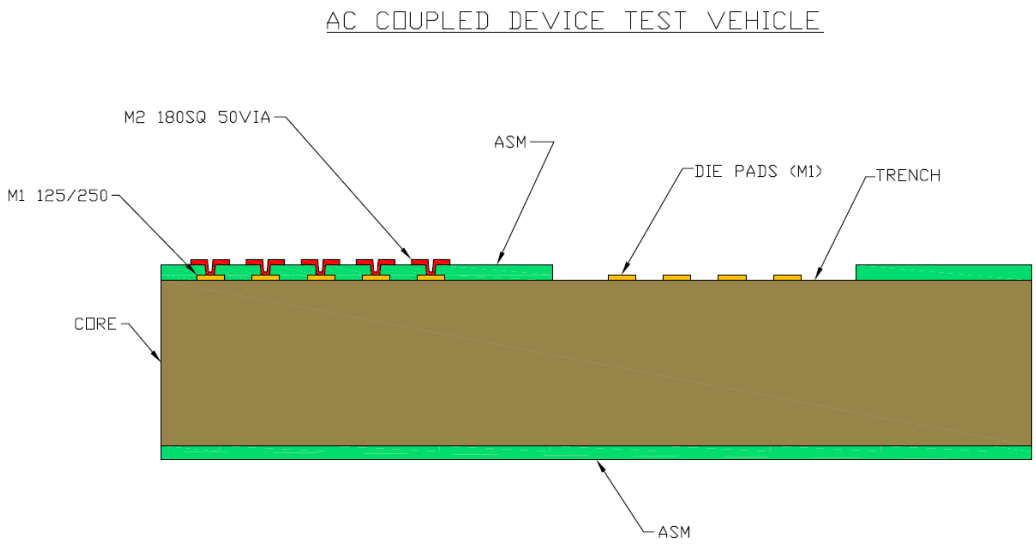

Figure 2 AC Coupled Device Test Vehicle

Our initial test samples from EIT were promising. All the vias we tested worked, leading us to choose 65-μm via diameters for the final run. In matching the separation limit of the samples, we also reduced the bump array sizes in our final run by 1-mm. Figure 3 shows a photomicrographic close-up of the top layer metal on one of the samples:

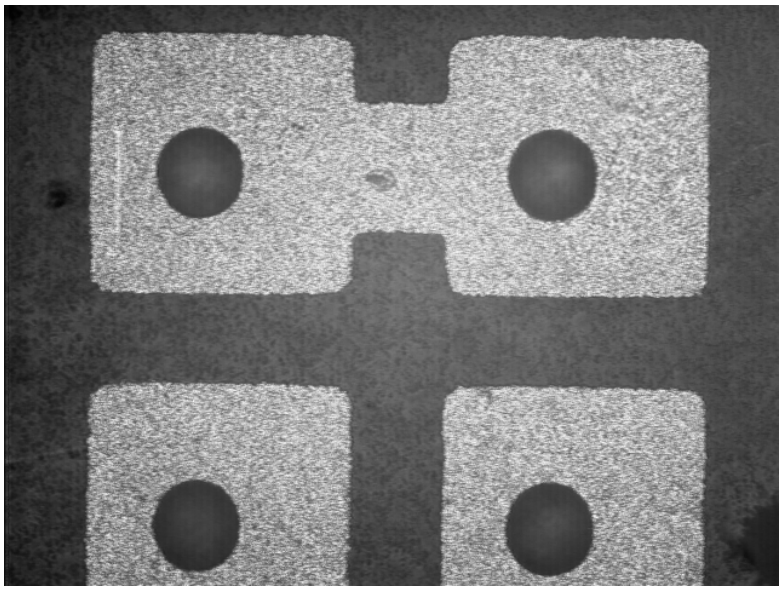

Figure 3 Top layer metal of EIT test sample

We observed an acceptable level of surface roughness in the samples, approximately 2 to 3 microns variation across a 5mm scan. We consistently observed approximately 1μm tilt in the surface flatness per millimeter across 5mm scans.

After providing these early samples it was clear that EIT was not interested in building the full structure shown in Figure 1. Unfortunately, commercially driven vendors are driven by large purchase orders, not modest R&D subcontracts. After numerous meetings with no progress, we terminated EIT's subcontract and switched to turn-key printed circuit board (PCB) providers. This meant accepting coarser features and increased surface roughness.

At this point a testbed was created. Our PCB vendor (Sunstone) was able to provide processes that allowed us to create capacitor plate features down to our smallest previously-agreed-upon size of 200-um, including wiring out to probe pads. We only needed a single metal layer for this board. The mating chip-side process was similarly reduced to a single-layer metal, patterned with similar feature sizes and deposited onto a glass slide in our cleanroom. Incorporating a transparent substrate such as simple glass allowed us to easily simulate a flip-chip alignment and die-attach process as well.

For all of our experiments the high-K layer was patterned onto a board using a screen printing method. This allowed us to cover entire quadrants of the board, while preventing the material from getting into the underfill ports and central solder-bump well. This simplified the patterning process; however, the material volume as it flowed through the mask, and the resulting drift, required us to characterize the necessary mask bias. The first-run screen mask included four test sites as shown in Figure 4.

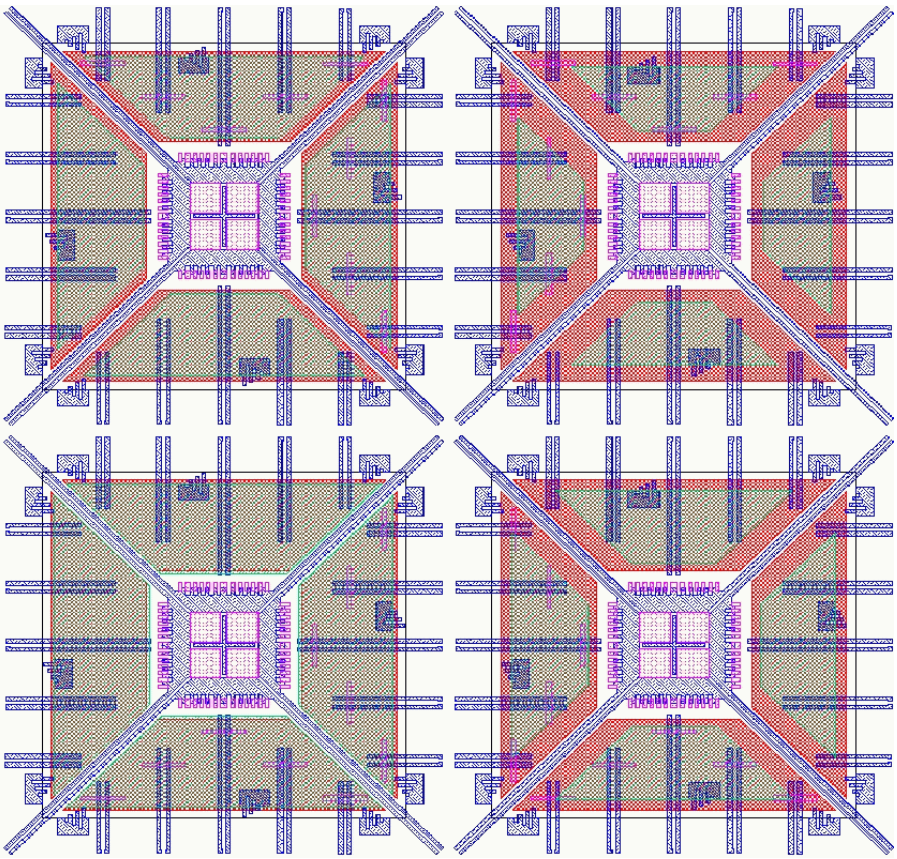

Figure 4 First run screen mask

The screen mask patterns are represented in green, while the desired resultant trapezoids (all the same size) are shown in red. This patterning test was complicated by our necessary move to a thicker board metal: a 1-ounce copper strip is about a 3.5 times larger step than we had patterned over previously. As a result of this run, we were able to establish that a negative 10% bias would be appropriate for future screen printing. We also determined that the mask feature area should be reduced to 40 X 30 cm, and features for alignment and characterization should be improved.

Unfortunately, the capacitance measurement results were not in line with our initial target of 40-pF/mm^2. We suspect much of this was due to variations in thickness from run-to-run, as is evident in deviations (from left graph to right graph) in the results shown in Figure 5. While we attempted to minimize the variations by using a consistent screen aperture and thickness, aging of the mixture caused it to print differently as time went on.

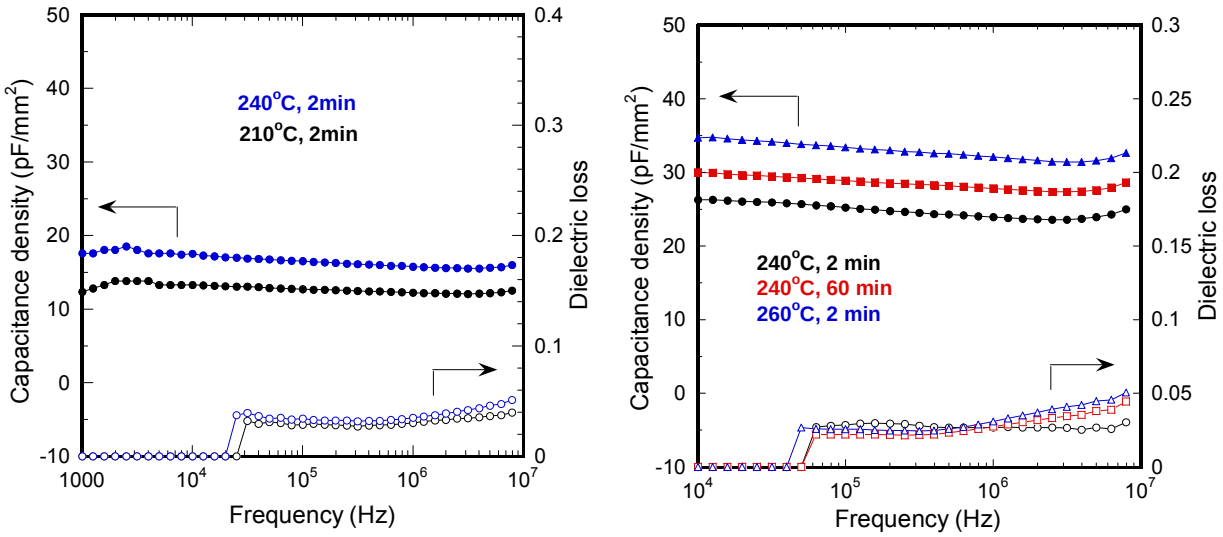

Figure 5 Results of two runs showing inconsistency in capacitance density

These graphs also show our attempts to vary the hard-bake cycle, with either extended time or extended temperature, in order to again move the capacitance toward our target. Unfortunately, this apparent insufficiency of the old slurry recipe means it could not meet the capacitance goal while retaining compatibility with a standard solder reflow process, including lead-free processes. We decided to add BYK-W-9010 to the slurry in order to stabilize the layer prior to the solder-reflow stage. We chose BYK-W-9010 because it improves wettability and improves dispersion by resisting the tendency of the slurry to bind on edges and corners.

We then created PCB designs which were submitted to Sunstone, a PCB manufacturer. These were combined with glass "chips" fabricated at NCSU. Figure 6 shows the layout of these glass chips for our second run:

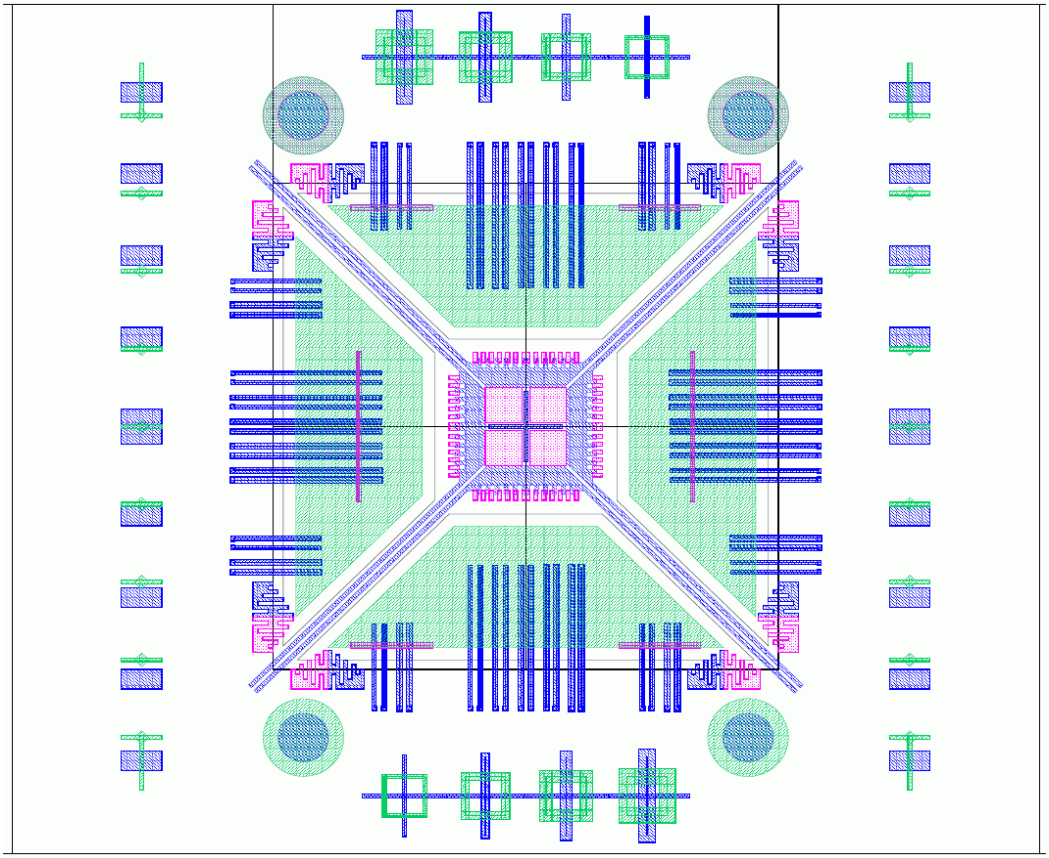

Figure 6 Layout for the second run

This includes screen-mask (green), board traces (blue), and chip features (red). Capacitance tests are grouped together, sharing a common chip-side cross-member. These range from 200 to 300 μm in dimension. They are duplicated on opposing chip quadrants, and include open cases for de-embedding on the other two quadrants. Since we were trying to mimic a large-die experiment, the short-term testbed's size was kept to 25 mm square; however, since we did not have the capability of performing ball bumping in-house, the central cavity was filled with structures for alignment. Alignment and post-alignment evaluation features include naked eye corner spots, microscopic verniers, and capacitive difference regions. The experiments were designed to yield even with misalignments approaching +/- 1 mm; as a benefit, simple chip-edge alignment proved sufficient (chip, slide, and board outlines are shown in black).

We then did a first run of the fabrication of our in-house chip on simple 1 X 3-inch glass microscope slides, with enough yield to allow us a handful of die-attach tests. This was largely a process learning experiment that led us to realize that these masks needed rework.

In early May 2008, we reworked the screen-mask and glass slides in order to improve fabrication. Screen-mask features were reduced to a 40 x 30 μm sweet-spot where screen tension would be flat and even. Similarly, glass slide features were moved away from the outermost edges, where photoresist spin-on was discontinuous. Upon receiving our PCB testbed boards back from fabrication at Sunstone, we immediately began assembly, using these new masks. This became the primary testbed for exploring changes to the high dielectric constant material formulation, so we ordered twenty two-sided boards for testing.

Using our original material recipe (cured at 240 ^{o}C for 2 minutes), we patterned the high dielectric constant material onto the boards. (This process wass thermally compatible with solder reflow.) Employing a microscope and several micromanipulators, we then flip-mounted our diced glass slides over the top, and affixed them with tape. On the board, pairs of probe pads extended beyond the high dielectric constant material patterns. On the slide, lines were drawn crossing the probe lines, forming capacitor pairs across the dielectric. These sites were probed and measured with the Agilent E4980A Precision LCR Meter. Unfortunately, values taken from every test site were far lower than expected, ranging from 218 to 541 fF. Assuming our 40 pF/mm^2 target, a value of 1.2 pF was expected from the 200 x 300 μm sites. Furthermore, we measured the back-side probe sites (bare board; identical to front-side, but lacking high dielectric constant material and slide coverings), we saw only slightly smaller values, ranging from 188 to 387 fF.

The test pattern was symmetrical, and had four identical sets of probe lines. Half of these probe lines intersected with circuit-completing lines on the slide while the other half were left open. We attempted to precisely calibrate the capacitor circuits against these 'open' cases. These measurements were made while downward pressure was applied to the slide, by either a steel weight or by maximum pressure by a Quarter XYZ500TIM manipulator with magnetic base, in order to ensure that any gap was minimized. Nevertheless, these still returned low values; a difference of only 38 fF was measured on the 200 x 300 μm capacitor-pair site, relative to the identical but uncovered site, much lower than the 1.2 pF we were expecting.

Since we had confirmed knowledge of the dielectric constant of our material from a non-testbed scenario done in work prior to this effort, we suspected the problem of low capacitance was due to surface topologies. (The high dielectric constant material dries and becomes a hard, non-sticky surface, to which our glass slides will not even adhere.) Thus, we assume air gaps existed between the dielectric and slides. We proceeded to measure the surfaces using a Dektak profilometer. The board wires measured approximately 30 μm in thickness, with only 3 μm or so roughness from the board surface. We saw very little roughness in the board or wire surfaces. We measured a bare wire trench (280 μm space between two wires) against an area covered with the high dielectric constant material, and found little difference in the dielectric thickness above the wires relative to the difference we found between them. We did notice, however, that the thickness did drop off severely at the metal edges, actually leaving some metal exposed. This observation was made visually, and confirmed by profilometry. Finally, by running the profilometer lengthwise along the wires, we measured the dielectric thickness to be somewhere between 8 and 12 μm. The major finding, however, was that it had a rough surface, with many 7 to 13 μm bumps, spaced roughly every 60 to 80 μm apart. This proved that an air gap did exist.

This is the unfortunate side-effect of losing access to EIT's HyperBGA technology – that was much smoother.

We attempted to repeat our capacitance measurement with a flexible substrate in place of the patterned hard glass slide. The flexible substrate was cut in the form of a wire, 36 to 38 mils wide, and placed metal-side down over the high dielectric constant material layer, and held down by a piece of Gelpak membrane, followed by a clear (unpatterned) glass slide. The goal of this experiment was to reduce this air gap, and it did the job partially. Our capacitance measurements improved to between 6 and 7 pF/mm^2. Nevertheless, this was far less than we expected.

Finally, we decided to replicate the original high dielectric constant material characterization structure with the modification of putting our board in place of the patterned platinum probe sites. While this experiment no longer emulated a chip-over-board topology, it did eliminate any air gap, and allowed us to measure the high dielectric constant material as it existed patterned over a board topology--or so we thought.

The construction was as follows: High dielectric constant material was patterned over our board, then a square region of platinum (approximately 1500 Angstroms) was patterned over that. The overlap distance between the platinum edge and the board-level probe wires was subsequently measured (these were roughly 100 μm), in order to later extract our impedance density. Unfortunately, all of the impedance measurements made during this experiment proved to be unstable; the capacitance values would vary from measurement to measurement, and impedances would come up effectively shorted, either immediately, or after a few tens of seconds. We concluded that the dielectric was thinning at the wire edges, due to its pre-bake settling (too viscous).

After several attempts, we were unable to replicate the dielectric properties and viscosity desired in order to meet the dielectric constant target while keeping to a solder reflow compatible temperature regime.

At this point, both the faculty (Kingon) and Postdoc (Kim) left NCSU for other positions before the experiments were complete. They were not able to complete the experiments in their new location. We evaluated doing the experiments in our own lab but did not have the safety equipment for the hazardous chemicals involved. Hence we had to abandon the flow of experiments that would have led to attempts to reduce the air gaps.

In the end, our failure to find a suitable replacement for EIT, coupled with the loss of our high-K dielectric team forced us to focus our attention on the other tasks described below. We were only able to fully complete the first three of the four elements of this task.

2.2 Task 2. Circuit Demonstrator with Laminate Package

In prior work, we demonstrated that power efficient short range links can be built using capacitive coupling for relatively short (~15 cm) links. For longer links some form of equalization is needed in order to compensate for the frequency characteristics of the channel. We realized that one very efficient way to build the equalizer would be to use capacitive coupling. **In this case, the capacitors are placed on-chip, not between the chip and the package**. This section describes this experiment and the successful outcome.

ACCI utilizes non-contacting capacitor plates between silicon and package as passive equalizers (EQ). The series capacitor de-emphasizes the low frequency component of the signal and, when combined with a low-pass channel, generates a flat band-pass response to eliminate inter-symbol-interference (ISI). The size of the capacitor in traditional ACCI, however, needs to be carefully selected to match the channel loss. That is, the fixed capacitor value creates constraints on the channel length and significantly increases the difficulty of designing for the rest of the system, including package design and PCB floorplaning. In the presence of longer channels, smaller coupling capacitors with higher frequency peaking are required to compensate for the increased attenuation.

Such a system usually is swing-limited; i.e., the signal swing is too small (due to a smaller coupling capacitor) for the receiver to detect. An easy way to mitigate the high-frequency loss in longer channels is to implement an active EQ in the system, such as feed-forward equalizer (FFE) at the transmitter (TX), a continuous-time linear equalizer (CTLE) or decision-feedback equalizer (DFE) at the receiver (RX). This task investigates a package structure to be integrated into the TX-side FFE and shows 90% improvement in power efficiency by enabling voltage-mode summing.

The FFE in high-speed chip-to-chip communication consists of two major blocks: a digital delay block and an analog summation block. The delay block is straightforward to implement because the signal is still in the digital domain, while the analog summation block is more difficult to design. Conventionally, analog summation is realized by current-mode summing, or by using a current-mode digital-to-analog converter (DAC). The advantage of this type of summation block is the relative simplicity of the topology. The channel impedance is matched at the TX side by the load resistor of the current-mode driver. The load resistor also sums all the tail currents modulated by the current/previous bits and their respective tap weights. The driver output swing is subsequently determined by the tail current. This topology has been popular over the years but there are several disadvantages associated with current-mode summation which make it less attractive: First, it consumes a significant amount of power due to the nature of the summation (or subtraction) of the current. Secondly, a large parasitic capacitor (dominant pole) exists on the node where all the current summing branches are joined together, thus limiting the frequency response of the FFE and the number of total branches in the DAC. Thirdly, the output currents are summed non-linearly due to channel-length modulation. A tap weight setting of X, in reality, might be less than X times the least-significant-bit (LSB) swing. Finally, current summation has limited dynamic range because of constraints on output swing to keep tail transistors in saturation.

A multi-capacitor (MultiCap) structure, as shown in Figure 7, and compatible with ACCI technology, can solve all the problems of current-mode summation at once. This passive structure facilitates a simpler yet more efficient way to perform the summing function. The coupling capacitors in traditional ACCI are formed by the top metal plates on the flipped chip and the package (left of the figure). CMOS technology has the advantage of much higher fabrication resolution than the package; therefore, it is straightforward to utilize this advantage toward creating the MultiCap structure (center of the figure) for signal coupling as well as voltage summing. The difference between the traditional coupling capacitor and the MultiCap structure is that the former has a single top plate, while the latter has a matrix of discrete top metal plates.

Each individual capacitor of the MultiCap, as illustrated in Figure 8, is designed to be connected to separate branch (1 LSB) of the voltage-mode output driver, replacing both the voltage-to-current and current summing blocks used in conventional FFE.

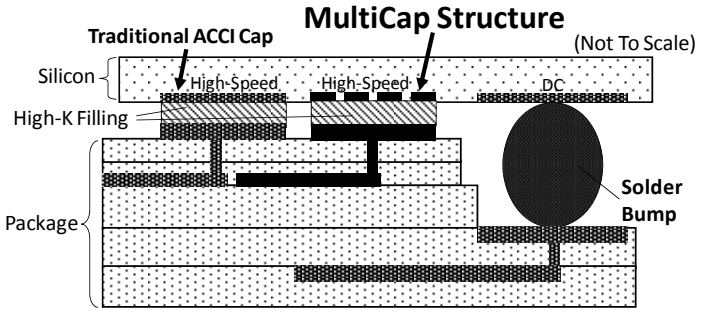

Figure 7 Multi-capacitor structure

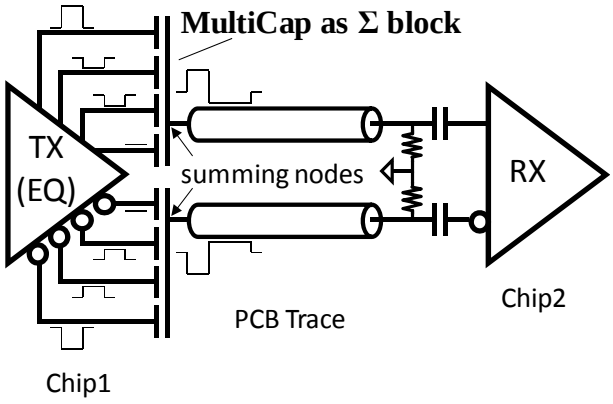

Figure 8 Summing the multi-caps

Instead of current being summed on a single branch, voltage is summed on the shared bottom plate. This structure solves the problems of the conventional current summation at once. First, simple voltage-mode circuits can be used throughout the TX, and the need for steady tail currents is eliminated. Secondly, each branch of the FFE is independently coupled to one common bottom plate, eliminating the presence of a single dominant parasitic capacitor. Thirdly, the tap weight is

selected by turning ON/OFF each identical sub-branch. Turning ON more branches would directly result in more signal swing but keeps the same edge rate. As an added benefit, there is no concern regarding the linearity in the voltage summation. Finally, unlike current-mode summing, the MultiCap sums the voltage regardless of the signal amplitude. It does not have constraints on the dynamic range of the output swing or the need of a common-mode feedback circuit.

The purpose of this task was to explore MultiCap implementation options. How can the MultiCap be realized with various substrates such as silicon (from RTI), and laminate package (EIT)? By going through the design phase for the whole system, we were able to make better estimates of the realistic size and cost of the whole system.

At an early stage of the project, much work was done developing the circuitry to work with the capacitively-coupled interconnect, mainly using the simulators HFSS, Sonnet, ADS, HSPICE, and Spectre to model the channel. The top view of the MultiCap structure, as shown in Figure9, is an array of individual capacitances (C_C) connected to the FFE voltage-mode driver.

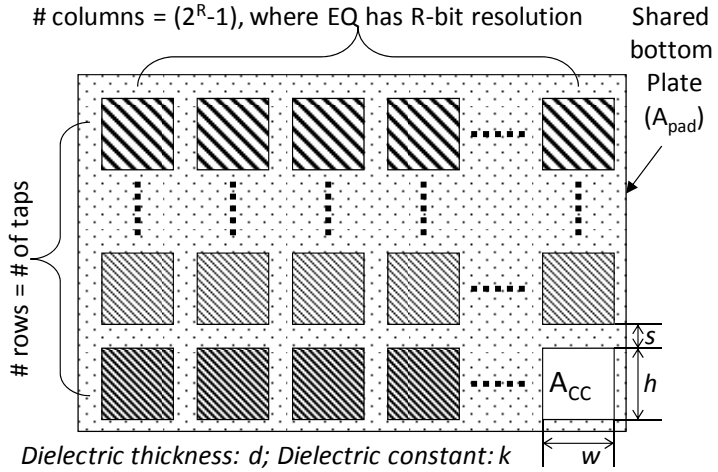

Figure 9 Top view of the multi-cap structure

The number of rows equals the number of taps of the FFE (N_{tap}), and the number of columns equals 2^R, where each tap weight has R-bit resolution. All individual capacitances are of identical geometries; therefore, the tap weights of the FFE are set by whether each branch has a signal or not. The nominal design values of the individual C_Cs are all identical to keep the circuit simple and linear. The following sections provide equations to help in choosing useable values of the MultiCap. The variables used in the following equations are the plate geometries (w, h, and s), individual plate area (A_{CC}), shared bottom pad area (A_{pad}), and dielectric constant (k). Note that the spacing (s) is neglected in most of the equations when calculating the MultiCap geometries.

The value of C_C, like all parallel-plate capacitors, is determined by the size of the overlapping plate area (A_{CC}), the dielectric thickness (d), and the dielectric constant (k):

$$C_C = k\varepsilon_0 \frac{A_{CC}}{d} \tag{1}$$

The upper-bound of the area of each top plate, $A_{CC,max}$, is limited so that the total area of the MultiCap is smaller than the bottom plate:

$$A_{CC,\max} \le \frac{A_{pad}}{N_{tap} \times (2^R - 1)} \tag{2}$$

In some special cases, A_{CC} may be more than indicated in (2). For instance, the total number of top plates is no longer (2^R-1) when the FFE has a different maximum tap weight for each tap. Nonetheless, $A_{CC,max}$ can still be calculated by assuming that each individual top plate is identical, and the equations discussed below can be modified accordingly.

The minimum area of each top plate, $A_{CC,min}$, is implicitly limited by the minimum RX sensitivity. That is, the total capacitance of one row of the MultiCap (connected to one tap), $C_{C,row,min}$, should be large enough such that the RX is able to detect the non-equalized signal after attenuation over the longest channel. The value of $C_{C,row,min}$ is treated as a given variable and should be estimated from a preliminary simulation, which includes an ideal driver and an ACCI channel. The total equivalent area of the entire row of top plates ($A_{CC,row,min}$) can be calculated:

$$C_{C,row,\min} = k\varepsilon_0 \frac{A_{CC,row,\min}}{d} \Rightarrow A_{CC,row,\min} = \frac{d}{k\varepsilon_0} C_{C,row,\min} \tag{3}$$

Thus, the corresponding minimum area of each top plate, $A_{CC,min}$, can be derived

$$A_{CC,\min} \ge \frac{A_{CC,row,\min}}{(2^R - 1)} = \frac{d}{k\varepsilon_0 (2^R - 1)} C_{C,row,\min} \tag{4}$$

A larger A_{CC} is preferable because, unlike the non-EQ version where larger than necessary C_C causes ISI, in this case the EQ is capable of either tuning down the driving power for shorter channels or generating more high-frequency components for longer channels, thereby eliminating ISI.

The range of the MultiCap area is bounded mainly by two variables, A_{pad} and $C_{C,row,min}$, from (2) and (4), respectively. The former is related to input/output (I/O) density and the later to circuit capability. By respectively substituting (2) and (4) into (1), then combining the results, the upper-bound and lower-bound of C_C can be calculated:

$$C_{C,\min} = \frac{C_{C,row,\min}}{(2^R - 1)} \le C_C \le \frac{k\varepsilon_0 A_{pad}}{d \times N_{tap} \times (2^R - 1)} = C_{C,\max} \tag{5}$$

Equation (5) is essential in estimating the available value of C_C. For example, given a set of variables; A_{pad} = 175x175μm^2 (200μm pitch and 25μm spacing on package), $C_{C,row,min}$ = 500fF (from preliminary simulation with TRX in 0.13μm standard CMOS and 80cm microstrip on FR4 PCB), d = 1μm (assumed), k = 18, N_{tap} = 4, and R = 3, substituting into (5) results in:

$$71.5fF = \frac{500 \times 10^{-15}}{(2^3 - 1)} \le C_C \le \frac{18 \times 8.85 \times 10^{-12} \times (175 \times 10^{-6})^2}{1 \times 10^{-6} \times 4 \times (2^3 - 1)} = 174.2fF \tag{6}$$

Further analysis can be achieved by equating both sides of (5):

$$C_{C,row,\min} \le \frac{k\varepsilon_0 A_{pad}}{d \times N_{tap}} \tag{7}$$

Equation (7) shows the trade-offs between three of the most dominant groups of variables for the MultiCap structure: the geometries (A_{pad} and d), the property of the dielectric filling (k), and the circuit complexity (N_{tap} and, implicitly, $C_{C,row,min}$). In order to satisfy (7), when driven toward higher-density I/O, the area of the bottom plate is limited; thus, a higher dielectric constant is needed. Alternatively, the same goal can be achieved by decreasing the dielectric thickness, decreasing the circuit complexity (such as lower N_{tap}), or improving the RX sensitivity (hence lowering $C_{C,row,min}$). Take the same example used in (6), if all the parameters are fixed except for k,

$$500\times10^{-15} \le \frac{k\times8.85\times10^{-12}\times(175\times10^{-6})^2}{1\times10^{-6}\times4} \Rightarrow k \ge 7.4 \tag{8}$$

Equation (8) shows that k needs to be at least 7.4 for a useable MultiCap structure. If the high-k dielectric (k =18) used in the example is not available, and an ordinary oxide (k=4) is used instead, (8) will not be satisfied unless the geometric constraints are relieved ($A_{pad} \ge$ 238x238μm^2 or d ≤ 0.54 μm), circuit complexity is decreased ($N_{tap} \le 2$), or the RX sensitivity is improved ($C_{C,row,min} \le$ 271fF).

As shown in Figure 10, there are two types of parasitic capacitances in the MultiCap structure: vertical parasitics ($C_{p,TG}$, and $C_{p,BG}$, coupled to ground) and horizontal parasitics (C_p, coupled to next row of MultiCap).

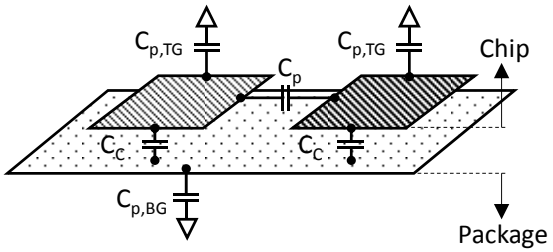

Figure 10 Two types of parasitic capacitances in the Multi-cap structure

The vertical capacitances, $C_{p,TG}$ and $C_{p,BG}$, are between the top plate and the on-chip ground plane, and between the bottom plate and the package ground plane, respectively. They are treated as constants from manufacturing process, due to the fixed height between plates and ground planes, and can be minimized by placing a cutout in the ground plane where the MultiCap is located. The other significant parasitic component is the horizontal fringe capacitance, C_p, between two rows of the MultiCap. The ratio between C_p and C_C, η, can be estimated (lower η is better):

$$\begin{cases} C_C = k\varepsilon_0 \dfrac{wh}{d} \\ C_p = k\varepsilon_0 \dfrac{wd}{s} \end{cases} \Rightarrow \quad \eta = \frac{C_p}{C_C} = \frac{d^2}{sh} \tag{9}$$

The capacitance ratio η is independent of *w* in (9) because of the shared width. The parasitic capacitance between columns can be calculated in the same manner, but is ignored due to the way the MultiCap connection is arranged (both terminals of the capacitor being connected to the

same tap of FFE). Notice that η is independent of k only when the dielectric filling is homogeneous – the dielectric filling is not patterned, nor layered; otherwise, η will be smaller than indicated in (9). Under normal circumstances, when d ≈ s, the ratio η is about d/h. If h>>d, C_p is negligible.

Figure 11 shows a comparison between calculated results using (9) and electromagnetic (EM) simulation results using Sonnet Suites (3D planar and model extraction). It is shown from both curves that C_p is less than 10% for a horizontal spacing larger than 1μm. EM simulation shows more parasitic capacitance because it takes into account the fringe capacitance.

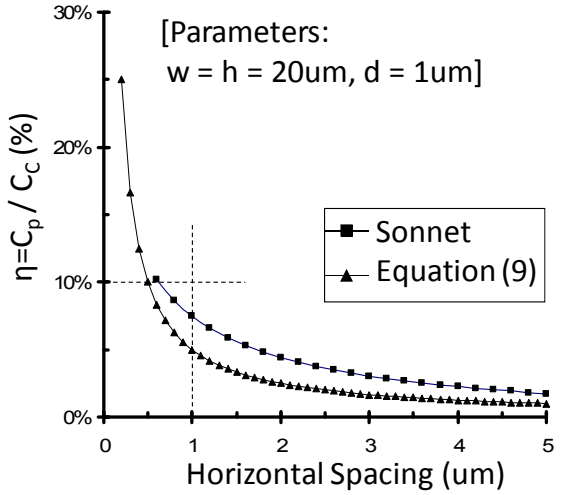

Figure 11 Sonnet simulations compared with closed-form solution

The original plan for this task was to design:

Active circuitry (NCSU via MOSIS)

- Die will have 3mm circuit area
- 1500 μm additional border area is required for High-K deposition

Modified HyperBGA (EIT)

- Panel size: 14”x22”
- Trench depth and port width to be determined
- Ball bumps by RTI, for case of large-die testbed

Back-end High-k build-up layer (NCSU)

- Panel will be cut to 4” prior to build-up (for silk-screening)
- Patterned over modified HyperBGA (~80um thick trench)
- Approximately 10um thick (+/- a few microns, possibly)

The channel includes two capacitors in its path - one very close to the transmitter; the other very close to the receiver. The capacitors are formed by the top metal of the silicon chip and the top metal of the carrier/package/PCB (in this case, the EIT board.) Between the two metal plates is a high-K dielectric developed by the NCSU Materials Science department.

Figure 12 shows the block diagram of the circuit designed for this task. The small grey blocks at the top of the diagram are simple delay inverters. The larger blocks near the center of the diagram (labeled "T") are multiplexor-like NFET-only drivers which work with low supply voltage.

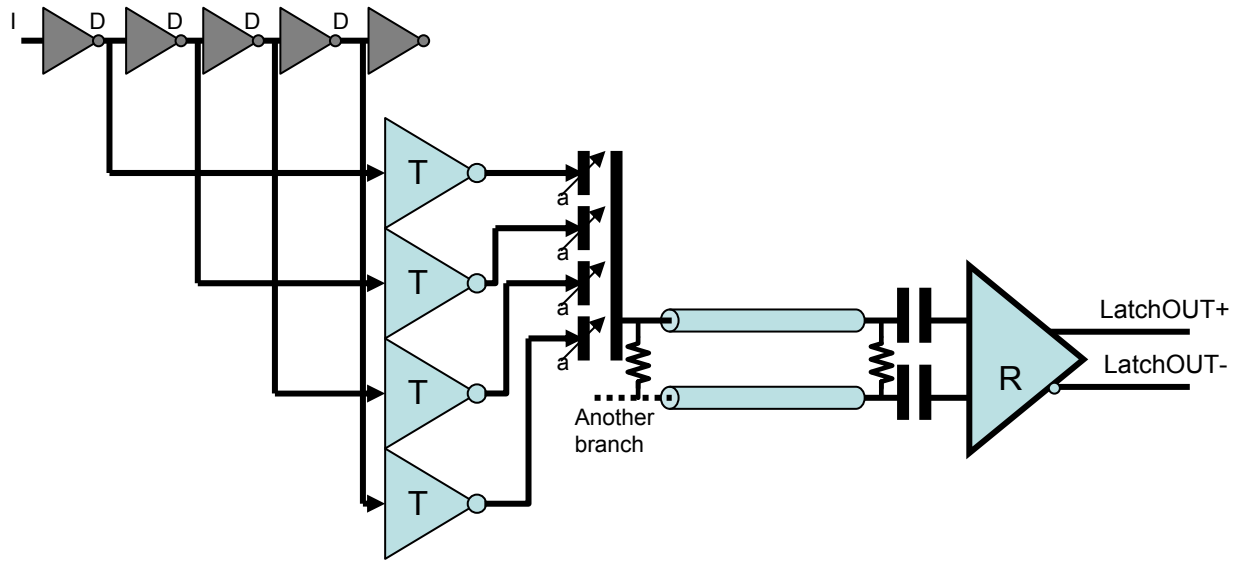

Figure 12 Circuit configuration for Task 2

The block diagram of the transmitter circuit is shown in Figure 13.

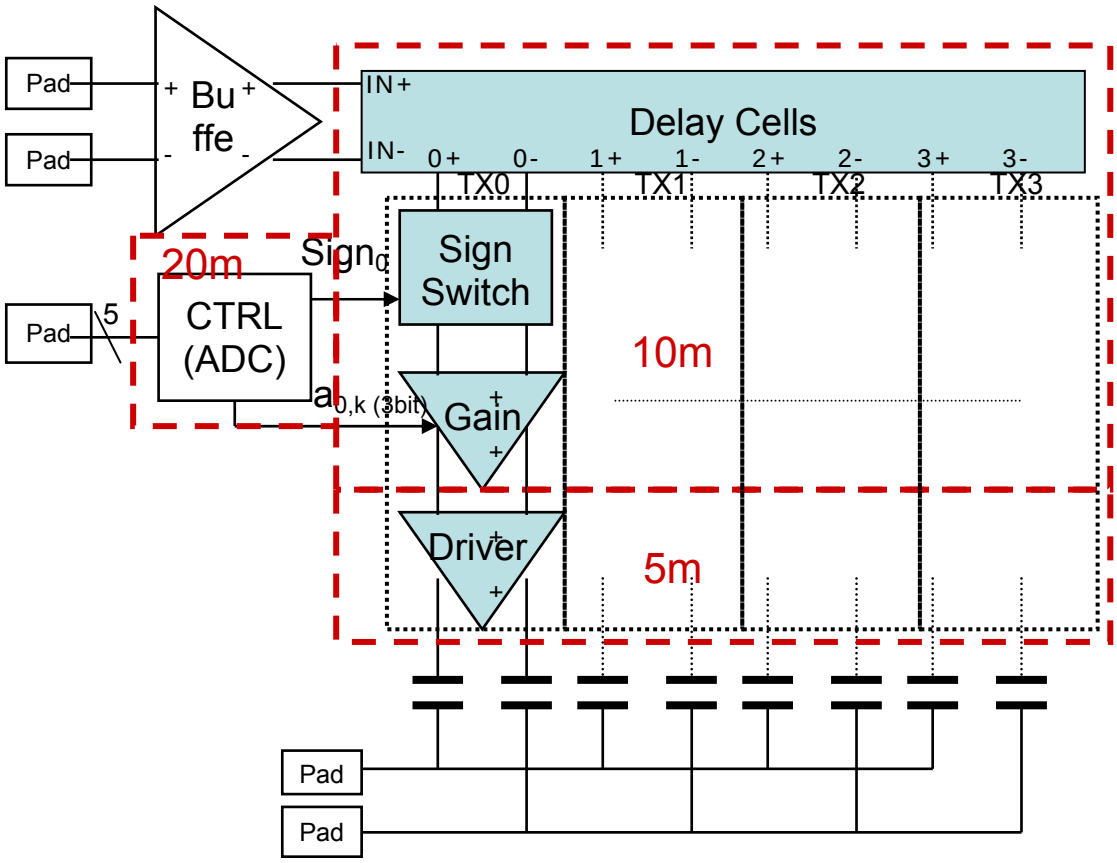

Figure 13 Transmitter circuit

The layout of the transmitters is shown in Figures 14 and 15. The transmitters themselves are shown in the first diagram. The second diagram shows a closeup of the metal-insulator-metal (MIM) on-chip capacitors that are used to connect the transmitters to the summing node.

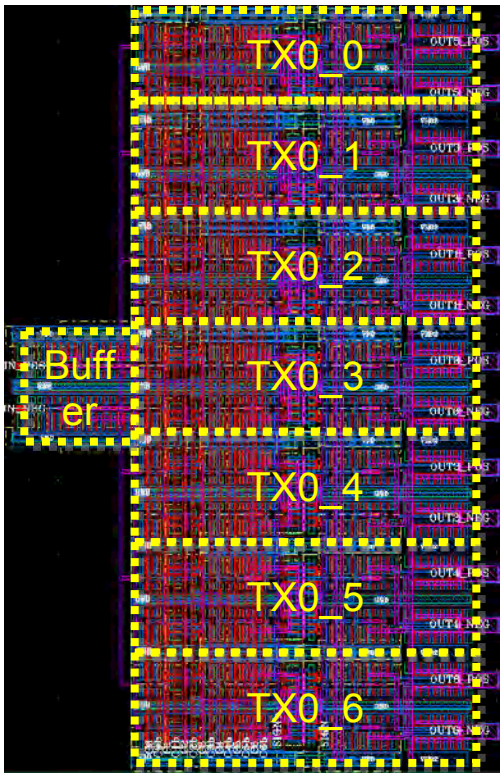

Figure 14 Overall transmitter layout

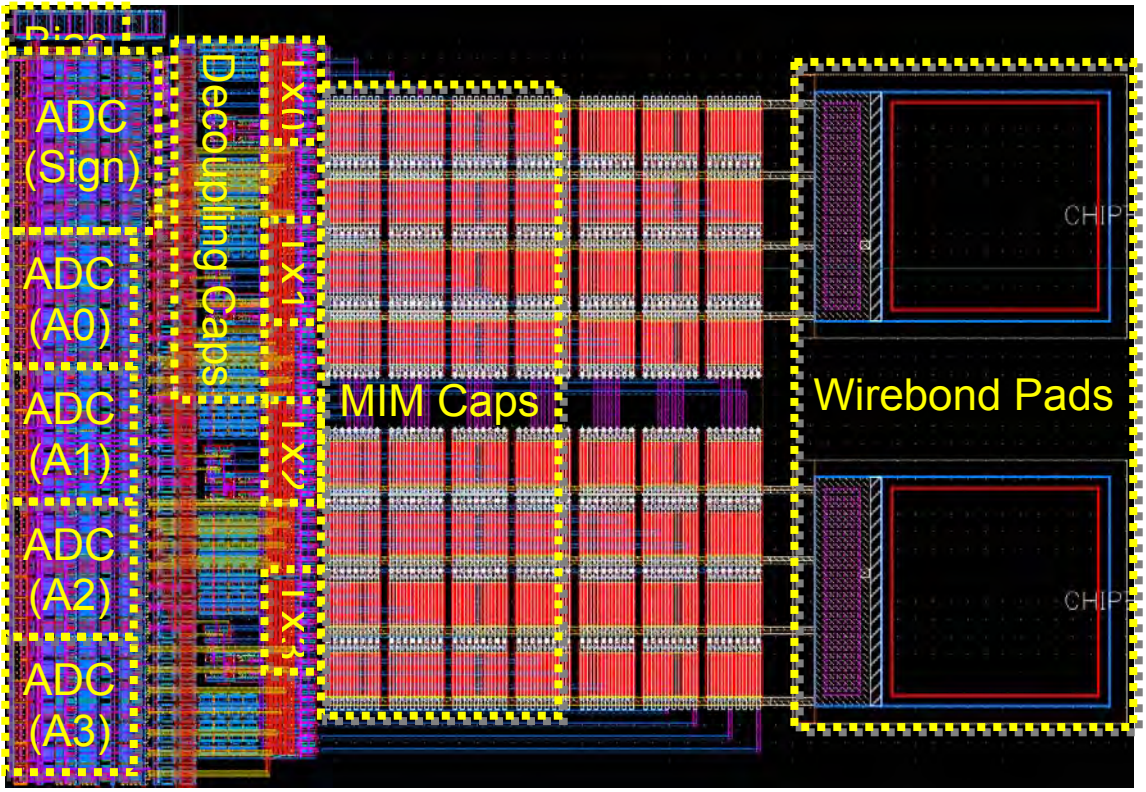

Figure 15 Closeup of the MIM capacitors

The receiver is a simple current-mode logic (CML) regenerative latch with one buffer stage and one latch stage, as shown in Figure 16. The layout is shown in Figure 17.

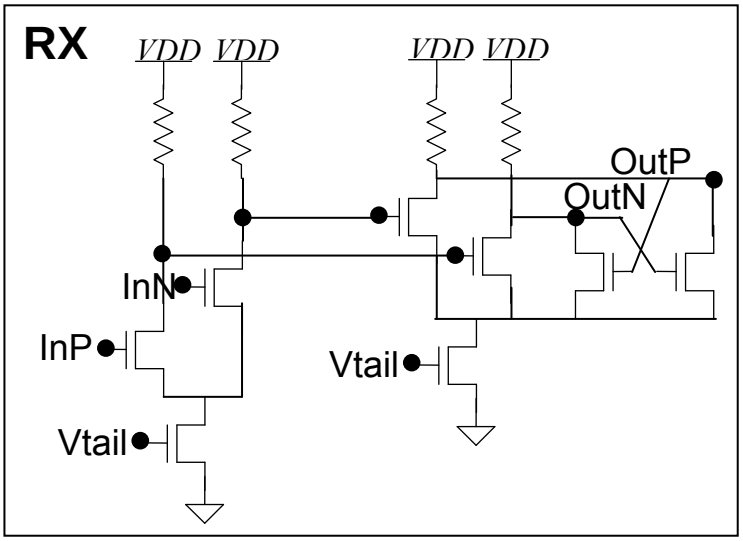

Figure 16 Receiver

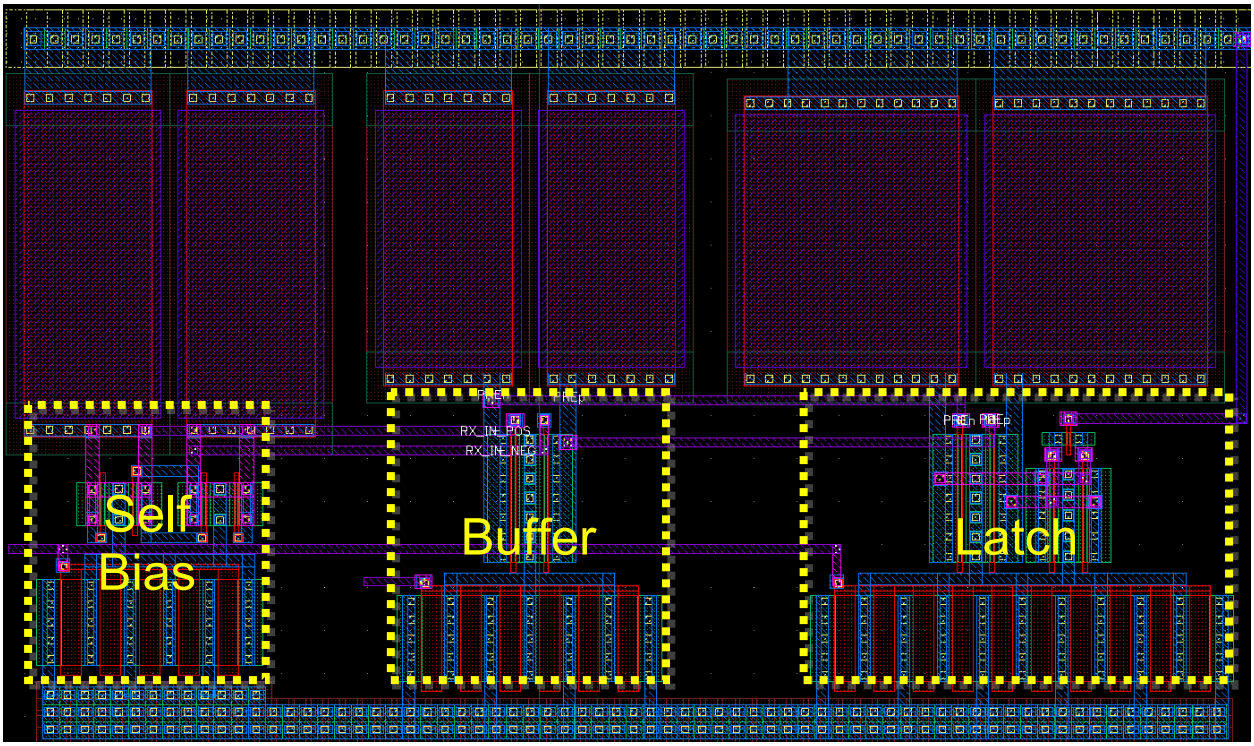

Figure 17 Receiver Layout

The CMOS chip is designed for lower power consumption while maintaining high performance. The total power consumption is down from 20 mW to 15 mW at worst case – 10 Gbps on a 1 m long channel. A control (CTRL) unit (five 3-bit flash ADCs) for coefficient adjustment was

added as the external control interface. The CTRL unit is not included in the total power consumption calculation because this is assumed to be available in a real chip.

Post-layout simulation results are shown in Figure 18. The RX input, instead of a traditional non-return-to-zero (NRZ) signal, is a pulse signal. Therefore the input eye has a y=0 trace. An amplifier followed by a latch can recover the NRZ signal from the pulse signal. 15 to 45 mV clearance is required to get a clean eye at the latch output. Notice that no clock is required in this scheme.

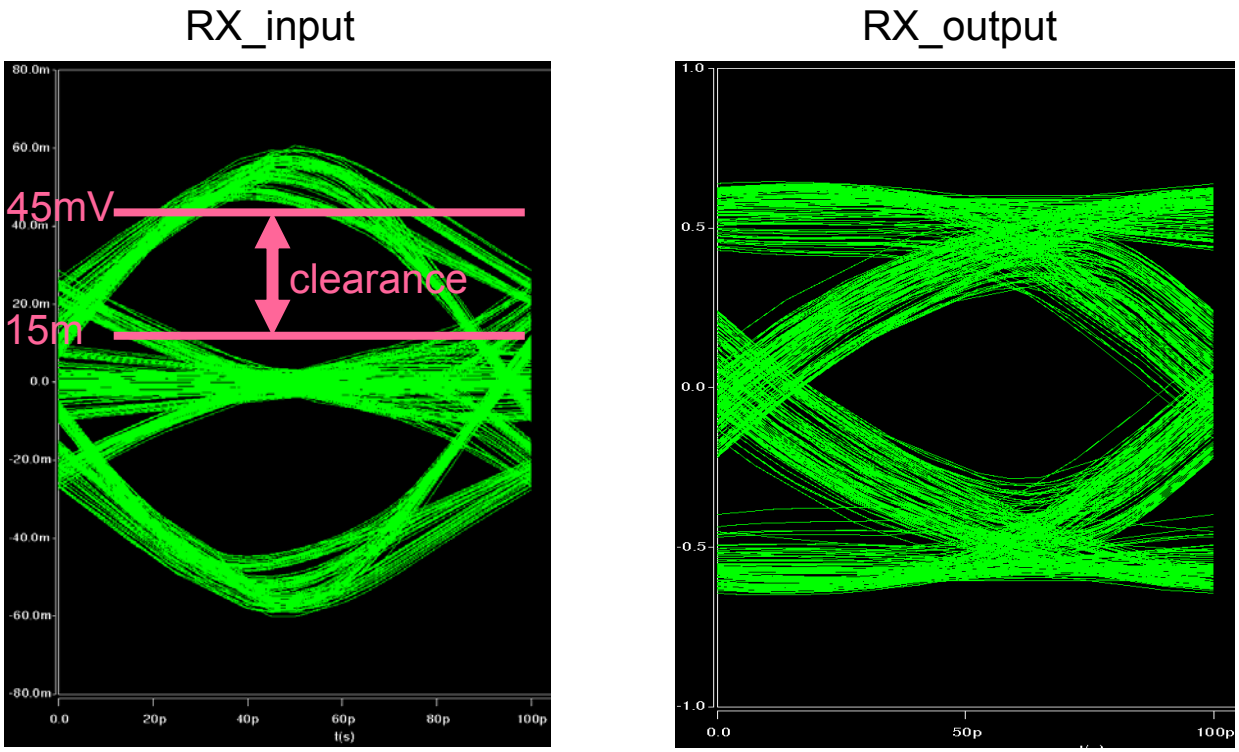

Figure 18 Post-layout simulation results

Unfortunately, during the project, we suspended cooperation with EIT because of the lack of interaction. However, the type of laminate package that EIT manufactured was above our minimum requirement. As this time, this type of laminate package was gaining its popularity in mobile consumer electronics – the so called high-density interconnect (HDI). Therefore, we are confident that the capacitively coupled interconnect on laminate package will be physically realizable with reasonable cost.

In the end, we were able to implement a realistic system with existing PCB technology, manufactured by Mid-Atlantic Circuits. There are two variations on the channel length and an additional main board for CTRL signal routing as shown in Figures 19-21.

Table 1 : PCB and corresponding test setup

#	Channel Length
1	40cm
2	60cm

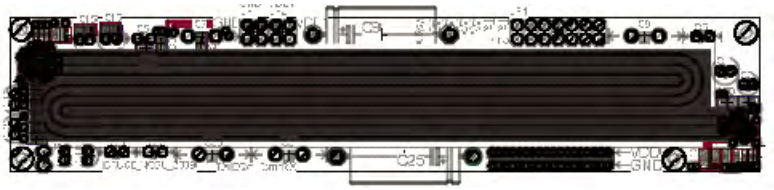

Figure 19 PCB #1 top view (40 cm channel)

Figure 20 PCB #2 top view (60 cm channel)

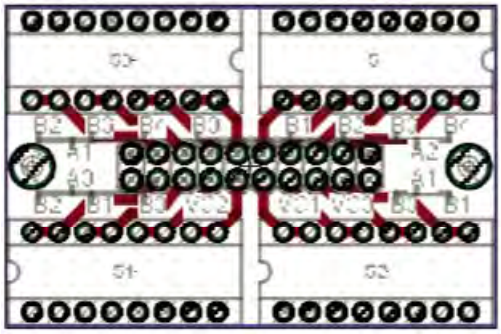

Figure 21 Control board

Assembly steps:

Step 1. Place SMD capacitors on top side by using solder paste and reflow. If possible, use hand soldering.
Step 2. Hand solder SMD capacitors on back side, insert all the DIP components from the back, then make a support platform.
Step 3. Epoxy the chip directly on board then do wire-bonding according to the bonding diagrams shown below

There were two types of transceiver chip manufactured through MOSIS. One has a slightly more efficient transmitter design. The transmitter and receiver wire bonding for the May 2008 0.013

micron run chips are shown in Figures 22 and 23; the bonding for the August run is shown in Figures 24 and 25.

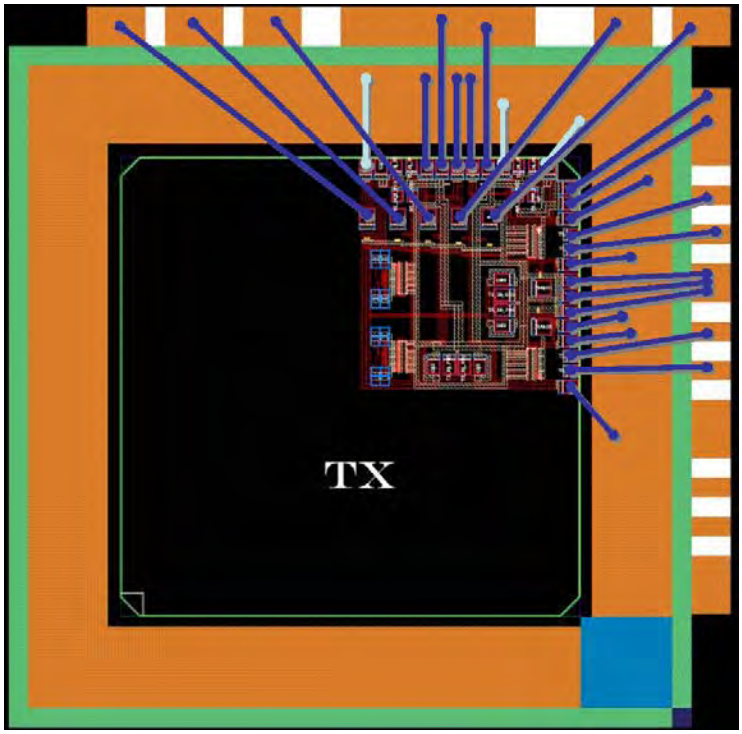

Figure 22 May 2008 transmitter bonding

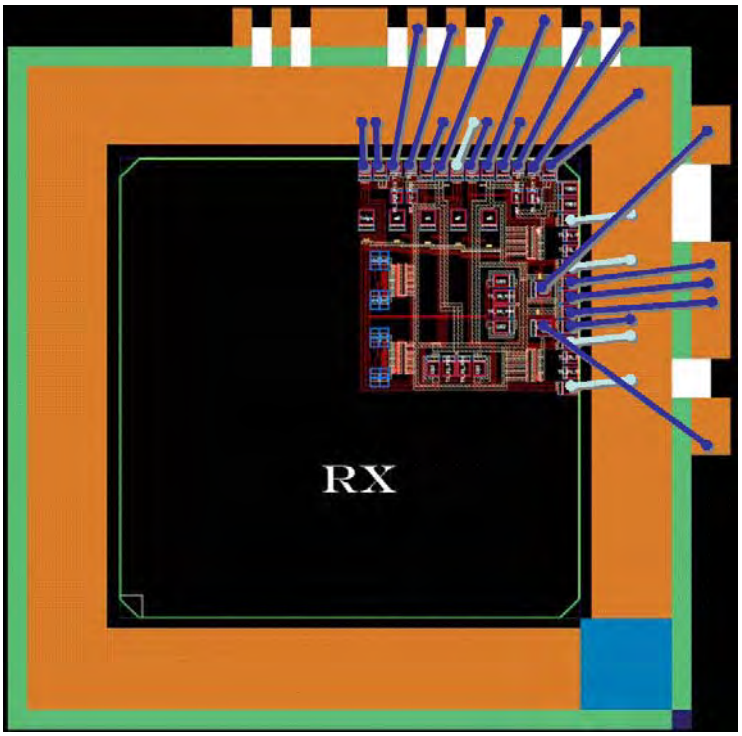

Figure 23 May 2008 receiver bonding

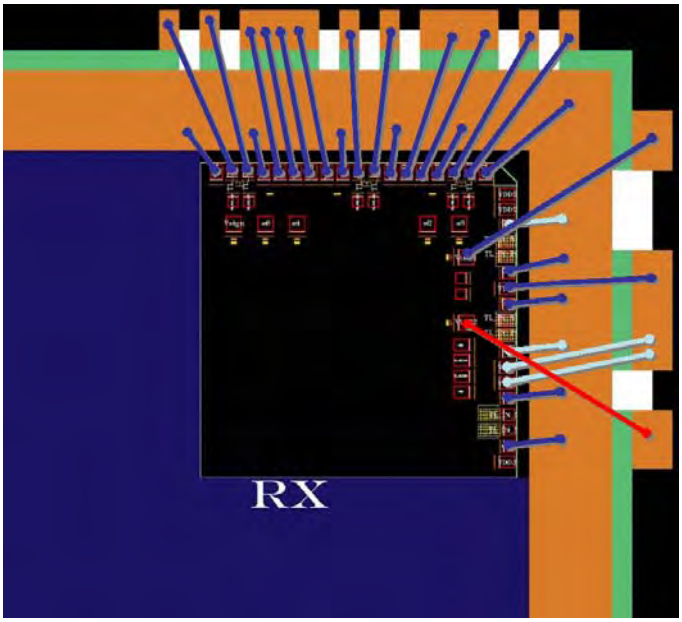

Figure 24 August 2008 receiver bonding

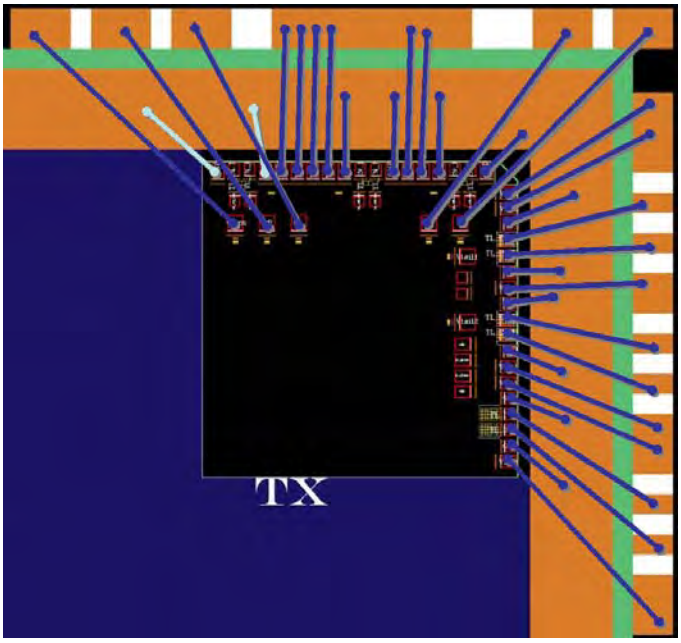

Figure 25 August 2008 transmitter bonding

Figure 26 shows the PCB configuration and silicon area of the TX and RX, including the MultiCap located underneath the bond pads:

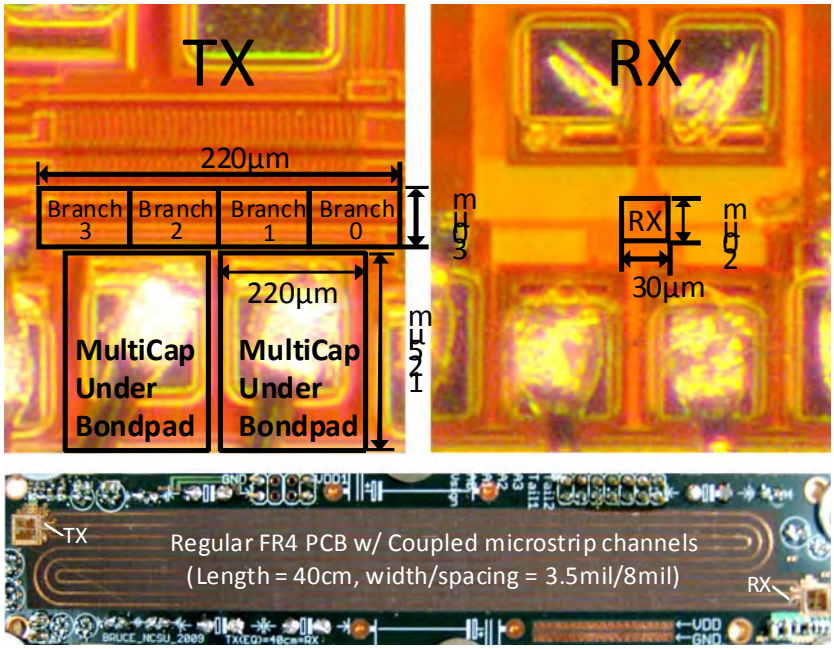

Figure 26 TX and RX PCB configuration

The delay cell is made of a pair of inverters having a measured delay of 84 ps. Figure 27 shows the effects of the changing tap weights by comparing RX inputs before and after the continous-time fractional-space equalizer (CT-FSE) is enabled:

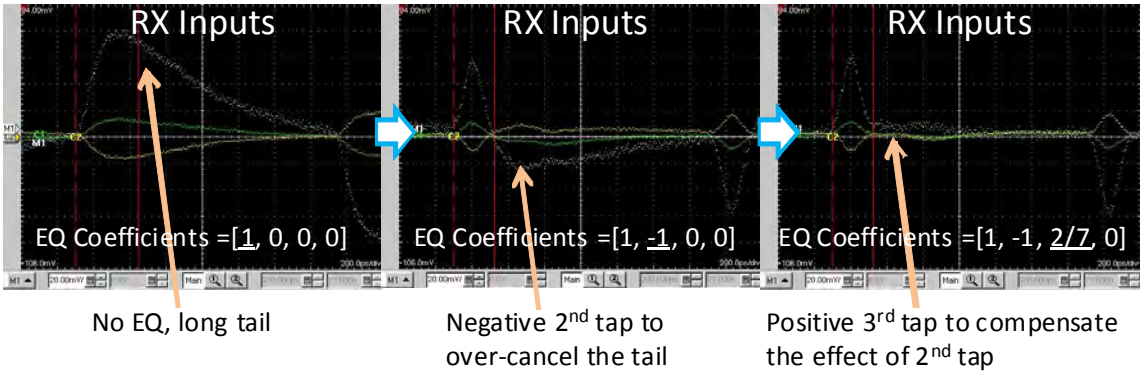

Figure 27 Effects due to tap weights

Figure 28 shows the 5 Gb/s NRZ data recovered by the receiver:

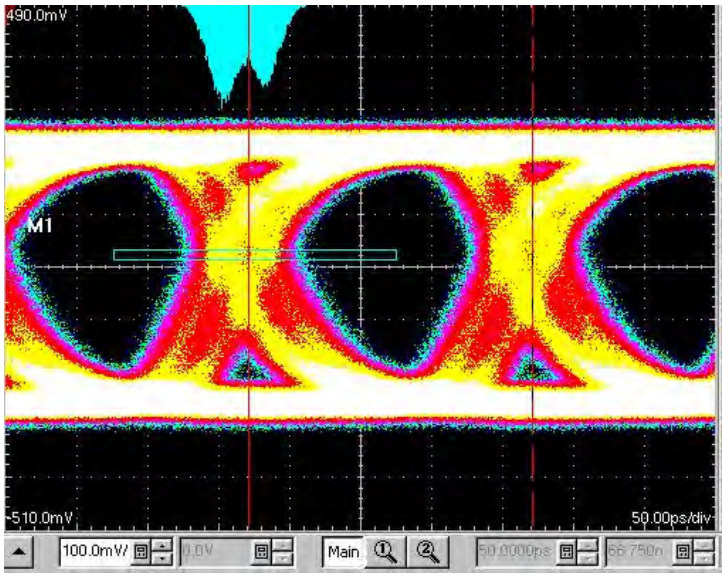

Figure 28 Data recovery by receiver

The system has a bit-error-rate (BER) of less than 10^{-12} for a pseudo-random bit sequence (PRBS) input pattern of length 2^{23}-1. For each channel, the TX and RX consume 6.5 mW and 1.1mW, respectively. Table 2 compares the TRX performance with state-of-the-art designs. It is demonstrated that the TRX, benefitting from the CT-FSE scheme combined with use of the MultiCap, has an area of only 0.007 mm^2 and achieves a low power requirement of 1.5 mW/Gb/s.

Table 2: Comparison of Performance

	This work[*1]			[palmer07]	[Uchiki08]	[Liu09]
Process	0.13 µm			90nm	90nm	65nm
Channel Length	10cm[*2]	40cm	80cm[*2]	75cm[*3]	38cm	127cm
Speed	5Gb/s			6.25Gb/s	6Gb/s	8.9Gb/s
TX Power	5.2mW	6.5mW	7.6mW	4.9mW	-	11.6mW
RX Power	1.1mW	1.1mW	1.1mW	8mW	10.9mW	5.4mW
TRX Total Area	.007mm^2			.307mm^2	.054mm^2	.023mm^2
TRX Power Efficiency (mW/Gb/s)	1.3	1.5	1.7	2.0	>1.8[*4]	1.9
TRX Config.	TX CT-FSE (k=4,R=3) Passive EQ			RX half-rate Cap. Degen.	RX Cap. Degen. ×4	RX half-rate DFE-IIR

[*1]No multiplexer incorporated. [*2]Simulation result. [*3]Estimated from loss. [*4]Only RX was reported.

We demonstrated that a MultiCap structure compatible with ACCI technology has performance as good as other transceivers implemented at better CMOS technology nodes, while operating at power consumptions lower than competing approaches even when the latter are built in more advanced circuit nodes. The parameters of the MultiCap can be calculated from a set of equations, which also allows quantitative trade-offs when designing with MultiCap embedded packages. We also found that the MultiCap structure is able to replace the current-summing block of a TX FFE when combined with proper FFE design. This type of system configuration supports high-speed, low-power (1.7mW/Gbps), chip-to-chip communication and the flexibility of active equalization.

2.3 Task 3. Capacitive Connections using embedded board capacitors

The primary objective of this task was to demonstrate circuit designs that support the scenario where the series capacitor is NOT immediately adjacent to the driver and/or receiver (such as was the case in our prior work and the tasks above). There are two scenarios where this is desirable:

1. When supporting an interconnect standard that forbids DC connection. Numerous interconnect standards, such as FiberChannel, require AC coupling but large capacitors are used. These capacitors have high parasitic, which degrade the potential performance. If smaller capacitors could be used, then high performance is possible and the capacitance cost (and parasitic) can be further reduced by placing the capacitor in the PCB as an embdedded structure.
2. For capacitively coupled sockets and connectors (task 4 below). The connector will not be immediately adjacent to the driver and/or receiver.

In this task we demonstrated the circuit principles required to support such channels.

This task involves the design of 0.18 micron CMOS circuits for testing capacitive connections made using PCB-embedded passive capacitors. The capacitors are formed between two close

layers inside the PCB with high-K dielectric. This type of system configuration could save area for on-chip or on-board capacitors. IBM and Sanmina are potentially interested in this structure.

Figure 29 shows the cross section of the structure of the embedded capacitor used for both decoupling and signaling using ACCI. It has been demonstrated that the embedded capacitor, built in a laminate package by adding at least one additional high dielectric constant dielectric layer, can be used in power supply decoupling. The signal channels have to be routed through that capacitor layer anyway, thus it is straightforward to utilize the coupling function and implement an ACCI channel: a transmission line with small series capacitor.

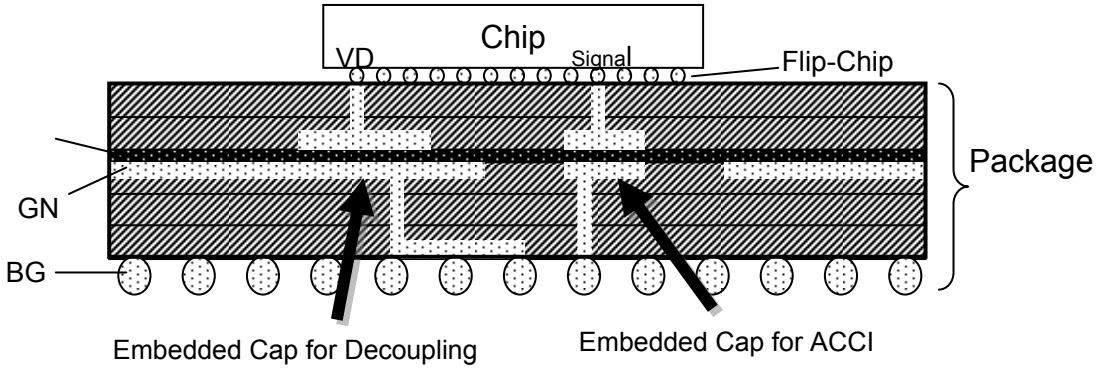

Figure 29 Cross section of structure showing embedded capacitors

A simplified model of the whole ACCI channel is shown in Figure 30. The stubs represent the routing channel inside the laminate package. The parasitics on the interface of the embedded capacitors -- including series inductance and resistance -- are not fully shown in the diagram but are included for more accurate simulation. The channel lengths are set to be a maximum of 10 cm for stubs and 1 m for channels. The coupling capacitor sizes are swept and analyzed using simulations.

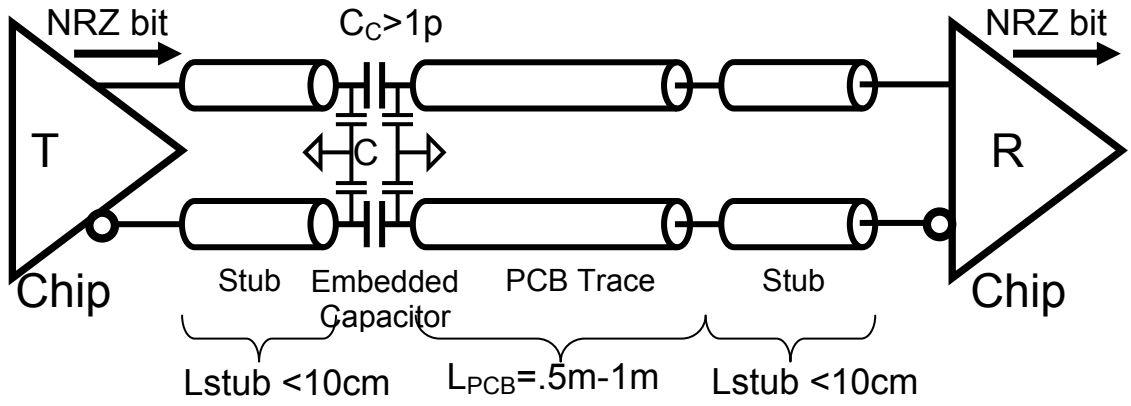

Figure 30 Simplified model of ACCI channel

The simulated transmission coefficient (S_{21}) with varying coupling capacitance (CC in the diagram) is shown in Figure 31. The location of the peak of the S_{21} shifts towards lower frequency as coupling capacitance increases. The magnitude of S_{21} increases as coupling capacitance increases but eventually reaches an upper limit when channel attenuation starts to dominate.

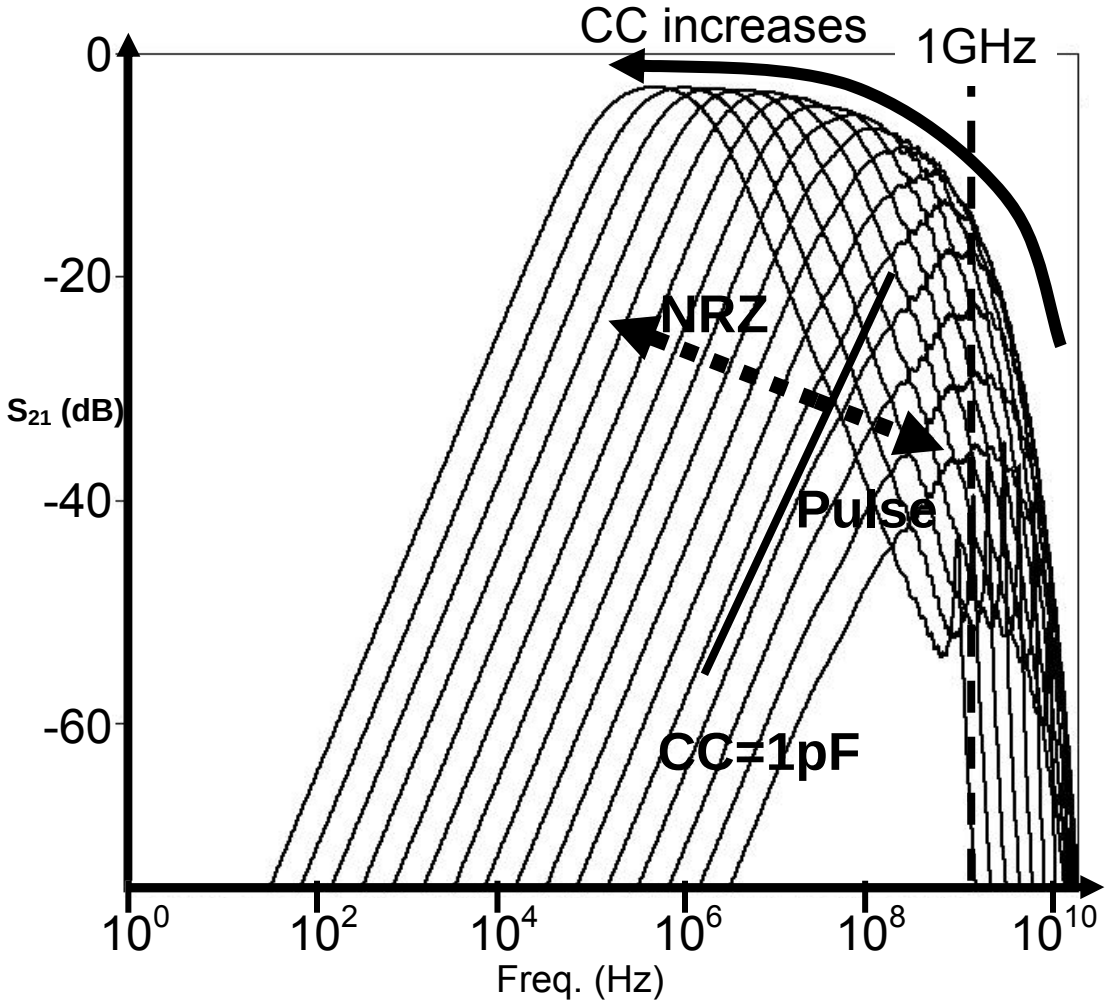

Figure 31 Simulated S_{21} with varying coupling capacitance

The whole ACCI channel can be treated as a band-pass filter (BPF) with center frequency and bandwidth which change as the coupling capacitor varies. Figure 32 shows that the center frequency is inversely proportional to coupling capacitance. The bandwidth, however, peaks when the coupling capacitance is around 1 pF. The length of the channel can also affect the response but is treated as a given constraint. This type of analysis allows for selection of the most suitable capacitors.

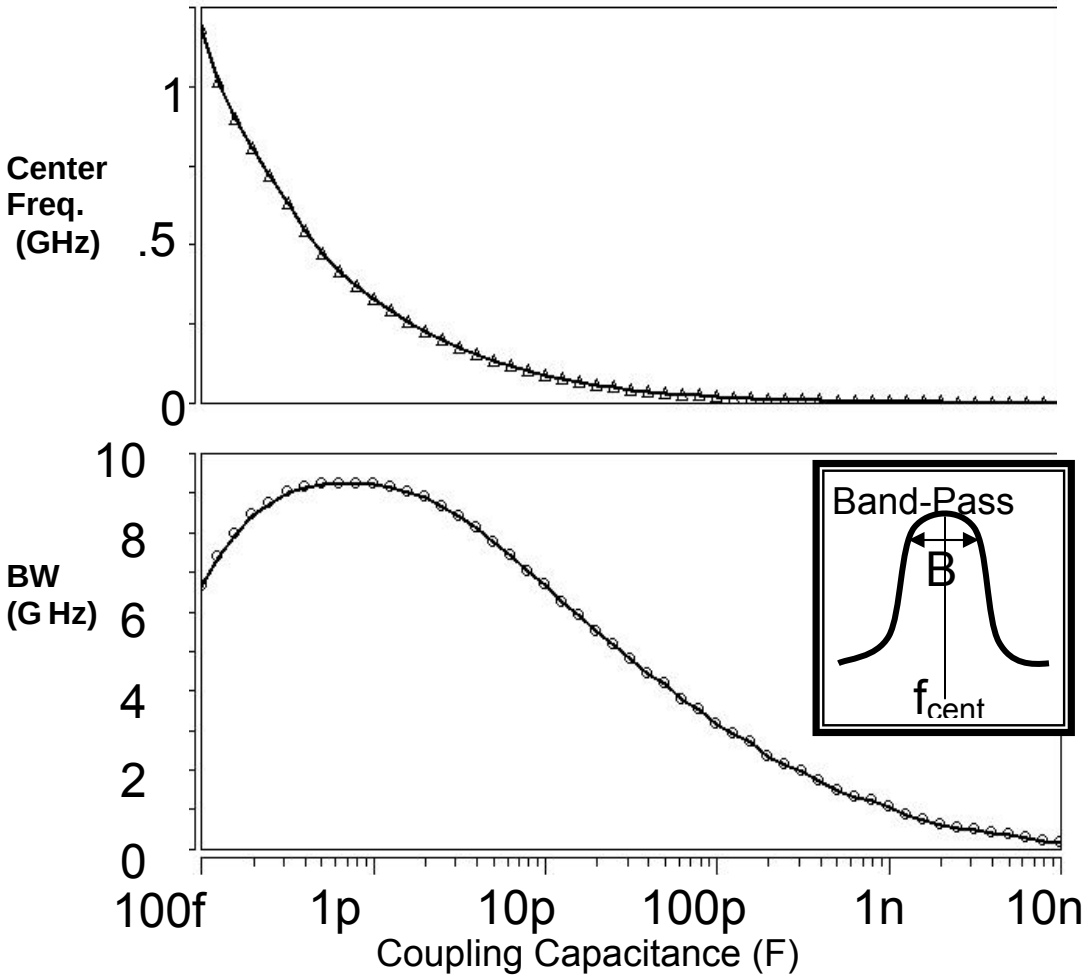

Figure 32 Center frequency and bandwidth as a function of coupling capacitance

ACCI utilizes the coupling capacitor to facilitate low-swing pulse signaling and a voltage mode transmitter is required. In our design, a simple pair of complementary progressively-sized inverters is used to transmit the NRZ bit stream.

It is possible to select the capacitor size and allowable channel length to be matched so that the passive equalization alone is enough to eliminate the ISI at specific operating frequencies. The circuit, however, would have very limited flexibility. Thus, a non-clocked version of the receiver-side FSE scheme is introduced to accommodate a wider range of parameter variation, such as channel length ranging from 10cm to 100cm and capacitor size ranging from 1pF to 10pF. The fractional delay permits equalization within the bit period, thus giving the agility to equalize a number of pulse and NRZ scenarios. The comparison of time-domain waveforms in Figure 33 shows that the FSE is able to cancel out the signal tails even when CC varies a lot.

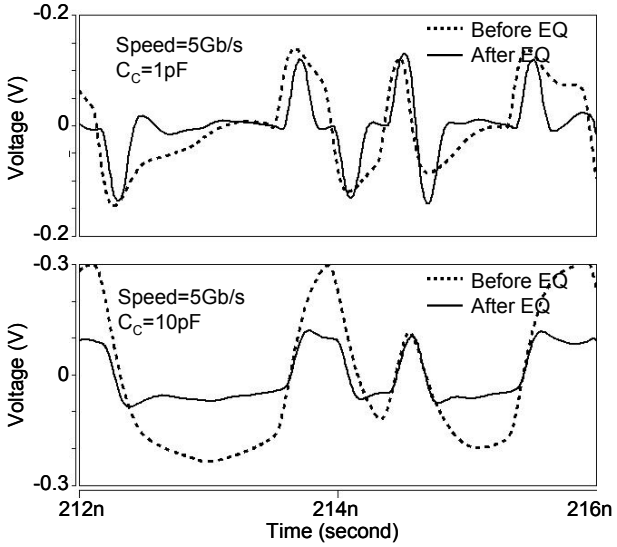

Figure 33 Comparison of time-domain waveforms

There are three major advantages gained by using this FSE receiver. First, the analog nature of the circuit makes the receiver work for both pulse signaling and NRZ signaling. Secondly, there is no clock involved in the transceiver thus saving the effort of designing the clock distribution network and saving power in the clock tree. Finally, the power consumption is low, compared with a digital filter at comparable speed with similar channel characteristics.

The simulated eye diagrams shown in Figure 34 demonstrate that a fractional finite input response (FIR) pulse receiver is capable of recovering the received pulse signals as well as NRZ signals. The transceiver is implemented in standard 0.18 CMOS technology and can send a 5Gps bit stream through a worst-case channel of [10cm]=[1pF]=[100cm]=[10cm]. The power consumption is 5mW for the transmitter and 15mW for the receiver.

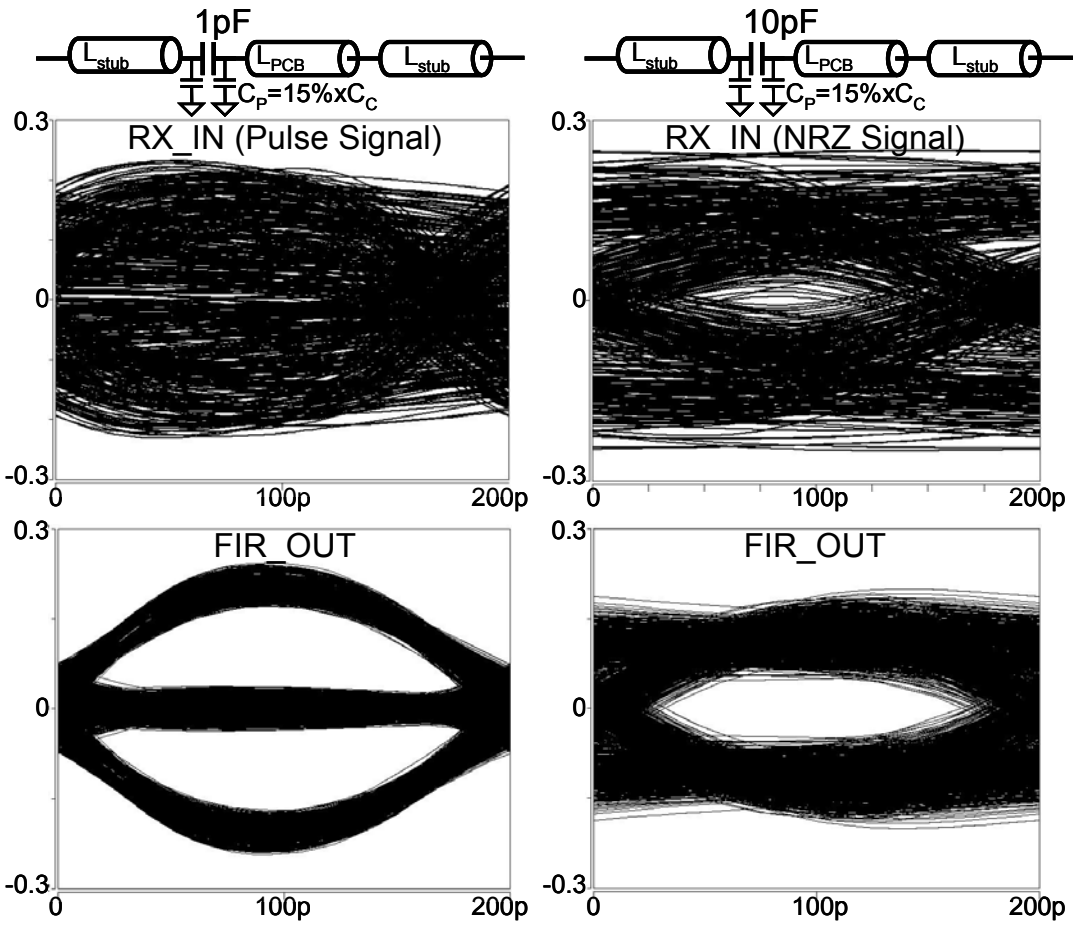

Figure 34 Simulated eye diagrams

Embedded capacitors combined with ACCI circuitry provide the advantages of low power, high speed, high flexibility, and relaxed circuit constraints in addition to the advantages that naturally come with embedded capacitors. The trade-offs between capacitor size, capacitor parasitics, channel length, and constraints on circuit performance are used to determine the most appropriate value for the embedded series capacitor. However, the performance is not highly sensitive to the value chosen. A fractional FIR pulse receiver is introduced to enable the channel to operate correctly across multiple length and capacitance ranges. When using this receiver, a series capacitance of around 1 pF is a good choice, giving an area of 1000 um^2, based on a density of 100 nF/cm^2. This combination is a good candidate for replacing SMT series capacitors with buried capacitors in backplane applications.

2.4 Task 4. High Density Capacitively Coupled Socket

The purpose of this task is to implement a socket that supports capacitively coupled structures. The system cross section is shown in Figure 35. The socket can be realized using HDI manufactured by EIT. The socket consists of pins protruding on the underside of a laminate package which can be mated to a PCB substrate. The green 1-mil spacer indicates a thin film developed by Sanmina to act as the dielectric in an array of capacitors formed by the interface of the underside of the laminate socket and the top of the board.

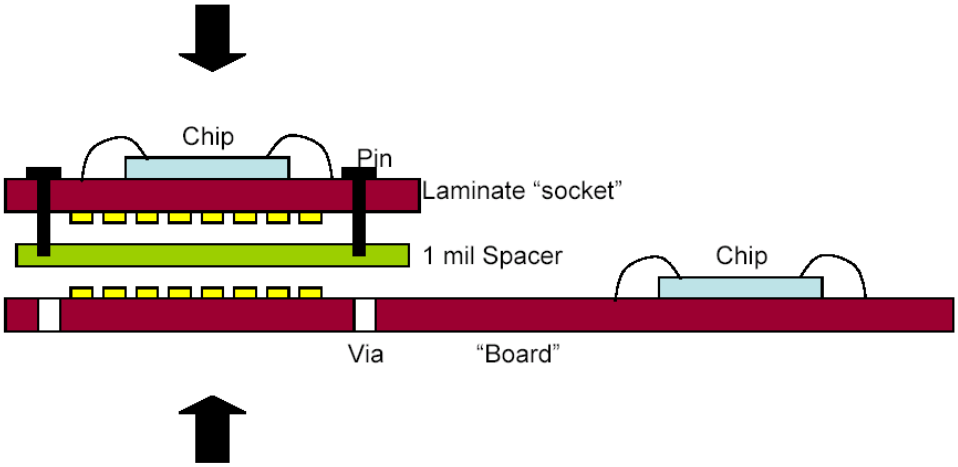

Figure 35 System cross section

Based on the channel characteristics, the transceiver circuitry for Task 4 and Task 2 can be consolidated into one design, given there is enough flexibility within the single chip. Therefore, a high-speed transmitter with transmitter side CT-FSE can be implemented in 0.13 CMOS technology. The detail of this circuitry is explained in Section 2.2.

Unfortunately, we suspended cooperation with EIT and Sanmina because of lack of interaction. Thus, in the end, we were unable to implement a real system. However, based on the solid results demonstrated in Task 2, we know it is possible to apply our theory to the capacitively coupled socket.

2.5 Task 5. High Density Inductively Coupled Socket

Because we were never able to achieve constructive interaction with Sanmina (they were very interested when the proposal was being written and then did not respond after we secured funding), we were not able to completely integrate, assemble and test the sockets, the last element of this task. However, similar to Tasks 2 and 4, we recognize that the difference in the channel models between this task and Task 6 were minimal, so the PCBs and CMOS circuitry developed for Task 6 actually fulfill the requirements of this task as well.

2.6 Task 6. Inductively Coupled Connector

Inductive coupling requires that a coil be built on both sides of the separable interface. The signaling principles are different than for capacitive coupling. The potential advantage of inductive coupling is that a larger air gap can be tolerated than with capacitive coupling. Thus our conclusion is that inductive coupling is more suited to separable connectors and sockets than capacitive coupling. This section describes the results of experiments that demonstrate the useful functioning of inductively coupled connectors.

Inductively coupled systems rely on pulse signaling due to the high-pass filter response of the transformer in the frequency domain acting as a differentiator in the time domain (Figures 36-38). Binary Non Return to Zero (NRZ) signals sent into the transformer are output as pulse signals with their amplitude and decay time determined by various parameters of the inductor, such as the diameter, the number of turns, turn width and spacing, and the spacing between the inductors comprising the transformer.

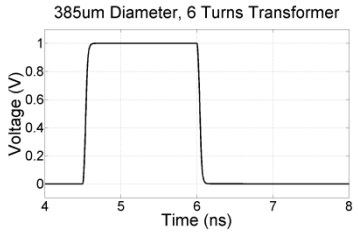

Figure 36 NRZ input for a 385 mm transformer

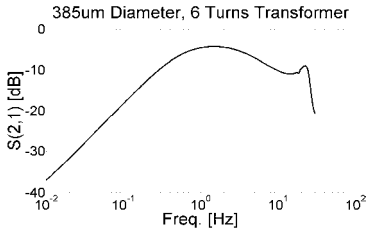

Figure 37 S_{21} for a 385 μm transformer

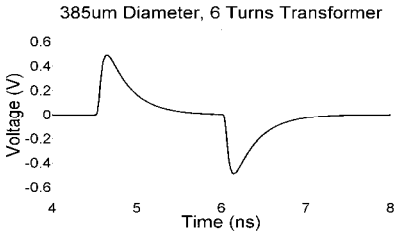

Figure 38 Output from a 385 μm transformer

When lossy transmission lines are used before or after the transformer in a complete inductive channel, the transmission lines act as low-pass filters, which combine with the inductor to form a band-pass response. In the time domain the result is similar to that of a transformer alone with a reduction in the magnitude of the resulting pulse while the decay time of the pulse is unaffected. Additionally, when the gap spacing between inductors in the transformer are increased, the effect in the time domain is also a reduction in the peak amplitude of the coupled pulse (Figure 40). Therefore to successfully signal over larger spacing between inductors or over lossy transmission lines, larger transformers that produce higher amplitude pulses are required.

As seen in Figure 39, the amplitude of the coupled pulses in the time domain can be increased by signaling over larger diameter transformers and/or using a transformer with more turns. However, increasing the peak amplitude of the pulse through transformer sizing also increases the time required for the pulse to decay back to the zero state. For slow speed signaling, this decay time may not be an issue, but as signaling speed is increased the slow decay time results in

ISI. Attempting to send in a '0' while the pulse is still decaying from the previous '1', can result in a reduction in the peak amplitude of the '0' pulse, closing the eye diagram as seen in Figure 41

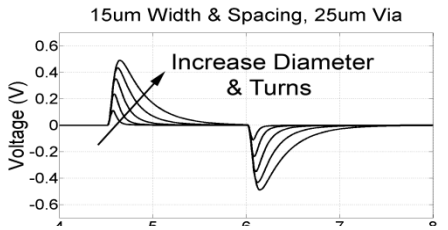

Figure 39 Time domain output for a range of transformer sizes

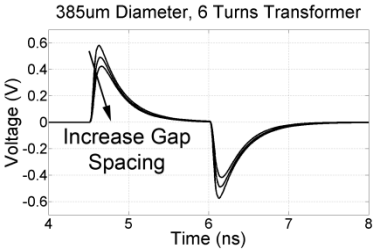

Figure 40 Time domain comparison of gap spacing between conductors

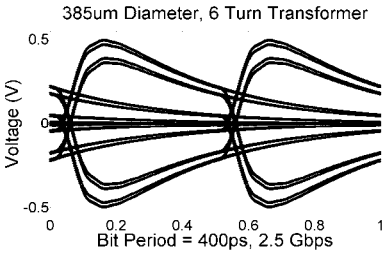

Figure 41 Eye diagram illustrating ISI due to the natural decay of the pulse tail

To enable high-speed signaling over larger transformers used for large gap spacing or in conjunction with lossy transmission lines by reducing ISI, we proposed using driver-side multi-bit fractional equalization. By using a digital FIR filter to de-emphasize the non-transition bits of a data stream, ISI was reduced. For inductive signaling this is only slightly beneficial since the transformer removes the DC content of the input signal and only the edges of the NRZ signal produce pulses. Extending Dally & Poulton's concept, we proposed using multiple levels within a single bit and over multiple bits, which we call multi-bit fractional equalization.

Using multi-bit fractional equalization, the optimal signal input to a transformer is created by preserving the rising and falling edges of the NRZ data, which produce the pulses output by the transformer. The slowly decaying tail of the pulse can be removed by de-emphasizing the DC component of the NRZ input. Using too little de-emphasis fails to adequately reduce the tail, while too much de-emphasis results in a reduction of the amplitude of the pulse produced by the next transition bit since the amplitude of the input signal was reduced. To equalize out longer tails, multiple bits after a transition may have to be equalized.

A multi-bit fractionally equalized input bit stream for a 385 µm (6 turns, 15 µm width & spacing) transformer is shown in Figure 42. In this example there are four possible signal levels for each of the two bits after a transition (0 → 1 or 1 → 0) bit. The peak amplitude is preserved for the first quarter of the transition bit, then the amplitude of the remaining portions of the transition bit and some of the second bit after a transition are reduced in order to minimize the decaying tail of the coupled pulse. Finally the input signal returns to full swing for the final half of the second bit after a transition and all subsequent bits in order to preserve the full swing from 1 → 0 or 0 → 1 and thus produce a coupled pulse with maximum amplitude. The NRZ input and fractionally equalized input for series of bits can be seen in Figure 43, followed by the resulting output from the 385 µm transformer in figure 44. The fractionally equalized input results in a significant reduction in the tail and a uniform amplitude of each pulse.

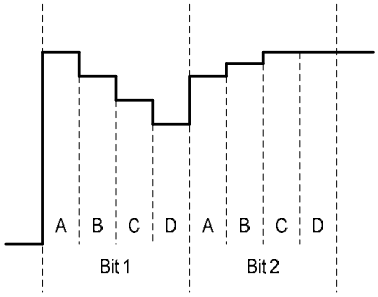

Figure 42 2-bit, 4-phase fractional equalization

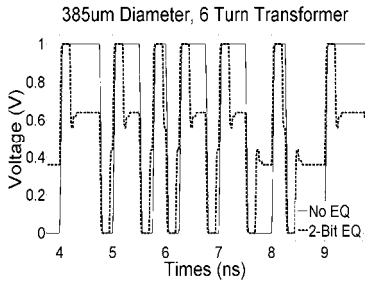

Figure 43 NRZ and multi-bit fractionally equalized input to a 6-turn transformer at 4 Gbps

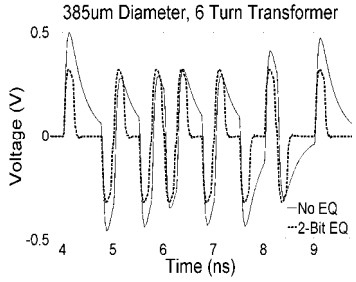

Figure 44 Output from a 6-turn transformer with and without fractionally equalized input at 4 Gbps

The reduction in ISI possible using mult-bit fractional equalization is apparent when viewing the eye diagrams for both the equalized (Figure 45) and the non-equalized NRZ input (figure 46) to 385 µm transformer. The fractionally equalized input stream allows for the transformer to be used at signaling speeds far greater than those possible without equalization.

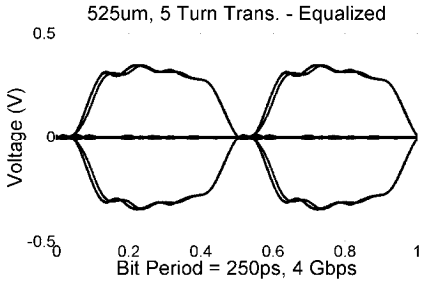

Figure 45 Multi-bit 4 Gbps signaling over the same transformer as Fig. 47

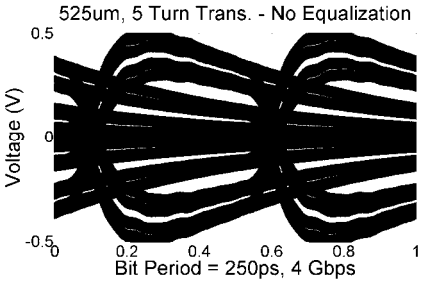

Figure 46 Gbps signaling over a transformer with too slow a natural decay

Additionally, a channel with two transformers and a transmission line, Figure 47, can be equalized for using the same multi-bit equalization scheme as for a single transformer. When two transformers are present in the channel, the first integrates the NRZ data resulting in a pulse, which is then integrated on both the rising and falling edges of the pulse by the second transformer, resulting in a heavily attenuated pulse with a small opposite magnitude pulse following. This "double pulse" produces ISI when signaling at higher data rates and can be minimized with multi-bit fractional equalization as shown in figures 48-50 where the black indicates signaling without equalization and the red uses equalization.

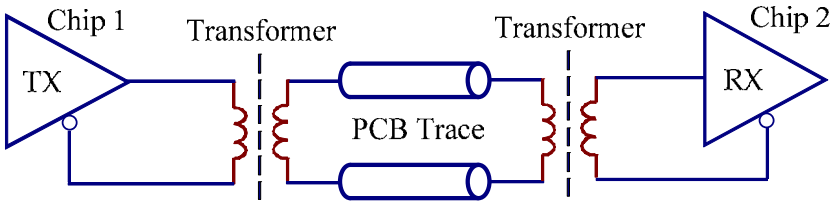

Figure 47. An ACCI channel with two transformers and a transmission line

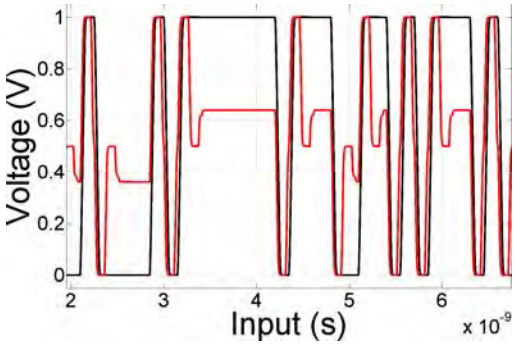

Figure 48 Equalized and non-equalized input onto a channel with two transformers

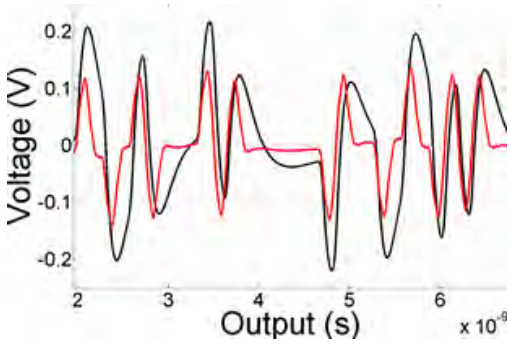

Figure 49 Output from a two transformer channel for the data input in Fig. 50

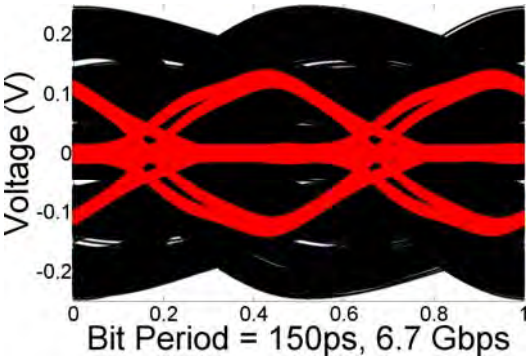

Figure 50 Eye Diagram for the channel output from Fig. 51

Figure 51 shows the layout of the inductively coupled circuitry:

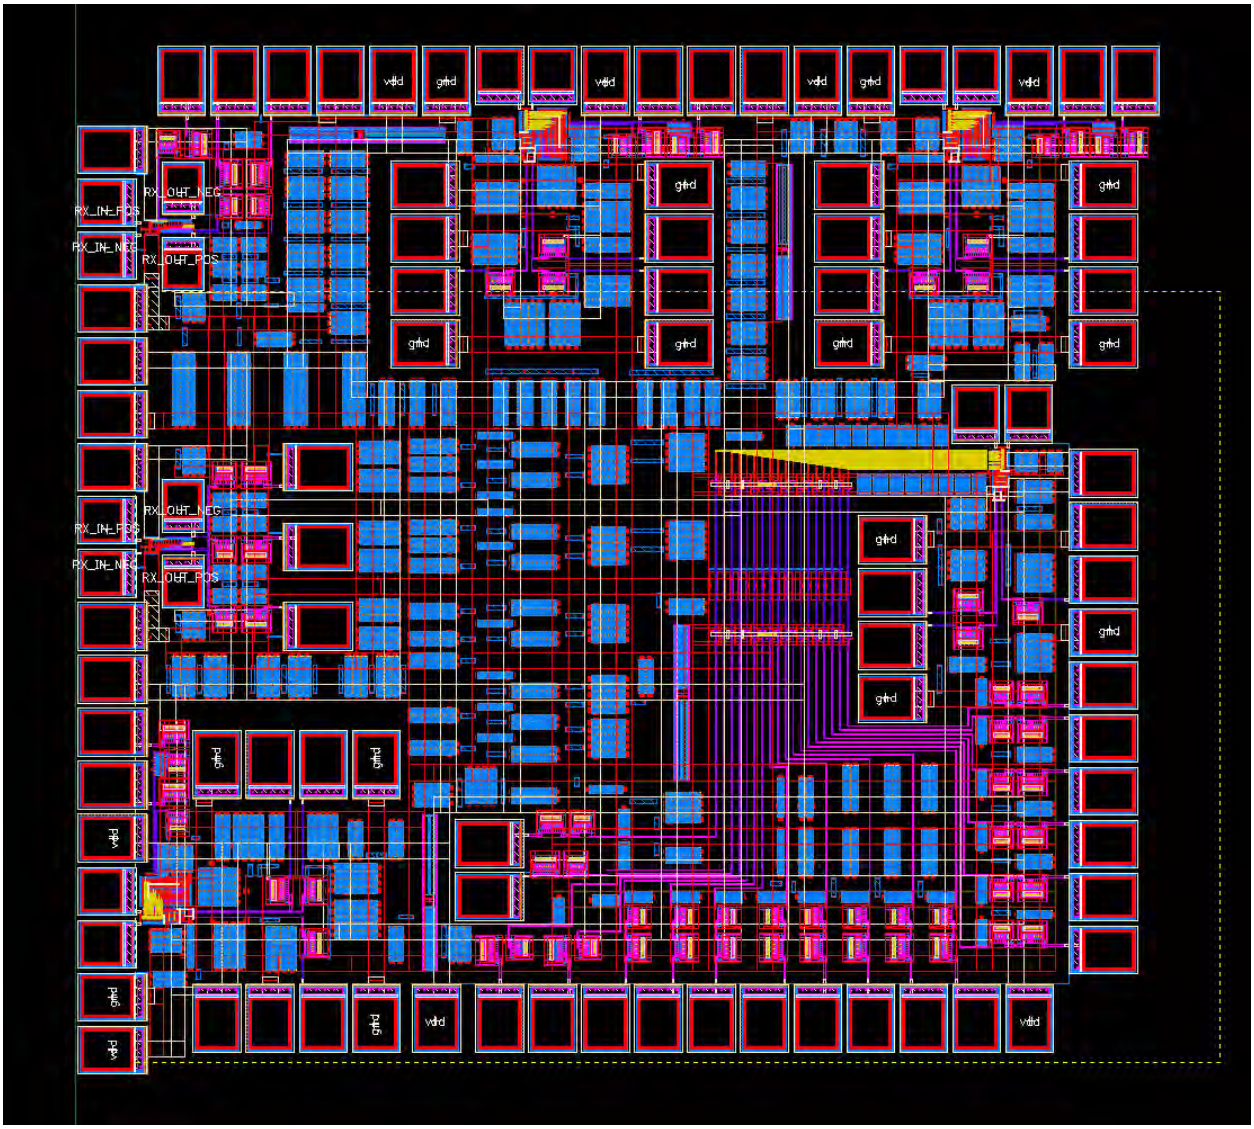

Figure 51 Die design for inductive coupled designs

2.6.1 Inductive Channel Design

In order to test the multi-bit fractional equalizing driver design, an inductive channel was created using two PCBs, each with an inductor, flipped on top of each other. By preventing the inductors from touching and thus shorting out, a transformer is created over which signaling can be achieved. The PCBs are connected using an array of small screws, overlapping just in the portion of each containing the inductors. Signals are input and output using SMA connectors and sufficient area for decoupling capacitors is provided (Figure 52).

Figure 52 Unpopulated driver and receiver printed circuit boards

Variations of the driver / receiver PCB pair were designed to produce 4 distinct inductive channels listed below. A complete flipped transformer during testing is shown in Figure 53.

1) 3 Turn, 4 mil W / 4 mil S, 2134 µm Diameter, In Phase
2) 3 Turn, 4 mil W / 4 mil S, 2134 µm Diameter, Out Phase
3) 2 Turn, 6 mil W / 6 mil S, 1727 µm Diameter, In Phase
4) 1 Turn, 4 mil W / 4 mil S, 914 µm Diameter, Out Phase

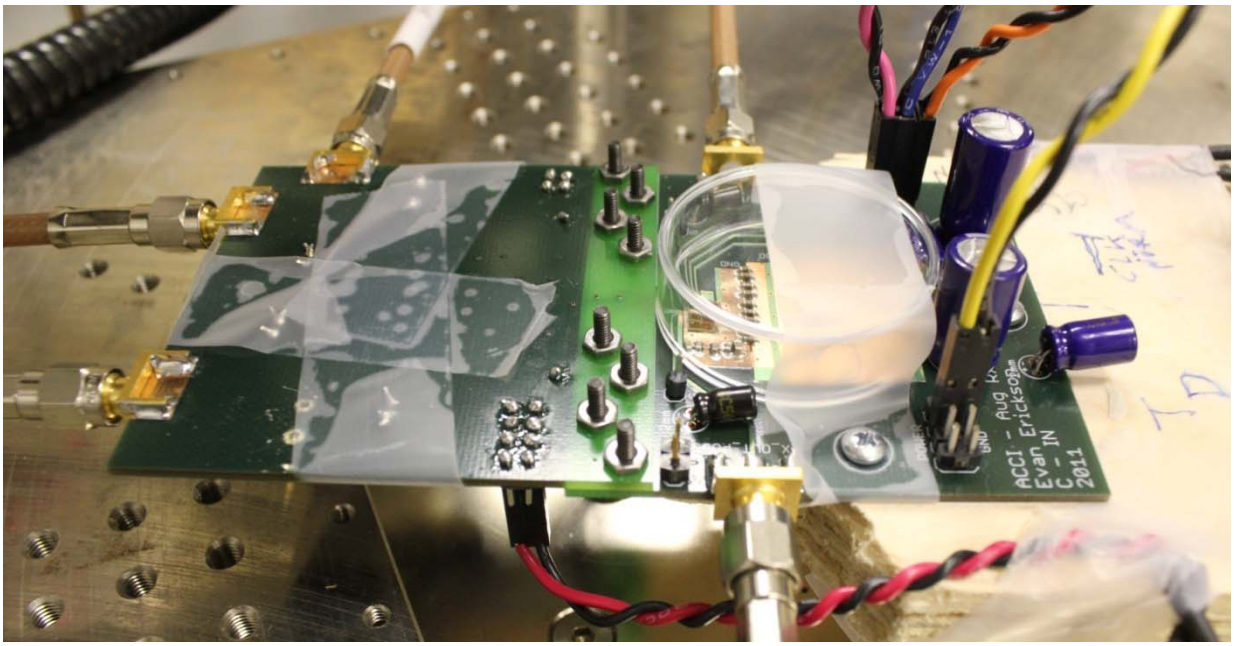

Figure 53 Complete transformer during testing

The driver die is upside down on the left, the receiver is visible through the protective cap on the right.

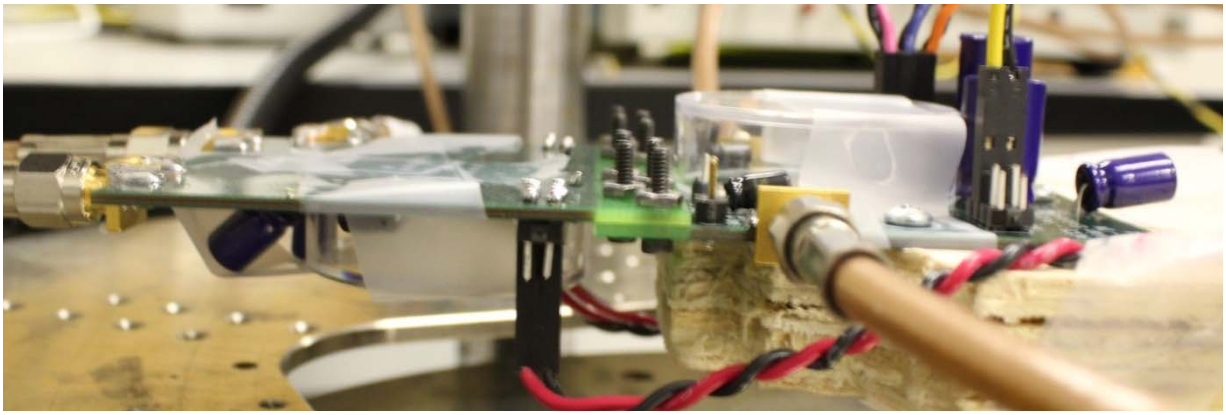

Figure 54 Side view of the inductive channel during testing

Due to the presence of wirebonds, plastic caps on the left and right protect the driver and receiver respectively. A 2 turn, 6 mil width and spacing, 1727 µm in phase transformer is between the screw

Initial simulation of the transformers were completed in Ansoft's HFSS, a 3-D electromagnetic field solver (Figure 55), which produced S-parameter data for characterization of the inductive channels in the frequency domain (Figure 56). This initial simulation data helped us determine the parameters of the transformers we'd like to use to investigate further in PCB testing.

Mathworks Matlab was used to fit the frequency data from HFSS and model signaling in the time domain. This technique enabled of multi-bit fractional equalization by enabling us to experiment with various equalization schemes on a high level (Figure 57).

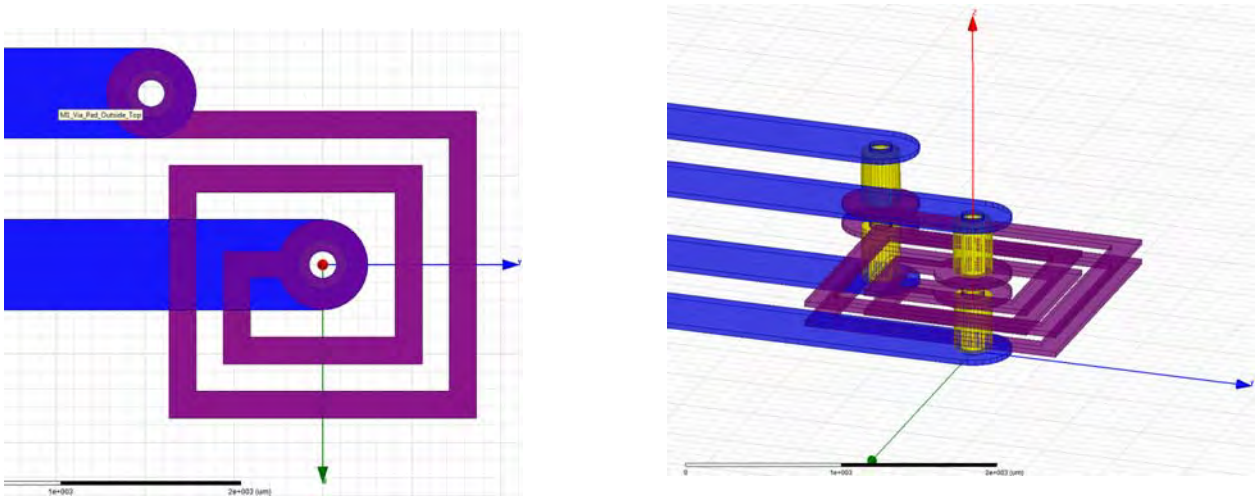

Figure 55 HFSS model of a 1727 µm transformer composed of two inductor on PCBs with a small air gap separating them

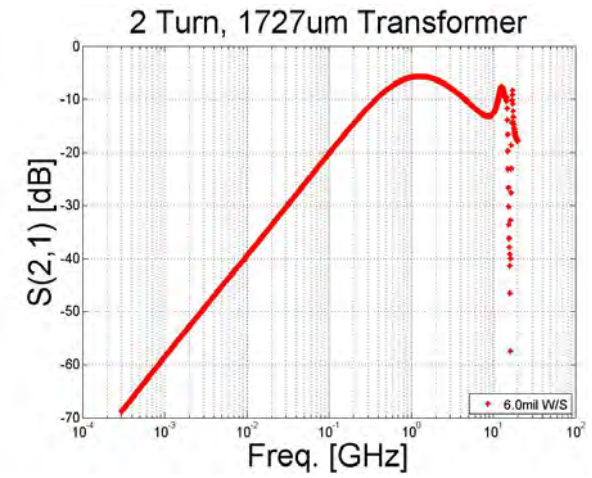

Figure 56 Frequency domain (S_{21}) for a 1727 µm transformer on PCB

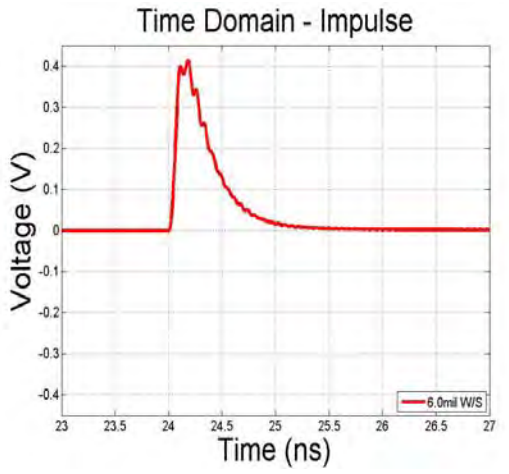

Figure 57 Time domain response of 1727 µm transformer

2.6.2 Fractionally Equalizing Inductive Driver Testing

Testing of the fractionally equalizing driver was completed to demonstrate signaling over larger size transformers created with a simple PCB process. The driver is designed to achieve up to 5 Gbps signaling rates, but since the larger transformers produce pulses with slow decay time, speeds above 2-3 Gbps experience significant inter-symbol interference as evidenced by the eye diagrams that follow. Data and a differential clock signal were input using a bit-error rate tester (BERT), control signals to setup the output current of the driver using a digital pattern generator, separate power supplies for the driver and receiver were used, and finally a digital sampling oscilloscope was used to measure the received data. The test setup is pictured in Figure 58.

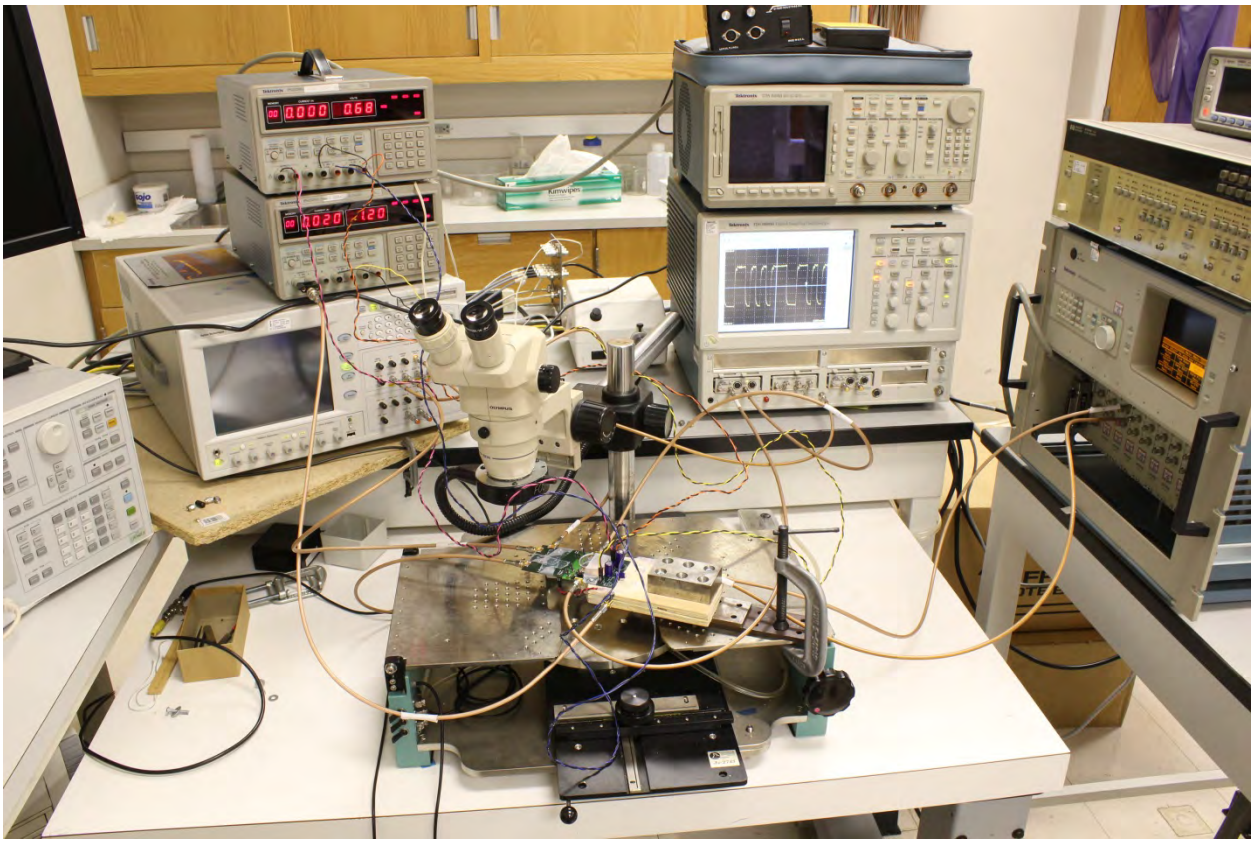

Figure 58 Setup for fractional equalizing inductive current-mode driver

Figures 59 through 64 show the output of the multi-bit fractionally equalizing driver, however the driver is current-mode and all the measured data below is voltage sent into a 50 ohm system.

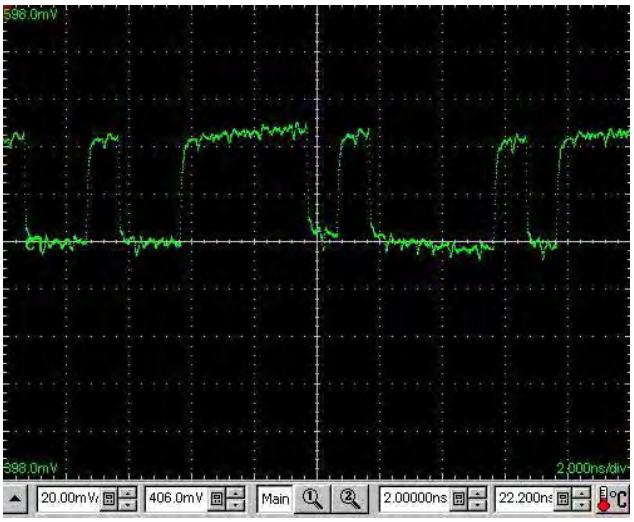

Figure 59 Random data stream showing minimum driver output of 0.8mA

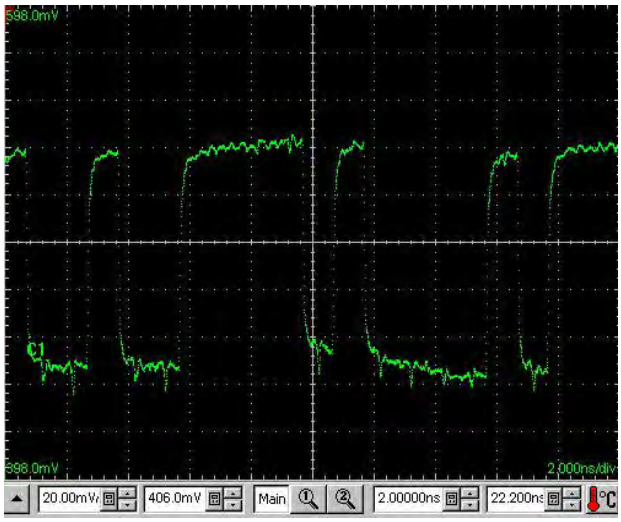

Figure 60 Random data stream showing maximum driver output of 2.0mA

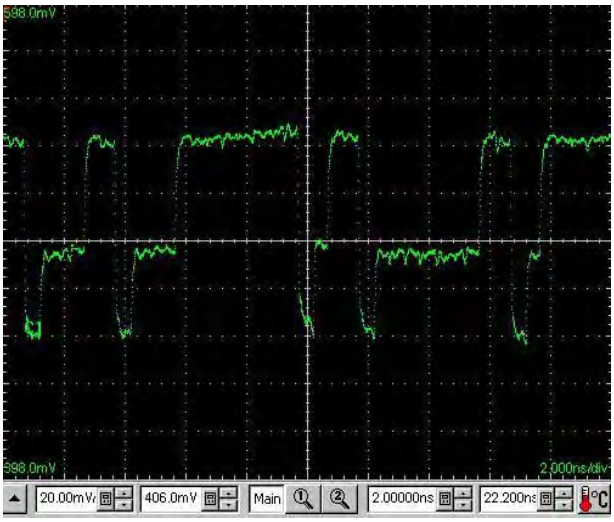

Figure 61 Random data stream showing the first half of tap 1 turned on and all others off

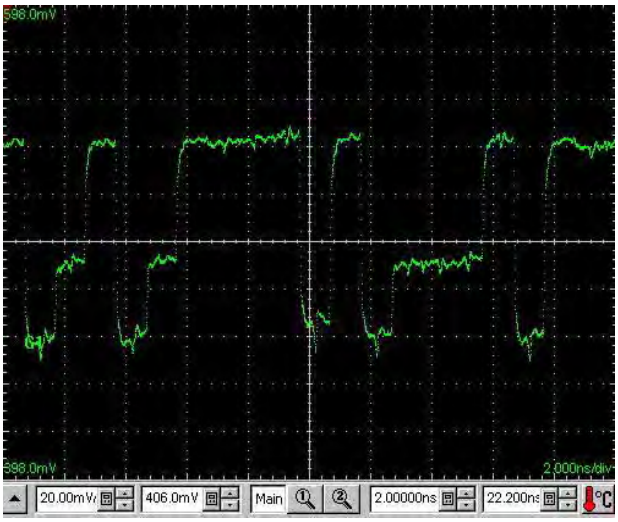

Figure 62 Random data stream showing the first and second halves of tap 1 turned on and all others off

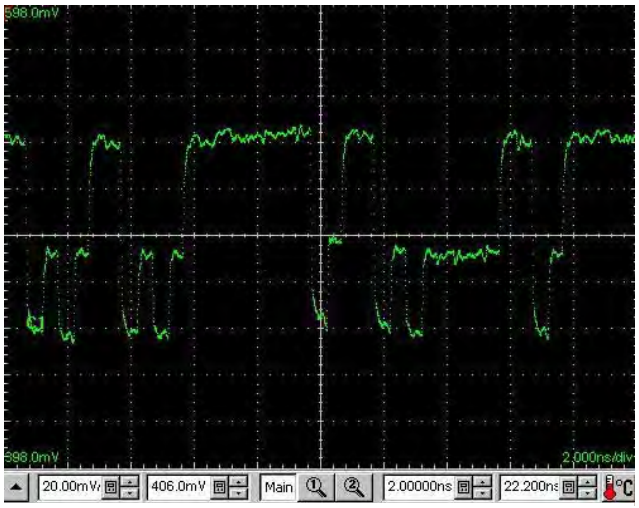

Figure 63 Random data stream showing the first halves of tap 1 & 2 on and all others off

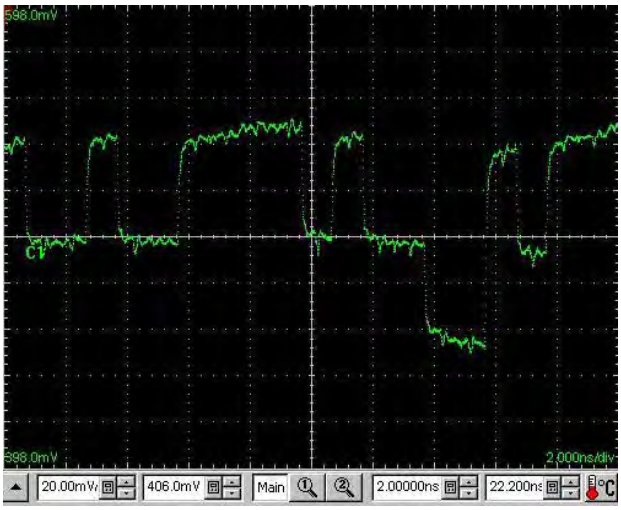

Figure 64 Random data stream showing taps 1 & 2 off and the final remaining tap on

The first two figures show the minimum and maximum output of the driver, by setting the baseline tap on and all other differential drivers off in the minimum case and all drivers on in the maximum case. Figures 59-64 show various configurations of the fractional taps being enabled or disabled to prove the driver is working correctly. Due to the driver being current-mode and only voltage being possible to test, the driver runs out of headroom and only displays correctly on the falling edge of the bit stream. The output is valid however because the correct output is seen when the output is measured across the transformer. Using this driver, the optimal equalization scheme for an inductive channel can be found.

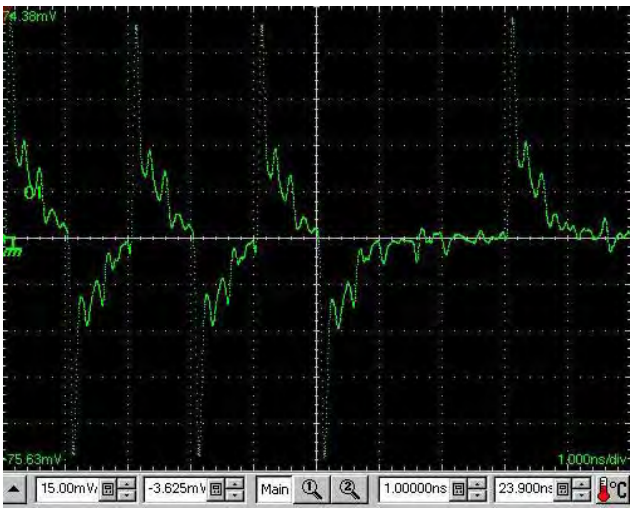

Figure 65 Channel output showing a full swing NRZ input at 1 Gbps; power is 20mA

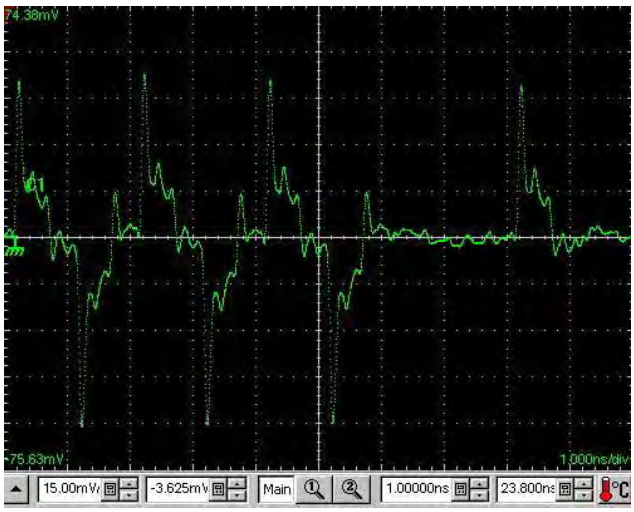

Figure 66 Channel output for full swing input for the transition bit only at 1 Gbps; power is 17 mA

For these experiments, the two-turn 1727 µm transformer described previously was used. The output from the channel is shown below as pulses in the time domain and as the corresponding eye diagram in Figures 65 through 71. Figures 65 and 66 show the effects of driver side equalization on the channel output. As the data rate increases the amount of ISI increases and thus the eyes close in the eye diagrams. A transformer this large cannot support signaling greater than 3 Gbps.

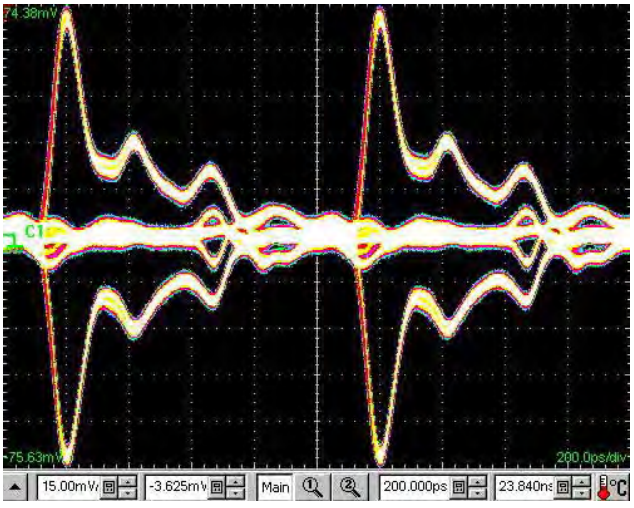

Figure 67 Eye diagram for channel output using an NRZ input signal at 1 Gbps

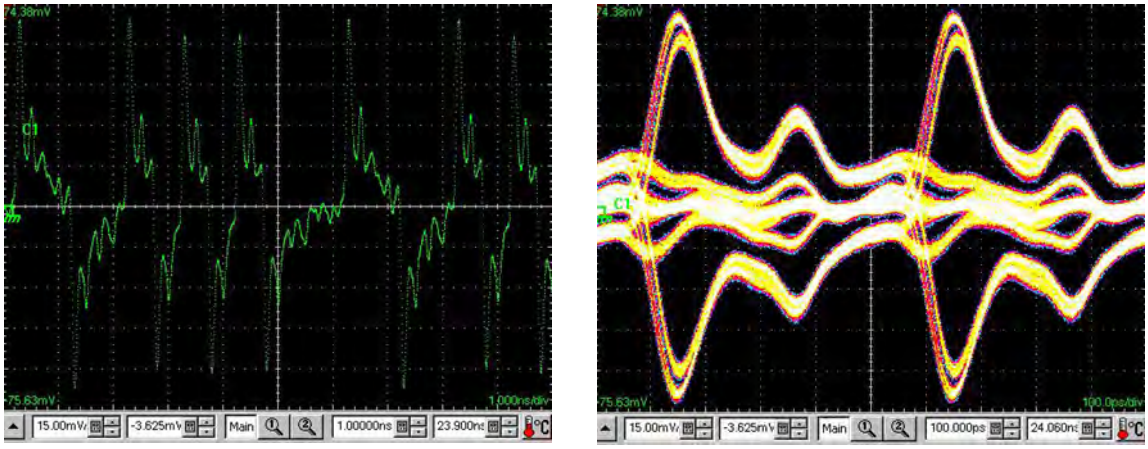

Figure 68 Channel output and eye diagram using an NRZ input signal at 2 Gbps

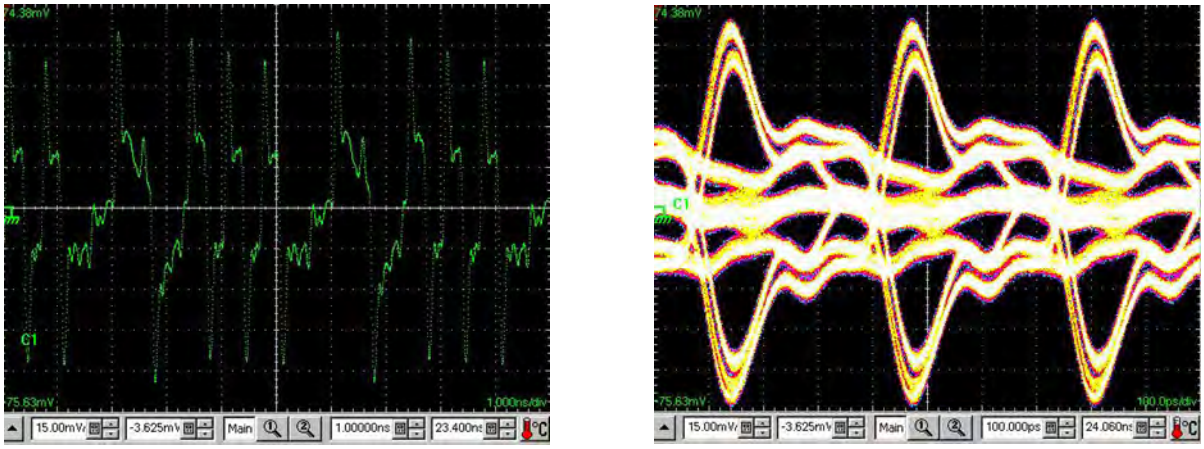

Figure 69 Channel output and eye diagram using an NRZ input signal at 3 Gbps

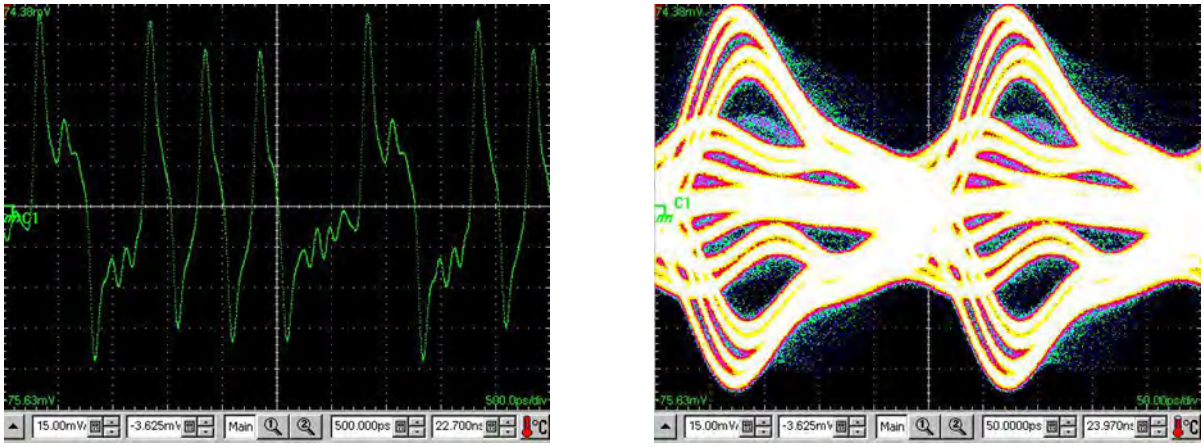

Figure 70 Channel output and eye diagram using an NRZ input signal at 4 Gbps

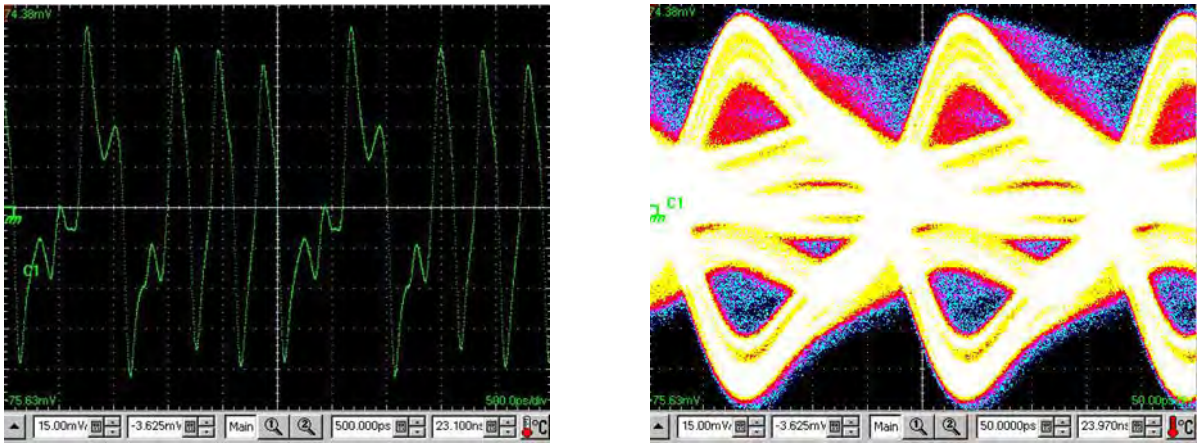

Figure 71 Channel output and eye diagram using an NRZ input signal at 4.95 Gbps

The current output from the channel is received by the current-mode pulse receiver, converted to voltage, amplified, returned to full-swing NRZ data, and finally amplified again. The receiver outputs 100mV data while driving a 50 ohm load. Figures 72 through 75 show the received data and the corresponding eye diagram, when the data rate reaches 4 Gbps the receiver is no longer able to recover the data as evident in the last figure.

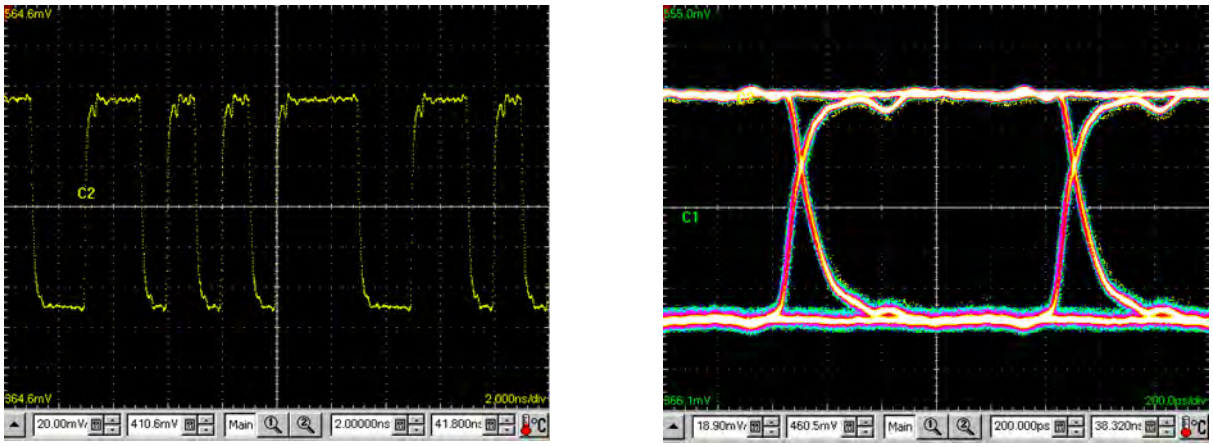

Figure 72 Receiver output and eye diagram using an NRZ input signal at 1 Gbps

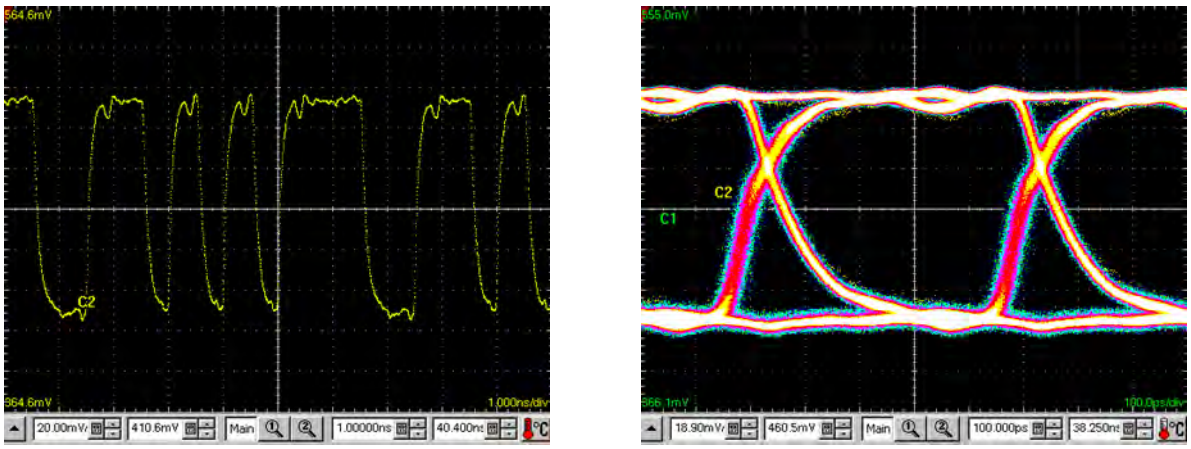

Figure 73 Receiver output and eye diagram using an NRZ input signal at 2 Gbps

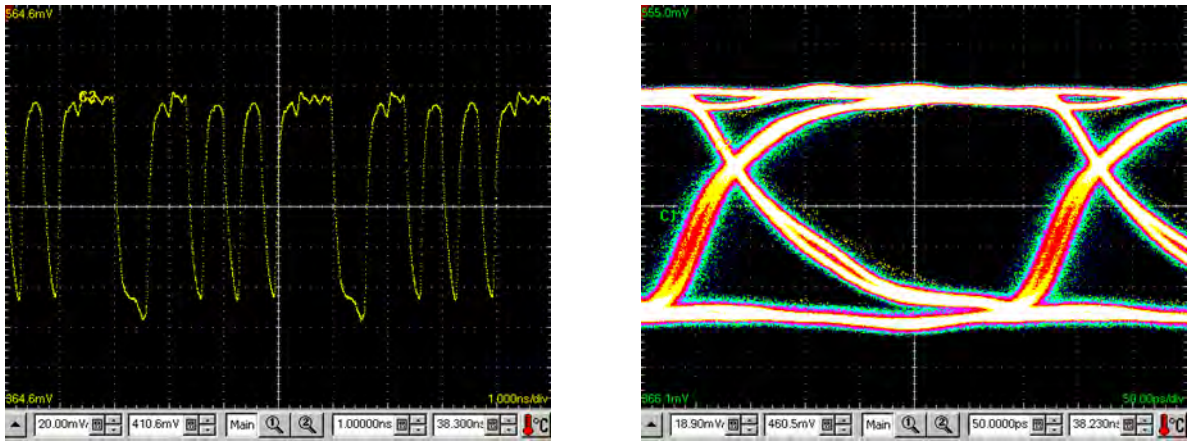

Figure 74 Receiver output and eye diagram using an NRZ input signal at 3 Gbps

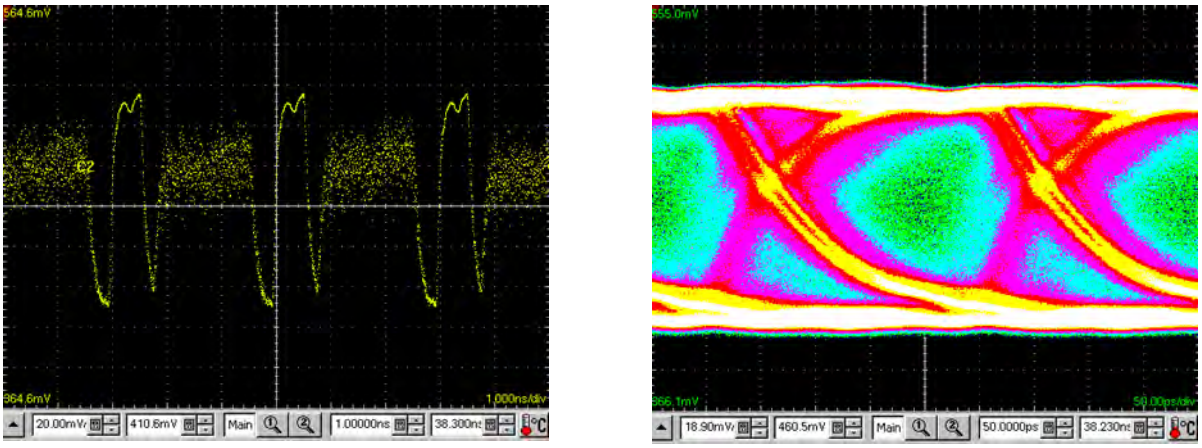

Figure 75 Receiver output and eye diagram using an NRZ input signal at 4 Gbps

Signaling over an inductive channel composed of two inductors on PCBs which are flipped onto each other has been demonstrated. A current-mode fractionally equalizing driver and current-mode pulse receiver have been developed and tested and speeds greater than 3 Gbps signaling over large transformers have been achieved.

In conclusion, we demonstrated that high quality signaling can be delivered over inductively coupled connectors consisting of just two turns on each side of the connector or socket structure. These inductors could be built at a pitch as tight as 8 mil (0.2 mm) using a high density laminate structure such as EIT's.

2.7 Task 7. Space Experiment

The Space Experiment was designed to demonstrate the use of the Appliqué Sensor Interface Module (ASIM) to collect in-flight data from an FPGA-based bit BERT.

Our initial intention was to characterize a data path formed using a high-speed Xilinx Virtex II RocketIO data path to feed data to an existing ACCI test chip. Figure 76 shows the original idea:

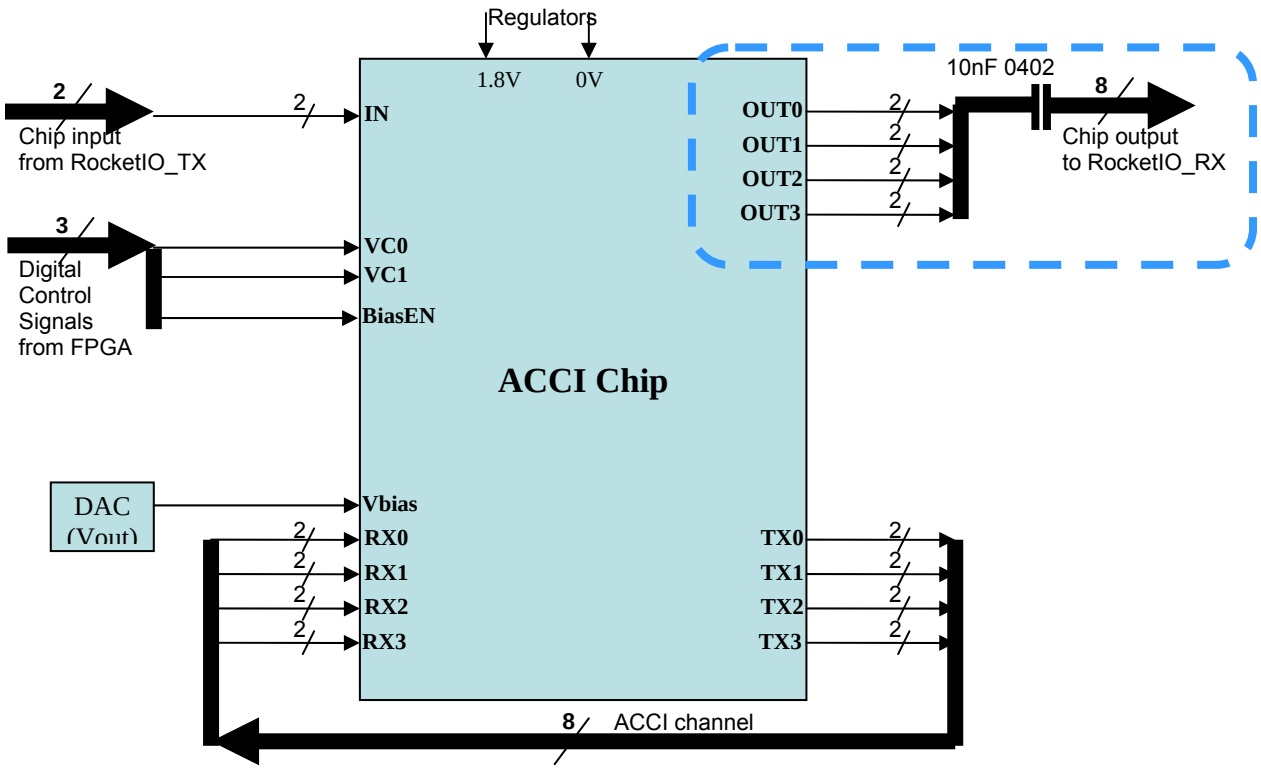

Figure 76 Original conception of the in-flight module architecture

Figure 77 shows the motherboard (the red board on the floor of the module) and the ASIM (the green board in the center):

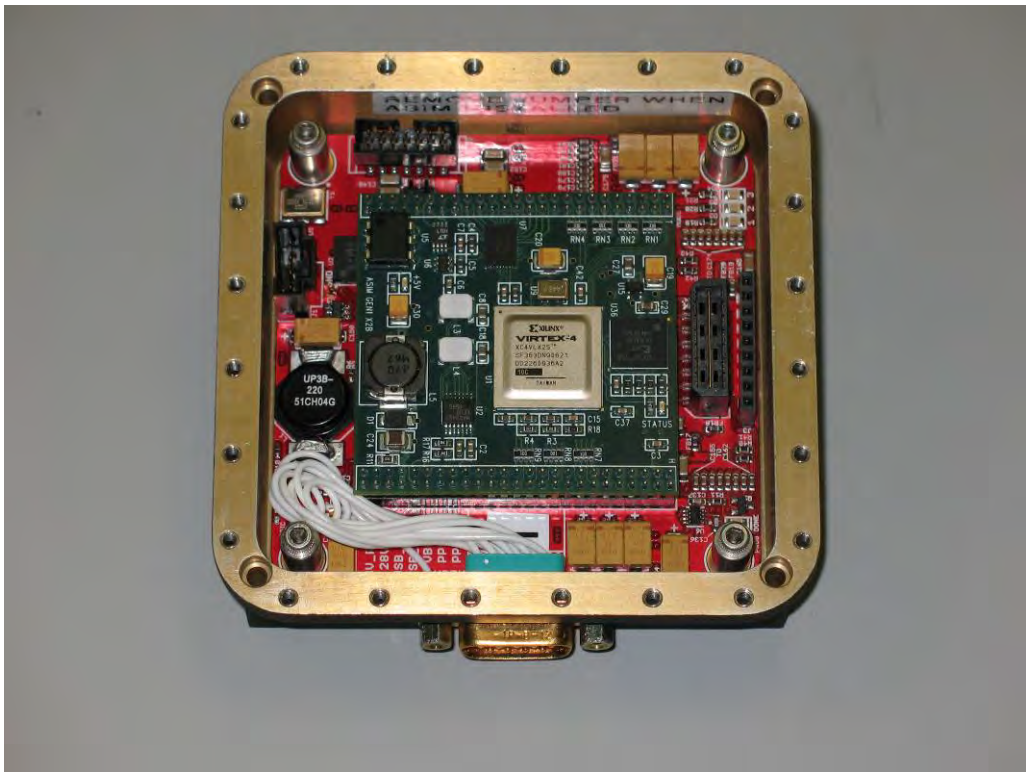

Figure 77 In-flight module showing motherboard (red) and ASIM (green)

The motherboard contains the Virtex II FPGA (which is underneath the ASIM in the picture.).

The cover of the module (Figure 78) is a circuit board that has the ACCI test chip as a wire bonded component:

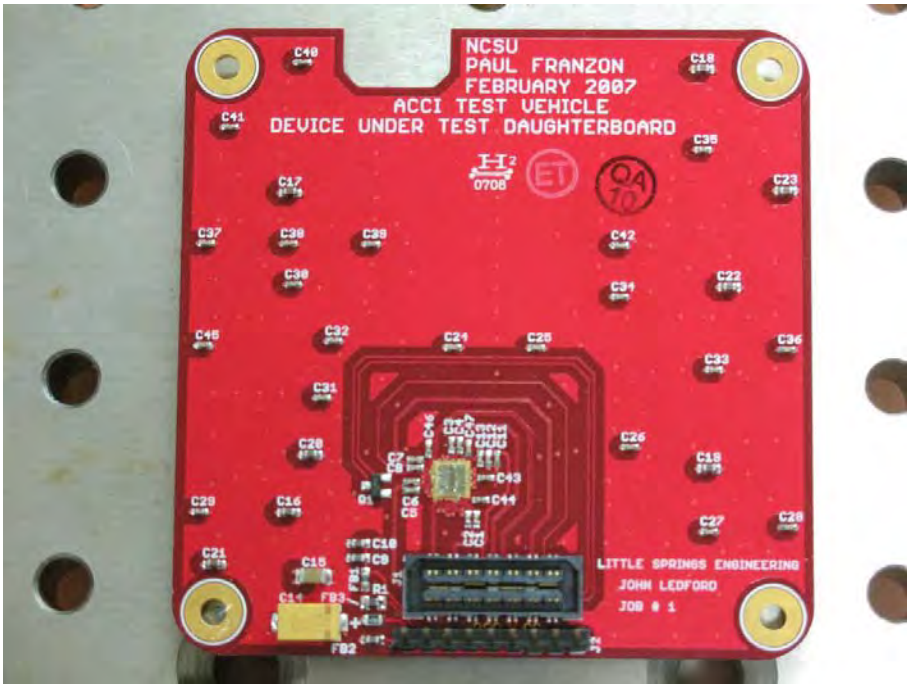

Figure 78 In-flight module cover

The ACCI chip is lower center of the picture. Figure 79 shows a close-up of the chip:

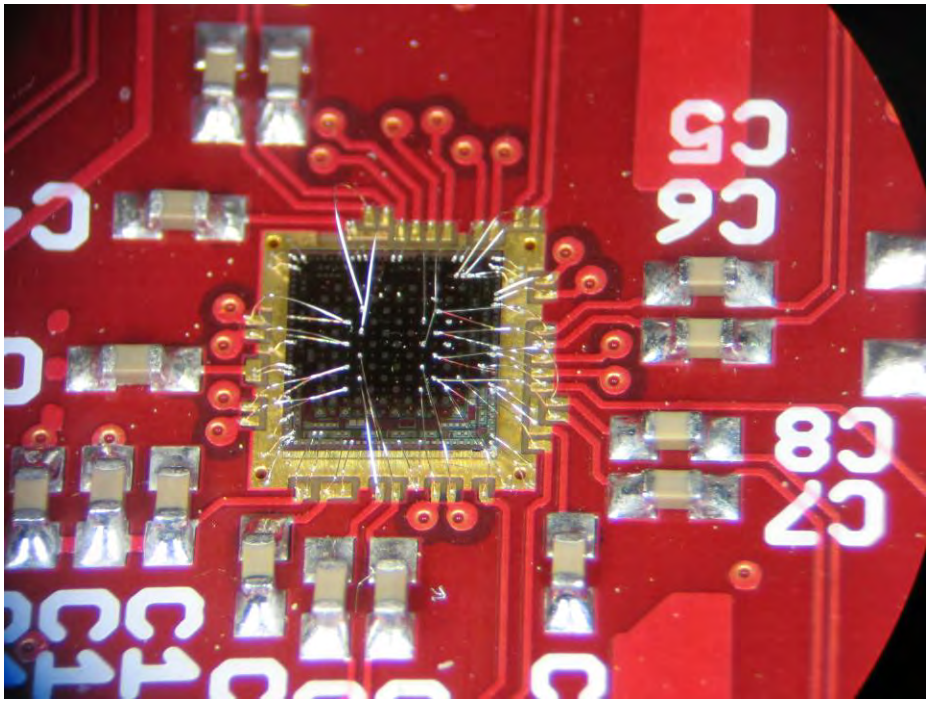

Figure 79 Close-up of the ACCI chip mounted on the in-flight module cover

Figure 80 shows the module with daughterboard installed, undergoing testing in our laboratory:

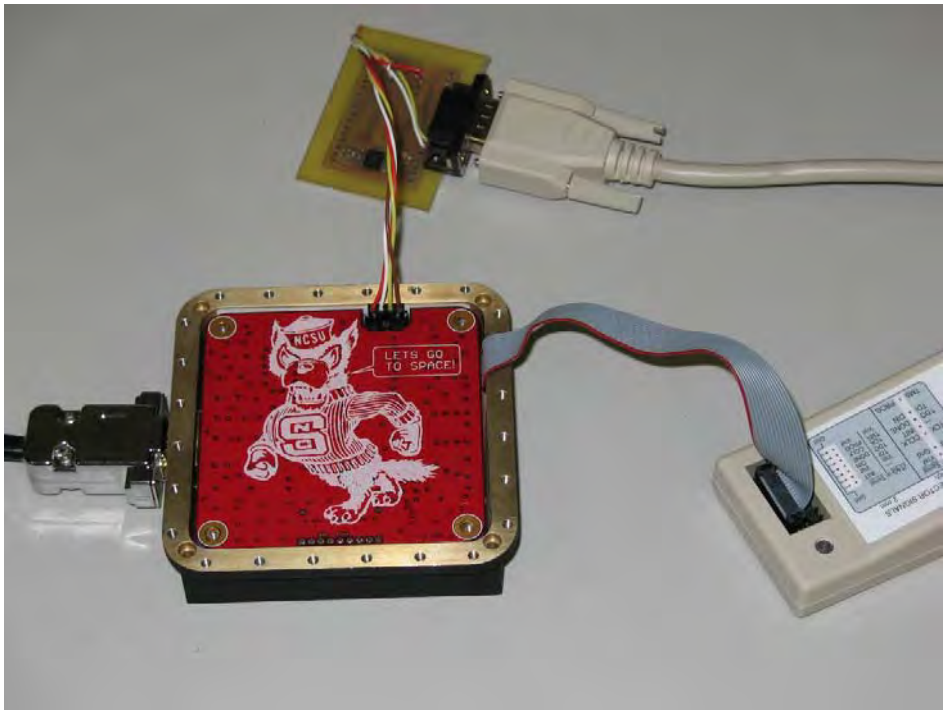

Figure 80 The in-flight module undergoing testing at NCSU

The module worked fine in our laboratory, but when tested in the real system by MicroSat there was a failure mode that we had not seen. We had to make an engineering change to the module to use a different ACCI chip than originally planned. We think that this was necessary because signal integrity issues caused voltage level mismatches in the RocketIO between the ACCI chip and the FPGA on the module's motherboard. We came up with a novel solution to this problem. We used a different ACCI chip that did not require RocketIO. In order to meet the schedule we had to work out a way to use the replacement ACCI chip without redesigning the module's motherboard or daughterboard. We were able to find a way to bond the replacement ACCI chip to the existing daughterboard in a way that allowed the system to work properly.

Figure 81 shows the replacement chip bonded to the original daughterboard:

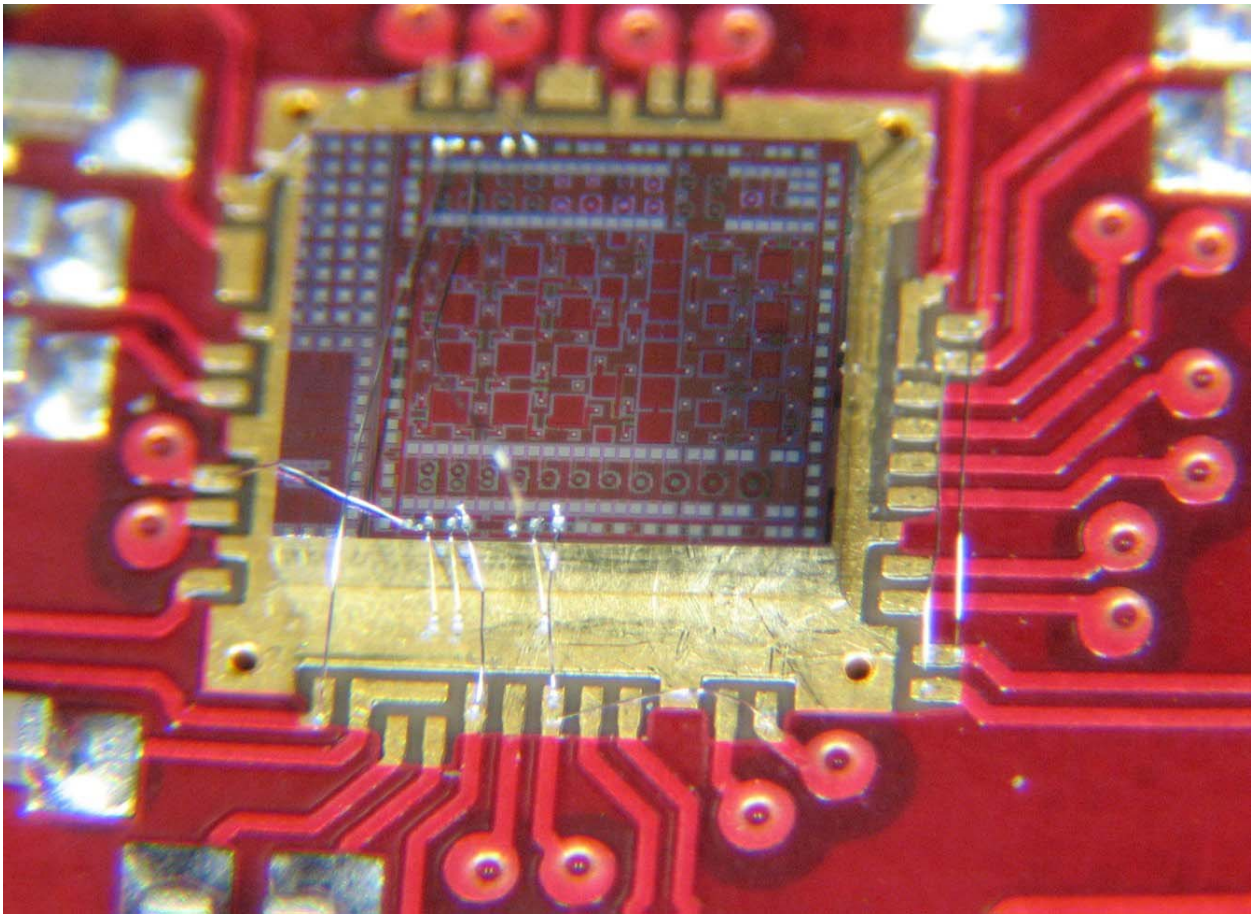

Figure 81 The replacement chip as bonded on the in-flight module's cover

After shipping the module we had to do some firmware updates, but we were easily able to get the module ready to launch before the May 19, 2009 liftoff. To date we have gotten no feedback from MicroSat Systems on the ultimate performance of the module.

3.0 CONCLUSION

The overall goals of the project, as stated in the proposal were :

1. Demonstrate, to the level of commercial acceptance, a capacitive coupled chip-package integrated structure, operating at 6 Gbps per pin and with a pitch of 65 um.
    a) 6 Gbps/65 um in a 0.18 um die on a silicon package.
    b) On a laminate package, our objective is to demonstrate 3 Gbps with a 100 um pad pitch.
2. Demonstrate, to the level of commercial acceptance, an inductively coupled socket, operating at over 1 Gbps and with a pitch of 1.3 mm.
3. Build a space-capable demonstration of a system using ACCI, capable of being flown in a USAF test rocket.

Goal 1a was achieved even before we began this project as a result of project F29601-03-2-0135.

Goal 1b was partially achieved in that we were easily able to exceed the target data rate over the channel. However, due to vendor issues outside our control we were not able to meet the pad pitch requirements.

Similarly to Goal 1b, we were able to easily meet the target data rate for Goal 2 over an equivalent channel, but could not implement the socket of that specific density without cooperation from EIT or Sanmina.

Goal 3 was completely achieved.

3.1 Academic Acceptance.

Several academic groups have published papers on this topic since we started work in this area. These include:

- Keio University
- University of Tokyo
- Technical University of Munich
- University of California, Los Angeles
- University of California, Davis
- University of California, Berkeley
- ARCES-University of Bologna, Bologna, Italy
- Yale University
- Johns Hopkins University

3.2 Commercial Acceptance

The main outcomes likely to be commercially relevant are (1) circuit structures that support embedded capacitors for AC coupled channels, and (2) inductively coupled connectors. Item (1) is likely to be adopted when the circuit speeds required for these channels reach the data rates that the parasitic of discrete components will become unacceptable. Item (2) is currently being explored with two commercial entitities. The discussions are confidential in nature.

Appendix A: Journal Papers

The following publications resulted from the work in this project:

1. K. Chadrashekar, J. Wilson, E. Erickson, Z. Feng, S. Mick, J Xu, and P.D. Franzon, “Inductively Coupled Connectors and Sockets for Multi-Gbps Pulse Signaling,” in IEEE. Trans. Adv. Pack., Vol. 31, No. 4, Nov 2008, pp. 749-758.

2. L. Luo, J.M. Wilson, S.E. Mick, J. Xu, L. Zhang, P.D Franzon, “3 Gb/s AC Coupled Chip-to-Chip Communication Using a Low Swing Pulse Receiver,” IEEE JSSC, Vol. 41, No. 1, Jan 2006, pp. 287-296.

3. S. Mick, L. Luo, J. Wilson, P. Franzon, “Buried Bump and AC Coupled Interconnection Technology," IEEE Trans. Adv.Packaging, 27(1), Feb 2004, pp. 121-125.

Appendix B: Conference Papers

1. B. Su, W. S. Pitts, P. D. Franzon, J. Wilson, "A zero power consumption Multi-Capacitor structure for voltage summing in high-speed FFE," Proc. IEEE Electrical Performance of Electronic Packaging, pp. 13-16, Oct. 2010.

2. E. Erickson, J. Wilson, K. Chadrasekar, P.D. Franzon, “Multi-bit fractional equalization for multi-Gb/s inductively coupled connectors,” Proc. IEEE EPEPS, Oct. 2009, pp. 121-24.

3. B. Su, P. Patel, S.W. Hunter, M. Cases, and P.D. Franzon, “AC coupled backplane communication using embedded capacitor,” IEEE EPEP 2008, Oct. 2008, pp. 295-298.

4. L. Luo, J. Wilson, S. Mick, J. Xu, L. Zhang, E. Erickson, P. Franzon, *"A 36Gb/s ACCI Multi-Channel Bus using a Fully Differential Pulse Receiver,"* Proceedings of the IEEE, Custom Integrated Circuits Conference, 2006, Sept. 2006, pp. 773-776.

5. J. Xu, J. Wilson, E. Erickson, P. Franzon, "Pulse Signaling in Inductively Coupled Sockets and Connectors," Paper 13.2, SRC Student Symposium, Oct. 2006.

6. J. Wilson, L. Luo, S. Mick, B. Chan, H. Lin, P. Franzon, *“AC Coupled Interconnect using Buried Bumps for Laminated Organic Packages,”* Proc. Electronic Components and Technology, ECTC, May/June 2006.

7. J.M. Wilson, S. E. Mick, J. Xu, L. Luo, E. L. Erickson, P.D. Franzon, *"Considerations for Transmission Line Design on MCMs using AC Coupled Interconnect with Buried Solder Bumps,"* Proc. 10th Annual Workshop on Signal Propagation on Interconnects, Berlin, Germany, May 2006, pp. 281-282.

8. L. Luo, J. Wilson, J. Xu, S. Mick, P. Franzon, “Signal integrity and robustness of ACCI packaged systems,” Proc. IEEE EPEP, Oct. 2005, pp. 11-14.

9. J. Xu, J. Wilson, S. Mick, L. Luo, "2.8 Gbps inductively coupled interconnect for 3D ICs," 2005 Symposium on VLSI Circuits, June 2005, pp. 352-355.

10. K. Chandrasekar, Z. Feng, J. Wilson, S. Mick, P. Franzon, “Inductively Coupled Board to Board Connectors,” ECTC’05, 31 May – 3 June, 2005, pp. 1109-1113.

11. J. Xu, J. Wilson, S. Mick, L. Luo and P. Franzon, “2.8 Gb/s Inductively Coupled Interconnect for 3-D ICs,” Japan VLSI Symposium, June 2005.

12. K. Chandrasakar, Z. Feng, J. Wilson, S. Mick, P. Franzon, “Inductively Coupled Board-to-Board Connectors,” ECTC 2005, May 2005.

13. L. Luo, J. M. Wilson, S.E. Mick, J. Xu, L. Zhang, and P.D. Franzon, "A 3 Gb/s AC Coupled Chip-to-Chip Communication," 2005 International Solid State Circuits Conference, San Francisco, February 2005.

14. Paul Franzon, Angus Kingon, Stephen Mick, John Wilson, Lei Luo, Karthik Chandrasakhar, Jian Xu, Salvatore Bonafede, Alan Huffman, Chad Statler, Richard LaBennett, Invited Paper, "High Frequency, High Density Interconnect Using AC Coupling," Fall MRS Conference, Boston, MA, December 2003.

15. Paul Franzon, Stephen Mick, John Wilson, Lei Luo, Karthik Chandrasakhar, Invited Paper, "AC Coupled Interconnect for High-Density High-Bandwidth Packaging," Proc. SPIE, Microlectronics: Design, Technology and Packaging, Perth, Australia, December 2003, pp. 67-69.

16. Paul Franzon, Stephen Mick, John Wilson, Lei Luo, Karthik Chandrasakhar, Invited Paper, "AC Coupled Interconnect for High-Density High-Bandwidth Packaging," Japan SSDM, Tokyo, Japan, September 2003.

17. Jian Xu, Stephen Mick, John Wilson, Lei Luo, Karthik Chandrasakhar, Paul Franzon, "AC Coupled Interconnect for Dense 3-D Systems," Proc. IEEE Conference on Nuclear Science and Imaging, Seattle, Washington, October 2003.

18. S. Mick, L. Luo, J. Wilson, P. Franzon, "Buried solder bump connections for high-density capacitive coupling," IEEE Electrical Performance of Electronic Packaging, 2002, pp. 205-208.

19. S. Mick, P. Franzon, A. Huffman, Packaging Technology for AC Coupled Interconnection, IMAPS Flip-Chip Conference, July 2002.

20. S.E. Mick, L. Luo, J.M. Wilson, P.D. Franzon, "Buried Solder Bump Connections for High-Density Capacitive Coupling," IEEE Electrical Performance on Electornic Packaging, October 2002.

21. S. E. Mick, J. M. Wilson, and P. Franzon, "4 Gbps AC Coupled Interconnection," Invited Paper, IEEE Custom Integrated Circuits Conference, May 12-16, 2002, pp. 133-140.

LIST OF SYMBOLS, ABBREVIATIONS AND ACRONYMS

3D	three dimensional
A	amperes
AC	alternating current
ACCI	AC-coupled-interconnect
A_{CC}	top plate area
ADC	analog-to-digital converter
ADS	a radio-frequency circuit analysis program
A_{PAD}	pad area
ASIM	Appliqué Sensor Interface Module
BCB	benzocyclobutene
BER	bit error rate
BERT	bit error rate tester
BPF	band-pass filter
C4	Controlled Collapse Chip Connection
C	Celsius
C_C	generic capacitance variable
CMOS	complementary metal oxide semiconductor
C_p	horizontal parasitic fringe capacitance
C_{pBG}	horizontal parasitic fringe capacitance from bottom plate to ground
C_{pTG}	horizontal parasitic fringe capacitance from top plate to ground
CTE	coefficient of thermal expansion
CT-FSE	continuous-time fractional spaced equalization (or equalizer)
CTLE	continuous-time linear equalizer
CTRL	control
DAC	digital-to-analog converter
DC	direct current
DFE	decision-feedback equalizer
d	thickness
EIT	Endicott Interconnect Technologies
EM	electromagnetic
EQ	equalizer
F	Farads
FFE	feed-forward equalizer
FIR	finite impulse response
FPGA	field programmable gate array
FR4	a printed circuit board dielectric material
h	height
HDI	high-density interconnect
HFSS	an electromagnetic field solver program
HSPICE	a circuit simulation program
IBM	International Business Machines
I/O	input/output
ISI	inter-symbol interference
k	dielectric constant

K	dielectric constant
LCR	inductance capacitance resistance
LSB	least significant bit
MCM	multi-chip module
MCNC	Microelectronics Center of North Carolina
MIM	metal-insulator-metal
m	meters
MOSIS	a semiconductor fabrication company
MultiCap	multi-capacitor structure
NCSU	North Carolina State University
NFET	n-type field effect transistor
NRZ	non-return-to-zero
N_{tap}	number of taps
PCB	printed circuit board
PRBS	pseudo-random bit sequence
Q	quality factor
RoHS	restriction of hazardous substances
R	tap-weight resolution
R&D	research and development
RTI	Research Triangle Institute
RX	receiver
RZ	return-to-zero
S_{21}	forward transmission coefficient
SMA	a type of connector (SubMiniature version A)
SMD	surface-mount device
SMT	surface-mount technology
SPICE	a circuit simulation program
s	spacing
S	spacing
TX	transmitter
USAF	United States Air Force
V	volts
w	width
W	width

DISTRIBUTION LIST

DTIC/OCP
8725 John J. Kingman Rd, Suite 0944
Ft Belvoir, VA 22060-6218 1 cy

AFRL/RVIL
Kirtland AFB, NM 87117-5776 2 cys

Official Record Copy
AFRL/RVSE/Cheth Ouch 1 cy