Exploration of FPGA Interconnect for the Design of Unconventional Antennas

Abhay Tavaragiri
abhay@vt.edu

Jacob Couch
jacouch@vt.edu

Peter Athanas
athanas@vt.edu

Bradley Department of Electrical and Computer Engineering
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061-0111 USA

ABSTRACT

The programmable interconnection resources are one aspect that distinguishes FPGAs from other devices. The abundance of these resources in modern devices almost always assures us that the most complex design can be routed. This underutilized resource can be used for other unintended purposes. One such use, explored here, is to concatenate large networks together to form pseudo-equipotential geometric shapes. These shapes can then be evaluated in terms of their ability to radiate (modulated) energy off the chip to a nearby receiver. In this paper, an unconventional method of building such transmitters on an FPGA is proposed. Arbitrary shaped antennas are created using a unique flow involving an experimental router and binary images. An experiment setup is used to measure the performance of the antennas created.

Categories and Subject Descriptors

B.4.1 [**Input/Output and Data Communications**]: Data Communications Devices—*Transmitters*

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Antenna Design, Geometric Routing, Embedded Transceivers, Hidden Transmitter

1. Introduction

Traditionally, circuits constructed on FPGAs are designed using Hardware Description Languages (HDL) like Verilog or VHDL. Digital and analog signals are presented to the FPGA at the boundary of the FPGA fabric. All signals are then converted to digital signals for use within the FPGA. This practice leads to the logical conclusion that the target circuits on an FPGA are typically digital in nature.

FPGA'11, February 27–March 1, 2011, Monterey, California, USA.

Modern FPGAs have highly complex, programmable, interconnection structures, fabricated using many metal layers. For example, the Virtex-4 FPGA is fabricated utilizing an 11-metal-layer process [14]. Typically, 80% of the configuration bits in a given bitstream control the wire segments and programmable switches that make up the programmable interconnect network [5]. This network is made up of long segment routing, dedicated routing, local routing, and switch matrices, which are in turn made of wire segments and programmable interconnect points (PIPs). There are instances where the programmable interconnect can be used for other things beyond connecting two or more wires. Such a process would not involve the CLBs in the FPGA; clearing the path for a non-digital circuit.

The creation of non-digital circuits on an FPGA is an interesting concept with the potential for extending the opportunities for FPGAs. It does, however, seem improbable since non-traditional circuits are neither supported by HDLs nor the FPGA primitives.

One non-traditional use of FPGAs is the creation of geometric shapes and filled patterns, by using the programmable interconnect as a pallet. A *shape* is defined here as a single two-dimensional curve that consists of a long network of interconnected wire segments forming a single net that completely fills the interior of the curve. Such structures serve no obvious function in a typical digital circuit, yet can provide other auxiliary unintended functionality such as the ability to create various shaped antennas. For the purposes of this paper, an antenna is defined as a structure that radiates energy in a controlled manner. This energy could then be used to transmit a digital signal. One theory investigated here is that these antennas in turn act as conventional antennas, by resonating at various frequencies depending upon the shape, and radiating energy when properly excited. Due to the abundance of the programmable interconnect, these geometric structures can reside *transparently* on top of a normal (digital) design without interfering or changing the functionality of the digital circuit. A signal or a set of signals from the digital circuit can be used to modulate a *carrier signal* whose frequency is selected based upon the resonant frequency of the geometric structure. The end result is an on-chip circuit capable of radiating selective information off-chip.

Xilinx FPGAs are used in this study since the underlying configuration bitstream can be readily modified at abstraction levels lower than HDL (XDL in this case). These low-level manipulations, which are not readily available on other FPGA platforms, allow the creation of these unconventional

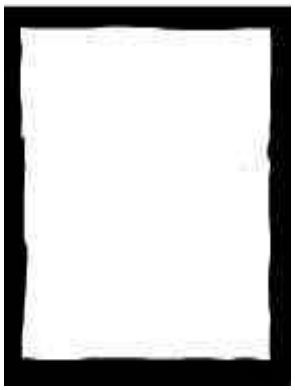

Figure 1: A sample binary bitmap file in Microsoft Paint

routing designs. This use-case utilizes the programmable interconnect as something that was never intended: as a medium for radiating energy off a device.

Section 2 of this paper describes the process of creating 2D shapes, or geometric structures using the programmable interconnect and a special router. The resulting shapes can serve a useful function. Section 3 outlines the process of transforming the shapes into antenna, and the experiments used to measure the effectiveness of the antenna. Section 4 presents applications of these techniques. And finally, Section 5 concludes the paper.

2. Creating Geometric Shapes

This section presents a way of creating 2D structures using the abundance of underutilized wires within the FPGA fabric. In the tool flow presented here, shapes are initially described as a simple binary bitmap image that can be created using any drawing tool (Figure 1). FPGAs contain many resources that are arranged on a grid where each cell is known as a *tile*. The crucial aspect is that the dimensions of the bitmap file provide a one-to-one mapping to the tile resources of the FPGA being targeted. With the use of some software manipulations, the bitmap file is mapped into a new structure of '0"s and '1"s with '0"s representing the white portion of the image and '1"s representing the black portion of the image. Therefore, a '1' in the mapping selects a tile in which all of its free routing resources are merged with all adjacent '1' tiles, and the '0"s map to tiles that are left alone.

It should be noted that only the routing resources of the FPGA are being utilized in this application. There are many other resources on the FPGA that do not allow routing such as SLICEL, SLICEM, DSP, BRAM, and IOB sites. With the exception of a few special resource columns in the FPGA, every other row contains an interconnect on the switch matrix that allows for precise structured routing. Although non-switch matrix resources are represented in the bitmap, the enablement of these resources allows for the inputs of the resource to be saturated while the resource itself is not utilized. This prevents any horizontal or vertical routing expansion through these non-routing resources [10].

The process of creating geometric structures on a FPGA is divided into two major phases: a static design phase, and a geometric shape design phase. The static design phase is the normal process of creating a digital circuit for an FPGA. There are no restrictions on placement or resources used in this design. The second phase is the creation of the geometric shape using the routing resources. These phases are then merged together. An electrical connection between the two phases is accomplished through the use of a "dummy" bus-macro. This is similar to how a partial region is connected to a static region in partially re-configurable designs per the Xilinx Partial Reconfiguration toolchain [2]. This single connection provides a means of actively driving the single-net geometric structure.

2.1 Design Merging

As stated above, once the static design is defined, it must provide a mechanism to be connected to the geometric structure. This is accomplished by placing a "dummy" bus-macro on the net which excites the geometric structure. Bus-macros are merely slices on the FPGA that pass a signal through the LUT and only buffer the logic value while not changing the logic value. Any pass-through logic element would work for the interconnection between the two phases, and is used only for its anchoring abilities. It also provides a new source which regenerates the signal for the geometric structure net. The bus-macro, and the interconnect tile directly to the west of the bus-macro slice, must be within the black region in the bitmap. If this does not occur, then the driving line for the geometric structure cannot be brought from the bus-macro to the actual routed geometric structure. Once the bus-macro is placed and added to the proper nets, the net is re-routed utilizing the standard Xilinx router [2].

After the static design is created and the bus-macro is placed on the FPGA, all wire segments on the FPGA that are utilized by the static design must be found and marked as used. This prevents the geometric structure design from interfering with the static design. With the Xilinx FPGA routing schema, routing conflicts are not always detectable based on PIP names. The router must assure that it does not attempt to place two nets on the same indivisible segment on the FPGA.

ADB is a connectivity database originally created at Virginia Tech that provides wiring data for most modern Xilinx FPGAs [6]. Within ADB, a concept known as wire segments provides a unique hash value for each possible indivisible wire grouping. By utilizing the ADB as a tool to identify the segment membership of each PIP, wire segment conflicts can be identified and removed from the design.

2.2 Fill Router

Typically, routers use algorithms such as Pathfinder that go through several routing iterations to determine the best possible routing for two or more nodes on a given network [4]. However, for the given application, an unconventional router is used. This router is agnostic with respect to critical paths and, its goal is to saturate the geometric structure within a specified region. This special router that saturates regions with routes, called the *Fill Router*, has been created specifically for this task [8].

The Fill Router works by having a collection of wire segments that is initially seeded with a known allowable wire segment. For the static design, the wire segment that is the output from the bus-macro is adequate to seed the Fill Router. From these wire segments, all possible fan-out paths

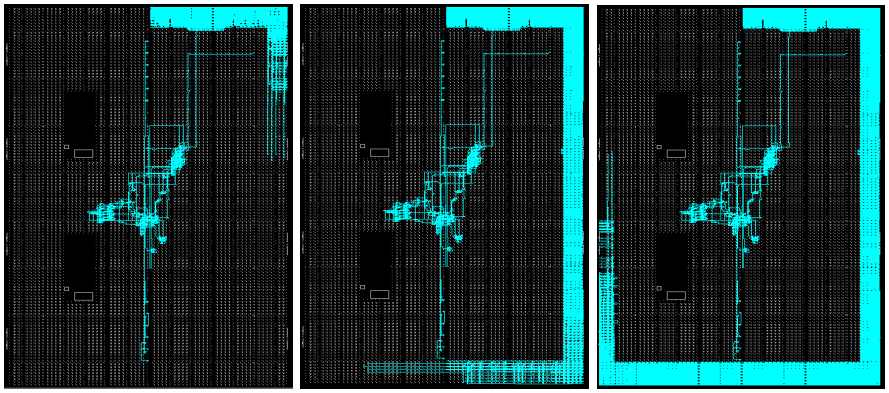

Figure 2: Fill Router coverage after six, eleven and fifteen wavefronts using the seed image in Figure 1.

are discovered by querying ADB. Each PIP that is found through the fan-out process is then evaluated to determine if it satisfies the following constraints:

- The newly merged segment does not conflict with the static design.
- The newly merged segment is currently not being utilized in the existing geometric structure.
- The newly merged segment is in a location that is within an allowable region based on the bitmap image.

Each new wire segment that is discovered and approved is then placed into a collection for the next round. Furthermore, each PIP that passes the constraint evaluation is added to the geometric structure net.

An iteration is completed after all PIPs discovered from the wire segment collection are evaluated, and the PIP is either added to the geometric structure net or discarded. This process continues until there are no additional PIPs added in a round. The iterative process is captured at multiple stages in the FPGA_Editor screen shots shown in Figure 2. Although this process does not have a control parameter to determine how much to grow each round, because the goal is to saturate the region and the region is bounded, the process will eventually complete after all available and allowable resources on the FPGA are exhausted. Figure 3 depicts the final design where the loop structure of Figure 1 is iteratively filled, and then layered onto a simple existing design.

The resulting design has the characteristic of being a single long trunk with many short branches. This structure should provide an almost uniform propagation path along the design by preventing multiple paths from the beginning to the end of the geometric structure path. Furthermore, the net retains its tree like parameters and is unlike a directed graph in which there can be a convergence of multiple paths.

2.3 Integration into Xilinx Toolchain

All manipulations of the design are conducted at the XDL level [11]. XDL is an intermediate data format that Xilinx provides that allows for the manipulation of modules, instances, and nets within a FPGA design. To get a design into the XDL format, the output file after the static place and route process is captured as a NCD file. The NCD file is

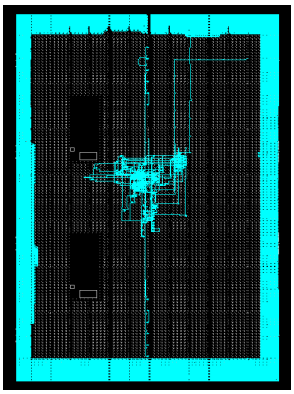

Figure 3: The consolidated geometric structure

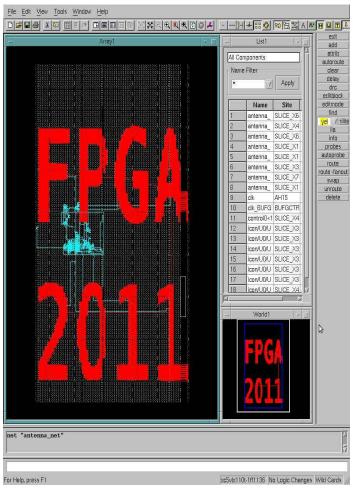

Figure 4: Another geometric structure implemented on a XCV5LX110T

then converted to a XDL file utilizing the XDL2NCD utility [11]. The XDL file is first evaluated and the custom bus-macro and geometric structure net are then inserted into the XDL representation of the design. Although the NCD file could be captured before the static translating, mapping, placing and routing, a risk remains that the static design could then not be inserted into the design if the geometric structure net is inserted first. Once all manipulations are added to the XDL file, the XDL can be converted back to a NCD file and re-inserted into the Xilinx toolchain to generate a bitstream.

3. Shapes as Antennas

In this paper, the sole purpose of creating the geometric structures is to evaluate their ability to resonate and radiate energy off-chip. Unlike a traditional radio transmitter process, there are no readily available amplifiers, pulse shapers, or analog filters. One must make-do with the limited digital circuitry available within the FPGA. In the experiments performed here, the output of a numerically-controlled oscillator is directly connected to the geometric shape through the dummy bus macro. Various shapes were created and tested, and radiated energy was measured. The theory tested here is that the shape of the antenna effects the radiated energy. It is possible that some other phenomenon is influencing the observed behavior, yet the experiments conducted appear to support the initial hypothesis. Control experiments were created where no geometric shapes were deployed to verify the effectiveness of the antenna structures.

The performance of the antenna structures created is measured by an experimental setup. Using a spectrum analyzer, connected to a signal detection coil placed under the target FPGA board, radiated energy can be measured.

It should be emphasized that the usage of the term "antenna" may not directly correspond to antennas in conventional radios. The transmission characteristics are quite different and the power efficiencies are worse than a conventional "antenna". The geometric structures created in this flow may indeed be driven in-phase by a single driver, yet is not necessarily an equipotential region. Note that there is no explicit knowledge of the underlying VLSI design beyond what is exposed in FPGA_Editor [10]. For typical antenna designs, designers are interested in the geometric shape and characteristic impedance of all segments. This design is unable to evaluate these characteristics because the physical characteristics are not released by Xilinx, thus there is no knowledge of which physical metal layers are being utilized for each net [16]. It is likely that the geometric structure net is passing through many metal layers of the FPGA, some of which may be entirely shielded. Because these nets are designed as digital networks, it is not only likely, but probable that the antenna net passes through resources such as digital switches, MUXes, keepers, terminators, and buffers. Interestingly, the antenna structures on the FPGAs have enough portions of the net that possess good radiating characteristics such that a corresponding receiver outside the FPGA can detect the radiation from the FPGA.

3.1 Signal Detection

The radiated energy transmitted by the antenna on the FPGA is captured using a signal detection coil. The coil is constructed to be a square of six turns with dimensions of $2cm \times 2cm$. The trivial nature of the detection circuitry

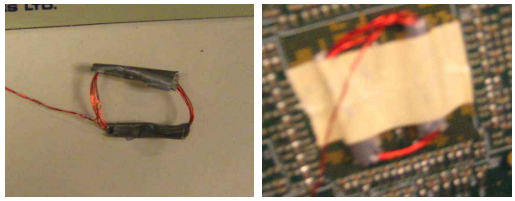

Figure 5: The detection coil taped to the FPGA board

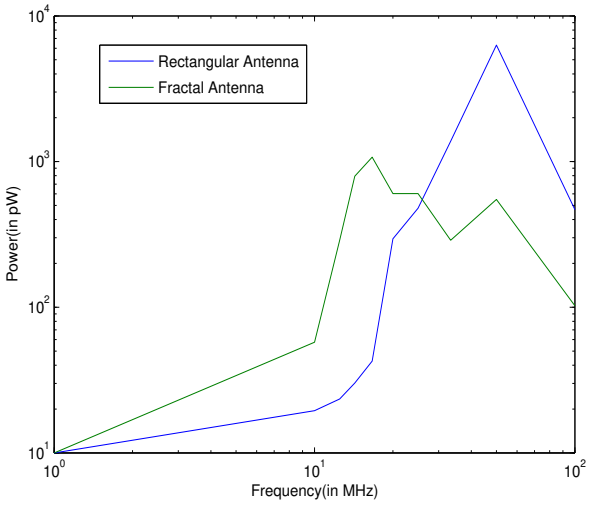

Figure 6: The antenna energy radiation graph

implies that the energy captured might be slightly lesser than what is actually being radiated. A more sophisticated system for detecting the radiation might lead to a better radiation efficiency.

The FPGA board used to conduct the initial experiments is the Xilinx ML410 (Revision E) which has an on-board Virtex-4 FPGA with the part number XC4VFX60-CES4S-FFG1152-11. The ML410 board is equipped with two crystal oscillator sockets (X6 and X10) each wired for standard LVTTL-type oscillators [15]. The X6 is populated with a 100MHz oscillator that provides the system clock which is used as an input to the user controlled excitation source previously described.

The Virtex-4 FPGA has a metal casing on top, which is in direct contact with a metal heat sink. This combination is likely quite effective in blocking radiated energy. Further, the Virtex-4 is packaged in a Flip-Chip BGA. Unlike traditional packaging in which the die is attached to the substrate face up and the connection is made by using wire, the solder bumped die in flip-chip BGA is flipped over and placed face down, with the conductive bumps connecting directly to the matching metal pads on the laminate substrate [13]. Thus, the maximum expected radiation would be if the detection is done below the FPGA chip rather than above it. In the experiments performed, the coil is attached to the board right beneath the FPGA chip (Figure 5) to capture the radiation from the FPGA through the PCB (also likely to have ground and power planes).

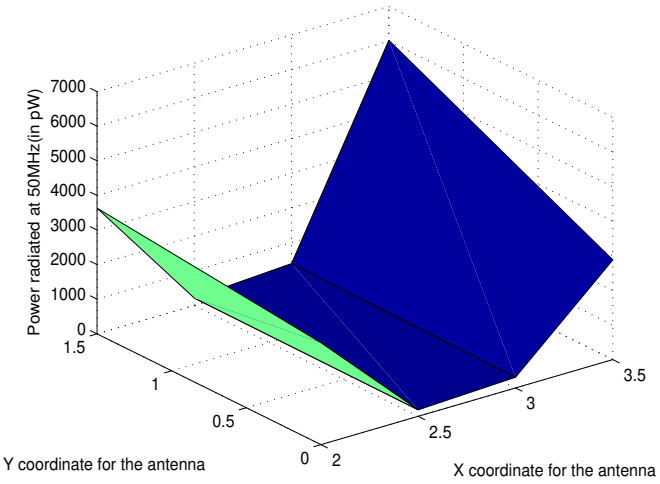

Figure 7: Effect of detection coil position on energy radiation

3.2 Methodology

The FPGA is initially powered on and the response of the signal detection coil is noted on a spectrum analyzer. The on-chip programmable oscillator is then turned on using Xilinx's ChipScope [12] software to the desired frequency and the response is measured again. As the frequency is swept, the peaks are observed for a given geometric shape. The frequency at which the highest peak is observed is assumed to be the resonant frequency of the geometric shape. When the oscillator is turned off, the peak reduces back to its original value, indicating that the energy dissipation is indeed caused by the excitation of the antenna in the FPGA. The energy radiation graph of the rectangular antenna is compared with that of a fractal antenna (discussed in Section 3.5) in Figure 6.

3.3 Sensing the Radiated Energy

An important aspect to be considered is the position of the detection coil under the FPGA chip. The detection coil is approximately $2cm \times 2cm$ whereas the FPGA chip is a square of approximately $3.5cm \times 3.5cm$. It is observed that the best performances are achieved when two of the sides of the square coil coincide with those of the FPGA. In contrast, the worst performance is when the coil is placed in the middle of the chip dimensions with no edges coinciding (Figure 7).

3.4 Antenna's Contribution to Radiation

An experiment was performed to verify that the antenna was actually contributing to the radiated energy. This was verified by making a slight change to the design. Instead of connecting the oscillator output to the antenna structure, it was connected to a slice just outside of the antenna net (Figure 8) and the procedure detailed in Section 3.2 was repeated. The radiated energy in this case was nearly zero confirming that it was indeed the antenna structure causing the radiation.

3.5 The Fractal Antenna

In order to illustrate the robustness and the system independence of the above mentioned approach, a wideband fractal antenna is also implemented. The fractal antenna structure used in this experiment has a maximum number of corners in addition to a maximal path length [1].

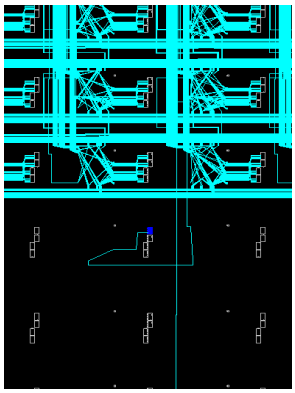

Figure 8: Oscillator connected just outside of the antenna net

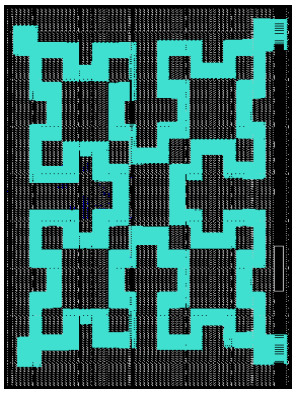

Figure 9: The fractal antenna

The XUPV5 (Revision A) board, consisting of the Virtex-5 FPGA XC5LX110T-FF1136-1, is used for this experiment. The experiment setup and the mechanism for measuring the radiated energy are the same as the ML410 experiment, except the XUPV5 is the FPGA under test. The energy radiation graph of the fractal antenna confirms that radiation is emanating from the FPGA in this experiment. An FPGA_Editor snapshot of the fractal antenna (Figure 9) is shown. As illustrated in Figure 6, although the total power radiated is less than the rectangular antenna, the frequency peak is much wider in comparison to the rectangular antenna.

3.6 Hidden Transmitter

An interesting application of the structures created is that of a hidden transmitter. The transmitter antenna designed can be embedded onto a dense logic design covering almost the entire FPGA. As mentioned before, these structures can be merged with a dense design while remaining electrically isolated and without impacting the operation of the intended circuit, all due to the abundance of programmable interconnect. The design of such a system would make it nearly impossible to detect the transmitter through visual inspection or by netlist extraction. To the untrained observer, the

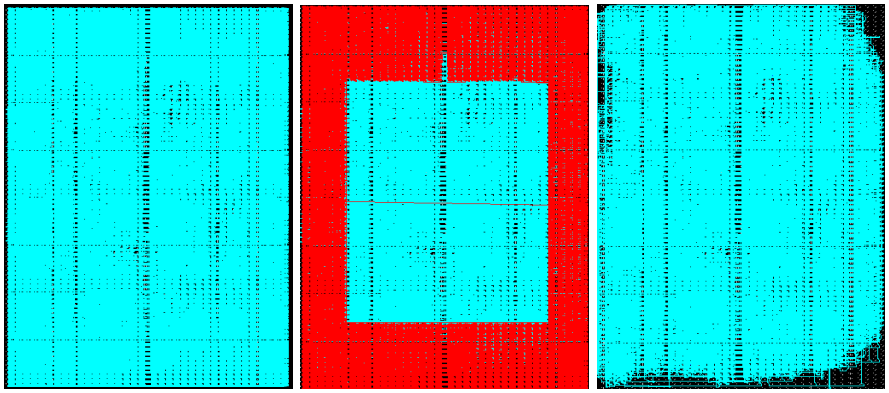

Figure 10: Dense design illustration in FPGA_Editor (a) design with antenna overlayed, (b) antenna net highlighted, (c) design without antenna

FPGA provides the intended digital operation, yet behind the scene, a hidden transmitter is actively producing a signal. The transmitter could also be further hidden by running the signal through bus-macros and route-throughs throughout the design. Figure 10 clearly illustrates this concept. The FPGA_Editor snapshot on the left shows the original dense design with the non-discernible embedded antenna structure. The middle snapshot shows the actual antenna structure when highlighted using FPGA_Editor. The snapshot on the right shows the untampered design without the antenna structure.

It is important to distinguish the concept of a hidden transmitter from that of side-channel power analysis attacks. A broad family of statistical power analysis techniques can be used to infer data from within a device. These can be further sub-categorized into Simple Power Analysis (SPA), Differential Power Analysis (DPA) and Correlation Power Analysis (CPA) [7]. SPA and DPA attacks directly use the power consumption information of a signal to break a cryptographic algorithm by monitoring the voltage and current fluctuations of a circuit [3]. CPA, on the other hand, is a technique of withdrawing the correct key using correlation coefficient of statistics [9].

Both the hidden transmitter and the power analysis techniques retrieve non-exposed data from a FPGA. However, the hidden transmitter uses it for radiating energy rather than capturing information about the underlying circuit.

4. Future Research

The results obtained from the experiments have been encouraging; however, there is still room for enhancement in certain aspects of the design. The initial experiments were performed on the ML410 (Revision E) board with a Virtex-4 FPGA and a XUPV5 board with a Virtex-5 FPGA. Both of these FPGAs have a metal cap as well as a heat sink on the chip. Although a radiated signal can be detected below the FPGA, much of the radiated energy is being absorbed into the metal cap and the heatsink. This may be affecting the radiation being captured by the spectrum analyzer. The PCB power planes are also an effective means of shielding which may further affect the energy captured. Preliminary experiments were performed on a Spartan 3E FPGA, which has a plastic cap. After initial observations this has resulted in higher power numbers.

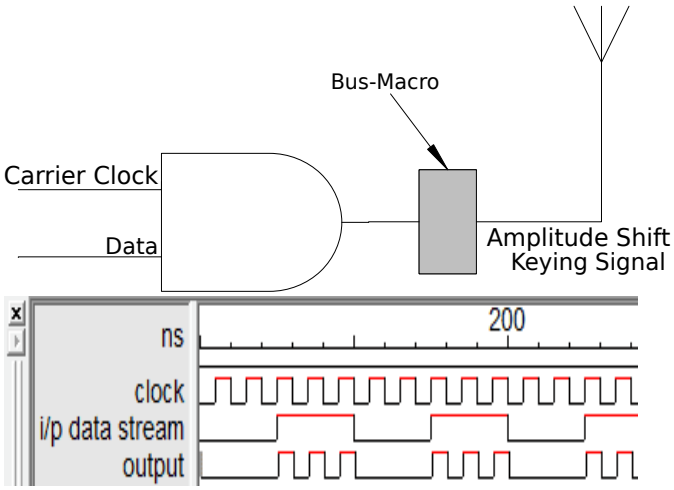

Figure 11: Amplitude shift keying setup

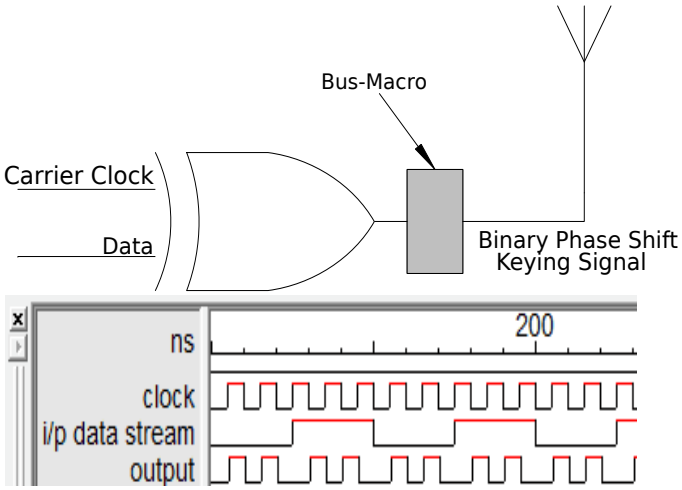

Figure 12: Binary phase shift keying setup

Another potential application for this work is the creation of a complete transmitter system. There are many possible modulation techniques that could be applied here. Two common modulation techniques, Amplitude Shift Keying (ASK) and Binary Phase Shift Keying (BPSK), are outlined below. These techniques can be easily implemented on an FPGA using basic logic gates.

ASK is a modulation technique where a carrier frequency is either present or absent indicating a '1' or a '0' for the transmitted data. A logical AND gate can be used to implement ASK. The clock signal is connected to one input of the AND gate with the other input connected to the input data stream. The ASK setup as well as a waveform illustrating the operation of the circuit are shown in Figure 11.

Phase-shift keying (PSK) is another modulation technique in which the carrier frequency's phase is altered to indicate the transmitted data. The simplest form of this technique is BPSK. BPSK works by locking onto a carrier frequency and then monitoring when it becomes inverted. There are multiple coding techniques for placing the data into this modulation technique, but a simple 0 for in-phase and 1 for out-of-phase could be implemented utilizing a XOR gate. The clock signal is connected to one input of the gate while the input data stream is connected to the other gate. The BPSK setup and the corresponding waveform are shown in Figure 12. This modulation scheme has the advantage in that it has a nearly flat power signature.

The modulated signals generated can then be transmitted using the antenna structures created to implement a comprehensive transmitter system.

5. Conclusion

A technique for the construction of transmitters on FPGAs has been discussed. The transmitter antenna demonstrated good energy radiation characteristics and responded well to slight variations. Since the structures generated on programmable interconnect are in no way similar to classical antennas, some ambiguity remains as to the exact mechanisms behind the radiated energy. But as long as there is radiation being detected outside the FPGA, it is plausible that with a modified design, data could be radiated off a net and decoded by antennas outside the FPGA. A more detailed study of the results could help understand the usefulness of the approach in the bigger scheme of things and the road ahead. It was also discussed that this technique is distinct from and is in no way related to the side-channel attacks involving power monitoring.

6. Acknowledgments

The authors would like to thank USC-ISI East and the Harris Corporation for supporting this research and for providing insight into the antenna structures. The authors would also like to thank the Brigham Young University Configurable Computing Lab for their assistance and insightful conversations. The authors would like to convey a special thanks to Neil Steiner and Aaron Wood from ISI for providing excellent tools for the routing and segment identification.

7. REFERENCES

[1] S. R. Best. A comparision of the resonant properties of small space-filling fractal antennas. In *IEEE Antennas and Wireless Propagation Letters*, 2003.

[2] C. Claus, B. Zhang, M. HÃijbner, C. Schmutzler, J. Becker, and W. Stechele. An xdl-based busmacro generator for customizable communication interfaces for dynamically and partially reconfgurable systems. 2007.

[3] J. Jaffe, P. Kocher, and B. Jun. Differential power analysis. In *Proc. 19th International Advances in Cryptology Conference CRYPTO99*, 1999.

[4] L.McMurchie and C.Ebeling. Pathfinder: a negotiation-based performance-driven router for fpgas. In *Proc. of the 1995 ACM Third Int. Symp. on Field Programmable Gate Arrays*, 1995.

[5] L.Wang, C.Stroud, and N.Touba. *System On Chip Test Architectures: Nanometer Design for Testability.* Morgan Kaufmann, 2007.

[6] N.Steiner. A standalone wire database for routing and tracing in xilinx virtex, virtex-e, and virtex-ii fpgas. Master's thesis, Virginia Polytechnic Institute and State University, 2002.

[7] M. Z. Rahaman and M. A. Hossain. Side channel attack prevention for aes smartcard. In *Proc. 11th International Conference on Computer and Information Technology ICCIT2008*, 2008.

[8] A. Sohanghpurwala. A preview of the openpr open-source partial reconfiguration toolkit for xilinx fpgas. 2010.

[9] K. Wu, H. Li, B. Peng, and F. Yu. Correlation power analysis attack against synchronous stream ciphers. In *Proc. 19th International Conference for Young Computer Scientists ICYCS2008*, 2008.

[10] Xilinx Inc. *FPGA Editor Guide 3.1i*, 2000.

[11] Xilinx Inc. *Xilinx Integrated Software Enviroment 4.2i Documentation*, 2001.

[12] Xilinx Inc. *Chipscope VIO, Datasheet DS 284*, 2004.

[13] Xilinx Inc. *XAPP426: Implementing Xilinx Flip-Chip BGA Packages*, 2006.

[14] Xilinx Inc. *Virtex4 Family Overview, Datasheet DS112*, 2007.

[15] Xilinx Inc. *ML410 Embedded Development Platform, User Guide*, 2008.

[16] Xilinx Inc. *PlanAhead User Guide 11.4*, 2009.