

A Real Time Cognitive Radio Testbed for Physical and Link Layer Experiments

Shridhar Mubaraq Mishra*, Danijela Cabric*, Chen Chang*, Daniel Willkomm[†], Barbara van Schewick[†], Adam Wolisz[†] and Robert W. Brodersen*

*School of Electrical Engineering and Computer Science
University of California, Berkeley, California 94704
Email: smm@eecs.berkeley.edu

[†]Telecommunication Networks Group, Faculty of Electrical Engineering and Computer Science,
Technische Universität Berlin, Berlin, Germany
Email: willkomm@tkn.tu-berlin.de

Abstract—Cognitive Radios have been advanced as a technology for the opportunistic use of under-utilized spectrum. Cognitive Radio are able to sense the spectrum and detect the presence of Primary Users. However, Primary Users of the spectrum are skeptical about the robustness of this sensing process and have raised concerns with regards to interference from Cognitive Radios. Furthermore, while a number of techniques have been advanced to aid the sensing process, none of these techniques have been verified in a practical system. To alleviate these concerns, a real time testbed is required, which can aid the comparison of these techniques and enable the measurement and evaluation of key interference and performance metrics. In this paper we present such a testbed, which is based on the BEE2, a multi-FPGA emulation engine. The BEE2 can connect to 18 radio front-ends, which can be configured as Primary or Secondary Users. Inherent parallelism of the FPGAs allows the simultaneous operation of multiple radios, which can communicate and exchange information via high speed low latency links.

I. INTRODUCTION

Over the past years, traditional approaches to spectrum management have been challenged by new insights into the actual use of spectrum. In most countries, spectrum at frequencies below 3 GHz has been completely allocated to specific uses. For example, the Federal Communication Commissions (FCC) frequency allocation chart indicates multiple allocations over all of the frequency bands [1]. Thus, within the current regulatory framework, spectrum is a scarce resource, at least at the frequencies below 3 GHz, which are particularly valuable due to their favorable propagation characteristics. On the other hand, actual measurements taken at the BWRC (see spectrogram in Figure 1) indicate low utilization especially in the 3 to 6 GHz bands. This view is supported by various recent measurements in the US and elsewhere. For example, studies by the FCCs Spectrum Policy Task Force reported vast temporal and geographic variations in the usage of allocated spectrum with utilization ranging from 15% to 85% [2]. These measurements seriously question the efficiency of the current regulatory regime.

As the measurements show very clearly, those who have been allocated frequency bands by the regulatory agency (Primary Users) are not using it all the time. At the same




Fig. 1. Spectrum use of 0 – 2 GHz frequencies over 10 minutes

time, others would like to use spectrum, but do not have a right to use the corresponding frequencies. Therefore, one way of increasing spectrum efficiency is to enable these other, Secondary Users to get access to frequency bands already allocated to Primary Users while these are not using it. One of the mechanisms to accomplish this sharing is called “Opportunistic Spectrum Sharing”. Under the opportunistic sharing regime, Secondary Users are allowed to operate in certain frequency bands without the consent of the Primary Users of these bands, as long as they do not interfere with the Primary User. The FCC has legalized this type of sharing in the 5 GHz band and is considering whether to allow it in the TV broadcast bands [3].

Cognitive Radios are a technology that may be used to implement opportunistic sharing, e.g. as proposed in [4], [5]. Cognitive Radios are able to sense the spectrum to see whether it is being used by the Primary User. However this sensing operation may be rendered difficult due to a degraded wireless

channel, which has prompted concerns from Primary users of the spectrum.

A. Concerns of existing Primary Users

In general, Primary Users have not been very receptive at the idea of opportunistic spectrum sharing. In particular, they are concerned that Cognitive Radios will harmfully interfere with their operation [6]–[9]. This interference may result from the following:

- The Cognitive Radio may not be able to reliably detect a Primary User and therefore may start sending although the Primary User is using the frequency. For example, while receiver sensitivity of existing Primary receivers in the TV bands may be poor, users may boost gain by employing high gain antennas. As opposed to this, Cognitive Radios are limited by fixed antenna gains and hence may not be able to sense the TV transmitter. This is the classic “hidden terminal problem” in wireless networks where a receiver is unable to “hear” the transmitter and starts its own transmission thereby interfering with the intended receiver of the transmission.
- A Cognitive Radio that is using a frequency that was deemed free by the sensing process may not be able to reliably detect that a Primary user has reappeared and therefore may not vacate the frequency.
- Even if a Cognitive Radio has detected the Primary User, it may fail to vacate the frequency quickly enough and therefore continue to send creating harmful interference to the Primary User transmission.

B. The way forward: Testbed and test cases

As the above discussion indicates, there is considerable debate whether it is possible to build a Cognitive Radio that does not disturb Primary Users. However, this debate cannot be resolved on a theoretical basis. Ultimately, Primary User concerns can only be addressed by working systems that demonstrate that the amount of interference is sufficiently low to justify allowing their use. This demonstration is important to regulators whose confidence in Cognitive Radios is not sufficient enough to go forward with the regulatory framework that would allow the technology; it may also help to reduce opposition from Primary Users.

At the moment, there is no agreement on the kind of demonstrations that would be needed to convince regulators and Primary Users that a Cognitive Radio meets these requirements. As it is impossible to test all possible cases, it is necessary to commonly agree on a set of “representative” test cases that a Cognitive Radio must pass to “prove” that the amount of interference is sufficiently low to justify allowing Cognitive Radio technology. In addition, metrics will be needed to assess how well a Cognitive Radio performs in a test case.

At the same time, it is far from clear what mechanisms are best suited to implement Cognitive Radios, both with respect to preventing interference and with respect to efficiency and performance. There are a plethora of techniques (cooperative sensing, cyclostationary detectors [10], etc.) that have been

proposed to enhance detection. None of these techniques have been tested in real world scenarios and their performance has yet to be characterized. Thus, there is a need for experimenting with different techniques in a real system, using a set of test cases and metrics to compare different Cognitive Radio implementations. In this paper, we propose such an experimental setup based on the Berkeley Emulation Engine 2 (BEE2) platform to experiment with various sensing techniques and develop a set of metrics and test cases, which will allow us to measure the sensing performance of these techniques.

The rest of the paper is organized as follows: Section II discusses the requirements and metrics for experimental evaluation. These requirements are partially driven by Primary User concerns and motivated by the need for real time evaluation. Section III discusses the basic architecture and implementation of the testbed using the Berkeley Emulation Engine (BEE2) platform. Section IV explains the experiment setup using BEE2 and the proposed set of experiments. Finally, we conclude in Section V.

II. METRICS AND REQUIREMENTS FOR EXPERIMENTAL EVALUATION

In order to facilitate the deployment of Cognitive Radio technologies for the secondary usage of spectrum it is crucial to prove the reliable detection of Primary Users by Cognitive Radios. We have to show that Primary and Secondary Users can coexist without a degradation of the Primary User transmission in order to convince regulatory authorities as well as Primary Users to enable such technologies. In particular, algorithms that aid reliable Primary User detection need to be verified and their performance has to be quantified by metrics and test cases.

In this Section we introduce metrics and test cases for the evaluation of Cognitive Radio mechanisms to detect Primary Users. Based on the defined metrics and test cases we present a list of features important for Cognitive Radio testbeds.

A. Metrics and test cases for Primary User detection

There are three aspects of Primary User detection that need to be verified and quantified in order to define metrics for Cognitive Radio systems:

- 1) The time until detection of the Primary User.
- 2) The reliability of Primary User detection.
 - a) The probability of **not** detecting the Primary User (false negative).
 - b) The probability of false alarms (false positive).
- 3) The time needed to clear the spectrum once a Primary User has been detected.

Before using any spectrum of a Primary User the Cognitive Radio has to ensure that the Primary User is currently not using its spectrum. With varying conditions on the channel between a Primary User and a Cognitive Radio due to severe fading, shadowing and building penetration losses, spectrum sensing can become extremely difficult. Furthermore, given concerns of Primary Users, we expect stringent requirements

on the probability of detection, which have to be met under all conditions. Whereas the time until detection is not that critical – since the Cognitive Radio is not using the Primary User spectrum yet – it is important to keep the probability of false alarms low in order to achieve a better exploitation of unused spectrum.

During an ongoing communication using temporarily available licensed spectrum, the involved Cognitive Radios have to continuously monitor the used spectrum for reappearing Primary Users. In this case the time until detection of the Primary User is very critical, as it has a big influence on the time the Cognitive Radio interferes with the Primary User. In order to achieve very short detection times we may be more tolerant with false alarms, i.e. trade a short detection time against more false alarms. The time needed to vacate the spectrum after detection is another important factor in this scenario. All communication peers have to be informed about the detected Primary User as they might not have detected it themselves.

Any Cognitive Radio system has to be tested and verified considering the above defined aspects. The evaluation has to include different scenarios of Primary User and Cognitive Radio placements, Primary User signal strengths and various numbers of Primary Users and Cognitive Radios as well as different scenarios of concurrent Cognitive Radio transmissions (to distinguish a Cognitive Radio transmission from a Primary User transmission). However, as stated earlier, opponents may always come up with a specific scenario, under which they state that the probability of not detecting the Primary User is too high. That is why these measurements are specially important for pathological placements of the Cognitive Radios (as suggested in comments to the FCC by Primary Users), in which shadowing and multipath conditions are particularly severe. To make these measurements meaningful, it is crucial to commonly agree on a “representative” set of test cases that every Cognitive Radio has to pass.

B. Comparison of different algorithms and Cognitive Radio systems

A variety of techniques have been proposed to enhance the detection of Primary Users. Some examples are the detection using a pilot signal [11] and cyclostationary feature detection [10]. Also for the detection while using the primary spectrum different mechanisms have been proposed. Transmission and sensing could be interleaved or done concurrently using active cancellation of the own signal or only using some part of the primary spectrum for transmission while concurrently sensing on the other part. Another mechanism proposed to enhance reliable Primary User detection is cooperative sensing [10]. The influence of such techniques in isolation as well as in conjunction with other techniques need to be investigated.

Furthermore, the trade-off between the probability of not detecting the Primary User and other parameters need to be investigated. Some of these parameters are:

- the computational overhead needed for detection,
- the power dissipation,

- the wasted transmission time / bandwidth for detection,
- the hardware area, and
- the per unit cost needed.

C. Requirements on the testbed

The above defined metrics and test cases impose certain requirements on a testbed for Cognitive Radios. Following is a list of features any Cognitive Radio testbed should provide.

- The ability to support multiple radios, which can serve as Primary or Secondary Users. Primary Users must be controllable in order to make precise measurements.
- The ability to connect various different front-ends in order to be able to test in different frequency ranges (and thus with different Primary Users).
- The ability for physical / link layer adaptation and fast information exchange between multiple radios for sensing and cooperation.
- The ability to perform rapid prototyping in order to experiment with different sensing algorithms.

Figure 2 shows an abstract diagram of such an emulation platform. To implement multiple radios, the emulation platform must provide plenty of parallelism and mechanisms to connect to multiple front-ends. Furthermore, the latency to exchange information between the various radios should be small.

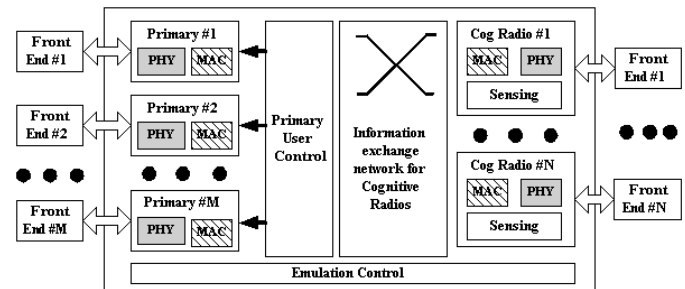


Fig. 2. Emulation platform for Cognitive Radios

III. TESTBED ARCHITECTURE AND IMPLEMENTATION

The testbed requirements stated earlier are met by the Berkeley Emulation Engine 2 (BEE2), which is a generic, multi-purpose, FPGA based, emulation platform for computationally intensive applications. Each BEE2 can connect to 18 front-end boards via multi-gigabit interfaces. The case for FPGAs, over DSPs and Microprocessors, has been argued in [12]. FPGAs offer rapid reconfigurability, exhibit rapidly increasing computational power per unit area and demonstrate the best computational performance per unit power consumed for key computational modules [12]. Furthermore, FPGAs provide plenty of parallelism, which can be used to implement multiple radios.

A. The BEE2 board

The BEE2 consists of 5 Vertex-2 Pro 70 FPGAs. Each FPGA embeds a PowerPC 405 core, which minimizes the

latency between the microprocessor and reconfigurable logic. These 5 FPGAs form a single compute module. Each FPGA can be connected to 4 GB of memory with a raw memory throughput of 12.8 Gbps. Four FPGAs are used for computation and one for control as shown in Figure 3. Adjacent FPGAs are connected via on-board low-voltage 40 Gbps parallel interfaces [12]. All computation FPGAs are connected to the control FPGA via 20 Gbps links. These high bandwidth, low latency links allow the five FPGA to form a virtual FPGA of five times the capacity.

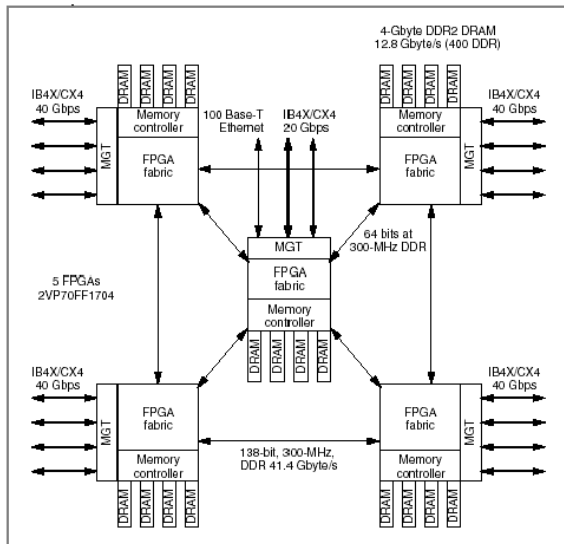


Fig. 3. BEE2 Compute Module

These FPGAs can connect to the external world using 10 Gbps full duplex Infiniband interfaces. There are a total of 18 such connectors per board. The Infiniband connectors allow the BEE2 compute module to connect to an Infiniband switch which enables multiple BEE2 compute models to communicate and exchange data. Figure 4 shows a picture of the BEE2 board. Details about BEE2 can be found in [12].




Fig. 4. BEE2 board

Each BEE2 board supports one 100 Base-T Ethernet interface, which is available on the control FPGA. The Power PC of the control FPGA can run Linux and a full IP protocol stack.

The board also contains USB and JTAG interfaces along with provisions for a flash card. The 100 Base-T interface allows remote management and control.

B. Modular front-end system

The front-end system has been designed in a modular fashion. The analog/baseband board contains the filters, ADC/DAC chips and a Xilinx Vertex-II Pro FPGA. Digital-to-analog conversion is performed by a 14-bit DAC running up to 128 MHz, while analog-to-digital conversion is performed by a 12-bit ADC running up to 64 MHz. The FPGA performs data processing and control, and supports 4 optical 1.25 Gbps links for transmitting and receiving data to/from BEE2. The optical link provides good analog signal isolation from digital noise sources and allows the front-end to be moved up to a third of a mile from BEE2 for wide range wireless experimentation. A separate RF modem module connects to the baseband board. The current RF modem module is capable of up/down converting 20 MHz RF bandwidth at 2.4 GHz. The RF frequency is fully programmable in the entire 80 MHz ISM band. A block diagram of a single RF modem is shown in Figure 5, while Figure 6 shows the RF and baseband boards.

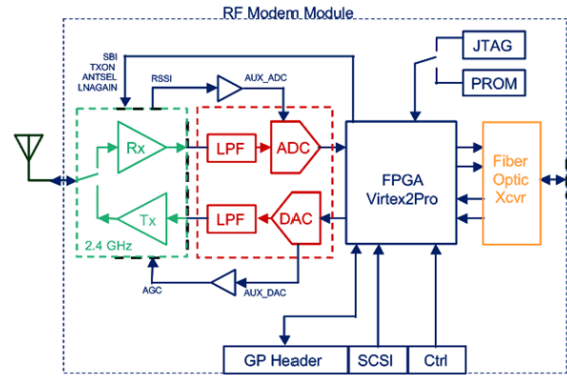


Fig. 5. RF Modem Module and Analog/baseband board




Fig. 6. Front-end boards

Scalability is achieved through parallel RF modem modules being provided with a common RF reference and clock signals. Two configurations are supported by this architecture:

- 1) All front-ends operate at the same radio frequency (the radios need to operate in Time Division Duplex (TDD) mode in a single 20 MHz band).
- 2) Groups of 4 or more antennas operate in different bands (the radios operate in Frequency Division Duplex (FDD) mode and occupy the entire 80 MHz band).

C. BEE2 programming model using Simulink

BEE2 can be programmed using Matlab/Simulink from Mathworks [13] coupled with the Xilinx system generator. The tool chain is augmented with BWRC developed automation tools for mapping high level block diagrams and state machine specifications to FPGA configurations. A set of parameterized library blocks have been developed for communications, control operators, memory interfaces and I/O modules.

IV. COGNITIVE RADIO SETUP

The BEE2 emulation platform supports all the features required for a Cognitive Radio testbed. This section describes how various Cognitive Radio functionalities can be mapped to the BEE2 platform and the key features of the BEE2 that facilitate this mapping process. Finally we describe a first experimental setup.

A. Multiple radios for Primary and Secondary User functionality

The BEE2 compute board can connect up to 18 front-ends, which can be split between Primary and Secondary User interfaces. This enables us to construct scenarios with multiple Primary Users exhibiting different channel use patterns. Additionally we can experiment with various setups of multiple Secondary Users collaborating in the detection of Primary Users.

The optical links from the BEE2 board to the front-end boards that reach up to 1/3 mile facilitate experimentation in different spatial configurations and different shadowing and multi-path environments.

Primary Users can be controlled directly and hence precise interference measurements can be performed.

B. Complex sensing algorithms

The BEE2 compute module can be used to perform complex signal processing functions. Each Virtex-2 Pro can support a peak rate of 81 GMACps (Multiply Accumulate (MAC) operations per second) enabling the implementation of cyclo-stationary feature detection and pilot tone detection algorithms. The parallelism offered by the FPGA cluster enables multiple radios to be operational simultaneously.

Since the FPGAs are connected via high speed links, cooperation among Cognitive Radios is easy to implement.

C. Management and control protocols

A user can connect to the BEE2 platform via standard Ethernet and control the Primary and secondary users. This Ethernet interface also provides a proxy for a real control channel needed to exchange setup and sensing information. Furthermore, the Power PCs on each FPGA can be used

to implement protocols with the possibility of moving time-critical functions to reconfigurable logic.

D. First experimental setup

Our first experimental setup consists of a pair of Primary Users and a Cognitive Radio network as shown in Figure 7. The Primary users will be laptops with standard IEEE 802.11 b/g cards controlled by BEE2. Each Cognitive Radio consists of a laptop computer with an IEEE 802.11 b/g card and an 2.4 GHz 80 MHz wide front-end. The front-end is only responsible for sensing in this setup while the laptop computers with the 802.11 cards are used for user data transmission. Information between the sensing radios and the transmission laptops is exchanged via the standard Ethernet interface which serves as the control channel in the first implementation.

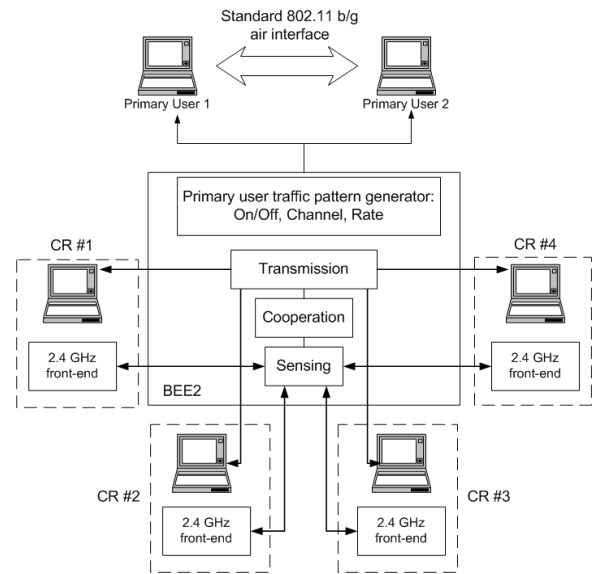


Fig. 7. Primary/Secondary User setup using the BEE2

We believe that the choice of the unlicensed 2.4 GHz ISM band for our first experimental demonstration is suitable for several reasons:

- 1) It is an unlicensed spectrum so the Cognitive Radio operation in this band is not a subject to an agreement with licensed users. Furthermore, it is considered as a very crowded spectrum with many unlicensed devices that are not able to intelligently control and avoid mutual interference.
- 2) Commercially available WLAN devices for the 2.4 GHz, such as IEEE 802.11 b/g cards within laptops, are quite programmable and allow users to control their transmission parameters. Therefore, they can be used for Primary User emulation in a controlled fashion as well as for Secondary User transmitters.
- 3) The ability to transmit standard compliant 802.11 b/g waveforms on the secondary links and coordinate control of transmission times, will allow us easy experimentation of protocols for medium access control.

- 4) All hardware and software support for the 2.4 GHz bands is already developed within BWRC to support Cognitive Radio experiments. Our BEE2 infrastructure supports multiple connections of laptop cards and 2.4 GHz front-ends that can be combined as a Cognitive Radio system capable of sensing and transmission. Furthermore, our 2.4 GHz front-ends are configurable to sense the whole 80 MHz spectrum instantaneously while commercial devices can only sense a single 20 MHz channel.
- 5) We believe that the performance of sensing algorithms for indoor 2.4 GHz experiments, if reported as function of input SNR, can be further extended to other frequency bands.

V. CONCLUSIONS

In this paper we present a testbed for experimenting with Cognitive Radios at the physical and link layer. The motivation for a testbed is provided by the need to validate various sensing algorithms to prove non-interference to licensed users and to evaluate their performance with well defined metrics. This testbed allows us to emulate Primary as well as Secondary Users and enables the evaluation of the performance of various spectrum sensing schemes. The 2.4 GHz spectrum was chosen for initial experimentation due to the availability of off-the-shelf transmission equipment and the ability to emulate Primary Users in a controlled manner. These 2.4 GHz radios are connected to the Berkeley Emulation Engine 2 (BEE2), which is a multi FPGA emulation platform. The FPGAs enable the implementation of complex signal processing functions and the inherent parallelism of the FPGAs supports concurrent operation of multiple radios. The Cognitive Radios can exchange sensing and setup information in a timely manner

since the BEE2 FPGAs are connected via high bandwidth low latency links.

ACKNOWLEDGMENT

The authors would like to thank the BEE2 team for their valuable comments.

REFERENCES

- [1] NTIA, "U.S. frequency allocations." [Online]. Available: <http://www.ntia.doc.gov/osmhome/allochrt.pdf>
- [2] "Spectrum policy task force report," Federal Communications Commission, Tech. Rep. 02-155, Nov 2002.
- [3] "Unlicensed operation in the tv broadcast bands and additional spectrum for unlicensed devices below 900 mhz in the 3 ghz band," Federal Communications Commission, NOTICE OF PROPOSED RULEMAKING 04-186, May 2004.
- [4] D. Cabric, S. M. Mishra, D. Willkomm, R. W. Broderon, and A. Wolisz, "A cognitive radio approach for usage of virtual unlicensed spectrum," in *14th IST Mobile Wireless Communications Summit 2005*, Dresden, Germany, June 2005.
- [5] T. Weiss and F. Jondral, "Spectrum pooling: An innovative strategy for the enhancement of spectrum efficiency," *IEEE Commun. Mag.*, vol. 42, pp. S8–S14, march 2004.
- [6] http://gulfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id.document=6516086573.
- [7] http://gulfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id.document=6516482880.
- [8] http://gulfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id.document=6516483135.
- [9] http://gulfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id.document=6516883484.
- [10] D. Cabric, S. M. Mishra, and R. W. Brodersen, "Implementation issues in spectrum sensing for cognitive radios," in *Asilomar Conference on Signals, Systems, and Computers*, 2004.
- [11] A. Sahai, N. Hoven, and R. Tandra, "Some fundamental limits on cognitive radio," in *Allerton Conference on Communication, Control, and Computing*, 2003.
- [12] C. Chang, J. Wawrzynek, and R. W. Brodersen, "Bee2: A high-end reconfigurable computing system," *IEEE Design and Test of Computers*, vol. 22, no. 2, pp. 114–125, 2005.
- [13] <http://www.mathworks.com/products/simulink/?BB=1>.