Picasso: Full Duplex Signal Shaping to Exploit Fragmented Spectrum

Steven Hong, Jeffrey Mehlman, Sachin Katti
{hsiying, jmehlman, skatti} @stanford.edu

ABSTRACT

Wireless spectrum is increasingly fragmented due to the growing proliferation of unlicensed wireless devices and piecemeal licensed spectrum allocations. Current radios are ill-equipped to exploit such fragmented spectrum since they expect large contiguous chunks of spectrum to operate on. In this paper we argue that future radios should provide full duplex signal shaping to the higher layers to systematically exploit fragmented spectrum. Such an architectural design would allow the radio to decouple the use of different spcetrum fragments. We present the design and implementation of Picasso, a system that provides such a general signal shaping abstraction. Picasso has two novel components: a self-interference cancellation technique and a programmable filter engine that enables it to simultaneously send and receive over different spectrum fragments. We provide an initial design and empirically evaluate the feasibility of both components.

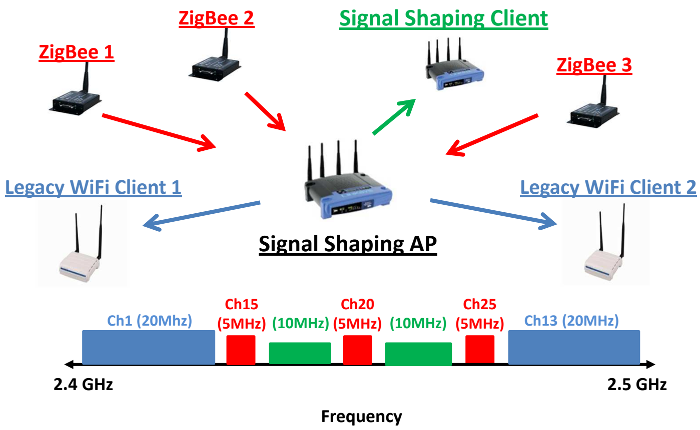

Figure 1: 3 ZigBee Interferers shown in red. Signal Shaping AP can simultaneously service multiple Legacy WiFi Clients and utilize fragmented spectrum for Signal Shaping Client simultaneously.

1. INTRODUCTION

Wireless spectrum is becoming increasingly fragmented. In the unlicensed ISM band it is not uncommon for users to carry multiple wireless devices each operating in their own contiguous, narrow bands of varying widths, which leads to fragmentation of the 100MHz ISM spectrum chunk. Similar situations exist in cellular bands. Due to piecemeal allocations by the FCC over the years, it is not uncommon for a single network operator to have fragmented, narrowband chunks spread over a large frequency range (e.g. AT&T owns nearly 40MHz of spectrum spread over 200MHz in the 700-900MHz range). Moreover, spectrum fragmentation varies over time and space; the set of available ISM bands depends on which devices are operating at a particular location at any given time. Similarly, region-wise spectrum allocations by the FCC as well as the existence of short term region specific spectrum leases [4] imply that cellular spectrum fragmentation is also variable in space and time.

Current radios are ill-equipped to take advantage of fragmented spectrum. Traditional design practice has been to build radios that operate on contiguous spectrum chunks, since it was quite complex and expensive to design configurable radios that could operate on tiny, discontiguous fragments of spectrum [3]. Coupled with the fact that contiguous spectrum chunks were invariably available, this conservative design approach has worked fairly well for the past few years.

However, we believe that this approach is untenable for designing future radios. The key reason is that the density of radios around us has dramatically increased over the past few years; it is fairly common for a single user to be carrying devices that have $\sim$ 10 radios operating in the ISM band alone [1]. Hence, large contiguous chunks of spectrum (e.g the 40MHz chunks that 802.11n needs) may be luxuries that are hard to find in the near future. Consequently to deliver high throughput, future radios will have to take advantage of whatever spectrum is available, even if it is fragmented.

In this paper, we argue that future radios should provide a general *full duplex signal shaping* capability to the higher layers to systematically exploit fragmented spectrum. By full duplex signal shaping, we mean that the radio can be configured to transmit on an arbitrary set of spectrum fragments and receive on a different arbitrary set of spectrum fragments, simultaneously. The key architectural consequence of such a radio is that it would *decouple the use of different spectrum fragments*,

Hotnets '11, November 14–15, 2011, Cambridge, MA, USA.

i.e. instead of designing one complex wideband PHY and MAC protocol that operates over the entire fragmented spectrum, the radio can simply run several independent, contiguous narrowband PHY and MAC instances on each spectrum fragment as shown in Fig. 2. Such narrowband protocols are well understood, already widely used, and simple to implement. Thus signal shaping radios can help preserve design modularity, and offer a general primitive that enables the reuse of well engineered higher layers.

To demonstrate the benefits of a full duplex signal shaping radio, we consider a concrete, commonplace example. An 802.11n access point (AP) is operating in the unlicensed 2.4GHz ISM band and serving 3 clients, as shown in Fig. 1. There are three ZigBee radios operating in 5MHz channels 15, 20 and 25 (around 2.425. 2.45 and 2.675 GHz respectively). With current radios, the AP cannot find a continuous 40MHz chunk and is forced to pick a smaller 20MHz band to operate on. However, with full duplex signal shaping, the AP can use all four available spectrum fragments and weave them together to access a total of nearly 80MHz of unused bandwidth if needed. Signal shaping also allows the MAC to flexibly designate different spectrum fragments to different clients. For example, if one of the clients is using a VoIP application and the other two are involved in file transfers, the AP can assign the 5MHz fragment to the VoIP client, and use the other larger fragments for the data clients.

Further, since full duplex signal shaping capability would allow the AP to simultaneously send and receive on different spectrum fragments, the *usage of each spectrum fragment is decoupled* since the AP does not have to worry about synchronizing transmissions and receptions across all the clients and can serve each one of them independently. Full duplex operation would also allow the AP to be backward compatible with legacy clients that do not have signal shaping capability, since it could simply use the 20MHz fragments to establish two independent WiFi networks for the legacy clients to connect, and use the non-standard 5MHz fragments to connect to newer signal shaping enabled clients.

As this example suggests, a full duplex signal shaping radio must solve two important challenges:

1. **Full Duplex Operation:** In order for a signal shaping radio to simultaneously transmit and receive on arbitrary but different fragments of spectrum, the key challenge it must overcome is *receiver saturation*, i.e. if a radio is transmitting, then the self interference saturates the receive circuit's ADC and zeroes out the received signal. Prior cellular radio designs [6] solve this problem by using *statically configured* analog notch filters to remove the self-interference. But a signal shaping radio has to be able to send and receive on arbitrary spectrum fragments, and designing programmable analog filters is expensive and complex.
2. **Programmable shaping:** Since the available spectrum fragments are dynamic and change over time, Picasso must provide a programmatic interface that allows the higher layers to dynamically specify which spectrum fragments to use, and then shape the transmitted signals to obey these higher layer directives.

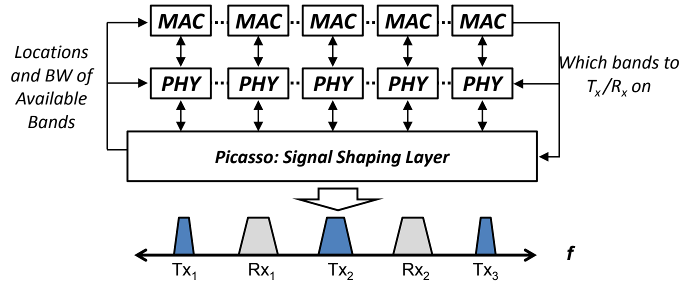

Figure 2: Picasso exposes a signal shaping API to the PHY/MAC

In this paper, we present ***Picasso***, our system design that realizes full duplex signal shaping. Picasso provides both programmable shaping and simultaneous TX/RX capabilities. It exposes an interface to the PHY/MAC layers where the upper layers either send or receive *digital, baseband sample streams* along with internal headers that specify what spectrum fragments these samples are to be transmitted or received on.

Picasso makes two key technical contributions. First, Picasso designs a self-interference cancellation technique to enable simultaneous TX/RX on separate bands. Here, the challenge is that even though the radio is simultaneously receiving while transmitting on different bands, the analog circuitry on the receive path cannot actually filter the transmitted signal out; it is extremely complicated and expensive to design highly configurable *analog filters*. The receive circuitry is overloaded by the transmitted signal, which is billions of times stronger than the received signal. Left unchecked, it saturates the ADC and any received signals are lost. We build on recent work in full duplex radios [6] and design a technique that cancels the transmitted signal in analog and ensures that the ADC is not saturated.

Second, Picasso designs a reconfigurable filter engine that removes any residual self-interference and provides efficient and programmable shaping using off-the-shelf components. The key tradeoff here is between complexity and efficiency, because the efficiency of practical filters (defined as how much they leak signals into adjacent bands and cause interference, less is of course better) is proportional to the implementation complexity. Picasso designs a filter engine with high-speed reconfigurable filters and resamplers that can dynamically and efficiently shape signals with reasonable complexity and shows that it can be implemented on existing hardware.

To demonstrate the feasibility of Picasso, we present a preliminary design on Xilinx Virtex-5 FPGA-based software radios. We briefly describe the design of the filter engine and the self-interference cancellation technique. We also evaluate the efficacy of the self interference cancellation technique using experimental results and show that it provides enough cancellation to prevent ADC saturation and enable simultaneous TX/RX on different bands.

2. DESIGN

Picasso's key contribution is that it decouples the use of different spectrum fragments, i.e. it allows one to use separate independent PHY layer encoding/decoding and MAC layer scheduling for individual spectrum fragments. Fig. 2 shows how Picasso differs from current radio designs - between the PHY/MAC layers and the actual transmission of RF signals on the antenna, there is an additional signal shaping layer. Picasso exposes an API to the PHY/MAC layers which consists of streams of complex digital baseband samples flowing between the signal shaping and the PHY layers for the spectrum fragments. Each stream carries an internal header which includes a tuple defining the spectrum fragment on which those digital complex samples should be transmitted or were received on. Picasso thus makes the architecture more amenable to evolution, since decoupling how fragmented spectrum is used from how packets are processed and scheduled allows the PHY and MAC can evolve independently and new innovations can be easily integrated without having to change the radio.

Next, we describe the two main components of Picasso, the self-interference cancellation block and the programmable filtering engine.

2.1 Simultaneous Transmit and Receive

To realize full duplex communication, Picasso has to handle the large self interference introduced by the transmitted signal to the received signal. Specifically, in Picasso a radio has two antennas, one for transmit and one for receive. The signal from the transmit antenna interferes with the receive signal of interest at the receive antenna and causes *receiver saturation.*

When the ADC samples the analog signal on the receive antenna, it converts each sample into a number corresponding to a voltage level. The value of each sampled point is stored as a fixed-length variable whose size is determined by the resolution, or dynamic range, of the ADC. If, for example, the ADC has n (typically $n \leq 12$) bit resolution, then the ADC can only hold values from 0 to 2^{n-1}. The problem is that the self interference is billions of times stronger than the received signal (e.g. for Wifi the self interference would be nearly 60-70dB stronger). The dynamic range of practical ADCs is not large enough to acquire the received signal in the face of such large self-interference, so the received signal is lost in quantization.

This is a problem even though the transmitted signal is on a different spectrum fragment than the received signal. The reason is that on the analog side of a Picasso radio, there is no programmable filtering and any signals received on the entire 100MHz ISM band are passed through. *Spectrum fragments over which a radio might transmit or receive are not fixed in advance and change over time and space, so we must be able to receive over the entire band.* Coupled with the fact that programmable analog filters are complicated and expensive, the analog front end of a Picasso radio has no choice but allow all signals across the ISM band through.

To solve this problem, Picasso makes a key observation: Instead of attempting to filter in analog, Picasso can sufficiently cancel the self-interference in analog such that the ADC has enough dynamic range to pass the received signal through without distortion, thus full duplex signal shaping is possible. The reason Picasso can cancel the transmitted signal is that *the self-interference is known* because it is coming from the same radio. We leverage recent work on single channel full duplex radio design [2][6] to design full duplex signal shaping over different spectrum fragments. Note that our requirements are not as onerous as those required for in-band full duplex because the transmission and reception occurs on different, discontiguous bands. Rather all Picasso need is enough cancellation of the transmitted signal such that the received signal can be captured within the dynamic range of the ADC. The rest of the filtering can then be handled by the filter engine.

Conceptually, self-interference cancellation is easy to describe - simply subtract an identical copy of the transmitted signal from the received signal to eliminate self-interference. However in practice this simple conceptual idea is hard to implement. First, analog circuits which subtract signals are much more difficult to design than circuits which add signals. Hence instead of trying to subtract, Picasso borrows a novel circuit design from [6] to first *obtain the exact inverse* of the transmitted signal, and then adds it to the received signal to eliminate the self-interference. Signal inversion is performed using a balun (balanced/unbalanced) transformer, which takes a signal as input and outputs the exact inverse of the signal. The left hand side of Fig. 3 shows a block diagram of the balun-based self-interference cancellation design. Baluns are frequently found in RF circuits for converting single-wire signals with a common ground into differential signals, but it was recently shown that they can also be used to negate self-interference using a technique called balun cancellation [6].

While a balun can provide a signal inverse, it's not sufficient to achieve self-interference cancellation. The first problem is that even though the transmitted signal is known exactly, once it leaves the transmit antenna and arrives at the receive antenna where it will be cancelled, the wireless channel has introduced delay and attenuation which distort the signal. Even if the distortion is slight, this precludes us from perfectly canceling the signal. To solve this problem, Picasso implements a calibration mechanism which uses feedback control to estimate the attenuation and delay so that we can compensate for the channel distortion. Due to space constraints the exact details are omitted, but the basic idea is that there exists a single optimum attenuation and delay pair, matching the effect of the wireless channel, that will maximize the self-interference cancellation. Because the relationship between attenuation/delay and the remaining energy of the canceled signal exhibits a convex conic structure, we can utilize a simple gradient descent algorithm which will quickly search different pairs until it has converged to the optimum point. We note that because the transmit and receive antennas are fixed and located in such close proximity, the channel between them will be fairly consistent and hence we would only need to run this calibration

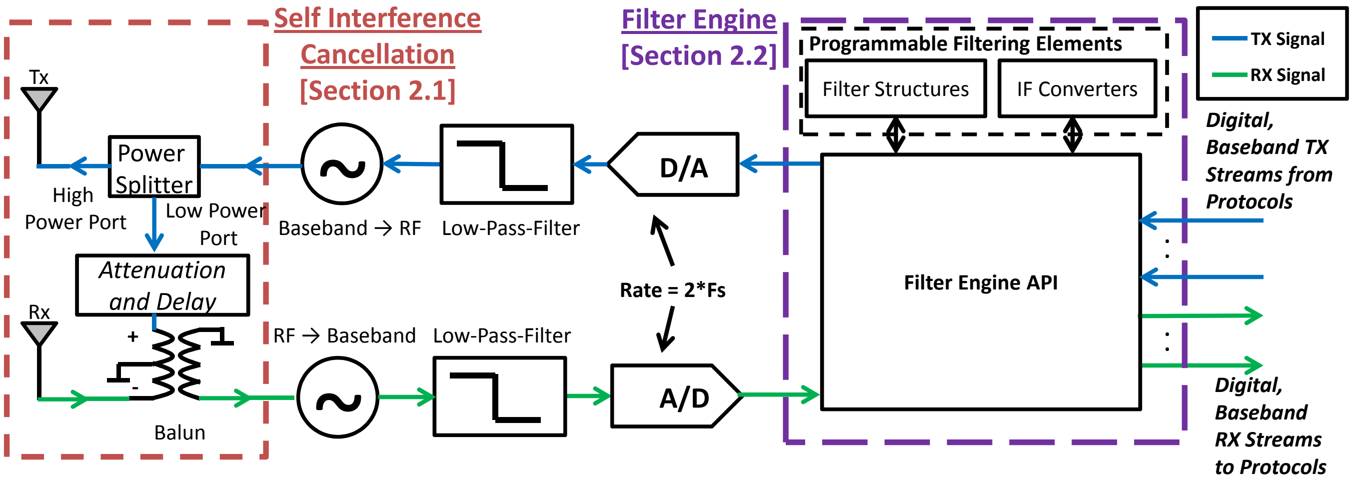

Figure 3: Picasso's solution sketch involves 2 main components: Balun inversion cancellation to enable simultaneous transmit and receive over different fragments and programmable shaping to utilize fragmented spectrum.

mechanism occasionally in order to achieve acceptable cancellation.

2.2 Filter Engine

The filter engine's job is ensure that the transmitted or received signals are shaped according to the MAC layer specification of which spectrum fragments to use. To illustrate this functionality more clearly, we re-visit the example from the introduction, where a modified WiFi AP was using two separate 20MHz spectrum fragments centered at 2.412 and 2.472GHz to simultaneously serve two legacy users. We assume that Picasso is doing signal shaping over the entire 100MHz ISM band centered at 2.45GHz. The MAC layer would inform Picasso about the two fragments, including their frequency locations and the intent to use both of them as independent TX/RX streams. When the node is transmitting, the PHY layer would send two 40 megasamples/sec (MS/s) (because you need to sample at 2× for Nyquist sampling) streams of digital baseband complex samples (representing the two encoded signals to be transmitted) to Picasso, which would shift the two streams to the appropriate 20MHz fragments as specified by the MAC. If instead the AP was receiving transmissions from the clients on either or both of the 20MHz fragments, Picasso would deliver one or two streams of digital samples corresponding to the received signal at 40MS/s to the PHY for decoding.

The filtering engine contains an ADC and DAC, both of which are capable of operating at 200MS/s (the required Nyquist rate to create signals that span the entire 100MHz ISM band). Further, on the analog RF side, we assume that there is a single oscillator at 2.45GHz which upconverts the shaped signal from Picasso to the ISM band. Continuing with this modified WiFi AP example, the filter engine has to perform three high level tasks to shape the signals for transmission (to receive a signal, the steps are reversed):

1. **Resampling:** Since the DAC expects an input signal at 200MS/s, first we must upsample the 40Msamples/sec streams to 200Msamples/sec. To accomplish this, the upsampler will interpolate (insert extra samples) to reach the 200MS/s mark.
2. **Filtering:** Upsampling creates aliases [7] that can cause interference. Hence the filter engine must low-pass filter both upsampled streams to remove any undesirable aliasing effects generated by (1) and retain only the upsampled-baseband version of each stream.
3. **Mapping to the appropriate spectrum fragments:** After the first two steps, the engine has two 200MS/s streams each occupying 20MHz at the center frequency. The final step is to move the 20MHz occupancies to the specified fragments in the 100MHz band, i.e. to -38MHz and 22MHz, respectively (corresponding to 2.412GHz and 2.472GHz at a center frequency of 2.45GHz).

Finally, these streams are added together and sent to the DAC, after which they are upconverted to the carrier frequency of 2.45GHz and transmitted over the air.

In Figure 3, the right side shows the block diagram of the current design of Picasso's filter engine. Corresponding to the three high level actions that the filter engine needs to perform, there are three components:

Reconfigurable Filter Structures: These are banks of programmable filters consisting of FIR, IIR, and resampling filter building blocks. These filters will be configured and sequenced to provide the capability of steps (1) and (2).

Intermediate Frequency Converters: These map the signal from incoming digital baseband to a digital intermediate frequency (IF), and provide the capability required for (3).

Filter Engine API: This component acts as the substrate that allows programmable interconnection of the reconfigurable filter and IF converter elements to obtain the desired shaping. It configures the filters, up/down samplers, and digital up/downconverters and also coordinates the movement of streams across these elements. Next, it collects all of the streams, adds them, and sends the final stream to the DAC. The analog output of the DAC is upconverted to 2.45GHz and transmitted. The process for receiving shaped signals is exactly the reverse.

3. FEASIBILITY

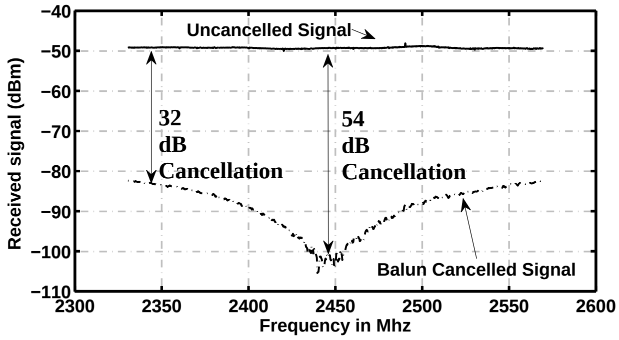

Figure 4: Received Signal Strength versus frequency for a wideband 240 MHz chirp centered at 2.45 GHz.

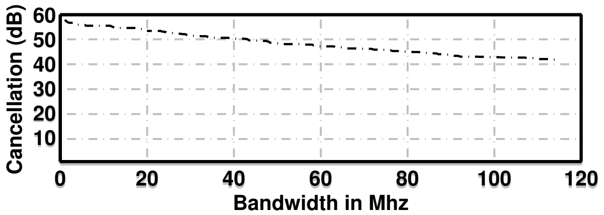

Figure 5: Average Cancellation versus signal bandwidth. We can still achieve more than 40dB cancellation even when canceling a wideband 100 MHz signal.

In this section, we evaluate the initial design of Picasso. We have designed a prototype using a Virtex-5 LX30 FPGA based software radios from National Instruments. This FPGA has a total of 19,200 basic LUT-FF and 32 DSP48E arithmetic unit resources. The FPGA is connected to an NI 5781 Baseband Transceiver, which features quadrature 100MS/s 14-bit ADCs and 100MS/s 16-bit DACs. Altogether, these can cover roughly 80MHz of total bandwidth, allowing us to provide signal shaping over almost the entire 2.4GHz ISM band. Next, we evaluate the feasibility of the two main components in Picasso: the self-interference cancellation and the filter engine.

3.1 Self-Interference Cancellation

The feasibility of a full-duplex signal shaping layer depends on three factors:

1. Dynamic range/resolution of ADC

2. Range of signal strengths expected

3. Amount of self-interference cancellation achievable

The first and third variables are design choices that are within our control. Depending on the second variable, we can determine how much dynamic range/self-interference cancellation Picasso will need, and determine whether or not such a system is feasible. To provide specific numbers, we'll focus on the requirements for an 802.11n system, but the same analysis applies to cellular systems too.

Dynamic range (DR) is defined as the ratio between largest and smallest possible values of a variable of interest. At the transmitter, the dynamic range of the DAC determines the maximum ratio between the powers of the strongest and weakest transmissions. At the receiver, the ADC's dynamic range defines the maximum ratio between the strongest and weakest received signal power. When the dynamic range is exceeded, the converter's quantization noise can bury the weaker signals. The dynamic range of the ADC can be calculated through the following formula [7]

$$\text{DR (dB)} = 6.02 \times n + 1.76\text{dB} \quad (1)$$

where n is the number of bits in the DAC/ADC Resolution. Higher dynamic ranges equate with better performance but also cost exponentially more. Typical WiFi systems use 8-bit DAC/ADCs which would provide 50 dB of dynamic range while 16-bit DAC/ADCs, which are considered high-end and are fairly expensive, would provide 98dB of dynamic range. With current technology, up to 12-bit ADCs are technically feasible and cost effective.

At the transmitter, the maximum ratio between transmit powers over different fragments will rarely exceed 30dB so DAC dynamic range is usually not a concern. On the other hand, if the transmitter is operating while the system attempts to receive, the dynamic range of the ADC at the receiver is critical because the transmitted signal is much stronger than the received signal. To determine the required ADC dynamic range, we calculate the second variable - the range of expected signal strengths. 802.11n is built to operate at SNRs as low as 5dB. Because the typical thermal noise-floor for WiFi systems is approximately -95dBm, the power of the weakest decodable signal is -90dBm. On the other side of the spectrum, the maximum output from a WiFi 2.4 GHz antenna is 23 dBm. Assuming that the transmit and receive antennas are reasonably separated, the attenuation between the two due to path loss can be calculated as follows:

$$\text{Path Loss(dB)} = 36.56 + 20\log_{10}(f) + 20\log_{10}(d) \quad (2)$$

where f is the carrier frequency in MHz and d is the distance in miles. If we assume that the transmit and receive antennas are separated by 15 cm, then the path loss between transmitter and receiver is approximately 25 dB. Thus, the amplitude of the strongest signal at the receiver is -2dBm. Thus our ADC would require more than 102dB (92dB+10dB margin) in dynamic range in order to simultaneously transmit and receive. Given these specifications, not even a top of the line 16-bit ADC would suffice.

The need for balun based self-interference cancellation now becomes apparent. The balun based cancellation technique provides just enough cancellation so that the dynamic range of off-the-shelf ADCs is sufficient to pass the received signal through without distorting it. In order to demonstrate the amount of cancellation, we've prototyped a receive system using this balun cancellation technique. We programmed a wideband signal generator to generate a 240 MHz chirp with a center frequency of 2.45 GHz. This signal is then split over 2 wires. One wire feeds into our balun inversion system, which also includes a variable attenuator and a variable delay element which can be tuned to match the distortion seen by the unmodified signal in the opposite wire. The signals from both wires are then combined and then attenuation is measured. While the balun should theoretically be able to cancel the transmitted signal perfectly, in practice the inverted signal is not an exact duplicate and cancellation is imperfect.

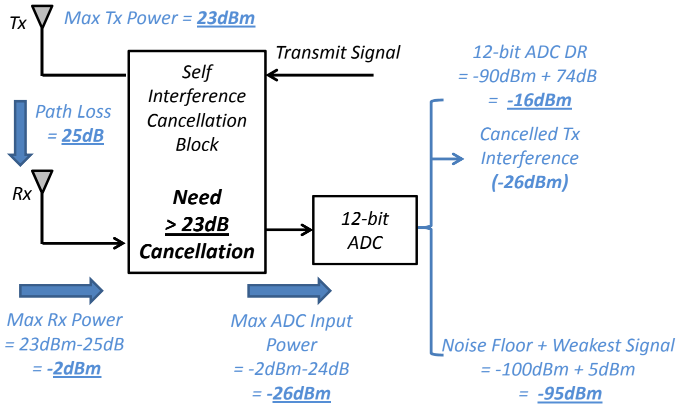

Figure 6: The maximum transmit power for WiFi is 23 dBm. Combined with pathloss, the self-interference cancellation block must provide at least 23dB of cancellation so that the cancelled transmit signal is less than -26dBm. The dynamic range of a 12-bit ADC is from -90 to -16 dBm, so this provides at least a 10dB margin. Our experiments show that up to 40dB of cancellation is possible, meaning that the cancelled transmit signal is well within the dynamic range.

Despite this imperfection, our technique still provides a significant amount of cancellation. Similar to the results seen in [6], we found that the technique can provide upwards of 40dB cancellation over a bandwidth of over 100 MHz. Figure 4 shows the cancellation across frequency for a wideband 240 MHz signal. We see that the technique provides greater cancellation at the center frequency but gradually tapers as we move further away from the center. To demonstrate the relationship between the bandwidth of the transmitted signal versus the amount of achievable cancellation, Figure 5 shows that we can still achieve an average of more than 40dB cancellation across a wide band of 100 MHz. Combining all the components together in Fig. 6, we see that balun signal inversion cancellation technique would allow us to build Picasso using fairly cheap commodity 12-bit ADCs while maintaining a comfortable margin in the dynamic range of the ADC of nearly 25dB.

3.2 Filter Engine

The filter engine relies upon two key atomic blocks to implement shaping: programmable filter structures and digital IF converters. For the filter structures, we utilize DSP48E slices available on the Xilinx LX30 [11]. Each of these slices is a highly configurable arithmetic logic unit which featured pipelined multiplier, adder and accumulator stages and can be clocked at up to 550MHz. Slices can be individually programmed and/or cascaded to implement FIR, IIR, and resampling filtering with relative ease.

The IF converters are implemented using the standard CORDIC approach for sin-cos generation [12] and a complex multiplier. In total, each IF converter requires ~1100 LUT-FF resources and 3 DSP48E slices, and can be clocked at up to 300MHz. Based upon stream requirements, both the CORDICs and DSP48E slices can be time-division multiplexed to support multiple streams. The initial Picasso prototype has 10 filter structures (10 DSP48E slices) and 3 IF converters (9 DSP48E slices and 3300 LUT-FFs). Thus the overall resource consumption on an off-the-shelf, mid-range FPGA is around 15% of the hardware resources. While we expect this fraction to go down further with newer Virtex-6 FPGAs, our initial prototype suggests that the filter engine is feasible using off-the-shelf hardware.

4. DISCUSSION & CONCLUSION

Current wireless devices generally offer very limited flexibility in terms of being able to choose what carrier frequencies and bandwidths they can use. While such a limited design may have sufficed in the past, this inflexibility is quickly becoming a handicap as spectrum becomes more and more scarce and fragmented. Picasso provides a general signal shaping layer that cleanly separates the concern of utilizing fragmented spectrum from the design of higher PHY/MAC layers.

However, we emphasize that Picasso is orthogonal to prior work [5][10][8] in designing MAC/coexistence protocols to dynamically share fragmented spectrum. Instead, Picasso aims to provide a generic programmable substrate on top of which all these prior as well as future, novel MAC designs can be easily implemented. Picasso thus decouples the PHY/MAC from signal shaping and allows them to evolve independently.

We believe that Picasso is a general architectural solution that is not just limited to operation in the unlicensed bands. For instance, cellular spectrum fragmentation is likely to remain an issue globally because of short-sighted regulatory planning. The problem is compounded by the fact that even the same service providers own different fragments of spectrum in different regions, forcing mobile chipsets to accommodate a wide frequency range of operation in order to support roaming. In this situation, Picasso would not only enable operators to utilize fragmented spectrum to support high throughputs, but also allow devices to work in multiple regional markets by providing the programmability to handle any fragmentation pattern. White-space devices operating in the TV bands would also greatly benefit from Picasso, as it would allow them to exploit white-space fragmentation, which varies dynamically depending on which TV channels are in use at any time.

5. REFERENCES

[1] P. Bahl, A. Adya, J. Padhye, A. Walman, "Reconsidering Wireless Systems With Multiple Radios," In *ACM SIGCOMM CCR*, 2004

[2] J. Choi, M. Jain, K Srinivasan, P. Levis, S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication", In *ACM MOBICOM*, 2010

[3] E. Coffman, P. Robert, F. Simatos, S. Tarumi, G. Zussman, "Channel Fragmentation in Dynamic Spectrum Access Systems - a Theoretical Study," In *ACM SIGMETRICS*, 2010

[4] FCC Table of Frequency Allocations, *http://transition.fcc.gov/oet/spectrum/table/fcctable.pdf*, May 6, 2011.

[5] S. Hong, S. Katti, "DOF: A Local Wireless Information Plane", In *ACM SIGCOMM*, 2011

[6] M. Jain, J. Choi, T. Kim, D. Bharadia, K. Srinivasan, P. Levis, S. Katti, P. Sinha, S. Seth, "Practical Real-Time Full Duplex Wireless", In *ACM MOBICOM*, 2011.

[7] A. Oppenheim, R. Schafer, "Discrete Time Signal Processing: 3rd Edition," *Prentice Hall Inc*, 2009

[8] H. Rahul, N. Kushman, D. Katabi, F. Edalat, and C. Sodini. "Narrowband friendly wideband radios." In *ACM SIGCOMM*, 2008

[9] J. Tourrilhes, "Fragment Adaptive Reduction: Coping with Various Interferers in Radio Unlicensed Bands," In *IEEE ICC*, 2001

[10] L. Yang, W. Hou, L. Cao, B. Zhao, H. Zheng, "Supporting Demanding Wireless Applications with Frequency-agile Radios", In *USENIX NSDI*, 2010

[11] UG193: XtremeDSP User Guide, *http://www.xilinx.com/support/docu-mentation/user_guides/ug193.pdf*, June 1, 2010.

[12] DS249: LogiCORE IP CORDIC v4.0 Data Sheet, *http://www.xilinx.com/support/documentation/ip_documentation/cordic_ds249.pdf*, March 1, 2011.