

Picasso: Flexible RF and Spectrum Slicing

Steven Hong
Stanford University
hsiyong@stanford.edu

Jeffrey Mehlman
Stanford University
jmehlman@stanford.edu

Sachin Katti
Stanford University
skatti@stanford.edu

ABSTRACT

This paper presents the design, implementation and evaluation of Picasso, a novel radio design that allows simultaneous transmission and reception on separate and arbitrary spectrum fragments using a single RF front end and antenna. Picasso leverages this capability to flexibly partition fragmented spectrum into multiple slices that share the RF front end and antenna, yet operate concurrent and independent PHY/MAC protocols. We show how this capability provides a general and clean abstraction to exploit fragmented spectrum in WiFi networks and handle coexistence in dense deployments. We prototype Picasso, and demonstrate experimentally that a Picasso radio partitioned into four slices, each concurrently operating four standard WiFi OFDM PHY and CSMA MAC stacks, can achieve the same sum throughput as four physically separate radios individually configured to operate on the spectrum fragments. We also demonstrate experimentally how Picasso's slicing abstraction provides a clean mechanism to enable multiple diverse networks to coexist and achieve higher throughput, better video quality and latency than the best known state of the art approaches.

Categories and Subject Descriptors: C2.1 [Computer Systems Organization]: Computer Communication Networks

General Terms: Algorithms, Performance, Design

Keywords: Radio Virtualization, Interference Cancellation

1. INTRODUCTION

Many applications can benefit from the capability to simultaneously and independently use arbitrarily sized but separate spectrum fragments with a single radio and antenna. By this capability we mean that the radio can simultaneously transmit, simultaneously receive, or simultaneously transmit (TX) and receive (RX) on arbitrary but separate spectrum fragments. For example, we can use it for spectrum aggregation in fragmented ISM bands as shown in Fig. 1(A). A WiFi AP can run independent OFDM PHY and CSMA MAC protocols on two WiFi channels to simultaneously serve two legacy WiFi clients assigned to different channels and achieve significantly higher throughput than a legacy AP that is restricted to the use of only one channel at a time. Similarly, a WiFi client radio with such a capability can simultaneously connect to multiple WiFi APs

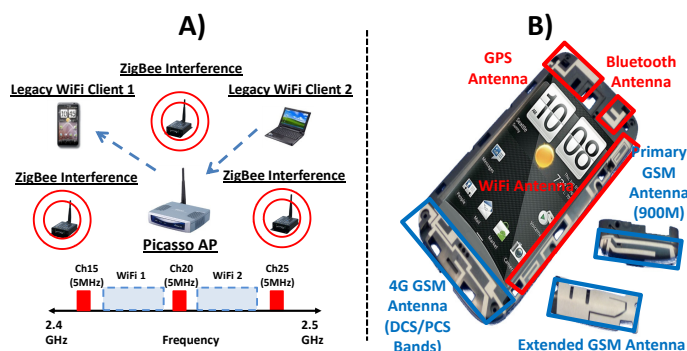


Figure 1: A) Fragmented Spectrum is common in ISM band: Picasso AP can aggregate multiple fragments and simultaneously service multiple clients on different channels. **B) Samsung Galaxy Printed Antenna Layout:** Instead of having an antenna per protocol, Picasso enables multiple protocols to use the same radio and antenna, saving space on the device.

on different channels and obtain a much higher aggregate throughput than current radios that can transmit or receive on only one channel at a time.

Such a capability could also be used for radio sharing and coexistence. Portable consumer devices such as smartphones must accommodate a growing list of separate ISM band protocols such as WiFi, WiFi-Direct, Zigbee, NFC, and Bluetooth. Current practice is to use a separate radio and antenna for each protocol, shown in Fig. 1(B), but as the number of radios increases it becomes difficult to find enough space to separately place all the antennas these radios would need (e.g., the iPhone 4 “antenna-gate” was caused by antennas placed too closely [30]). Instead, a radio with a single antenna that allows simultaneous TX/RX on arbitrary spectrum fragments could be shared among all the above protocols. WiFi would use one fragment, Zigbee would use another, and so on, saving valuable real estate on space-constrained devices. They can operate their own independent PHY/MAC protocols on the shared radio without interfering with each other.

However, it is non-trivial to design such a radio. The key obstacle is that current radios cannot simultaneously transmit and receive on different arbitrary spectrum fragments with a single, shared Radio Frequency (RF) front end and antenna. The reason for this is that the transmitted signal causes high-powered self-interference, which saturates the RX chain and Analog-to-Digital Converter (ADC), consequently nulling the received signal. While the standard solution is to utilize static, analog RF filters to eliminate the self-interference, such an option is infeasible because spectrum fragmentation is dynamic—available spectrum in the ISM band varies in space and time, depending on the presence of other wireless networks. Consequently, if a radio wants to leverage all the available spectrum and be able to simultaneously transmit and receive on different fragments, the shared analog front end would need programmable analog filters that can

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'12, August 13–17, 2012, Helsinki, Finland.

Copyright 2012 ACM 978-1-4503-1419-0/12/08 ...\$15.00.

be dynamically configured to let only the received signals through and filter out the self-interference. Analog filters however are typically statically configured and programmable analog filters that can be changed dynamically are expensive and impractical to deploy in current radios [18].

In this paper, we present the design and implementation of Picasso, a novel full duplex circuit design that sufficiently cancels (instead of filters) the self-interference in analog and prevents RX front end and ADC saturation, enabling the radio to cleanly recover the received signal. Our key contribution here is a circuit design that (1) isolates TX and RX signals at a single antenna by incorporating a circulator [23], and (2) exploits the fact that the self-interference signal travels through the fixed, known circulator channel to design a passive self-interference cancellation circuit. This allows a *radio to simultaneously transmit and receive on arbitrary spectrum fragments even while using a single RF front end and antenna*.

Our design improves on all prior related work on full duplex [9, 14, 5] wireless since they require at least two antennas (one for TX, one for RX), and these need to be separated by 15–20cm, which is untenable for small personal gadgets such as smartphones or tablets.

Picasso leverages this full duplex capability to build an abstraction that allows one to flexibly slice a single radio and antenna into separate independent slices operating on different spectrum fragments. Each slice is associated with a specific spectrum fragment in the ISM band (whose width/position can be programmatically specified). The key property is that the operation of each slice is decoupled from the other slices, i.e., the slice is free to run whatever narrowband PHY and MAC protocols it chooses, and the protocol behavior is not impacted by any other slice that may be present on the shared radio and antenna. Thus, in the above scenarios, the AP would have two slices corresponding to the two spectrum fragments and run two independent WiFi OFDM/CSMA protocols on the two slices in parallel. Similarly, a radio could be shared amongst multiple protocols (e.g. WiFi, Bluetooth, and NFC) by assigning independent slices to the appropriate spectrum fragment and running the corresponding protocol.

Further, to ensure that each slice can use existing, well-engineered narrowband PHY and MAC protocols on each slice, Picasso includes a reconfigurable filter engine that transparently takes signals spread over different spectrum fragments, and efficiently filters and resamples them so that the higher layers just see a simple sample stream consisting of narrowband digital samples. The higher layers are then free to process these samples with any narrowband PHY technique they choose, and schedule access to the slice with a MAC protocol of their choice. The slice, for all intents and purposes, appears as their own piece of spectrum centered at zero, operating on their own radio. Picasso thus completely abstracts out the complexity of spectrum fragmentation.

We design and implement a prototype of Picasso on Xilinx Virtex-5 FPGA-based software radios [32]. Our implementation consists of both the radio design that provides the slicing abstraction, as well as a WiFi-like contiguous OFDM PHY and CSMA MAC to operate on top of the slices. We show that Picasso’s implementation of the slicing abstraction provides strong decoupled operation, i.e. there’s no SNR loss because of either programmable filtering or simultaneous TX/RX. In other words, a Picasso radio achieves the same throughput as one would have achieved by using several independent radios statically configured to operate on individual slices.

To demonstrate the generality and applicability of Picasso, we demonstrate that Picasso provides clean abstractions to solve coexistence problems in dense deployments by evaluating a scenario where a number of diverse networks have to coexist with each other while sharing radios and fragmented spectrum (shown later in Fig. 16).

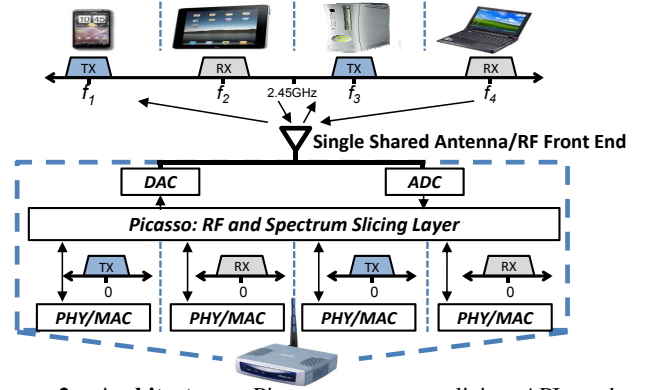


Figure 2: Architecture: Picasso exposes a slicing API to the PHY/MACs of the different slices and allows them to independently operate.

Each coexisting network has different objectives: the WiFi network wants to maximize throughput, the gameplay network wants to maximize video quality, and the game controller network wants to minimize latency. We compare Picasso with a state-of-the-art technique for exploiting fragmented spectrum, SWIFT [24], and show that Picasso achieves nearly $2\times$ higher throughput, a 10dB higher PSNR in video quality, and lowers latency by a factor of $2\times$. Further, the improvements are achieved with simple narrowband OFDM PHY and CSMA MAC implementations running on top of each slice, leading to a modular design that reuses existing, well-engineered components.

Picasso builds on our recent workshop paper [12] that sketched out the basic approach. In this paper, we present a novel full duplex design, as well as present a full design and prototype implementation. Finally, Picasso’s focus is the mechanism to implement slicing of spectrum and RF front ends. An equally interesting question is given the set of coexisting networks and available fragmented spectrum, what is the best allocation of slices, and the distributed signaling algorithm needed to inform all devices their slice allocations? While this question is orthogonal to Picasso’s design, we show in Sec. 7 that Picasso simplifies this problem considerably compared to a design that uses traditional radios which are incapable of slicing.

2. OVERVIEW

Picasso’s key contribution is that it allows one to flexibly slice a radio into multiple slices each operating on different, arbitrary spectrum fragments with their own independent PHY/MAC protocols. Fig. 2 shows how Picasso differs from current designs—between the PHY/MAC layers and the actual transmission of RF signals on the antenna, there is an additional slicing layer. Picasso exposes an API to the PHY/MAC layers which consists of streams of complex digital baseband samples flowing between the slicing and the PHY layers for the slices. Each stream carries an internal header which includes a tuple specifying the spectrum slice on which those digital complex samples are transmitted or received. Picasso thus makes the architecture more amenable to evolution. By decoupling how spectrum and radios are sliced from how packets are processed and scheduled, the PHY and MAC can evolve independently and innovations can be easily integrated without having to change the radio. To realize this abstraction, Picasso needs to overcome two key challenges.

1. RF Isolation: First, Picasso must provide complete isolation between the operation of multiple slices. Slice isolation is relatively easy in the scenario where the device is simultaneously transmitting (or receiving) on all adjacent slices. But when the radio is transmitting on one spectrum slice while receiving in another, isolation is significantly harder. In order to provide this functionality in practice, the radio has to cope with high-powered self-interference introduced

by its own transmitted signal at the antenna and the RX chain, which causes *receiver saturation* [14]. Since our design utilizes a single antenna, the transmitted signal directly leaks into the RX chain.

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 80+dB stronger). As we will show in Section 3, 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 receiver saturates and the received signal is lost in quantization. Note that this is a problem even though the transmitted signal is on a different band than the received signal. Spectrum slices 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 and have no choice but to allow all signals across the ISM band through.

A naive way to prevent receiver saturation is to utilize multiple, parallel receive RF chains. However, supporting multiple RF chains on mobile devices is expensive [11] in terms of space and power, and does not scale well with an increasing number of slices. Another plausible solution to prevent receiver saturation is to use statically configured analog filters to filter out the transmitted signal and only pass the received signal. But again, since spectrum slices assigned vary over time and space, the analog filters would have to be dynamically reconfigurable. However, analog filters are typically optimized for static operation [18]. Hence in order to support dynamic configurations, we would require a bank of static analog filters of varying widths, with at least as many filters as the combinatorial number of spectrum slices, possible slice widths, and center frequencies.

2. Programmable Abstraction: Picasso’s second challenge is to provide the abstraction of a single narrowband radio on each allocated spectrum slice to the higher layer protocols. To do so, Picasso must ensure that the digital samples being sent down from the PHY for transmission occupy a bandwidth equal to the slice’s spectrum allocation w , and are centered at DC. Picasso must also ensure that the samples are shifted to the correct center frequency f_c that is associated with that slice, and manage this abstraction across multiple MAC/PHY instances. The opposite is performed on reception.

Picasso’s goal is to design a radio with slicing capability using only one TX and RX RF chain, one antenna, and one simple wide-band analog filter. Solving these challenges in aggregate enables Picasso to virtualize the radio, allowing multiple independent MAC/PHY protocols to share the same RF front end. The next two sections elaborate on how to solve these challenges.

3. Picasso: RF FRONT END SLICING

In this section, we describe how Picasso slices a single RF front-end. We focus on the challenging case where the radio is simultaneously transmitting and receiving over different slices. Picasso’s key insight is to cancel the self-interference in analog instead of filtering it. Picasso subtracts the self-interference signal from the received signal so that its effect is eliminated and receiver saturation does not occur. We describe the design of Picasso’s novel self interference cancellation technique below, beginning with a discussion of how much cancellation is needed for Picasso to prevent receiver saturation.

How much self-interference cancellation is needed? The amount of cancellation required depends on two factors:

1. Dynamic range/resolution of ADC
2. Range of the expected signal strengths

To provide specific numbers, we focus on the requirements for an 802.11 WiFi system, since it is the protocol with the highest transmit power in the ISM band. Dynamic range (DR) is defined as the

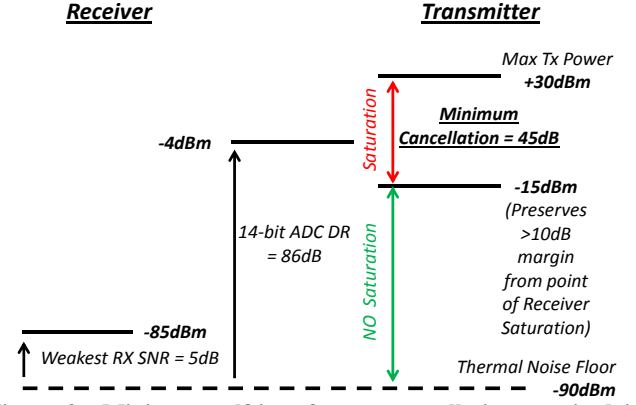


Figure 3: Minimum self-interference cancellation required is 45dB: The maximum transmit power for WiFi is 23 dBm, with a PAPR of 7dB, and so the maximum peak power is 30dBm. The dynamic range of a 14-bit ADC ranges from -90 to -4 dBm, but to preserve >10dB to receive the weakest detectable signal (5dB is the minimum but we preserve at least 10dB to provide some buffer), the maximum power of the canceled transmit signal must be less than -15dBm.

ratio between largest and smallest acceptable values of a variable of interest. At the transmitter, the dynamic range of the Digital-to-Analog Converter (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 buries the weaker signals. The dynamic range of a DAC or ADC is calculated as follows [20]

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

where n is the number of bits in the DAC/ADC Resolution. For example, Picasso uses 14-bit ADCs, which provide about 86dB of dynamic range.

At the transmitter, the maximum ratio between transmit powers over different slices 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 receiver ADC 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. 802.11 is built to operate at SNRs as low as 5dB. Because the typical thermal noise-floor for WiFi systems is approximately -90dBm, the power of the weakest decodable signal is -85dBm. On the other side of the spectrum, the maximum output from a WiFi 2.4 GHz antenna is 23 dBm. Further, OFDM exhibits relatively high PAPR (Peak to Average Power Ratio) of around 7dB [3], so the highest transmit power a signal will have is around 30dBm.

Commodity 14-bit ADCs [33] that can operate at a rate of 200 MS/s (enough to cover the 100MHz 2.4GHz ISM band) are commercially available. A 14-bit ADC has 86dB of dynamic range. However, with a noise floor of -90dBm and a max transmit power of 30dBm, we would need 120dB of dynamic range, far beyond the capabilities of commodity ADCs. Hence in order to ensure receiver saturation does not occur, Picasso must ensure that self-interfering transmitted signal power is reduced by at least 45dB. Fig. 3 demonstrates the required interference power reduction in order to comfortably prevent receiver saturation.

Note that the above analysis is for the *worst-case* scenario with maximum transmit power and PAPR for OFDM signals. In practice, transmit powers used are smaller than 23dBm and a PAPR of




Figure 4: A) In Band Cancellation - QHX-Balun Cancellation achieves sufficient in-band cancellation; **B) Out of Band Cancellation Leakage - QHX-Balun cancellation** leaks a significant amount of signal power into adjacent bands, raising the noise floor by as much as 9dB.

7dB is quite high [3]. Further, other ISM PHYs such as Zigbee and Bluetooth do not have high PAPRs. Hence, if Picasso can design a technique that provides 45dB of self interference reduction, it should be more than sufficient.

3.1 Prior Full Duplex Techniques

There are three recent full duplex designs [14, 9, 15] that provide sufficient self interference cancellation (45dB). However, none of this work is directly applicable for Picasso because of the following reasons:

- First, all prior designs use two or more antennas to achieve full duplex, i.e. they use TX and RX antennas that are separated by distances ranging from 15–20 cm. However, in a mobile device such as a smartphone or a small tablet where circuit real estate is at a premium, such multiple and distant antenna placement is infeasible (e.g., the iPhone 4 “antenna-gate”). Picasso’s design must use a single antenna for both TX/RX to ensure it is applicable to small consumer devices.
- Second, as we show below, while prior work provides sufficient in-band cancellation, they do not provide isolation across slices. Specifically, since the interference cancellation technique uses active circuits that leak interference to neighboring spectrum, they harm the received signal’s SNR and hurt performance.

To cancel self-interference, the best performing prior design [14] utilized separate TX/RX antennas separated by 15cm, so the transmitted signal is received after some variable attenuation and delay at the RX antenna. A cancellation circuit tries to obtain the *inverse* of the transmitted signal using a a balanced/unbalanced (balun) transformer, dynamically adjusts the attenuation and delay of the inverse signal to match the over-the-air self interference, and combines it with the signal received on the RX antenna. The prior design used active QHX220 chips to mimic programmable attenuation and delay, thus we call this technique *QHX-balun cancellation*. We implemented and evaluated the amount of cancellation one could get with QHX-balun cancellation, with TX/RX antennas spaced 15cm apart. We found that the QHX-balun cancellation techniques provide a significant amount of cancellation, 35dB on average for a 40MHz transmitted signal (Fig. 4(A)) which, when combined with the signal attenuation of 20dB due to the TX/RX antenna separation, is comfortably sufficient for our receivers that use 14-bit ADCs. The number is consistent with the prior reported results.

However, there is a catch. Even though the technique manages to cancel the self-interference, it ends up leaking a lot of interference into adjacent spectrum where there is no transmitted signal. We plot the amount of interference leakage in adjacent spectrum as a function of the spectrum distance from the edge of the transmit band. As Fig. 4(B) shows, the QHX-balun based technique leaks a significant

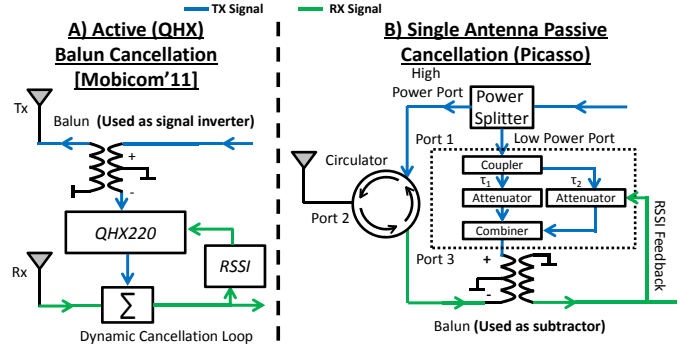


Figure 5: A) QHX-Balun Cancellation Diagram; B) Picasso Single Antenna, Passive Cancellation: Picasso’s cancellation design is *passive*, preventing interference leakage into adjacent bands, and *requires only a single antenna*.

amount of interference, as high as 10dB in adjacent bands that are separated by 1 MHz.

Apart from the fact that it requires separate TX and RX antennas, *interference leakage in adjacent spectrum* is a show stopper for Picasso since the adjacent spectrum may be used to receive concurrently, and the received signal’s SNR will be reduced. To put the impact of interference leakage in perspective, a 10dB reduction in SNR can result in an order of magnitude decrease in throughput. The culprit is the active QHX chip used in the prior design. Even though it does provide self-interference cancellation in the transmit spectrum fragment, it distorts and leaks the transmitted signal power into neighboring spectrum fragments.

3.2 Picasso’s Self-Interference Cancellation Technique

Picasso designs a novel self-interference cancellation technique that makes two significant contributions over all prior full duplex designs: (1) Single antenna full duplex and (2) Passive, leakage-free cancellation. We describe each in more detail below.

Single Antenna Full Duplex: Fig. 5(B) shows the block diagram of Picasso’s single antenna, full duplex circuit. There are two key differences. First, the two separate antennas are replaced by a single antenna and a circulator (which we describe below). Second, instead of the active programmable QHX-220 component, there is a passive cancellation circuit which consists of two fixed delay lines and programmable, resistive switched attenuators. The key takeaway is that all the components in Picasso’s design are *passive*, whereas prior approaches used active components. We briefly describe the circuit components and how they work together to provide the required 45dB of cancellation.

A circulator [23] is a three-port device—labeled 1, 2 and 3 (Fig. 5(B))—that provides limited isolation between any two ports, and has historically been used in ranging systems such as radar [23]. A common circulator design contains an arrangement of conductor striplines and ferrite blocks. When the ferrites are biased, the circulator produces magnetic fields which enhance or oppose any RF energy flowing across the striplines. This biasing causes directional RF signal flow, routing the energy in a circular fashion from port 1 to 2, 2 to 3, and 3 to 1, but not in the reverse direction. Thus the power scattering matrix which represents the power flow allowed by the circulator between any pair of ports, S , for an ideal circulator that has perfect isolation quantified by the factor $\Gamma = 0$ can be written as

$$S = \begin{pmatrix} \Gamma & 1 - \Gamma^2 & \Gamma \\ \Gamma & \Gamma & 1 - \Gamma^2 \\ 1 - \Gamma^2 & \Gamma & \Gamma \end{pmatrix} \quad (2)$$

where the row indicates the source port and the column indicates the destination port. Microwave circulators are available at ISM band




Figure 6: A) Picasso Cancellation vs. Delay Spread For delay mismatches larger than those introduced by 1mm copper traces (≈ 4.76 ns), Picasso will not be able to achieve sufficient cancellation. **B) Picasso Cancellation vs. Attenuation Spread:** Picasso requires attenuation spreads of less than .5dB.

frequencies off the shelf [34] and their operation is similar in principle. Our design utilizes a 3-port circulator to concurrently route the TX and RX signals to and from the antenna as shown in Fig. 5(B). The TX signal is first passed through a 8dB power splitter, which splits the input signal into a high-power and low-power output. The high-power output has 85% of the input power and is connected to port 1 of the circulator. Hence, very little transmit signal power is wasted. The low-power output is used as a reference signal for self interference cancellation, which we describe after the next paragraph.

The antenna is connected to port 2 and the RX chain is connected to port 3. An ideal circulator would ensure that the signal on port 1 (the TX signal) would never leak to port 3 (the RX signal). However, in practice, the impedance of the circulator input ports are not perfectly matched with the outputs resulting in less than perfect isolation, $\Gamma \neq 0$, and there is leakage from port 1 to port 3. Despite this mismatch, we show in Section 6 that the circulator can still provide around 15dB of isolation between the transmit and receive chains, i.e. the self-interference we see on the receive port is 15dB lower than the transmitted signal power. However, remember that we need to provide 45dB of self-interference reduction for Picasso—15dB is not enough. Further, the cancellation should not leak any additional interference to neighboring spectrum bands.

Leakage-Free, Passive Cancellation: Our cancellation technique builds on two key observations. First, the main cause of interference leakage is the active QHX220 chip which serves as the programmable attenuation/delay component. Hence, to avoid interference leakage, we have to use passive components which do not introduce such distortion. However, even though passive and programmable attenuators are available off-the-shelf, passive delay lines are not.

Our key insight here is that the transmitted signal is going through static circuit components—it never travels through the air. Hence, the variance in the expected delay is significantly smaller, and Picasso can measure this limited range via a simple calibration procedure (these devices are designed to provide predictable behavior for a given frequency band). Picasso discards the active components required to amplify the signal to compensate for the larger possible range of delays and attenuation.

While the circulator limits the possible range of variation, environmental factors such as temperature and multipath reflections still cause some spread in the delay and attenuation of the self-interference. To simulate what effect these distributions have on cancellation, Fig. 6 shows a simulation, where we send a 40MHz OFDM signal at 2.4GHz carrier, and pass it through a single fixed delay and attenuation using measurements from our prototype. Next, we simulate varying delay and attenuation spreads by varying the wire length and attenuation components. As expected, we see that cancellation declines with increasing variance in the spread of delay and attenuation. While the circuit can achieve 30dB+ of cancellation when the length and at-




Figure 7: Reference Signal shown in blue (X), interference signal shown in green (Y), residual signal shown in red (Z). **A) Effect of delay mismatch on signal cancellation:** Delay offsets result in imperfect cancellation. **B) 2 fixed delay paths with tunable attenuators:** Different attenuator values compensate for different delay values. **C) Helps mitigate effects of multipath due to reflections:** Picasso is tuned to cancel the aggregate response (Y_{Tot}), not just the LOS component.

tenuation deviations are small (less than 1mm in length and .5dB in attenuation spread), the performance suffers at larger mismatches.

To get a sense for what reasonable spreads in delay and attenuation can be, we look at what happens when a hand touches the antenna, a situation which commonly occurs in a mobile handset application. The hand changes the impedance of the antenna, which in turn causes reflections through the circulator to manifest at different delays. Our Vector Network Analyzer measurements showed that touching a standard WiFi Dipole antenna resulted in delay and attenuation spreads as large as 20ns and 3dB respectively. In copper traces, this would correspond to delay offsets of more than 4.2mm. As predicted in Fig. 6, we saw the cancellation drop precipitously from 40dB+ to less than 20dB. These possible perturbations underscore the need for a certain degree of tunability in order to ensure that the system is robust. While tunable delay elements are not readily available, tunable attenuators are. We leverage these devices in order to provide a margin for disturbances.

Our cancellation circuit utilizes fixed delay lines and programmable attenuators to compensate for a range of disturbances, including those which result in variable delay spreads seen in the example above. The trick is to realize that a change in delay manifests itself as a change in phase in the frequency domain. Analog signals of a particular frequency can be represented by a phasor/vector [19] so to aid our illustration, we sketch the vector representation of signal cancellation with delay mismatch ($\hat{\tau}$) in Fig. 7A. As you can see, the greater the delay mismatch, the larger the residual signal becomes.

A simple solution here would be to simply change the phase of the reference signal as well. The QHX220 does exactly this, but unfortunately requires integrated LNAs to inject additional power in order to change the phase by controlling the I and Q voltages. As we showed, these active components leak significant power into adjacent bands and thus should be avoided if possible.

We observe that we can achieve the desired phase shift if we simply utilize two delay paths instead of one. Recall that while programmable delay elements are not readily available, programmable attenuators are. By provisioning a programmable attenuator on each of our delay paths and then taking the aggregate response of both paths, we are able to mimic phase offset and adjust to variable delay spreads caused by perturbations. Again, let's take a look at the vector representation. As you can see in Fig. 7B, the aggregate response of two delay lines with programmable attenuators enables us to change the phase of the signal. Notice that in a single path, a programmable attenuator would not help at all with delay spread matching; the addition of the additional degree of freedom provided by the second path is essential.

Lastly, this approach enables us to combat internal reflections within the circulator and multipath in the surrounding environment. Because of imperfect isolation, the output signal from the circulator is actually spread in time. For instance, the power from port 1 leaking

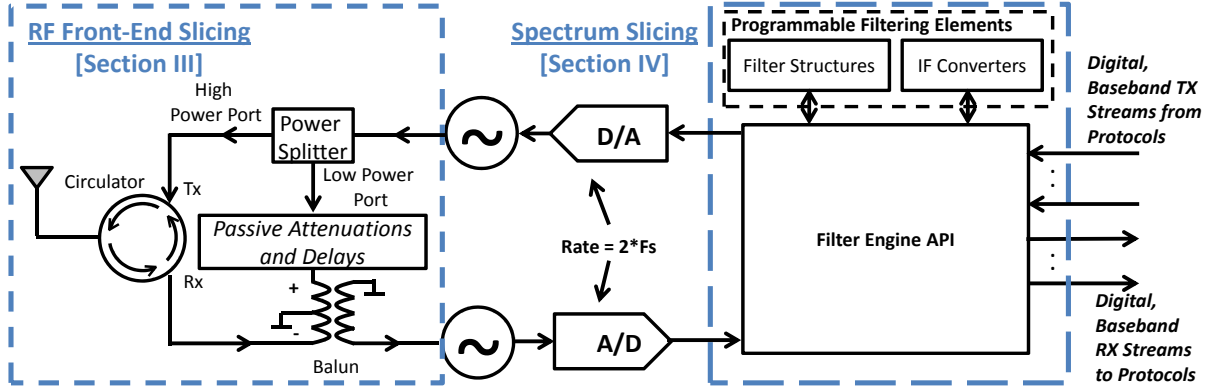


Figure 8: Picasso Block Diagram: The self-interference cancellation setup uses passive components for matching attenuation and delay. The filter engine setup processes digital baseband streams and appropriately resamples, filters and maps each stream to the appropriate fragment.

into 3 which we are trying to cancel, actually leaks to 2 as well. It then subsequently leaks back into 1 and ends up at our receive port again, albeit at a lower power (45dB lower with our circulator). If we simply match for the first delay path, we would be limiting our cancellation at 45dB due to this reflected path. Similarly, the multipath from the environment is typically on the order of 30-50dB lower as well. As you can see in Fig. 7C, the multipath components slightly affect the LOS component. With our cancellation circuit, the attenuation values can be fine-tuned to cancel the aggregate signal and achieve higher cancellation.

In summary, our cancellation circuit consists of two programmable attenuators and fixed delay line lengths which are statically configured to roughly straddle the delay the transmitted signal experiences through the circulator and any connecting wires (e.g., one delay line has less delay while the other has a larger delay than the transmitted signal). The attenuator values for each delay path are adaptively adjusted once every millisecond using a simple gradient descent algorithm [14] in order to compensate for changing environmental factors. Note that despite the fact that the attenuator values are adaptively changed, the attenuators themselves as well as the delay lines are passive. Thus, Picasso does not leak interference into adjacent bands as the QHX220 nonlinear active components do [14].

Reducing Transmit Insertion Loss: A final wrinkle is that instead of using the balun as a signal inverter, we use it as a subtractor circuit. Fig. 5(A) shows how a balun in the prior design [14] takes an input signal on the unbalanced tap and produces two output signals which are inverses of each other on the balanced taps. The drawback of this approach is that the input power is split between the two inverses, and thus half the power is spent in obtaining the reference signal for cancellation. Our insight is that the balun's operation can be modeled in reverse as taking two inputs on the balanced side and producing the subtraction of the input signals as the signal on the unbalanced side, shown in Fig. 5(B). Hence if the two inputs are replicas of each other, the output will be a zero signal. By using the balun to subtract rather than obtain the inverse, unlike prior designs, Picasso saves nearly half the transmit power from going to waste.

Picasso passes the output of the cancellation circuit to one of the balun inputs, and the other input is from port 3 of the circulator. We show in Sec. 6 that coupled with the circulator, this technique provides sufficient cancellation ($>45\text{dB}$).

Picasso's design thus eliminates the need for separate TX and RX antennas, preserving valuable real-estate in mobile devices. Further, unlike all prior work [14] which need to accommodate the large variance of the wireless channel between the TX and RX antennas with active cancellation circuits, Picasso's technique *requires no active components*. Because the circulator scattering characteristics are relatively fixed, once Picasso is calibrated at production time only small adjustments are required. Since the tuning scale is relatively small,

all the components used in the design can be passive, enabling Picasso to guarantee that there is no interference leakage.

4. Picasso: SPECTRUM SLICING

Picasso provides transparent spectrum slicing, i.e., it allows each slice to have its own separate PHY and MAC protocol. The higher layer protocols are presented with the abstraction of a single narrowband radio that is exclusively for their use on the allocated spectrum slice. Picasso designs a digital filter engine that provides such spectrum slicing. Specifically, recall that digital samples being sent down from the PHY for transmission occupy a bandwidth equal to the slice's spectrum allocation w , and are centered at zero. The filter engine's job is to ensure that the samples are shifted to the correct center frequency f_c that is associated with that slice. On the receive side, the filter engine is doing the reverse,

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, there is a single oscillator at 2.45GHz which upconverts the signal to the ISM band. The filter engine performs 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 Picasso upsamples the 40MS/s streams to 200MS/s. To accomplish this, the upsampler will interpolate (insert extra samples) to reach the 200MS/s mark.
2. **Filtering:** Upsampling creates aliases [20] that can cause interference. Hence the filter engine must low-pass filter both upsampled streams to remove any undesirable aliasing effects generated by the first step and retain only the upsampled-baseband version of each stream. To accomplish this, Picasso designs a reconfigurable filter bank, which consists of banks of programmable filters consisting of FIR, IIR, and resampling filter building blocks. These filters will be configured and sequenced to both resample and remove the aliasing effects.
3. **Mapping to the appropriate spectrum slices:** After the first two steps, the filtering 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). For this Picasso uses intermediate frequency converters, which map the signal from incoming digital baseband to a specified digital intermediate frequency.

Finally, these streams are added together and sent to the DAC, then upconverted to the 2.45GHz carrier frequency and transmitted. Fig. 8, the right hand side shows the block diagram of the current design of Picasso's filter engine.

Note that the filtering also automatically takes care of any residual self interference. The analog self-interference cancellation prevents the receiver ADC from saturating, but by itself, it is insufficient to fully cancel out the self-interference. But now with the ADC dynamic range no longer saturated, Picasso's filtering engine can digitally remove the remaining self-interference. To summarize, Picasso's design consists of 2 main components, the self-interference cancellation block and the filtering engine. Individually, each component is inadequate but together, they enable Picasso to provide the full slicing abstraction.

5. IMPLEMENTATION

Picasso (Fig. 9) is implemented in two parts – the self interference cancellation is realized as a discrete board level prototype (shown in yellow) and the spectrum slicing portion is realized using a Virtex-5 LX30 FPGA based software radio from National Instruments[32] (shown in red). The FPGA is connected to an NI 5781 Baseband Transceiver which uses the USRP XCVR2450 as the radio front end. The radio front ends are then connected to Picasso's self-interference cancellation board.

5.1 Self-Interference Cancellation

ADC resolution is an important specification because it determines the receiver's dynamic range. The NI 5781 baseband transceiver features quadrature 200MS/s 14-bit ADCs and 200MS/s 16-bit DACs and provide approximately 86dB of dynamic range and coverage over the entire 100MHz ISM band.

Analog components have an active power range of operation over which they act linearly and predictably. Beyond these points, performance is often uncharacterized and starts to break down. Let's take a look at the XCVR2450, an off-the-shelf radio front end utilized in Picasso, to see which components we need to keep an eye on and make sure are not saturated.

One of the XCVR2450's active components we have to watch out for is the low noise amplifier (LNA). The figure of merit for the LNA is the input 1dB compression point, which represents the maximum power that this device can handle before it saturates. The XCVR2450 LNA [36] has a 1dB compression point of 0dBm. As long as we maintain a sufficient level (6dB is a good rule of thumb) below this point, the LNA performance will not suffer. With a maximum self-interference signal of 30dBm including PAPR, Picasso must provide at least 36dB worth of cancellation. Another component in the XCVR2450 to keep an eye on is the mixer [37], which has a Third Order Intercept (IIP3) of 11.5dBm. The IIP3 is a measure of the nonlinearity of the component and specifies the maximum handleable power before nonlinearities are no longer negligible. Because it comes after the LNA which has a maximum gain of 17dBm, we require 41.5dB of cancellation to limit the input power at the LNA ($11.5\text{dBm} - 17\text{dB} - 6\text{dB} = -11.5\text{dBm}$), allowing the gain control access to its full range. Fortunately, we've already obtained enough cancellation for the ADC Dynamic Range which turns out to be the most sensitive part of our system. Both of these cancellation numbers (36dB and 41.5dB) are lower than the 45dB needed for the ADC, so as long as Picasso can achieve 45dB of self-interference cancellation, RF performance should not suffer.

5.2 Filter Engine

The filter engine is implemented on two main processing modules in the software radio, a Virtex-5 95T FPGA and a real-time processor (NI PXIe-8133 RT Module). The Virtex-5 95T FPGA has 14,720 slices and nearly 9Mb of random access memory (RAM). The slices are used for almost every operation on the FPGA, while the RAM use is limited to storing data while averaging and providing temporary

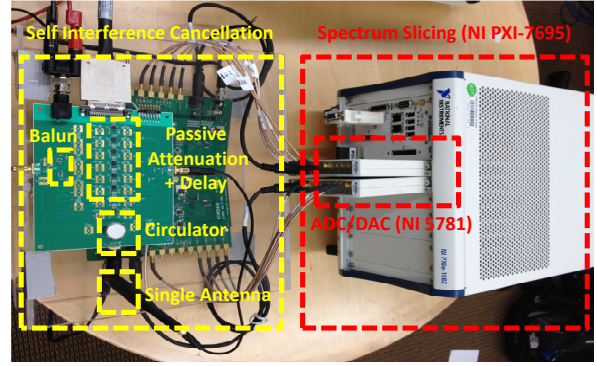


Figure 9: Picasso Radio Physical Setup: Built with off the shelf components and software radios from National Instruments.

storage when restructuring the output of the definition of FFTs from bit order to natural order. For the filter structures, we utilize DSP48E slices available on the Xilinx LX30 [38]. Each of these slices is a highly configurable arithmetic logic unit which feature 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.

5.3 PHY/MAC Implementation

We have also implemented a WiFi-style OFDM PHY and a CSMA MAC split across the FPGA and the realtime OS. The OFDM PHY can be configured to operate over different bandwidths (from 20MHz to 5MHz), and runs in realtime on the Software Defined Radio platform. It supports all the WiFi constellations (from BPSK to 64-QAM) as well as channel coding rates (1/2, 2/3, and 3/4 convolutional coding). We have not implemented an automatic rate adaptation algorithm, but in the experiments we empirically pick the combination of constellation size and channel coding rate that maximizes the link throughput. Thus, our prototype implements all the relevant features from a WiFi PHY/MAC. Note that we can run multiple instances of the WiFi PHY and MAC, one for each slice on the device.

6. EXPERIMENTAL EVALUATION

In this section we evaluate the performance of Picasso's prototype in an indoor testbed of NI Virtex-5 software radios that implement Picasso. As discussed before, we have implemented a WiFi style OFDM PHY and multiple instances (up to 4 due to FPGA hardware constraints) of these can be run on top of the slices in a single Picasso radio node. We benchmark Picasso's overall throughput performance, the efficacy of its self interference cancellation and programmable filtering components.

Compared Approach: Our main claim is that Picasso can slice a single RF front end and available fragmented spectrum into multiple slices, each of which can operate independent and decoupled PHY protocols. Hence the natural comparison is to a system that uses *multiple, separate radios*, one for each slice, with each radio configured to work on the corresponding specific spectrum slice, using its own PHY protocol. Unless stated otherwise, our comparison metric is total throughput.

An artifact of our current prototype is that our NI FPGA software radios use the USRP XCVR2400 as the RF front end. The XCVR's widest analog filter is only 40MHz, hence even though our ADCs and DACs have enough speed to cover the entire 100MHz ISM band, the RF front end limits us to 40MHz. Hence, in all our experiments the maximum spectrum that is available, either contiguous or the sum of the fragments, is 40MHz. However, if a wider front end were available, all of the results would carry over proportionally since none of our techniques are bandwidth-limited.




Figure 10: Aggregate Throughput for 4 Links: Picasso provides nearly the same throughput over 4 links as 4 independent radios would, and significantly outperforms the throughput a single legacy radio could provide over the same set of spectrum fragmented.

6.1 Picasso: Overall Performance

Method: We statically place one Picasso node in the testbed and allow it to operate over four spectrum slices with separate WiFi PHY instances for each slice. For each of the four spectrum slices, we create a software radio client which is implementing a contiguous narrowband WiFi PHY on top of the spectrum slice. The four clients are placed at random locations in the testbed. The experiments are run in a 40MHz piece of spectrum, which is divided into four non-contiguous fragments randomly. The fragment sizes vary from 5-10MHz. We then measure the maximum possible PHY throughput by choosing the bitrate and coding combination for each of the spectrum fragments and measure the PHY throughputs achieved. To compare against the multiple radio setup, we replace the Picasso node with four physically separate nodes, each of them running a single narrowband contiguous PHY instance configured to work on one of the spectrum slices. Once again we measure PHY throughputs by picking the highest bitrate and coding combination. We also compute the throughput achieved by legacy radios which can only transmit on one spectrum fragment at a time, hence they always use the largest fragment. The experiment is repeated for each location and each setting of spectrum fragmentation. We plot the CDF of these throughputs in Fig. 10.

Analysis: Picasso achieves almost the same throughput with one radio as the physically independent setup with four physically separate radios, validating the claim that slicing allows decoupled operation of each slice. As expected, it outperforms the legacy radio approach by nearly a factor of 4 \times .

At the higher throughputs, there is a 10% dip in Picasso's performance relative to the multiple radio configuration. The reason is that at these throughputs, the transmission is likely using a very dense constellation and high coding rates (e.g. 64 QAM and 3/4 rate convolutional coding). When there are two adjacent slices very close to each other, separated only by the minimum default value of 500KHz, there is a small amount of signal power leakage from Picasso's self interference cancellation and programmable filtering. Hence if the Picasso node is transmitting while it is simultaneously receiving a densely coded stream, the small leakage can cause a slight performance drop, which we see in the above experiment. The loss is relatively small however, around 10%, and can be completely avoided by increasing the guard band between slices a bit higher than the default value of 500KHz. We leave that tradeoff as a tunable parameter which the network designer can choose depending on his needs. The next experiment sheds more light on this tradeoff.

6.1.1 Impact of Slice Spectrum Separation

Method: In this experiment we use Picasso nodes that are using 10MHz spectrum fragments each, but the separation between the fragments is an experimental parameter. We conduct two sets of experiments, one where a Picasso node is using two spectrum slices,




Figure 11: Picasso Throughput for Varying Guardbands: For guardband separations as small as 500 KHz, Picasso throughput does not suffer. At smaller separations, the Rx/Rx case maintains its performance but interference leakage causes a 20% degradation in performance for the Tx/Rx case.

one for TX and one for RX; and another where the Picasso node is either receiving or transmitting on both spectrum fragments. The first configuration tests whether Picasso's self interference cancellation technique leaks interference, and the second tests of Picasso's programmable filtering component leaks interference. We place two Picasso nodes in the testbed such that they can barely use the densest constellation and highest coding rate, since this is the configuration which suffers the most from interference leakage. The separation between the two spectrum slices is varied in both experiments and the PHY throughput achieved is measured as before. Fig. 11 plots the throughput achieved vs. the spectrum slice separation for Picasso along with the results for filtering and self interference cancellation individually.

Analysis: Picasso's throughput doesn't suffer much even for separations as small as 1MHz for either experiment. At 500KHz (the minimum default slice spectrum separation in Picasso), the simultaneous TX/RX case starts to see a small throughput loss of around 5 – 10% which progressively increases to 20% for separations of 100KHz. The two RX case suffers relatively minor throughput loss. The reason is that Picasso's self-interference cancellation, despite using passive analog components to ensure frequency flat behavior and minimize leakage, still leaks a small amount of interference to the close neighboring band. Perfectly flat, non-leaking analog components are hard to obtain off-the-shelf, we suspect performance can be further improved with better analog components.

6.1.2 Impact of Relative Signal Powers at Different Slices

Another practical concern is whether the throughput achieved in each slice depends on the relative powers of the transmissions and receptions on other slices. For example, suppose a Picasso radio is operating two slices. While one slice is transmitting at 23dBm, the other slice is receiving a signal which after attenuation has a power of -85dBm. In other words there is a gap of 108dB between the two signals. One might wonder if throughput is affected because of any dynamic range issues with Picasso. We evaluate this scenario experimentally below.

Method: Here we use a Picasso node that is utilizing two spectrum slices of width 10MHz each, and the two slices are connected to two separate client radios configured to work on the respective slices. The Picasso radio as well as the clients are configured to transmit with a power of 23dBm on each slice. We vary the locations of the two radios, and for each location measure the difference in signal power at the two slices. If the Picasso radio is transmitting on one slice and receiving on the other, then the difference is between the transmit power of 23dBm and the received signal power on the other slice. If the Picasso radio is receiving on both slices, the difference is between the received signal powers on both slices. We plot the




Figure 12: Impact of Relative SNRs of Different Slices: 14-Bit ADC provides sufficient dynamic range for all possible SNR gaps. 12-Bit ADC may suffice in certain scenarios if performance loss at high SNR gaps are acceptable.

throughput achieved by each slice when both slices are active simultaneously normalized by the throughput when only that slice is active vs. the difference in signal powers between the two slices at the Picasso radio in Fig. 12.

Analysis: As we can see, with Picasso’s current prototype, even very high signal power differences between slices have no impact on throughput. The reason is that Picasso’s self-interference cancellation coupled with a 14-bit ADC provides sufficient dynamic range so that throughput does not suffer in spite of large gaps between signal powers in concurrently active slices. To check the impact of the ADC resolution, we repeat the same experiment with 12-bit and 10-bit ADCs. Recall from Sec. 3 that lower resolution ADCs have lower dynamic range and hence are more likely to get saturated if the self-interference is too strong. As Fig. 12 shows, there is a slight loss of performance with 12-bit ADCs, and more with 10-bit ADCs. This is expected, since the lower dynamic range leaves less ADC resolution for the received signal and consequently introduces higher quantization noise when the difference in signal powers is very high. Nonetheless, the performance is still reasonable, and suggests even a 12-bit ADC might suffice for Picasso in most cases.

6.2 Picasso’s Self Interference Cancellation

We experimentally evaluate the performance of Picasso’s cancellation technique. From a single Picasso node, we transmit a 40 MHz OFDM signal and measure the amount of cancellation Picasso’s technique achieves. We plot the cancellation achieved vs. the bandwidth of the transmitted signal Fig. 13(A).

As we can see the cancellation achieved on average is around 45-50dB, which is sufficient for Picasso’s purposes. Further, note the almost complete lack of interference leakage as Fig. 13(B) shows. Even when the adjacent spectrum fragment is as close as 1MHz, there is almost no interference leakage from the transmitted signal. At the minimum slice guard band value of 500KHz there is a small leakage which shows up as the slight loss of throughput we discussed in the throughput experiment in Sec. 6.1.1. And note that even though the experiment was run for several minutes and across different days, we could use the same statically configured setting to consistently obtain the above cancellation. Prior techniques need to adapt on a per packet basis.

Throughput Impact of using Passive Components: In Section. 3.1 we discussed why the prior QHX-balun based self-interference cancellation techniques used for full duplex are not a good fit for Picasso. Here we evaluate experimentally the throughput impact if the QHX-balun based cancellation technique is used compared to Picasso’s technique.

Method: We run an experiment with two nodes which have their transmit and receive antennas spaced 15 centimeters apart and implementing QHX based balun cancellation. Each radio has two slices of 10MHz width, and the spectrum separation between the slices is varied. For each slice spectrum separation setting, we apply both cancel-




Figure 13: A) In Band Cancellation - Picasso achieves sufficient in-band cancellation ($> 45\text{dB}$) to prevent receiver saturation; **B) Out of Band Cancellation Leakage** - While QHX-Balun cancellation leaks significant interference into adjacent sidebands, Picasso cancellation does not.

lation techniques separately using the same power and gain settings. We once again measure PHY throughput for both techniques and plot it vs. the separation between the transmit and receive spectrum fragments in Fig. 14.

Analysis: Fig. 14 quantifies the throughput impact of the out-of-band interference leakage the prior QHX-balun cancellation technique introduces. Picasso’s technique, because it does not leak self-interference, achieves a throughput that is nearly $5\times$ better when the separation is around 500KHz.

6.3 Programmable Filter Engine

In this section, we discuss the efficacy of Picasso’s programmable filtering. We focus on two metrics that can be traded off against each other: out-of-band interference rejection and the complexity of the filtering implementation. Complexity is measured as the number of coefficients used in the linear digital filter. The logic utilization for the filter engine is a function of the filter length, coefficient precision, coefficient symmetry, and input data precision. While the filter engine can be configured to utilize any parameters, Picasso utilizes 16-bit symmetric coefficients with an input sample precision of 16 bits. Thus, the only remaining variable parameter is the filter length. The complexity of the filter must also be weighed against the configuration speed as the number of clock cycles required to reconfigure the filter engine is directly proportional to the number of coefficients used in the filter.

6.3.1 Complexity vs. Interference Rejection

To determine how different filter lengths impact performance, we conduct the following experiment using 2 Picasso nodes: node 1 transmits two 10 MHz OFDM signals in two spectrum slices which node 2 must decode. We vary the separation between the two spectrum fragments. Performance is measured in terms of the achievable PHY throughput between nodes 1 and 2. We plot the throughput achieved for different filter complexities and varying separation between spectrum fragments.

Analysis: From Fig. 15, we can see that as filter lengths increase, the throughput between nodes 1 and 2 increase. This is expected, a longer filter length does a much better job of rejecting out-of-band interference since its frequency response starts approximating the ideal square shape and roll-off is reduced. In Picasso’s implementation, we use 20th order filters, which consume modest resources and exhibit very good interference rejection even when the adjacent slice is as close as 500KHz.

6.3.2 Complexity vs. Reconfiguration Speed

The filter engine facilitates on the fly filter configuration by reloading new coefficient vectors. The limitation is that the filtering operation must be halted while the new coefficient values are loaded




Figure 14: Throughput Comparisons with different self interference cancellation schemes: Active cancellation introduces sideband leakage which reduces system throughput. Picasso’s passive cancellation does not, which allows for near optimal throughput.

and some of the internal data structures are subsequently initialized. While longer filters invariably take a longer time to load, they provide better performance as shown in Fig. 15. The worst-case load time for the longest filter is less than $4.5\mu\text{s}$, which is less than one WiFi slot length.

7. APPLICATIONS

Picasso’s slicing capability is a general primitive that can be used for a variety of applications. In this paper we show how Picasso can help diverse networks coexist gracefully in dense deployments and dynamic fragmented spectrum. More broadly though, Picasso can be applied to designing WiFi APs, multi-channel mesh networks, and even adaptive duplexers in phones. We discuss these applications briefly in Sec. 9.

We demonstrate how Picasso helps tackle coexistence by using the example deployment shown in Fig. 16. The deployment paints a typical home network: a WiFi AP aims to provide Internet access to a laptop, a tablet, an Xbox and a smartphone. Concurrently, there is a multiplayer game going on, and the Xbox wants to stream high definition individual gameplay video to the smartphone and tablet, and common gameplay to a TV. The smartphone and tablet also double up as wireless gameplay controllers. Further, with the emergence of P2P technologies such as WiFi-Direct [31] (a new WiFi standard for directly connecting two devices without incurring the overhead of an AP), except for the WiFi network, all the other connections will likely be connected via independent WiFi-Direct links.

As we can see, the scenario has a number of diverse networked applications that have to coexist while sharing a single radio on each device in the fragmented ISM band spectrum. Further, their requirements are diverse. The WiFi network wants to maximize throughput. The gameplay streaming group (Xbox to smartphone, tablet and TV) wants to maximize video quality. The gameplay controller group have relative low throughput requirements but low latency is a critical requirement for seamless gameplay.

Picasso provides a clean architecture for multiple such networked applications to coexist by allowing the applications to decouple their operation from each other using slicing. Specifically, Picasso enables us to slice each device’s radio into three slices. The first slice would be assigned to the WiFi network, which is now free to operate its own optimized PHY and MAC protocols to maximize throughput. The second slice would be for the gameplay video streaming application (Xbox, smartphone, tablet and TV), and if it so chooses, would operate independent PHY and MAC protocols optimized to maximize video quality. Finally, a third slice is for the gameplay controller application (Xbox, smartphone, tablet), with an independent PHY/MAC optimized for low latency. Thus Picasso *decouples the operation of the networked applications from each other*, and enables them to gracefully coexist yet independently design their protocols to optimize their respective objectives.

Compared Approaches: We compare the performance of Picasso with the following approaches:




Figure 15: Guardband separation between adjacent fragments vs. achievable throughput for varying filter lengths: Out of band rejection performance increases for higher order filters.

1. **Multiple Radios (MR) [1]:** This is the optimal scheme which uses separate dedicated radios for each network group. In other words each device would have multiple radios, one for each networked application.
2. **SWIFT [24]:** This is the state-of-the-art practical technique that can weave multiple spectrum fragments into a single logical connection via a modified OFDM PHY. Medium access is based on CSMA. However, standard CSMA is not optimized to provide QoS, since it uses the same backoff counter for all nodes. Hence we modify SWIFT’s CSMA to use 802.11e’s MAC protocol [6], a modification to the standard 802.11 MAC protocol for providing QoS. Specifically, WiFi network traffic is classified as best effort, and uses a congestion window range of 31 – 1023. Video streaming traffic has higher priority and when it is contending, it uses a congestion window range of 15 – 31, and finally 7 – 15 for the latency-sensitive traffic which has the highest priority. These are standard parameters from the 802.11e specification. Finally, note that a SWIFT radio has to either transmit or receive on all fragments at the same time, it cannot operate in full duplex mode.

Traffic Model: In order to emulate a realistic network, each device’s traffic pattern is modeled using empirical trace based models [10]. Specifically for the WiFi network group, we use a traffic model that models Internet use behavior on WiFi networks [10]. For the video streaming group, we use a traffic model that a video encoder would generate in order to guarantee 720P quality video. Note that contrary to perception, guaranteeing a consistent video quality does not imply a consistent data rate, due to variations in the video frames the actual data rate required to maintain a fixed video quality varies over time. Finally for the gameplay controller group, we assume that short packets are generated at random intervals with an average period of 50 milliseconds between them [21], which is slightly lower than the median human response times of 200 milliseconds.

Slice Allocation: A natural question is how are slices allocated for Picasso as well as the multiple radio approach. We use a simple greedy algorithm. For the latency-sensitive group, we allocate a small spectrum slice of 1MHz that is sufficient for the low rate latency-sensitive traffic. The next slice is allocated to the video streaming group, with sufficient spectrum to provide 720p streaming (which requires roughly 3Mbps per stream, and thus nearly 10Mbps). Left-over spectrum is allocated to the WiFi network group. Note that this implies that the WiFi network might be allocated multiple slices corresponding to whatever spectrum fragments are leftover. The network will run separate narrowband OFDM PHY and CSMA MAC instances for these slices. We assume that a higher layer mechanism such as multi-path TCP [27] (MPTCP) is used to aggregate the throughput.

An important point is that this is just one allocation; a different policy might use a different allocation (e.g. a parent may prioritize WiFi network throughput over the gameplay streaming!). However, the key point is that Picasso decouples policy from mechanism, and enables a wide range of policies.

Another issue is how slice allocation decisions are coordinated and



Figure 16: Coexistence in Dense Environments: Picasso enables devices operating in multiple networks to flexibly share fragmented spectrum and RF front ends, yet operate in a decoupled manner.

communicated to all the gadgets. To achieve this, we allocate a slice of 1MHz as a control channel, that all devices access to announce their QoS needs, and coordinate to figure out the best slice allocation. Note that such a simple decoupled control channel functionality is also because of the slicing capability, non-Picasso radios would have to co-ordinate in band using mechanisms such as RTS/CTS.

The SWIFT approach of course does not have any slice allocation issues. All devices operate over the entire available spectrum, including fragmented spectrum.

Optimized PHY/MACs for the individual slices: Each networked application uses an optimized PHY/MAC for their slices. The WiFi network group has two choices, it could run CSMA with OFDM where all the devices contend, or it could use FDMA like prior work has suggested [28] by taking advantage of Picasso’s slicing capability to allocate slices for each individual device in the group. We chose the latter approach since FDMA has much less overhead compared to CSMA. For the video streaming slice, we use a simple OFDM PHY and CSMA MAC, since all the transmissions originate from the Xbox and thus CSMA causes no contention overhead. For the gameplay controller group, we use CSMA since that provides the lowest latency access.

Experimental Method: We begin by evaluating the performance of Picasso along three axes - WiFi Network’s Aggregate Throughput, Video Streaming Group PSNR, and Gaming Group Average Latency - and compare it with the performance of the multiple radio and SWIFT approaches.

We randomly vary two independent parameters in our network. First, we vary the relative channel conditions between the devices. Because the varying channel conditions inevitably lead to a disparity in absolute performance, we normalize all three axes relative to the values achieved by the MR approach to demonstrate the relative performance gains of implementing Picasso over SWIFT independent of the channel conditions. Second, we vary the fragmentation pattern of the network. We assume that there are other narrowband interferers in the ISM band (e.g. Zigbee devices) which are transmitting on random frequencies, and hence the spectrum available to the Networks is fragmented. The Picasso and MR approaches use the slice allocations from the algorithm described above. SWIFT devices of course always use all the available spectrum.

Analysis: Fig. 17(A) plots the first tradeoff between latency and throughput. Picasso provides nearly $2\times$ higher throughput and $2\times$ lower latency than SWIFT. With SWIFT, there is a positive correlation between the two variables, the higher the throughput achieved by the WiFi network group, the higher the latency for the controller group. The reason behind this effect is that higher throughput implies latency-sensitive packets have to wait longer to obtain channel access and in turn incur higher latency using SWIFT-CSMA. On the other hand, Picasso has none of these issues, since the latency-sensitive traffic has a separate slice, they obtain a consistent low latency. Further, note that Picasso is able to obtain consistently higher throughput than SWIFT. The reason is that since SWIFT mediates access to the channels in time, there is significant overhead associated with idle slots, collisions, and backoffs. Because Picasso slices in frequency,

the only overhead it pays is in the lengths of the guardbands, which enables it to make more efficient use of the spectrum.

Fig. 17(B) plots the second tradeoff between latency and PSNR. Picasso provides nearly 10dB higher median PSNR than SWIFT for the video streams. Once again, there is a slight positive correlation here between PSNR and Latency for SWIFT with higher video qualities corresponding to higher latencies. The reason is the same as above: the video traffic occupies a significant fraction of the channel time to maintain high video quality, with the latency-sensitive traffic starved for channel access in spite of the aggressive backoff counters it employs. Picasso again is able to maintain slice isolation and thus can guarantee that the latency-sensitive traffic obtains consistently low latency, and that the video streams also obtain a predictable rate to maintain a high video quality. However, sometimes SWIFT achieves higher PSNRs than Picasso, which is expected since SWIFT could allocate a majority of the channel access for video transmissions depending on the channel access. That increases video quality, but comes at the expense of latency and throughput for the controller and WiFi groups. Further, note that SWIFT cannot guarantee a predictable high rate, the video quality varies significantly across time, which causes a loss of video fidelity.

Fig. 17(C) plots the third tradeoff between video quality and throughput. Picasso, as expected, maintains consistent video quality and normalized throughput. SWIFT cannot guarantee consistent throughput for the video streams, hence video encoding adapts to a lower quality. SWIFT does not achieve higher throughput either, since as discussed above, SWIFT CSMA adds significant overhead due to contention and limits the actual useful throughput that can be achieved. With Picasso, in the WiFi network group, we can assign slices for each individual device that is part of the group and operate FDMA for medium access, and thus eliminate contention overhead. Consequently, Picasso achieves both higher throughput and video quality.

8. RELATED WORK

Picasso bridges and builds upon research in two previously unrelated areas of wireless: in-band full duplex and transceiver design for fragmented spectrum.

Picasso distinguishes itself from prior single-channel full-duplex cancellation work [5, 14, 9, 15] by achieving self-interference cancellation using a single antenna and passive attenuation/delay components which do not leak interference into adjacent spectrum. Furthermore, unlike [14], Picasso does not require constant tuning and as such would be much easier to manufacture on a chip.

The RF front end and spectrum slicing abstraction that Picasso provides differentiate it from prior transceiver designs for fragmented spectrum because it provides strong decoupled operation between slices. Because prior works [24, 7, 22, 26] are unable to transmit while simultaneously receiving, the usage and performance on all spectrum fragments are inevitably tied together, greatly complicating the design of higher-layer coexistence protocols.

The virtue of Picasso is that because it enables higher layers to systematically exploit each fragment of spectrum independently, coexistence protocols can be greatly simplified. Thus, Picasso’s slicing capability provides a key building block for coexistence protocols [25, 2, 28], facilitating the allocation of available spectrum fragments amongst competing radios. Picasso in turn builds on top of prior spectrum sensing frameworks such as [13, 16, 28], which provide information about the RF environment and inform Picasso regarding availability of spectrum slices.

9. DISCUSSION & CONCLUSION

Picasso’s ability to simultaneously use arbitrary spectrum fragments with a single radio/antenna allows it to virtualize a single radio




Figure 17: A) Throughput vs. Latency: Picasso eliminates the contention latency incurred by the gaming system when competing with larger WiFi Network packets; **B) PSNR vs. Latency:** Picasso can provide latency and deadline guarantees whereas SWIFT cannot; **C) Throughput vs. PSNR:** Picasso achieves higher aggregate throughput and PSNR by eliminating the idle slots inherent in SWIFT CSMA due to contention and random access.

into separate independent slices operating on different spectrum fragments. This slicing abstraction provides a clean solution for many wireless applications, of which fragmented spectrum exploitation, heterogeneous network coexistence, and radio sharing stand out.

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. Not only would a system like Picasso enable handset manufacturers to save costs by replacing the disparate chipsets with a single integrated solution, it would also facilitate global roaming and liberate consumers to more easily switch network operators, potentially driving improved quality of service due to increased competition between service providers.

Implementing Picasso within a mobile chipset would be challenging, but not impossible. Isolation devices such as circulators [40], are readily available in miniaturized form, but provide slightly less isolation. Consequently, due to the limited power handling capability of miniaturized microwave components, Picasso would likely require additional self-interference cancellation on the order of 55dB or so in order to ensure that the LNA, mixers, and lower resolution ADCs operate within the specified power ranges. We believe that such performance is feasible and will be realized in the near term.

Acknowledgements: We would like to thank Dave Levin, Ratul Mahajan, Asaf Cidon, and the anonymous SIGCOMM reviewers for their insightful comments. This research was made possible through the financial support of the NSF POMI Expedition Grant, as well as NDSEG and NSF Graduate Research Fellowships.

10. REFERENCES

- [1] P. Bahl, A. Adya, J. Padhye, A. Walman, "Reconsidering Wireless Systems With Multiple Radios," In *ACM SIGCOMM CCR*, 2004
- [2] P. Bahl, et. al "KNOWS: Kognitiv Networking Over White Spaces," In *IEEE DySPAN*, 2007
- [3] M. Briggs, J. Martinez, D. Bare, "Power Measurements of OFDM Signals," In *IEEE Symposium on Electromagnetic Computability*, 2004.
- [4] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, P. Bahl, "A Case for Adapting Channel Width in Wireless Networks," In *ACM SIGCOMM*, 2008
- [5] J. Choi, M. Jain, K. Srinivasan, P. Levis, S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," In *ACM MOBICOM*, 2010
- [6] S. Choi, J. Prado, N. Shankar, S. Mangold, "IEEE 802.11 e contention-based channel access (EDCF) performance evaluation", In *IEEE ICC*, 2003.
- [7] E. Coffman, et. al, "Channel Fragmentation in Dynamic Spectrum Access Systems - a Theoretical Study," In *ACM SIGMETRICS*, 2010
- [8] FCC Table of Frequency Allocations, <http://transition.fcc.gov/oet/spectrum/table/fcc-table.pdf>, May 6, 2011.
- [9] E. Everett, M. Duarte, C. Dick, A. S. Bhargwal, "Empowering Full-Duplex Wireless Communication by Exploiting Directional Diversity," *IEEE Asilomar*, 2011.
- [10] A. Gember, A. Anand, A. Akella, "A Comparative Study of Handheld and Non-Handheld Traffic in Campus Wi-Fi Networks," *Passive and Active Measurement Conf*, 2011.
- [11] T. Gill, "RF performance of mobile terminals - a challenge for the industry", <http://www.cambridgewireless.co.uk/Presentation/Trevor%20Gill%20-%20Radio%20SIG%20310311.pdf>, 2010.
- [12] S. Hong, J. Mehlman, S. Katti, "Picasso: Full Duplex Signal Shaping to Exploit Fragmented Spectrum," In *ACM HotNets*, 2011.
- [13] S. Hong, S. Katti, "DOF: A Local Wireless Information Plane", In *ACM SIGCOMM*, 2011
- [14] 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.
- [15] M. Khojastepour, et. al, "The Case for Antenna Cancellation for Scalable Full Duplex Wireless Communications" In *ACM HOTNETS*, 2011.
- [16] K. Lakshminarayanan, S. Sapra, S. Seshan, and P. Steenkiste. "Rfdump: An Architecture For Monitoring the Wireless Ether." In *ACM CoNEXT*, 2009.
- [17] T. Lamprecht, et. al, "Passive alignment of optical elements in a printed circuit board", In *Electronic Components and Technology Conference*, 2006.
- [18] G. Matthaei, et. al, "Microwave filters, impedance-matching networks, and coupling structures", *Dedham, Mass: Artech House Books*, 1980.
- [19] T. Moon, W. Stirling, "Mathematical Methods and Algorithms for Signal Processing," *Prentice Hall*, 1999.
- [20] A. Oppenheim, R. Schaffer, "Discrete Time Signal Processing: 3rd Edition," *Prentice Hall Inc*, 2009
- [21] C. Palazzi, S. Ferretti, S. Cacciaguerra, M. Rocchetti, "A RIO-Like Technique for Interactivity Loss-Avoidance in Fast-Paced Multiplayer Online Games," In *ACM Computers in Entertainment Magazine*, 2005.
- [22] J. Poston, W. Horne, "Discontiguous OFDM Considerations for Dynamic Spectrum Access in Idle TV Channels," In *IEEE DySPAN*, 2005.
- [23] D. Pozar, "Microwave Engineering: 3rd Edition", Wiley, 2004.
- [24] H. Rahul, N. Kushman, D. Katabi, F. Edalat, and C. Sodini. "Learning to Share: Narrowband friendly wideband radios." In *ACM SIGCOMM*, 2008
- [25] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, Y. Zhang, "Fine Grained Channel Access in Wireless LAN," In *ACM SIGCOMM*, 2010.
- [26] K. Tan, et. al, "Spectrum Virtualization Layer," *MSR Tech Report*, <http://research.microsoft.com/apps/pubs/default.aspx?id=154410>, 2011.
- [27] D. Wischick, C. Raiciu, A. Greenhalgh, M. Handley, "Design, Implementation and Evaluation of Congestion Control for Multipath TCP," In *USENIX NSDI*, 2011
- [28] L. Yang, W. Hou, L. Cao, B. Zhao, H. Zheng, "Supporting Demanding Wireless Applications with Frequency-agile Radios", In *USENIX NSDI*, 2010
- [29] L. Yang, B. Zhao, H. Zheng, "The Spaces Between Us: Setting and Maintaining Boundaries in Wireless Spectrum Access", In *ACM MOBICOM*, 2010
- [30] Gizmodo, "iPhone 4 Antenna-Gate", <http://gizmodo.com/5846638/giz-explains-whats-so-smart-about-the-iphone-4ss-antenna>, 2011.
- [31] WiFi Direct Industry White Paper, <http://www.wi-fi.org/discover-and-learn/wi-fi-direct>, 2010.
- [32] NI PXIe-8133 User Manual, <http://www.ni.com/pdf/manuals/372870b.pdf>, 2011.
- [33] NI5781 Datasheet, <http://sine.ni.com/ds/app/doc/p/ids-212/lang/en>, 2011.
- [34] MECA Circulator, <http://www.e-meca.com/rf-circulator-isolator.php>
- [35] UHD Daughterboard Application Notes, http://files.ettus.com/uhd_docs/manual/html/dboards.html, 2011.
- [36] Avago MGA-82563 17dBm GaAs MMIC Amplifier www.avagotech.com/docs/AV02-1985EN.
- [37] Avago AD8347 Direct Conversion Quadrature Demodulator, www.analog.com/static/imported-files/data_sheets/AD8347.pdf.
- [38] UG193: XtremeDSP User Guide, http://www.xilinx.com/support/documentation/user_guides/ug193.pdf, June 1, 2010.
- [39] DS249: LogiCORE IP CORDIC v4.0 Data Sheet, http://www.xilinx.com/support/documentation/ip_documentation/cordic_ds249.pdf, March 1, 2011.
- [40] 3.2 mm Miniaturized Circulator Data Sheets, http://www.hitachimetals.com/product/rf_microwave/isolators/small/3_2mmSmallIsolators.html, 2012.