

Towards Battery-Free HD Video Streaming

*Saman Naderiparizi, Mehrdad Hesar, Vamsi Talla, Shyamnath Gollakota and Joshua R. Smith
University of Washington*

Abstract – Video streaming has traditionally been considered an extremely power-hungry operation. Existing approaches optimize the camera and communication modules individually to minimize their power consumption. However, designing a video streaming device requires power-consuming hardware components and computationally intensive video codec algorithms that interface the camera and the communication modules. For example, monochrome HD video streaming at 60 fps requires an ADC operating at a sampling rate of 55.3 MHz and a video codec that can handle uncompressed data being generated at 442 Mbps.

We present a novel architecture that enables HD video streaming from a low-power, wearable camera to a nearby mobile device. To achieve this, we present an “analog” video backscatter technique that feeds analog pixels from the photo-diodes directly to the backscatter hardware, thereby eliminating power-consuming hardware components, such as ADCs and codecs. To evaluate our design, we simulate an ASIC, which achieves 60 fps 720p and 1080p HD video streaming for 321 μW and 806 μW , respectively. This translates to 1000x to 10,000x lower power than it used for existing digital video streaming approaches. Our empirical results also show that we can harvest sufficient energy to enable battery-free 30 fps 1080p video streaming at up to 8 feet. Finally, we design and implement a proof-of-concept prototype with off-the-shelf hardware components that successfully backscatters 720p HD video at 10 fps up to 16 feet.

1 Introduction

There has been recent interest in wearable cameras like Snap Spectacles [16] for applications ranging from life-casting, video blogging and live streaming concerts, political events and even surgeries [18]. Unlike smartphones, these wearable cameras have a spectacle form-factor and hence must be both ultra-lightweight and cause no overheating during continuous operation. This has resulted




Figure 1: Target application. Our ultra-low-power HD streaming architecture targets wearable cameras. We achieve this by performing analog video backscatter from the wearable camera to a nearby mobile device (e.g., smartphone).

in a trade-off between the usability of the device and its streaming abilities since higher resolution video streaming requires a bigger (and heavier) battery as well as power-consuming communication and processing units. For example, Snap Spectacle, while lightweight and usable, cannot stream live video [16] and can record only up to one hundred 10-second videos (effectively less than 20 minutes) [16, 13] on a single charge.

In this paper, we ask the following question: Can we design a low-power camera that can perform HD video streaming to a nearby mobile device such as a smartphone? A positive answer would enable a wearable camera that is lightweight, streams high quality video, and is safe and comfortable to wear. Specifically, reducing power consumption would reduce battery size, which in turn addresses the key challenges of weight, battery life and overheating. Finally, since users typically carry mobile devices like smartphones that are comparatively not as weight and power constrained, they can relay the video from camera to the cloud infrastructure.

To understand this challenge, let us look at the different components of a video-streaming device: image

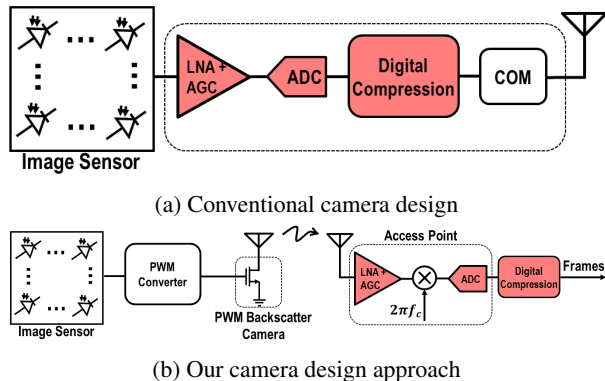


Figure 2: The amplifier, AGC, ADC and compression module consume orders of magnitude more power than is available on a low-power device. In our design, these power hungry modules have been delegated to the reader, eliminating their power consumption overhead from the wireless camera.

sensor, video compression and communication. Image sensors have an optical lens, an array of photo-diodes connected to amplifiers, and an ADC to translate analog pixels into digital values. A video codec then performs compression in the digital domain to produce compressed video, which is then transmitted on the wireless medium. Existing approaches optimize the camera and communication modules individually to minimize power consumption. However, designing an efficient video streaming device requires power-consuming hardware components and video codec algorithms that interface to the camera and communication modules. Specifically, optical lens and photo-diode arrays can be designed to consume as little as $1.2 \mu\text{W}$ [22]. Similarly, recent work on backscatter can significantly lower the power consumption of communication to a few microwatts [26, 23] using custom ICs. Interfacing camera hardware with backscatter, however, requires ADCs and video codecs that significantly add to power consumption.

Table 1 shows the sampling rate and data rate requirements for the ADC and video codec, respectively. HD video streaming requires an ADC operating at a high sampling rate of more than at least 10 MHz. While the analog community has reduced ADC power consumption at much lower sampling rates [46, 20], state-of-the-art ADCs in the research community consume at least a few milliwatts at the high sampling rates [34]. Additionally, the high data rate requires the oscillator and video codec to run at high clock frequencies, which proportionally increases power consumption. Specifically, video codecs at these data rates consume 100s of milliwatts to a few watts of power [2].

We present a novel architecture that enables video streaming on low-power devices. Instead of independently optimizing the imaging and the communication modules, we jointly design these components to significantly reduce system power consumption. Our architec-



Figure 3: Sample HD video frame streamed with our analog video backscatter design. Our prototype was placed 4 feet from the reader.

ture, shown in Fig. 2b, takes its inspiration from the Great Seal Bug [5], which uses changes in a flexible metallic membrane to reflect sound waves. Building on this idea, we create the first “analog” video backscatter system. At a high level, we feed analog pixels from the photo-diodes directly to the backscatter antenna; we thus eliminate power-hungry ADCs, amplifiers, AGCs and codecs.

Our intuition for an analog video backscatter approach is to shift most of the power-hungry, analog-to-digital conversion operations to the reader. Because analog signals are more susceptible to noise than digital ones, we split the ADC conversion process into two phases, one performed at the video camera, and one accomplished at the reader. At the video camera, we convert analog pixel voltage into a pulse that is discrete in amplitude but continuous in time. This signal is sent via backscatter from the camera to the reader. Avoiding the amplitude representation in the wireless link provides better noise immunity. The reader measures the continuous length pulse it receives to produce a binary digital value. Philosophically, the continuous-time, discrete amplitude pulse representation used in the backscatter link resembles that used in an extremely power-efficient biological nervous system, which encodes information in spikes that do not vary in amplitude but are continuous in time [45].

Specifically, our design synthesizes three key techniques. First, we show how to interface pixels directly to the backscatter hardware without using ADCs. To do, this we transform analog pixel values into different pulse widths using a passive ramp circuit and map these pulses into pixels at the reader. Second, we achieve intra-frame compression by leveraging the redundancy inherent in typical images. Intuitively, the signal bandwidth is proportional to the rate of change across adjacent pixels; since videos tend to be redundant, the bandwidth of the analog signal is inversely proportional to the redundancy in the frame. Thus, by transmitting pixels consecutively, we can implicitly perform compression and significantly reduce wireless bandwidth. Finally, to achieve inter-frame compression, we design a distributed algorithm that reduces the data the camera transmits while delegating most

Table 1: Raw digital video sampling and bitrate requirement

Video Quality	Frame Rate: 60 fps		Frame Rate: 30 fps		Frame Rate: 10 fps	
	Sampling Rate (MHz)	Data Rate (Mbps)	Sampling Rate (MHz)	Data Rate (Mbps)	Sampling Rate (MHz)	Data Rate (Mbps)
1080p (1920x1080)	124.4	995.3	62.2	497.7	20.7	165.9
720p (1280x720)	55.3	442.4	18.4	221.2	9.2	73.7
480p (640x480)	18.4	147.4	9.2	73.7	3.1	24.58
360p (480x360)	10.4	82.9	5.2	41.5	1.7	13.8

inter-frame compression functionality to the reader. At a high level, the camera performs averaging over blocks of nearby pixels *in the analog domain* and transmits these averaged values using our backscatter hardware. The reader compares these averages with those from the previous frame and requests only the blocks that have seen a significant change in the average pixel value, thus reducing transmission between subsequent video frames.

We implement a proof-of-concept prototype of our analog backscatter design on an ultra-low-power FPGA platform and a custom implementation of the backscatter module. Because no HD camera currently provides access to its raw pixel voltages, we connect the output of a DAC converter to our backscatter hardware to emulate the analog camera and stream raw analog video voltages to our backscatter prototype. Fig. 3 shows a sample frame from an HD video streamed with our backscatter camera. More specifically, our findings are:

- We stream 720p HD video at 10 frames per second up to 16 feet from the reader. The Effective Number of Bits (ENOB) received for each pixel at distances below six feet exceeds 7 bits. For all practical purposes, these results are identical to the quality of the source HD video.
- Our inter and intra-frame compression algorithms reduces total bandwidth requirements by up to two orders of magnitude compared to raw video. For example, for 720p HD video at 10 fps, our design uses a wireless bandwidth of only 0.98 MHz and 2.8 MHz in an average-case and worst-case scenario video, respectively.

We design and simulate an ASIC implementation for our system, taking into account power consumption for the pixel array. Our results show that power consumption for video streaming at 720p HD is 321 μW and 252 μW for 60 and 30 fps, respectively. Power consumption at 1080p full-HD is 806 μW and 561 μW at 60 and 30 fps, respectively. We run experiments with RF power harvesting from the reader which shows that we can support 1080p full-HD video streaming at 30 fps up to distances of 8 feet from the reader. *Our results demonstrate that we can eliminate batteries and achieve HD video streaming on battery-free devices using our analog video backscatter approach.*

2 A Case for Our Architecture

Fig. 2a shows the architecture of a traditional wireless camera. Photodiodes' output is first amplified by a low noise amplifier (LNA) with automatic gain control (AGC). The AGC adjusts the amplifier gain to ensure that the output falls within the dynamic range of the analog to digital converter (ADC). Next, the ADC converts the analog voltage into discrete digital values. The video codec then compresses these digital values, which are transmitted on the wireless communication link.

Unfortunately, this architecture cannot be translated to an ultra-low-power device. Although camera sensors consisting of an array of photodiodes have been shown to operate on as low as 1.2 μW of power [22] at 128×128 resolution, amplifiers, AGC, ADC and the compression block require orders of magnitude higher power. Furthermore, power consumption increases as we scale the camera's resolution and/or frame rate. A low resolution 144p video recorded at 10 fps requires an ADC sampling at 368 KSPS to generate uncompressed data at 2.95 Mbps. With the latest advancements in ADCs [33] and backscatter communication [26], uncompressed data can be digitally recorded using low-power ADCs and transmitted using digital backscatter while consuming below 100 μW of power, which is within the power budget of energy harvesting platforms. However, as we scale the resolution to HD quality and higher, such as 1080p and 1440p, the ADC sampling rate increases to 10-100 MHz, and uncompressed data is generated at 100 Mbps to Gbps. ADCs operating at such high sampling rates consume at least a few mW [34]. Further, a compression block that operates in real time on 100 Mbps to 1 Gbps of uncompressed video consumes up to one Watt of power [2]. This power budget exceeds by orders of magnitude that available on harvesting platforms. Thus, while existing architectures might operate on harvested power for low-resolution video by leveraging recent advancements in low-power ADCs and digital backscatter, as we scale the resolution and/or the frame rate of the wireless camera, these architectures' use of ADCs and compression block drives up power consumption to levels beyond the realm of battery-free devices.




Figure 4: Architecture of the PWM converter.

3 System Design

In the rest of this section, we first describe our analog video transmission scheme followed by the intra-frame compression and finally the interactive distributed inter-frame compression technique.

3.1 Analog Video Backscatter

Problem. At a high level, our design eliminates the ADC on the camera to significantly reduce power consumption. This limits us to working with analog voltage output from the photodiodes. A *naive* approach would leverage existing analog backscatter technique used in wireless microphones [49] to implement an analog backscatter camera. In an analog backscatter system, sensor output directly controls the gate of a field-effect transistor (FET) connected to an antenna. As the output voltage of the sensor varies, it changes the impedance of the FET which amplitude modulates the RF signal backscattered by the antenna. The reader decodes sensor information by demodulating the amplitude-modulated backscattered RF signal. However, this approach cannot be translated to a camera sensor. Photodiode output has a very limited dynamic range (less than 100 mV under indoor lighting conditions). These small voltage changes map to a very small subset of radar cross-sections at the antenna [19]. As a result, the antenna backscatters a very weak signal. Since wireless channel and receivers add noise, this approach results in low SNR signal at the receiver, which limits the system to both poor signal quality and limited range. One can potentially surmount this constraint by introducing a power-hungry amplifier and AGC, but this would negate the power-saving potential of analog backscatter.

Our solution. Instead of using linear amplifiers to amplify the analog video stream, we use comparators to pulse width modulate [24] camera sensor’s analog output and convert into digital domain as shown in Fig. 4. At a high level, pulse width modulation (PWM) converts analog input voltage to different pulse widths. Specifically, the output of PWM modulation is a square wave, where the duty cycle of the output square wave is proportional to the analog voltage of the input signal.

As we show next, a PWM converter can be implemented with passive RC components and a comparator, thereby consuming very low power. Fig. 4 shows the

PWM converter architecture. The input is a square wave operating at frequency f , which is determined by the camera’s frame rate and resolution. First, the square wave is low-pass filtered by an RC network to generate a triangular waveform, as shown in the figure. This waveform is then compared to the pixel value using a comparator. The comparator outputs a zero when the triangular signal is less than the pixel value and a one otherwise. Thus, the pulse width generated changes with the pixel value: lower pixel values have a larger pulse duration, while higher pixel values have a smaller pulse duration. We choose the minimum and maximum voltages for the triangular signal to ensure that the camera’s pixel output is always within this limit which also maximizes the gain (analog voltage to digital pulse width) of the PWM modulation operation.

In the frequency domain, in addition to the desired baseband signal, PWM modulation also generates additional harmonics which contain redundant copies of the input analog video signal [24, 47] and can result in electromagnetic interference (EMI) from the antenna at higher frequencies. However, since the typical power re-radiated by the antenna in a backscatter system is significantly lower than the power transmitted by an active radio, the resulting EMI from the harmonics of the baseband PWM signal in a backscatter system are also negligible (less than <-20 dBm in all scenarios). There are different ways one can suppress the minuscule EMI generated by harmonics of baseband PWM signal. A PWM conversion with full dynamic range generates harmonics which are about 4 dB lower than the desired baseband signal. However, if we increase the dynamic range by a factor of 2, the PWM harmonics reduce to 10 dB lower than the baseband. So, we can tradeoff harmonic suppression for slightly reduced performance. This is especially useful since EMI due to harmonics is a potential issue when the camera is close to the reader where it can re-radiate higher power EMI. However, at close distances we can safely tradeoff a bit of performance for EMI rejection. At farther distances, since the backscattered power is low, EMI is not an issue and we don’t need to tradeoff any performance. Finally, in addition to performance tradeoff, we can introduce additional low power pass filter after the sub-carrier modulation operation in the digital domain to suppress harmonics and transmit filtered signal using backscatter. However, the implementation of EMI suppression is beyond the scope of this work.

The final component in our hardware design is sub-carrier modulation, which addresses the problem of self-interference. Specifically, in addition to receiving the backscatter signal, the receiver also receives a strong interference from the transmitter. Since the sensor data is centered at the carrier frequency, the receiver cannot decode the backscattered data in the presence of a strong, in-band interferer. Existing backscatter communication

Table 2: Intra-frame compression for average/worst-case scenarios across 100 videos.

Video Resolution	Frame Rate: 60 fps		Frame Rate: 30 fps		Frame Rate: 10 fps	
	Raw Data Rate(Mbps)	Analog BW (MHz)	Raw Data Rate(Mbps)	Analog BW (MHz)	Raw Data Rate(Mbps)	Analog BW (MHz)
1080p	995.3	10.8/30.5	497.6	5.6/18.7	165.8	1.9/9.3
720p	442.3	6.3/15.6	221.1	3.1/8	73.7	0.98/2.8
480p	147.4	3/6.7	73.7	1.6/3.3	24.5	0.53/1.4
360p	82.9	1.9/3.9	41.4	0.94/2.3	13.8	0.32/0.78

systems, such as RFID and Passive Wi-Fi, address this problem using sub-carrier modulation. These systems use a sub-carrier to shift the signal from the carrier to a frequency offset Δf from the carrier frequency. The receiver can then decode the backscattered signal by filtering out of band interference. Another consideration in choosing sub-carrier frequency is to avoid aliasing; sub-carrier frequency should not be smaller than the effective bandwidth of the analog signal.

Our PWM-based design integrates subcarrier modulation. We implement this modulation with a simple XOR gate. The sub-carrier can be approximated by a square wave operating at Δf frequency. We input sub-carrier and PWM output to an XOR gate to up-convert the PWM signal to a frequency offset Δf . Sub-carrier modulation addresses the problem of self-interference at the reader, and as a result, the PWM backscattering wireless camera can now operate at a high SNR and achieve broad operating ranges. We show in §5.1 that our PWM backscatter wireless camera can operate at up to 16 feet for a 10 fps monochrome video signal in HD resolution. We also show in §6 that our camera system can work at up to 150 feet for a 50 KHz video signal in 112×112 resolution, over a $4\times$ improvement relative to the 3 kHz bandwidth analog backscatter wireless microphone [51]. Subcarrier modulation with a square wave generates high order harmonics copies of the baseband PWM signal but these higher order components can be eliminated via harmonic cancellation techniques introduced in prior work [48].

3.2 Intra-Frame Compression

There is significant redundancy in the pixel values of each frame of uncompressed video. Redundancy occurs because natural video frames usually include objects larger than a single pixel, which means that the colors of nearby pixels are highly correlated. At the boundaries of objects (edges), larger pixel variations can occur; conversely, in the interior of an object, the amount of pixel variation is considerably less than the theoretical maximum. The net result of pixel correlations is a reduction in the information needed to represent natural images to levels far below the worst-case maximum. Traditional digital systems use a video codec to reduce pixel value redundancy. Specifi-

cally, the ADC first digitizes the camera’s raw output at the Nyquist rate determined by the resolution and frame rate. The raw digital data stream is then fed to the video codec, which implements compression algorithms.

In the absence of the ADC, our wireless camera transmits analog video directly. However, we note that the bandwidth of any analog signal is a function of the new information contained in the signal. Inspired by analog TV broadcast, which transmits pixel information in a raster scan (left to right), we introduce and implement a zig-zag pixel scanning technique: pixels are scanned from left to right in odd rows and from right to left in even rows. The intuition here is that neighboring pixels have less variation, and the resulting signal would thus occupy less bandwidth.

We evaluate how well zig-zag transmission performs in terms of bandwidth reduction of the PWM modulated analog signal. We download one hundred 60 fps Full-HD (1080p) videos from [17] to use as our baseline. These videos range from slow to fast moving and contain movie action stunts, running animals, racing, etc. To create video baselines at different resolutions and frame rates, we resize and subsample these Full-HD video streams. For each of the resulting video resolutions and frame rates, we create an analog video by zig-zagging the pixels as described above. We then apply low-pass filters with different bandwidths on this analog signal and report the minimum bandwidth of the baseband component of the PWM modulated analog video stream at which the PSNR of the resulting signal exceeds 30 dB. The bandwidth requirement reported in Table 2 shows the average/worst-case scenario, i.e., the bandwidth that ensures the recovery of average/worst-case videos with minimal quality degradation. We outline the uncompressed digital data rate for reference. Compared to the digital uncompressed data sampled at the Nyquist rate, the analog signal occupies 17.7–32.6x less bandwidth for the worst-case and 43–92x for the average-case, demonstrating our intra-frame compression technique’s capability. Fig. 5a shows the CDF of effective bandwidth for our 30 fps 720p video dataset and demonstrates that an average-case 30 fps 720p video achieves up to a $71\times$ improvement compared to raw digital video transmission.




Figure 5: CDF of effective bandwidth for zig-zag pixel scanning (a), and improvement it provides over raster scanning (b).

Finally, we note that, compared to raster, a zig-zag scan faces less discontinuity in pixel values: instead of jumping from the last pixel in a row to the first pixel in the next row, it continues at the same column in the next row, thereby taking greater advantage of the video’s inherent redundancy. This further lowers bandwidth utilization in the wireless medium. As an example, on average, zig-zag pixel scanning occupies $\sim 120\text{KHz}$ and $\sim 60\text{KHz}$ less bandwidth than raster scanning in a 60 fps and 30 fps 720p video stream, respectively. Fig. 5b shows the CDF of bandwidth improvement for zig-zag scanning over raster in our one hundred 30 fps 720p videos dataset. The plot makes clear that use of zig-zag pixel scanning provides greater bandwidth efficiency than its raster counterpart.

3.3 Distributed Inter-Frame Compression

In addition to the redundancy within a single frame, raw video output also has significant redundancy between consecutive frames. Existing digital architectures use a video codec to compress the raw video prior to wireless transmission to remove inter-frame redundancy and reduce power and bandwidth. Our approach to address this challenge is to once again leverage the reader. Like backscatter communication, where we move power-hungry components (such as ADCs) to the reader, we also move compression functionality to the reader. Specifically, our design distributes the compression algorithm between the reader and the camera. We delegate power-hungry computation to the reader and leverage the fact that our camera system can transmit *super-pixels*. A super-pixel value is the average of a set of adjacent pixel values. A camera frame consists of an $N \times N$ array of pixels, which can be divided into smaller sections of $n \times n$ pixels. A super-pixel corresponding to each section is the average of all the pixels in the $n \times n$ frame section. The camera sensor, a photodiode, outputs a current proportional to the intensity of light. The camera uses a buffer stage at the output to convert the current into an output voltage. To compute the super-pixel, the camera first combines the current from the set of pixels and then converts the combined current into a voltage output, all in the analog domain. We note that the averaging of close-by pixels is supported by commercial cameras including the CentEye Stonyman camera [4].




Figure 6: Distributed compression. The sequence of frames and pixels transmitted by the camera to the reader.

Instead of transmitting the entire $N \times N$ pixel frame, the camera transmits a lower resolution frame consisting of $n \times n$ sized super-pixels (the average value of the pixels in the $n \times n$ block), called the low-resolution frame or *L* frame. Doing so reduces the data transmitted by the camera by a factor of $\frac{N^2}{n^2}$. The reader performs computation on *L* frames and implements a change-driven compression technique. At a high level, the reader compares the incoming *L* frame with the previous *L* frame. If a super-pixel value differs by more than a predetermined threshold between frames, then the super-pixel has sufficiently changed, and the reader asks the camera to transmit all pixels corresponding to it. If the difference does not cross the threshold, then the reader uses the pixel values corresponding to the previous reconstructed frame to synthesize the new frame without requesting new pixel values. We call the frame that contains the pixel values corresponding to the sufficiently changed super-pixels, the super-pixel frame or *S* frame. In addition to transmitting the *S* and *L* frames, the camera periodically transmits uncompressed frames (*I*) to correct for potential artifacts/errors that the compression process may have accumulated.

In streaming camera applications, the communication overhead of the reader requesting specific pixel values is minimal and is implemented using a downlink similar to prior 100-500 kbps designs [26, 44], where the receiver at the camera uses a simple envelope detector to decode the amplitude modulated signal from the reader. We note that prior designs can achieve Mbps downlink transmissions [42] as well as full-duplex backscatter [31], which can be used to remove the downlink as a bottleneck.

Fig. 6 shows the sequence of frames and pixels transmitted by our camera. Between two *I* frames, the camera transmits *M* low-resolution *L* frames and *M* super-pixel *S* frames which contain pixel values corresponding to super-pixels whose value differences exceed the threshold between consecutive frames. The number of *L* and *S* frames (*M*) transmitted between consecutive *I* frames trades off between the overhead associated with transmission of full resolution frames and the artifacts and errors the compression algorithm introduced. In our implementation of 10 fps HD video streaming, we transmit an *I* frame after a transmission of every 80 *L* and *S* frames.




Figure 7: **HD wireless camera prototype.** A laptop emulating an analog output camera feeds raw pixel data into our backscatter prototype. Then the PWM encoded pixels are backscattered to the reader.

4 Implementation

We built our wireless camera using off-the-shelf components on a custom-designed Printed Circuit Board (PCB). We use the COTS prototype to evaluate the performance of the wireless camera in various deployments. We then present the application-specific integrated circuit (ASIC) design of the wireless camera that we used to quantify the power consumption for a range of video resolutions and frame rates.

COTS implementation. Our wireless camera design eliminates the power-hungry ADCs and video codecs and consists only of the image sensor, PWM converter, a digital block for camera control and sub-carrier modulation, a backscatter switch and an antenna. We built two hardware prototypes, one for the high definition (HD) and another for the low-resolution version of the camera.

We built the low-resolution wireless camera using the 112×112 grayscale random pixel access camera from CentEye [4], which provides readout access to individual analog pixels. We implement the digital control block on a low-power Igloo Nano FPGA by Microsemi [7]. The analog output of the image sensor is fed to the PWM converter built using passive RC components and a Maxim NCX2200 comparator [10]. We set $R_1 = 83K\Omega$, $R_2 = 213K\Omega$ and $C = 78pF$ in our PWM converter design to support video frame rates of up to 13 fps. PWM converter’s output acts as input to the FPGA. The FPGA performs sub-carrier modulation at 1.024 MHz using an XOR gate and outputs the sub-carrier modulated PWM signal to the Analog Devices ADG919 switch which switches a 2 dBi dipole antenna between open and short impedance states. The FPGA injects frame and line synchronization patterns into the frames data before backscattering. We use Barker codes [3] of length 11 and 13 for our frame and line synchronization patterns, respectively. Barker codes have a high-autocorrelation property that helps the reader more efficiently detect the beginning of the frame in the presence of noise.

We use Verilog to implement the digital state machine for camera control and sub-carrier modulation. Verilog design can be easily translated into ASIC using industry standard EDA tools. We can further reduce our system’s power consumption by using the distributed compression technique. As described in §6, a camera deployed in a normal lab space can achieve an additional compression ratio




Figure 8: **Signal processing steps at the reader.** Recovering video from a PWM backscatter signal.

of around $30\times$, which proportionately reduces wireless transmissions.

To develop an HD-resolution wireless camera, we need access to the raw analog pixel outputs of an HD camera. Currently, no camera on the market provides that access. To circumvent this constraint, we download from YouTube HD-resolution sample videos lasting 1 minute each. We output the recorded digital images using a USB interface to an analog converter (DAC) to simulate voltage levels corresponding to an HD quality image sensor operating at 10 fps. Given the USB speeds, we achieve the maximum frame rate of 10 fps. We feed the voltage output to our PWM converter. Fig. 7 shows the high-level block diagram of this implementation. For the high-resolution version of the wireless camera, we set $R_1 = 10K\Omega$, $R_2 = 100K\Omega$ and $C = 10pF$ and use an LMV7219 comparator by Texas Instruments [8] in our PWM converter. The digital block and other system components were exactly the same as for the low-resolution wireless camera described above except that sub-carrier frequency is set to ~ 10 MHz here to avoid aliasing.

ASIC Design. As noted, our design eliminates the power-hungry LNA, AGC and ADC at the wireless camera by delegating them to the reader to reduce wireless camera power consumption by orders of magnitude. However, since commercially available Stonyman cameras (like CentEye) and components (such as FPGA) are designed for flexibility and ease of prototyping and are not optimized for power, our COTS implementation cannot achieve the full power savings from our design. Therefore, we analyze the power consumption of an application-specific integrated circuit (ASIC) implementation of our design for a range of video resolutions and frame rates. An ASIC can integrate the image sensor, PWM converter, digital core, oscillator and backscatter modulator onto a small silicon chip. We implement our design in a TSMC 65 nm LP CMOS process.

We use Design Compiler by Synopsis [14] to synthesize transistor level from the behavioral model of our digital core, which is written in Verilog. We custom-design the PWM converter, oscillator and backscatter modulator described in §3 in Cadence software and use industry standard simulation tools to estimate power. To support higher resolution and higher frame rate video, we simply increase the operating frequency of the oscillator, PWM converter and digital core. As an example, 360p at 60 fps requires a 10.4 MHz input clock, which consumes

Table 3: Power consumption

	Frame Rate: 60 fps	Frame Rate: 30 fps	Frame Rate: 10 fps
Video Quality	Power (μW)	Power (μW)	Power (μW)
1080p (1920x1080)	806.50	560.63	167.77
720p (1280x720)	320.94	252.10	78.31
480p (640x480)	126.88	106.78	36.71
360p (480x360)	75.63	65.68	25.11

a total of $42.4 \mu\text{W}$ in the digital core, PWM converter and backscatter switch; a 1080p video at 60 fps requires an ~ 124.4 MHz input clock, which consumes $408 \mu\text{W}$ in the digital core, PWM converter and backscatter switch. To eliminate aliasing in all cases, we choose a sub-carrier frequency equal to the input clock of each scenario. Note that sub-carrier frequency cannot be lower than the effective bandwidth of the signal reported in Table 2.

We use existing designs to estimate the power consumption of the image sensor for different video resolutions. State-of-the-art image sensors consume $3.2\text{pW}/(\text{frame} \times \text{pixel})$ [52] for the pixels-only image sensor, which results in $33.2 \mu\text{W}$ for 360p resolution; this increases to $398 \mu\text{W}$ for 1080p resolution video at 60fps. Table 3 shows the power consumption of the ASIC version of our wireless camera for different video resolution and frame rates. Note that these results show power consumption before inter-frame compression distributed across the reader and camera, which could further reduce wireless bandwidth and power consumption.

Reader Implementation. We implement the reader on the X-300 USRP software-defined radio platform by Ettus Research [15]. The reader uses a bi-static radar configuration with two 6 dBi circularly polarized antennas [1]. Its transmit antenna is connected to a UBX-160 daughterboard, which transmits a single-tone signal. USRP output power is set to 30 dBm using the RF5110 RF power amplifier [11]. The receive antenna is connected to another UBX-160 daughter board configured as a receiver, which down-converts the PWM modulated backscattered RF signal to baseband and samples it at 10 Msps. The digital samples are transmitted to the PC via Ethernet.

Fig. 8 shows block diagram of the signal processing steps required to recover the transmitted video. For example, for low-resolution video, the received data is centered at an offset frequency of the 1.024 MHz; therefore, we first filter the received data using a 500 order bandpass filter centered at 1.024 MHz. Then, we down-convert the signal to baseband using a quadrature down-conversion mixer. Next, we correlate the received data with 13 and 11 bit Barker codes to determine the frame sync and line sync. After locating frame and line sync pulses, we extract the time periods corresponding to a row of PWM-modulated pixel values and low-pass filter the signal with a 500 or-

der filter to remove out of band noise. We divide the row in evenly spaced time intervals that corresponded to the number of pixels in a single row of the image sensor. We recover the pixel value by calculating the average voltage of the signal, which corresponds to the duty cycle of the PWM-modulated signal. We sequentially arrange recovered pixel values into rows and columns to create video frames. Finally, we run a histogram equalization algorithm on the video to adjust frames intensity and enhance output video contrast [21].

5 Evaluation

We now evaluate various aspects of our wireless camera system. We start by characterizing the received video quality from the camera as a function of its distance to the reader. Next, we evaluate the performance of our wireless camera while it is worn by a user under different head positions and orientations. We then evaluate the power available by harvesting energy from RF signals transmitted by the reader to demonstrate the feasibility of a battery-free wireless camera. Finally, we evaluate the distributed interactive compression algorithm under two different scenarios.

5.1 Operational Range

We deploy our high-definition wireless camera device in a regular lab space. We use the USRP-based reader implementation (§4), which we set to transmit 23 dBm into a 6 dBi patch antenna. This is well below the 30 dBm maximum transmit power permitted by FCC in the 900 MHz ISM band. We vary the distance between the reader and the wireless camera prototype from 4 to 16 feet and configure the camera to repeatedly stream a 15-second video using PWM backscatter communication. The 720p resolution video is streamed at 10 fps, and the pixel values are encoded as 8-bit values in monochrome format. We record the wirelessly received video at the reader and measure the Signal to Noise Ratio (SNR). From that, we calculate the Effective Number of Bits (ENOB) [27] at the receiver.

We plot the ENOB of the video digitized by the reader as a function of distance between the reader and the wireless camera in Fig. 9. The plot shows that up to 6 feet




Figure 9: **ENOB of the Received Video.** ENOB of the received video versus distance of our wireless camera prototype from the reader.

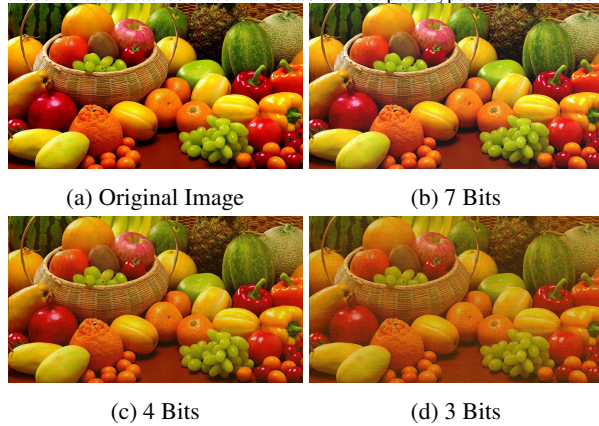


Figure 10: **Effective Number of Bits (ENOB) versus video quality.** The frame corresponding to an ENOB of 3 (d) shows video quality degradation.

from the reader, we achieve an ENOB greater than 7, which indicates negligible degradation in quality of the video, streamed using PWM backscatter. As the distance between the reader and the camera increase the SNR degrades, indicates a decrease in ENOB. Beyond the distance of 16 feet, we stop reliably receiving video frames. A separation of 16 feet between the reader and wearable camera is more than sufficient for wearable cameras, typically located a few feet away from readers such as smartphones. For reference, Fig. 10 shows frames from a 720p HD color video backscattered with our prototype at different distances from the reader; this resulted in different ENOB values.

We conclude the following: our analog video backscatter approach is ideal for a wearable camera scenario since the video is streamed to a nearby mobile device such as a smartphone. In this scenario, the SNR is high; hence, quality degradation due to an analog approach is not severe. Further, we gain significant power reduction benefits.




Figure 11: **ENOB of received video under different head motions.**

5.2 Effect of Head Pose and Motion

We ask a participant to wear the antenna of our wireless camera on his head and perform different head movements and poses while standing around five feet from a reader with fixed location on a table. The poses included standing still, moving head left and right, rotating head to the side, moving head up and down, and talking. Hence, our evaluation includes scenarios with relative mobility between the reader and camera; in fact, it also includes cases where no line of sight communication exists between the reader and camera. We next, evaluate our wireless camera for in-situ applications and assess how movements and antenna contact with a body affect video quality. We record the streaming video with the reader and measure the SNR and ENOB as we did in §5.1. Fig. 11 plots the ENOB of the received video for five different poses and movements. This plot shows that we can achieve sufficient ENOB while performing most of these gestures, resulting in a high-quality video compared to the original source video.

5.3 RF Power Harvesting

Next, we evaluate the feasibility of developing a battery-free HD streaming camera that operates by harvesting RF signals transmitted by the reader. We build an RF harvester for the 900 MHz ISM band based on a state-of-the-art RF harvester design [43]. The harvester consists of a 2 dBi dipole antenna and a rectifier that converts incoming RF signals into low-voltage DC output. The low-voltage DC is amplified by a DC-DC converter to generate the voltage levels to operate the image sensor and digital control block. We measure the power available at the output of the DC-DC converter.

We configure the reader to transmit a single tone at 30 dBm into a 6 dBi patch antenna and move the RF harvester away from the reader. Fig. 12 plots available power at the harvester as a function of distance. Based on available power, we also plot in Fig. 12 the maximum resolution and frame rate of the video that could be transmitted by an RF-powered ASIC version of our wireless camera. At close distances of 4 and 6 feet, we see sufficient power available from RF energy harvesting to operate the wireless camera at 60 fps 1080p resolution. As the distance increases, available power reduces, which lowers resolution of video being continuously streamed from the wireless camera. At 16 feet, the wireless camera continuously streams video at 10 fps 480p resolution; beyond this distance, the harvester does not provide sufficient power to continuously power the wireless camera. Note that Fig. 12 shows camera performance without using the distributed inter-frame compression algorithm described in §3.3. That algorithm, distributed across the wireless camera and reader, reduces camera transmissions, which lowers power consumption and consequently increases




Figure 12: **Power harvesting.** We plot the average power harvested by the battery-free hardware over different distances from the reader.




Figure 13: **Evaluation of distributed inter-frame compression algorithm for HD video.** The plots show PSNR at different compression ratio for two HD videos.

operating distances.

5.4 Distributed Inter-Frame Compression Evaluation

We consider two scenarios to evaluate our distributed inter-frame compression (subsection 3.3). We analyze HD video streamed from a fixed highway monitoring camera and from an action camera mounted on a user riding a motorcycle [9]. We evaluate the trade-off between the compression ratio and PSNR under both static and dynamic video feeds using our design. We measure the Peak Signal to Noise Ratio (PSNR) for different compression ratios by varying the threshold at which we consider the super-pixel (20 by 20 pixels) to have significantly changed. A higher threshold would result in higher compression ratios but at the cost of a degraded PSNR. Fig. 13 shows PSNR of the compressed video as a function of the compression ratio for our distributed inter-frame compression technique. For static cameras, we achieve a compression ratio of at least 35 \times while maintaining a PSNR above 30 dB. For dynamic videos recorded from a motorcycle, we achieve a compression ratio of 2 \times for a PSNR greater than 30 dB. This is expected since mobile scenarios significantly change majority of pixel values between frames, resulting in lower compression using our approach. We could address this by implementing a more complex compression algorithm at the reader to track moving objects in the frame and request specific pixel locations, achieving compression levels similar to video codecs. Implementing such complex compression algorithms, however, is beyond the scope of this paper.




Figure 14: **Prototype of our low-resolution, video streaming home security camera.** Image of the analog camera, FPGA digital core and pulse width modulated (PWM) backscatter, all implemented using COTS components. The overall board measures 3.5 cm by 3.5 cm.

6 Low-Resolution Security Camera

So far, we have demonstrated how high-resolution video offers a useful paradigm for a wearable camera system. However, various other applications, such as security systems and smart homes, do not require high-resolution video; lower resolution video would suffice for applications such as face detection. Specifically, wireless cameras are increasingly popular for security and smart home applications. In contrast to wearable camera applications, these cameras require low-resolution but much longer operating distances. To show that our design can extend to such applications, we evaluate the operation range at 13 fps 112 \times 112 resolution. Our IC design at this resolution consume less than 20 μ W without accounting for any distributed inter-frame compression saving.

To evaluate the range, we use a 13 fps 112 \times 112 resolution, gray-scale random pixel access camera from Cent-Eye [4] as our image sensor. The camera has a photodiode image sensor array, a trans-impedance buffer stage (to convert photodiode current to voltage) and a low-noise amplifier. It is extremely flexible, and the user can modify various settings, such as gain of the amplifier stage and, if desired, completely bypass the amplifier. We use this unique camera feature to prototype our wireless camera, which directly transmits analog values from the image sensor (sans amplification) using PWM backscatter. Fig. 14 shows photographs of our wireless camera prototype. The camera allows random access, i.e., any pixel on the camera can be accessed at random by setting the corresponding value in the row and column address registers. It can also be configured to output a single pixel, two adjacent pixels, or a super-pixel with sizes ranging from 2 \times 2 to 7 \times 7. We use the camera’s random access and super-pixel functionality to implement our distributed inter-frame compression algorithm. The power consumption of our low-resolution, off-the-shelf analog video streaming prototype is 2.36 mW. We emphasize that this off-the-shelf camera is used only to demonstrate operational range; to achieve the tens of microwatts power budget, we need to use our ASIC design.

Deployment results. We deployed our wireless camera system in the parking lot of an apartment complex. We




Figure 15: Operational range of low-resolution security camera.

use the USRP-based reader implementation set to transmit 30 dBm into a 6 dBi patch antenna. We vary the distance between the wireless camera prototype, and, the reader and at each separation, we stream 20 seconds of video from the camera to the reader. Simultaneously, we record camera output using a high input impedance National Instrument USB-6361 DAQ as the ground truth. We choose the popular PSNR metric commonly used in video applications to compare the video wirelessly streamed to the reader using PWM backscatter to the ground truth video recorded at the camera using an NI DAQ. PSNR computes the difference between the ground truth and wirelessly received video.

We measure the PSNR of the received video to evaluate the performance of the wireless camera under normal lighting conditions (below 300 lux) at a frame rate of 13 fps. To evaluate how much our sole wireless communication method affects the quality of received video, we consider PWM converter output as the ground truth for PSNR measurement. Also, to isolate the impact of AGC, which occurs at the reader, unaltered video received by the reader prior to applying any AGC is compared to the ground truth for PSNR measurement. Fig. 15 plots the PSNR of the received video at the reader as a function of the separation between the reader and the wireless camera. The plot shows that wireless camera streamed video at an average PSNR greater than 24 dB to a reader up to a distance of 150 feet away. Beyond 150 feet, the reader does not reliably decode the sync pulses, which limits the operating range of our wireless camera system. Thus, our analog backscatter approach achieves significantly longer range for low-resolution video compared to the HD version of the camera due to the trade-off between bandwidth/data rate and operating distances.

Applying our distributed inter-frame compression algorithm to low resolution videos. We deploy this security camera in a normal lab space and then implement our distributed inter-frame compression technique on the videos received. We evaluate the performance of our algorithm for three different super pixels sizes, 3×3 , 5×5 and 7×7 pixels, and plot results in Fig. 16. We achieve a $29.4\times$ data reduction using our distributed inter-frame compression technique while maintaining a PSNR greater than 30 dB.

Face detection accuracy. Next, we demonstrate that the




Figure 16: Distributed interactive compression with our low-resolution security camera. PSNR of the compressed video as a function of compression parameters in a typical lab setting.

quality of the video streamed from our low-resolution COTS implementation is sufficient for detecting human faces. Such a system can be used to detect human occupancy, grant access (such as Ring [12]), or set off an alarm in case of an intruder. To evaluate the system, we place the wireless camera at five different distances ranging from 16 to 100 feet from the reader. We ask ten users to walk around and perform gestures within 5 feet of the camera. We stream a 2 minutes video at each location at 13 fps and use the MATLAB implementation of the Viola-Jones algorithm to analyze the approximately four thousand video frames. Fig. 17 shows the accuracy of face detection as a function of the PSNR of the received video: as the quality (PSNR) of the video improves, the accuracy of face detection increases. We accurately detect up to 95% of human faces when the PSNR exceeds 30 dB.

7 Related Work

Prior work falls in two different categories.

Backscatter communication. An early example of analog backscatter was a gift by the Russians to the US embassy in Moscow, which included a passive listening device. This spy device consisted of a sound-modulated resonant cavity. The voice moved the diaphragm to modulate the cavity’s resonance frequency, which could be detected by analyzing the RF signals reflected by the cavity. [5]. A more recent example of analog backscatter is a microphone-enabled, battery-free tag that amplitude-modulates its antenna impedance using microphone output [51, 49]. In contrast, we design the first analog video backscatter system. Further, prior microphone designs had a low data-rate compared to video streaming. Our camera at 10 fps transmits about $9.2M$ pixels per second; for a microphone, a few kilo-samples of audio transmission is sufficient to fully recover the voice. In addition, our 13 fps $163K$ pixels per second camera operates at more than four times the range of the microphone in [51] due to pulse width modulation.

Ekhonet [53] optimizes the computational blocks between the sensor and the backscatter module to reduce the power consumption of backscatter-based wireless sensors. Our design builds on this work but differs from it in mul-




Figure 17: **Face detection with our low-resolution security camera.** We show the accuracy of face detection on the video streamed from our wireless camera.

multiple ways: 1) prior work still uses ADCs and amplifiers on the cameras to transform pixels into the digital domain and hence cannot achieve streaming video on the limited harvesting power budget. In contrast, we provide the first architecture for battery-free video streaming by designing an analog video backscatter solution.

Recent work on Wi-Fi and TV-based backscatter systems [25, 26, 23, 30, 42] can achieve megabits per second of communication speed using a backscatter technique. Integrating these designs with our video backscatter approach would prove a worthwhile engineering effort.

Low-Power cameras. [41] introduces a self-powered camera that can switch its photo diodes between energy harvesting and photo capture mode. Despite being self-powered, these cameras do not have wireless data transmission capabilities. [50, 38, 40, 35, 36, 37] show that using off-the-shelf, low-resolution camera modules, one can build battery-free wireless cameras that will capture still images using the energy they harvested from RF waves, including Wi-Fi and 900 MHz transmissions. Despite their ability to transmit data wirelessly, they are heavily duty cycled and cannot stream video. In particular, these designs can send a new frame at most every ten seconds when they are very close to the RF power source (within about a foot) and once every few tens of minutes at longer distances [50].

[32] presents a 90×90 pixels image sensor with pixels that are sensitive to changes in the environment. If a pixel receives sufficient illumination variation, the pixel address will be stored in a FIFO, thus compressing the image to the pixels that have significantly changed. Despite enabling image compression to occur at the image sensor, this system does not stream live video. In addition, at this low resolution, it burns about 3 mW of power when running at 30 fps. [28] introduces a 128×128 pixel event-driven image sensor that emphasizes low latency for detecting very fast moving scenes, so its power consumption is orders of magnitude higher than our system's.

[29] addresses the problem of conventional image sensors' power consumption not scaling well as their resolution and frame rate increases. In particular, the authors propose to change the camera input clock and aggressively switch the camera to standby mode based on de-

sired image quality. However, streaming video requires addressing the power consumption of multiple components, including camera, communication, and compression. Our work jointly integrates all these components to achieve the first battery-free video streaming design.

Finally, [39] shows that a regular camera on wearable devices burns more than 1200 mW, which limits the camera's operation time to less than two hours on a wearable device. They instead design a low-power wearable vision system that looks for certain events to occur in the field of view and turns on the primary imaging pipeline when those events happen. The power and bandwidth savings, however, are limited to the application and do not address communication. In contrast, we present the first battery-free video streaming application by jointly optimizing both backscatter communication and camera design and by eliminating power-consuming interfaces such as ADCs and amplifiers.

8 Limitations and Conclusion

This paper takes a significant first step in designing video streaming for battery-free devices. In this section, we discuss limitations and a few avenues for future research.

Security. Our current implementation does not account for security. However, to secure the wireless link between the camera and reader, we can leverage the fact that our digital core processes the PWM signal. Each wireless camera can be assigned a unique pseudo random security key. Based on this key, the camera's digital core can modulate the width of the PWM-encoded pixel value using an XOR gate. The reader, which knows the security key, can map the received data to the desired pixel values by performing the analogous operation.

ASIC versus off-the-shelf. While existing works on backscatter cameras focus on using off-the-shelf components, they treat cameras and backscatter independently and just interface the two. Thus, these works cannot achieve video streaming and the low-power demonstrated in this paper. Our key contributions are to make a case for a joint camera and backscatter architecture and to design the first analog video backscatter solution. However, this holistic architecture cannot be achieved with off-the-shelf components. The cost of building ICs in a research environment is prohibitively high. We believe that we spec out the IC design (in §4) in sufficient detail using industry standard EDA tools to take it from the lab to industry.

Mobile device as reader. To support our design for HD video streaming from wearable cameras, the smartphone must support a backscatter reader. We can use RFID readers that can be plugged into the headphone jack of the smartphone [6] to achieve this. In the future, we believe that backscatter readers would be integrated into smartphones

to support video and other applications [23].

Enabling concurrent streaming. Our analog backscatter camera design can use frequency division multiplexing techniques to share the wireless medium across multiple devices. The reader can coordinate communication by assigning different frequency channels to each camera. Multiple cameras can simultaneously backscatter on different channels by using different (assigned) frequency offsets in our sideband modulation. However, evaluating this approach is beyond the scope of this paper.

9 Acknowledgments.

We thank Kyle Jamieson and the anonymous reviewers for their helpful feedback on the paper. Also, we thank Ye Wang for his efforts in building a real-time demonstration of our low-resolution security camera. This work was funded in part by NSF awards CNS-1407583, CNS-1305072, CNS-1452494, CNS-1420654, Google Faculty Research Awards and Sloan Fellowship.

References

- [1] 6-dBi gain patch antenna. https://www.arcantenna.com/index.php/product_documents/get_document/id/492/. Accessed: March 2017.
- [2] Ambarella a9 ultra hd 4k camera soc product brief. <http://www.ambarella.com/uploads/docs/A9-product-brief.pdf>.
- [3] Barker code. <http://mathworld.wolfram.com/BarkerCode.html>. Accessed: September 2017.
- [4] CentEye Stonyman image sensor datasheet. <http://www.centeye.com/products/current-centeye-vision-chips/>. Accessed: March 2017.
- [5] The great seal bug. <http://counterespionage.com/the-great-seal-bug-part-1.html>. Accessed: March 2017.
- [6] Head jack pluggable rfid reader for smartphones. http://www.rfidtagworld.com/products/Ear-Jack-reader-UHF_1056.html.
- [7] IGLOO Nano FPGA datasheet. <https://www.microsemi.com/document-portal/docview/130695-ds0110-igloo-nano-low-power-flash-fpgas-datasheet>. Accessed: March 2017.
- [8] Lmv7219 comparator. <http://www.ti.com/lit/ds/symlink/lmv7219.pdf>. Accessed: September 2017.
- [9] Moving camera on a motorcycle video clip. <https://www.youtube.com/watch?v=sHj3xSG-RE&t=376s>.
- [10] NCX2200I low power comparator datasheet. http://www.nxp.com/documents/data_sheet/NCX2200.pdf. Accessed: March 2017.
- [11] RF5110 amplifier. http://www.rfmd.com/store/downloads/dl/file/id/30508/5110g_product_data_sheet.pdf. Accessed: March 2017.
- [12] Ring video doorbell. <https://ring.com/>. Accessed: March 2017.
- [13] Spectacles battery life. <http://www.barrons.com/articles/snapchat-spectacles-review-the-good-the-bad-the-revolutionary-1487846715>.
- [14] Synopsis design complier. <http://www.synopsys.com/Tools/Implementation/RTLSynthesis/DesignCompiler/Pages/default.aspx>.
- [15] USRP X-300. <https://www.ettus.com/product/details/X300-KIT>. Accessed: March 2017.
- [16] Wearable spectacles. <http://www.techradar.com/reviews/snap-spectacles>.
- [17] Youtube website. <http://www.youtube.com>.
- [18] H. R. Beard, A. J. Marquez-Lara, and K. S. Hamid. Using wearable video technology to build a point-of-view surgical education library. *JAMA surgery*, 151(8):771–772, 2016.
- [19] S. R. Best and B. C. Kaanta. A tutorial on the receiving and scattering properties of antennas. *IEEE Antennas and Propagation Magazine*, 51(5), 2009.
- [20] L. Chen, X. Tang, A. Sanyal, Y. Yoon, J. Cong, and N. Sun. A 0.7-v 0.6-uw 100-ks/s low-power sar adc with statistical estimation-based noise reduction. *IEEE Journal of Solid-State Circuits*, 52(5):1388–1398, 2017.
- [21] R. C. Gonzalez and R. E. Woods. *Image processing*, volume 2. 2007.
- [22] S. Hanson, Z. Foo, D. Blaauw, and D. Sylvester. A 0.5 v sub-microwatt cmos image sensor with pulse-width modulation read-out. *IEEE Journal of Solid-State Circuits*, 45(4):759–767, 2010.
- [23] V. Iyer, V. Talla, B. Kellogg, S. Gollakota, and J. Smith. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *Proceedings of the 2016 Conference on ACM SIGCOMM 2016 Conference*, SIGCOMM '16, pages 356–369, New York, NY, USA, 2016. ACM.

- [24] K. Kapucu and C. Dehollain. A passive uhf rfid system with a low-power capacitive sensor interface. In *RFID Technology and Applications Conference (RFID-TA), 2014 IEEE*, pages 301–305. IEEE, 2014.
- [25] B. Kellogg, A. Parks, S. Gollakota, J. R. Smith, and D. Wetherall. Wi-fi backscatter: Internet connectivity for rf-powered devices. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, 2014.
- [26] B. Kellogg, V. Talla, S. Gollakota, and J. Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In *Usenix NSDI*, 2016.
- [27] W. Kester. Understand sinad, enob, snr, thd, thd+ n, and sfdr so you don't get lost in the noise floor. *MT-003 Tutorial*, www.analog.com/static/importedfiles/tutorials/MT-003.pdf, 2009.
- [28] J. A. Leñero-Bardallo, T. Serrano-Gotarredona, and B. Linares-Barranco. A 3.6us latency asynchronous frame-free event-driven dynamic-vision-sensor. *IEEE Journal of Solid-State Circuits*, 46(6):1443–1455, 2011.
- [29] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong, and P. Bahl. Energy characterization and optimization of image sensing toward continuous mobile vision. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, pages 69–82. ACM, 2013.
- [30] V. Liu, V. Talla, and S. Gollakota. Enabling instantaneous feedback with full-duplex backscatter. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking, MobiCom '14*.
- [31] V. Liu, V. Talla, and S. Gollakota. Enabling instantaneous feedback with full-duplex backscatter. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 67–78. ACM, 2014.
- [32] U. Mallik, M. Clapp, E. Choi, G. Cauwenberghs, and R. Etienne-Cummings. Temporal change threshold detection imager. In *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, pages 362–603. IEEE, 2005.
- [33] F. Michel and M. Steyaert. A 250mv 7.5 μ w 61db snr cmos sc $\delta\sigma$ modulator using a near-threshold-voltage-biased cmos inverter technique. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 476–478. IEEE, 2011.
- [34] B. Murmann. Adc performance survey 1997-2017. <http://web.stanford.edu/~murmann/adcsurvey.html>.
- [35] S. Naderiparizi, Z. Kapetanovic, and J. R. Smith. Battery-free connected machine vision with wispcam. *GetMobile: Mobile Computing and Communications*, 20(1):10–13, 2016.
- [36] S. Naderiparizi, Z. Kapetanovic, and J. R. Smith. Wispcam: An rf-powered smart camera for machine vision applications. In *Proceedings of the 4th International Workshop on Energy Harvesting and Energy-Neutral Sensing Systems*, pages 19–22. ACM, 2016.
- [37] S. Naderiparizi, Z. Kapetanovic, and J. R. Smith. Rf-powered, backscatter-based cameras. In *Antennas and Propagation (EUCAP), 2017 11th European Conference on*, pages 346–349. IEEE, 2017.
- [38] S. Naderiparizi, A. N. Parks, Z. Kapetanovic, B. Ransford, and J. R. Smith. Wispcam: A battery-free rfid camera. In *RFID (RFID), 2015 IEEE International Conference on*, pages 166–173. IEEE, 2015.
- [39] S. Naderiparizi, P. Zhang, M. Philipose, B. Priyantha, J. Liu, and D. Ganesan. Glimpse: A programmable early-discard camera architecture for continuous mobile vision. In *Proceeding of the 15th annual international conference on Mobile systems, applications, and services*. ACM, 2017.
- [40] S. Naderiparizi, Y. Zhao, J. Youngquist, A. P. Sample, and J. R. Smith. Self-localizing battery-free cameras. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 445–449. ACM, 2015.
- [41] S. K. Nayar, D. C. Sims, and M. Fridberg. Towards self-powered cameras. In *Computational Photography (ICCP), 2015 IEEE International Conference on*, pages 1–10. IEEE, 2015.
- [42] A. N. Parks, A. Liu, S. Gollakota, and J. R. Smith. Turbocharging ambient backscatter communication. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, 2014.
- [43] A. N. Parks, A. P. Sample, Y. Zhao, and J. R. Smith. A wireless sensing platform utilizing ambient rf energy. In *Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireless), 2013 IEEE Topical Conference on*, pages 154–156. IEEE, 2013.

- [44] A. P. Sample, D. J. Yeager, P. S. Powladge, A. V. Mamishev, and J. R. Smith. Design of an rfid-based battery-free programmable sensing platform. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, 2008.
- [45] R. Sarpeshkar. Analog versus digital: extrapolating from electronics to neurobiology. *Neural computation*, 10(7):1601–1638, 1998.
- [46] M. D. Scott, B. E. Boser, and K. S. Pister. An ultra-low power adc for distributed sensor networks. In *Solid-State Circuits Conference, 2002. ESSCIRC 2002. Proceedings of the 28th European*, pages 255–258. IEEE, 2002.
- [47] J. Sun. Pulse-width modulation. In *Dynamics and control of switched electronic systems*, pages 25–61. Springer, 2012.
- [48] V. Talla, M. Hessar, B. Kellogg, A. Najafi, J. R. Smith, and S. Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):105, 2017.
- [49] V. Talla, B. Kellogg, S. Gollakota, and J. R. Smith. Battery-free cell phone. ACM UBIComp, 2017.
- [50] V. Talla, B. Kellogg, B. Ransford, S. Naderiparizi, S. Gollakota, and J. R. Smith. Powering the next billion devices with wi-fi. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies, CoNEXT '15*, pages 4:1–4:13, New York, NY, USA, 2015. ACM.
- [51] V. Talla and J. R. Smith. Hybrid analog-digital backscatter: A new approach for battery-free sensing. In *RFID (RFID), 2013 IEEE International Conference on*, pages 74–81. IEEE, 2013.
- [52] F. Tang and A. Bermak. An 84 pw/frame per pixel current-mode cmos image sensor with energy harvesting capability. *IEEE Sensors Journal*, 12(4):720–726, 2012.
- [53] P. Zhang, P. Hu, V. Pasikanti, and D. Ganesan. Ekhoneet: High speed ultra low-power backscatter for next generation sensors. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 557–568. ACM, 2014.