

SONIC: Connect the Unconnected via FM Radio & SMS

Ayush Pandey
ayush.pandey@nyu.edu
New York University Abu Dhabi
Abu Dhabi, United Arab Emirates

Rohail Asim
rohail.asim@nyu.edu
New York University Abu Dhabi
Abu Dhabi, United Arab Emirates

Khalid Mengal
kqm1@nyu.edu
New York University Abu Dhabi
Abu Dhabi, United Arab Emirates

Matteo Varvello
matteo.varvello@nokia.com
Nokia Bell Labs
New Jersey, United States

Yasir Zaki
yasir.zaki@nyu.edu
New York University Abu Dhabi
Abu Dhabi, United Arab Emirates

Abstract

As of 2022, about 2.78 billion people in developing countries do not have access to the Internet. Lack of Internet access hinders economic growth, educational opportunities, and access to information and services. Recent initiatives to “connect the unconnected” have either failed (project Loon and Aquila) or are characterized by exorbitant costs (Starlink and similar), which are unsustainable for users in developing regions. This paper proposes SONIC, a novel connectivity solution that repurposes a widespread communication infrastructure (AM/FM radio) to deliver access to pre-rendered webpages. Our rationale is threefold: 1) the radio network is widely accessible – currently reaching 70% of the world – even in developing countries, 2) unused frequencies are highly available, 3) while data over sound can be slow, when combined with the radio network, it takes advantage of its broadcast nature, efficiently reaching a large number of users. We have designed and built a proof-of-concept for SONIC, which shows encouraging initial results.

CCS Concepts

• **Networks** → **Network design principles**; *Wireless access networks*; • **Information systems** → **World Wide Web**.

Keywords

FM radio, World Wide Web, Developing regions

ACM Reference Format:

Ayush Pandey, Rohail Asim, Khalid Mengal, Matteo Varvello, and Yasir Zaki. 2024. *SONIC: Connect the Unconnected via FM Radio & SMS*. In *Proceedings of the 20th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '24)*, December 9–12, 2024, Los Angeles, CA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3680121.3697812>

1 Introduction

Access to the Internet is vital in today’s interconnected world. According to the International Telecommunication Union (ITU) about

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT '24, December 9–12, 2024, Los Angeles, CA, USA
© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1108-4/24/12
<https://doi.org/10.1145/3680121.3697812>



Figure 1: Pre-rendered webpage delivered as WebP via SONIC with no frame lost (left), 10% losses (center), and pixel interpolated on 10% losses (right).

2.9 billion people are still offline [25] (96% of them are from developing regions). Several initiatives attempted to extend Internet connectivity to remote areas, such as Google Loon [2] (via high-altitude balloons) and Facebook Aquila [4] (via solar-powered drones). Both projects were eventually halted due to high cost and performance issues (notably Aquila crash [3]). Starlink [7], or Internet access provided via low-earth orbit satellites, is another recent solution that is vastly growing (1.5 million users as of June 2023) but characterized by excessive costs for most people from developing regions.

The goal of this work is to design a novel technology to provide some form of *free* Internet access in developing regions. Our rationale is to leverage existing infrastructures that have broad adoption but were not intended for Internet data dissemination. We propose *SONIC*, a system that exploits the FM radio infrastructure – currently covering roughly 70% of the world according to the United Nations [46] – to enable the *download* of simplified webpages encoded over sounds. SONIC further relies on the SMS (Short Message Service) network, when available, as an uplink to interact with a webpage or request a new one. SONIC complements existing initiatives aimed at improving Internet availability. For instance, (expensive) satellite communication can be used to provide Internet connectivity to a few radio stations or towers. These stations can then leverage SONIC to extend access to end-users, making Internet connectivity more affordable and widespread.

SONIC consists of a *client* and a *server*. The server extends an FM transmitter with the capability to load webpages, render them, and then encode their appearances over sounds that are transmitted

along with regular radio programs. The client consists of a low-end mobile device and an FM receiver, often integrated. The client is equipped with an application to decode a webpage from sound. SONIC transmits images of rendered webpages for data savings (about 10x compression) and loss resilience, differently from Web objects like HTML and JavaScript. SONIC further uses nearest neighbor pixel interpolation [40] to recover from errors. Figure 1 shows that a webpage is still readable despite about 10% loss rate.

We have built a proof-of-concept of SONIC consisting of a Raspberry Pi acting as an FM radio transmitter [33] and a Xiaomi Redmi Go [10] acting as a client using its internal FM tuner as well as an external receiver. For encoding and decoding images of webpages, we rely on the Quiet library [37] and a chunking mechanism we have developed. Using this setup, we demonstrate that a rate of 10 kbps is sustainable with no losses when the FM receiver is integrated into the mobile device, and a Received Signal Strength Indication (RSSI) of -65 to -85 dB is achieved. Next, we demonstrate that this transmission rate is sufficient to broadcast 100 popular Pakistani websites with acceptable readability, as confirmed by a user study involving 151 participants. Multiple frequencies can be used to increase the rate, allowing for either broadcasting more content or supporting user requests via uplink.

2 Background and Related Work

Data Over Sound: Several research papers [11, 27, 30, 39, 41] and open source tools [21, 26, 37] have explored how to transmit data over both sound and ultrasound or inaudible audio frequencies, i.e., above 18kHz. Original studies [27, 30, 41] focus on the near-ultrasound band (i.e., 17-22kHz) and achieve overall low throughput at variable distances: 15bps (up to 25 meters) [30], 16bps (2.7 meters) [27], and 2.76kbps (1 meter) [41]. BackDoor [39] achieves higher speed (4kbps and up to 1 meter) by leveraging the non-linearity in microphone hardware to create a “shadow” in the audible frequency range. BatComm [11] further expands on the BackDoor idea by using orthogonal frequency division multiplexing (OFDM [31]) to transmit data over multiple orthogonal channels with an ultrasound frequency carrier. Batcomm demonstrates an overall much higher throughput than BackDoor: 47kbps (at 10 centimeters) and 17kbps (at 2.3 meters).

Among open source tools, AudioQR [26] works in the near-ultrasonic frequency band (17.5-19.5kHz) and can reach low speeds of about 100 bps but supporting long distances (up to 150 meters). GGwave [21] is a tiny data-over-sound library implementing a transmission protocol based on frequency-shift keying (FSK [48]) which can reach up to 128 bps over short distances. Finally, Quiet [37] enables data over both audible and inaudible frequencies, as well as different modulations such as 1024-QAM [49] and gmsk [19]. Quiet claims audible transfers at 7 kbps over the air, and up to 64kbps in cases where two devices are connected over an audio jack cable.

BatComm is an interesting approach for SONIC, but it is not open source and could not be shared by the authors upon our request. AudioQR is a solid product but not a good fit for SONIC since it sacrifices transmission speed for high distance, while we target very low “air distance”, e.g., some phones directly embed an FM receiver (Figure 3). With this respect, GGwave and Quiet




Figure 2: Visualization of the FM spectrum.

are better candidates for SONIC; we choose Quiet due to its higher performance and flexibility.

Innovative Usage of FM Radio: To the best of our knowledge, SONIC is the first system which uses FM radio to disseminate Web content. Further, SONIC combines the FM radio with the SMS infrastructure to provide both a downlink and an uplink. Few other systems have proposed innovative usage of FM radio as a downlink. RevCast [42] leverages the broadcast nature of FM radio for certificate revocation, i.e., the act of invalidating a TLS/SSL certificate before the validity period expires. RevCast relies on Radio Data System (RDS or RDBS in the US), a communications protocol for embedding small amounts of digital information in FM radio broadcasts. Both RDS and RBDS carry data at 1,187 bits per second on a 57 kHz subcarrier, as shown in Figure 2. The authors show, using a 3 kW commercial radio station and an RDS-enabled smartphone, that RevCast can deliver complete and timely revocation information, anonymously.

The work of [13, 14] use FM radio broadcasting to disseminate warning information to drivers. This solution leverages DSSS-CDMA transmissions [45], where each vehicle’s geographical position is used as a CDMA key to decode the relevant image. The authors validate their solution using a combination of simulation (MATLAB) and experimental setup (USRP devices [20]), and show data rates of up to 35 kbps. Finally, in 2003, Microsoft used FM subcarrier signals to turn ordinary gadgets into smart gadgets. MSN Direct [50] was a subscription network that sent short text updates over DirectBand (see Figure 2), i.e., the 67.65 kHz subcarrier leased by Microsoft from commercial radio broadcasters. This subcarrier delivers about 12 kbit/s of data per tower, for over 100 MB per day per city. MSN Direct has found some success in delivering traffic notifications and weather forecasts to Garmin GPS devices. However, it suffered from: 1) competition with cellular broadband, 2) downlink only communication. Microsoft ended the service in 2012.

3 SONIC Design

This Section designs SONIC (SOund NIC, or network interface controller), a technology which relies on FM broadcasting to provide access to simplified webpages encoded over sounds. The remainder of this section is organized as follows. We first provide an overview of SONIC. Next, we dive into the design of its key components.

3.1 Workflow

Figure 3 visualizes how SONIC works. Each mobile phone showcases a potential SONIC user with different capabilities. User-B owns a smartphone equipped with a Qualcomm chip mounting an




Figure 3: User-C requests a webpage via SMS (1). The SONIC server fetches and renders it into a WebP image (2), then passes it to an FM transmitter (3) for broadcast via sound (4). The webpage is delivered to user-C and any other device currently listening to the radio (5).

FM receiver, which is popular in developing regions [6]. User-A owns a smartphone with no usable FM receiver, but (s)he is in the proximity of a classic radio, e.g., at home or in the car. Both users can receive SONIC data, i.e., simplified webpages encoded over sounds, but cannot request a specific webpage due to their lack of SMS service.

User-C connects her phone to an FM radio via the audio jack, for higher performance (see Section 2), and also pays for Short Message Service (SMS). User-C can both *receive* and *request* webpages. In the figure, user-C is interested in news from `cnn.com`. This is achieved via an SMS to a SONIC provided number (1) which triggers the generation of a simplified webpage for `cnn.com` (2), e.g., the landing page, and its broadcasting over the FM radio (3)(4), so that user-C (as well as user-A and user-B) will eventually receive it (5).

SONIC Server: It is responsible for generating the simplified webpages to broadcast via FM radio. We assume that the FM radio infrastructure consists of multiple transmitters (and frequencies) at different locations, with some form of Internet access. Accordingly, the FM transmitters receive simplified webpages to be encoded via sound, and then transmit them. A central SONIC server produces such webpages and inform the respective transmitters.

The SONIC server produces simplified webpages according to users requests. Further, it maintains a list of the most popular websites in a region which are preemptively *pushed* to users in an attempt to improve their experience. For example, popular news sites can be pushed early in the morning.

SONIC Client: As visualized in Figure 3, users access SONIC using different hardware configurations. The minimum configuration implies a smartphone and access to an FM radio, either over the air or via an audio jack, which provides the *downlink*. An SMS service is needed for users interested in having access to the *uplink*, e.g., request a specific webpage or query a search engine. Software-wise, SONIC runs as a standalone application we describe next.

The SONIC application allows to *browse* webpages previously received via the FM radio, and request new ones. Accordingly, the app shows a catalog of available webpages, organized by content, popularity, and/or user interest. Within a webpage, a user might be interested in visiting some internal pages following classic hyperlinks. If the requested internal page is locally available, i.e., it was previously received via the FM radio, then the page would

instantly load. If not, an active uplink is required to request such page. Similarly, SONIC users with an active uplink can request additional webpages not previously received, as well as send queries to search engines (e.g., Google and Duckduckgo) and AI chatbots (e.g., chatGPT). Note that pages requiring authentication, such as online banking and social media, are not supported yet due to their privacy risk, as discussed in Section 3.4.

SONIC clients request *new* webpages, i.e., not previously received and locally cached, via SMS. Each request contains the URL, e.g., `cnn.com/index.html`, and geographic location of the user. The latter is needed by SONIC server to inform the proper transmitter along with its frequency, as discussed below. As the request is received, the SONIC server produces a *simplified* version of the webpage, either from its cache, e.g., if recently requested by another user, or by directly accessing it. The simplified page is then transferred to the radio transmitter which can physically reach the current user location. Meanwhile, SONIC server quickly responds to the user via SMS to acknowledge the request, and provide an estimate on when the page will be received. When the webpage is received, it is inserted in a cache with expiration date set according to a time indicated by the server.

3.2 Simplified Webpage Creation

According to [8], the average webpage size is around 2 MB for mobile sites. Given the low speeds of data over sound, i.e., tens of kbps at best (see Section 2), broadcasting 2 MB via SONIC might require tens of minutes. Further, receivers might not be able to reconstruct the received page, given that most Web components like JavaScript would break in presence of unrecoverable errors.

It follows that we need a mechanism to: 1) greatly compress webpages, 2) make webpages resilient to noise. Many solutions exist to reduce webpage sizes, such as compression proxies [5, 23], reader modes [9, 22], JavaScript cleaners [16–18] and redundant code removal [28, 32]. However, these solutions are still affected by noise and would thus require a large amount of FEC, which implies dimensioning the system for a worst case receiver. Instead, we would like a solution whose performance decreases as a function of the receiver’s RSSI, similarly to how the audio quality is reduced for users with bad reception.

We took inspiration from the work in [13, 14], which shows how the quality of images transmitted over RDS decreases as a function of a receiver’s RSSI. Accordingly, we propose to transmit *images* of loaded webpages, instead of actual Web files (html, js, css, etc.). Such images achieve both our goals since they compress a page, e.g., a 2MB page can be compressed as a few hundreds KB image, and they are loss resilient, as despite missing pixels users might still understand the content provided.

DRIVESHAFT (DS) [12] is a system to speedup the mobile Web by distributing crowdsourced screenshots of webpages. While the overall system is quite different, e.g., it involves complex mechanism to merge harvested screenshots while stripping them of private information, SONIC could leverage some of DS’s components, e.g., how to add interactivity to the screenshots. Unfortunately, DS’s code is not available, and we could not experiment and integrate it with SONIC. We here discuss SONIC’s mechanism to form images of webpages and its relationship with DS.




Figure 4: Preliminary performance evaluation of SONIC. (a) Frame loss rate as a function of the over the air distance between FM receiver and a SONIC client. “Cable” refers to a distance of zero, i.e., either an integrated FM receiver on the phone or a jack audio cable connection. (b) CDF of the size of images (WebP) of rendered webpages assuming variable image quality (Q) and pixels height (PH). (c) Evolution over time of the amount of data to be broadcasted as a function of transmission rates and number of webpages (N).

Interactivity: modern webpages allow user interaction either via simple hyperlinks or with menus, search boxes, etc. Images are instead static by definition. To allow some interactivity, [12] proposes to build *click maps*, i.e., maps containing a list of $\langle x, y \rangle$ coordinates of where an image is interactive. We adopt the same approach; as the user clicks on such coordinates, SONIC informs the server (via SMS, if available) and request the next image, e.g., an internal page accessed via an hyperlink, unless already available in the cache. Due to the more challenging network conditions of SONIC, e.g., potentially seconds in uplink and minutes in downlink, we limit the interactivity to hyperlinks only.

Images Format: DS requires lossless png due to its need to desensitize and merge crowdsourced screenshots. This is not a requirement for SONIC; we thus resort to WebP [1], a modern image format that provides superior lossless and lossy compression. We capture images of webpages as WebP with 10% quality, which significantly reduces the size while remaining acceptable, as we evaluate in Section 4. These images are created with 1,080 pixels width, and a maximum of 10k pixels height. The idea is to allow a user to “scroll” down on an image, as commonly done on a webpage, while avoiding to waste broadcasted data, given that user interest is lower for content below the fold. Depending on the mobile phone screen resolution, and using the scaling factor (i.e., mobile phone screen width / 1,080), the images are resized by multiplying both the width and height with the scaling factor. The scaling factor is also used to scale the coordinates of the click maps.

3.3 Sonic Data Transmission

Using the Quiet library [37], we create a new transmission profile inspired by their “audible-7k-channel”. The new profile uses Orthogonal Orthogonal Frequency-Division Multiplexing (OFDM) – a multi-carrier modulation system where data is transmitted as a combination of orthogonal narrow-band signals known as sub-carriers – with 92 sub-carriers. The data rates achieved by this profile reach 10 kbps. Upon transmitting a rendered page, we first divide the image vertically into multiple partitions, each with a width of 1 pixel. Each partition is then divided into fixed-sized frames of 100 bytes each. Each frame carries a partition and a sequence number

used to reassemble the image on the receiver end. We further leverage the error detection and correction schemes using: crc32 as the checksum, the inner Forward Error Correction (FEC) scheme (v29), and an outer FEC scheme (rs8). These allow the receiver to correct transmission errors.

In case of lost frames, the SONIC receiver uses nearest neighbor pixel interpolation [40] to correct missing pixels. Accordingly, missing pixels are replaced with the value of their adjacent pixel, prioritizing the left pixel given that webpage consists mostly of text read from left to right. In literature, more advanced and better performing techniques exist to recover missing pixels in images [15, 35, 44], leveraging deep neural network to learn patterns and structures of the image or leveraging sparsity and gradients in the data to fill in the missing regions. These techniques are both memory and CPU intensive, far beyond what a low-end device can support today.

3.4 Discussion

Incentives and Monetization: For users, the incentive is clear since they get access to a simplified Web in areas where such content is currently inaccessible. For SONIC providers, they could directly charge their users. This can be challenging since, for example, downlink users are completely passive leaving no trace to when any content is accessed. Alternatively, they could bundle their service with an SMS service, thus offering a paid service for users who can request content in real time, while being *free* for other users. Anyway, the cost of broadcasting over the radio is constant regardless of the number of “listeners”.

Finally, a more interesting monetization scheme arises when considering the current business model of radio stations: increasing their user-base to sell more advertisement. SONIC represents a novel service that radio stations can offer to attract more users, thus increasing their revenues via ads. Further, ads can now also be integrated with webpages, expanding from audio only to images.

Privacy Concerns: In essence, SONIC functions similarly to Web acceleration tools like Google AMP [23] and WebLight [5] which rewrite the content before sending it to a user. To do so, SONIC requires full visibility on the requested URL, as well as the content being currently accessed by a user. It follows that a SONIC server

could collect enough information to build user profiles and eventually violate user privacy. Differently from these solutions, SONIC relies on a broadcast medium (the FM radio) thus not knowing precisely how many users might be accessing a given content. Indeed, no privacy violation is possible for downlink-only users who can never request any content but just occasionally browse popular content requested by other users in their area.

Content Limitations: SONIC does not provide access to webpages behind a login, such as a personal bank website or a social media account. Such access is impossible for downlink-only users; for users with an uplink, they would have to share their username and password to allow the SONIC server to retrieve the required content. Further, the resulting page (e.g., showing the balance of the user) would be broadcasted to potentially many users adding on the privacy issue.

Next, SONIC lacks support for video, which is used at large in today’s Web not only by video streaming services but also by news websites, social media, etc. Simply put, SONIC does not have enough bandwidth to support video streaming; instead, videos are replaced by their thumbnails which are not “clickable”. Similarly, SONIC transmits simplified (pre-rendered) webpages with no support of complex visual effects normally achieved via either JS or CSS.

Scalability: SONIC is a best effort system, aiming to offer few webpages to a large number of users – anyone in the range of an FM radio – despite some very small bandwidth. This is achieved by bending the current strictness of webpages, and accepting some visual quality degradation as well as lack of full interactivity. Further, the catalog of webpages to be served has to be small (see Figure4(c)) and thus properly curated to offer interesting and useful content to its users.

4 Preliminary Evaluation

Methodology: We build an FM transmitter using a Raspberry Pi, a 20cm wire acting as an antenna connected to GPIO pin 4 of the Pi, and the code from [33], which uses the general clock output to produce frequency modulated radio communication in the range from 1MHz to 250MHz. We use 93.7MHz since it is unused at our location. For data transmission, we use the Mono channel of the FM baseband (30Hz to 15KHz, see Figure 2) with a carrier center frequency of 9.2KHz. We envision that other bands can be used to increase the data rate, e.g., using the left and right band of the Stereo channel, or even the DARC band. We left this exploration as future work. As FM receivers, we consider the scenario showcased in Figure 3. User-B leverages the internal FM tuner of a Xiaomi Redmi Go [24], and a pair of headphones as an antenna. User-A is another Redmi Go receiving sound from a nearby FM radio over the air. Finally, User-C is a Redmi Go connected to the same FM radio but using the audio jack. We ensure high RSSI between transmitter and receiver.

For encoding data over sound, we use Quiet [37] as discussed in Section 3.3. For decoding, we develop an Android application extending the sample application [36] from Quiet. We approximate user-B using the NextRadio app [34] to receive the FM radio signal and a jack cable to act both as an antenna and to share the received sound to a second device which runs the SONIC app. In the future, we will integrate FM radio receiving within the SONIC app.

For the content to transmit, we assume a deployment in Pakistan. Accordingly, we select the 25 most popular Pakistani websites from the Tranco list [29] filtered using the .pk domain name. For each landing page, we select three random internal pages, resulting in a total of 100 webpages (25 landing, and 75 internal). We rendered these 100 pages in Chrome on hourly basis for three days, while generating variable quality images of the rendered webpages.

General Results: Figure4(a) shows the frame loss rate, at the receiver, as a function of the distance between the phone and the receiving radio. “Cable” refers to a scenario where the receiver is directly connected to the FM radio, either using the internal FM tuner (User-B in Figure 3) or an actual cable (User-C in Figure 3). Each experiment is repeated 10 times, and we assume high RSSI (-70 db or higher) at the FM receiver. The figure shows no frame loss recorded over cable, and up to 10-20% frame losses (at the median) when considering about one meter or over the air communication between phone and FM receiver. The loss rate growth is not completely linear with the increased over the air distance, as the alignment between speaker and microphone also has a significant impact. We did not control for this variable, as SONIC users would likely not do so as well. We also observe a 100% loss rate at distances above 1.1m.

Figure 1 visualizes the impact of 10% losses. Overall, the impact is significant but *tolerable*, unless unlucky scenarios where losses are concentrated in important parts of the webpage, like title and text. In our early prototype, each portion of an image is transmitted equally; one optimization consists of adopting a dynamic scheme with higher error protection for important parts of an image/webpage. Nevertheless, we envision that SONIC users would mostly rely on their phone internal FM tuner – or a cable connection to a radio – to improve their performance and minimize lost frames.

Figure 4(b) shows the Cumulative Distribution Function (CDF) of the size of images (WebP) of rendered webpages assuming variable image quality (Q). WebP image quality is defined on a scale from 0 (worst) to 95 (best). The figure also reports the pixels height (PH), either “none” (i.e., maximum length of the webpage loaded) or cropped at 10K pixels (see Section 3.2). Assuming a low quality of 10, most webpages can be compressed in less than 200 KB, which instead requires 700 KB at high quality of 90. For SONIC, given the challenging network conditions of the FM radio, we opt for a low quality of 10 given that no significant image degradation is observed. The figure also shows that cropping to 10k pixels saves about 100 KB for 75% of the webpages. The tails of the CDFs show images with twice the size of the 90th percentile, e.g., 500 KB versus 250 KB when considering Q equal to 10 and PH equal to 10 k. These images can take significant time to be delivered, e.g., up to 6/7 minutes if assuming a single delivery frequency and a 10 kbps rate (see Section 3.3), and might thus require some further optimizations.

Finally, Figure 4(c) shows the evolution over time of the amount of data broadcasted with SONIC assuming the above 100 Pakistani webpages, and the current rate we have verified (10kbps) as well as 20 and 40 kbps which can be achieved via multi frequency (see Section 3). Note that we truncated the figure at two days – despite three days of data – since the pattern repeats and it helps improving its visibility.




Figure 5: Distribution of median user ratings (0-10 Likert scale) per top 50 Pakistani webpages, for question-b (plain) and question-a (hatched), with and without pixel interpolation (orange and blue, respectively).

Figure 4(c) shows that, with the current transmission rate (10 kbps), SONIC can only work in broadcast mode. In fact, the amount of data to be sent – due to variation in the supported webpages – rarely reaches zero (for which a rate of 20 or 40 kbps is required). However, SONIC is *scalable*, meaning that the amount of data to be sent does not grow indefinitely.

Variable RSSI: We evaluate SONIC’s frame losses across various levels of RSSI (Received Signal Strength Indication). RSSI measures signal strength in decibels, typically ranging from 0 (strongest) to -120 (lowest). FM requires RSSI comprised between -65 and -80 dB, e.g., radio stations in our area exhibit RSSI between -70 and -80 dB.

We utilize the TR508 FM Transmitter [38], which supports transmissions up to 1 km. We pair the SONIC client in “cable” mode (i.e., with no additional frame loss rate, as shown in Figure 4(a)) with the Real FM Radio app [47], which records RSSI values. We broadcast on an empty FM frequency and manually explore distances between transmitter and receiver to observe RSSI signals within the -65 to -90 dB range. At approximately 5 dB intervals, we transmit a single webpage up to 10 times and measure SONIC’s frame loss rate. For the RSSI range from -65 to -85 dB, we consistently observe no frame losses. For the -85 to -90 dB range, we record a fluctuating frame loss rate between 2 and 15%. As expected [43], for RSSI below -90 dB, we are unable to receive any frames.

User Study: To assess the impact of losses on webpage readability, we create screenshots of the top 50 Pakistani webpages (according to Tranco [29]) with synthetic variable losses (5%, 10%, 20%, and 50%). Missing pixels are either shown as dark or corrected using nearest neighbor pixel interpolation (see Section 3.3). This results in 400 screenshots (50 pages \times 4 loss rates \times 2 approaches, i.e., with and without pixel interpolation). We then recruit 151 students from a university in Pakistan to evaluate 20 randomly selected screenshots, thus averaging at least 7 ratings per screenshot.

Students rate screenshots on a Likert scale from 0 to 10 based on two questions: a) *Indicate your perception of content understanding in the above image (0 = extremely unclear, 5 = neutral, and 10 = extremely clear);* and b) *Rate the readability of the text in the above image considering the level of noise present (0 = very difficult to read, 5 = neither easy nor difficult, and 10 = very easy to read).* Figure 5 shows boxplots of the median rating per page for both questions

(plain for question-b and hatched for question-a), differentiating when pixel interpolation is used (orange) or not (blue). Overall, the figure demonstrates the benefits of pixel interpolation to enhance both text readability and content comprehension, improving the rating by at least one point regardless of the loss rate. For example, even at a loss rate of 20%, observed at a distance of 1.1 meters and above (see Figure 4(a)), pixel interpolation achieves a median content readability score of 7 (somewhat clear understanding). Text readability is more susceptible to losses, but still acceptable even with 20% losses.

5 Conclusion

This paper has presented SONIC, a novel connectivity solution which re-purposes the FM radio to broadcast prerendered webpages over sound frequencies. By serving prerendered webpages, SONIC supports graceful degradation of the user experience in presence of losses. Further, SONIC leverages the SMS infrastructure to provide an uplink, and allow its users to request specific websites, similarly to how in the 80s (lucky) radio listeners could call in to request a song. We built a proof of concept of SONIC and investigate its performance with respect to frame losses, quality of the pre-rendered webpages, and a combination of speed and volume of broadcasted pages. We find that most webpages can be compressed to 200KB images (using WebP and a quality level of 10) and still be “readable” even with a 10% loss rate. Further, a rate of 10kbps is attainable per frequency. With such rate, SONIC can effectively serve the 100 most popular webpages in Pakistan.

References

- [1] 2023. An image format for the Web. <https://developers.google.com/speed/webp>.
- [2] 2023. Expanding internet connectivity with stratospheric balloons. <https://x.company/projects/loon/>.
- [3] 2023. Facebook Halts Aquila, Its Internet Drone Project. <https://www.nytimes.com/2018/06/27/technology/facebook-drone-internet.html>.
- [4] 2023. High altitude connectivity: The next chapter. <https://engineering.fb.com/2018/06/27/connectivity/high-altitude-connectivity-the-next-chapter/>.
- [5] 2023. Load mobile pages faster with Web Light. <https://support.google.com/websearch/answer/9836344?hl=en>.
- [6] 2023. Snapdragon 835 Mobile Platform. <https://www.qualcomm.com/products/mobile/snapdragon/smartphones/snapdragon-8-series-mobile-platforms/snapdragon-835-mobile-platform>.
- [7] 2023. Starlink. <https://starlink.com/>.
- [8] 2023. Web Almanac. <https://almanac.httparchive.org/en/2022/>.
- [9] Apple 2023. Safari Reader Mode. <https://support.apple.com/guide/safari/hide-ads-when-reading-sfri32632/mac>.
- [10] GSM Arena. 2023. Xiaomi Redmi Go. https://www.gsmarena.com/xiaomi_redmi_go-9534.php.
- [11] Yang Bai, Jian Liu, Li Lu, Yilin Yang, Yingying Chen, and Jiadi Yu. 2020. BatComm: enabling inaudible acoustic communication with high-throughput for mobile devices. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 205–217.
- [12] Ketan Bhardwaj, Ada Gavrilovska, Moritz Steiner, Martin Flack, and Stephen Ludin. 2018. DRIVESHAFT: Improving Perceived Mobile Web Performance. *arXiv preprint arXiv:1809.09292* (2018).
- [13] Radu Gabriel Bozomitu and Florin Doru Hutu. 2023. Drivers’ Warning Application Through Personalized DSSS-CDMA Data Transmission by Using the FM Radio Broadcasting Infrastructure. *IEEE Access* 11 (2023), 11711–11731.
- [14] Radu Gabriel Bozomitu and Ștefan Corneliu Stoica. 2022. A Robust Radiocommunication System for FM Transmission Based on Software Defined Radio Technology. In *2022 IEEE 28th International Symposium for Design and Technology in Electronic Packaging (SIITME)*. IEEE, 78–81.
- [15] Xintao Chai, Hanming Gu, Feng Li, Hongyou Duan, Xiaobo Hu, and Kai Lin. 2020. Deep learning for irregularly and regularly missing data reconstruction. *Scientific Reports* 10, 1 (24 Feb 2020), 3302. <https://doi.org/10.1038/s41598-020-59801-x>
- [16] Mounema Chaqfeh, Russell Coke, Jacinta Hu, Waleed Hashmi, Lakshmi Subramanian, Talal Rahwan, and Yasir Zaki. 2022. Jsanalyzer: A web developer tool for

- simplifying mobile web pages through non-critical javascript elimination. *ACM Transactions on the Web* 16, 4 (2022), 1–31.
- [17] Moumena Chaqfeh, Muhammad Haseeb, Waleed Hashmi, Patrick Inshuti, Manesha Ramesh, Matteo Varvello, Lakshmi Subramanian, Fareed Zaffar, and Yasir Zaki. 2022. To Block or Not to Block: Accelerating Mobile Web Pages On-The-Fly Through JavaScript Classification. In *Proceedings of the 12th International Conference on Information and Communication Technologies and Development, ICTD 2022, Seattle, USA, June 27-29, 2022*.
- [18] Moumena Chaqfeh, Yasir Zaki, Jacinta Hu, and Lakshmi Subramanian. 2020. JSCleaner: De-Cluttering Mobile Webpages Through JavaScript Cleanup. In *Proceedings of The Web Conference 2020*. 763–773.
- [19] Melvin L Doelz and Earl T Heald. 1961. Minimum-shift data communication system. US Patent 2,977,417.
- [20] ettus 2023. Universal Software Radio Peripheral. <https://www.ettus.com/all-products/un210-kit/>.
- [21] Georgi Gerganov. 2023. GGwave. <https://github.com/ggerganov/ggwave>.
- [22] Mohammad Ghasemisharif, Peter Snyder, Andrius Aucinas, and Benjamin Livshits. 2019. Speedreader: Reader mode made fast and private. In *The World Wide Web Conference*. 526–537.
- [23] Google. 2019. AMP is a web component framework to easily create user-first web experiences - amp.dev. <https://amp.dev>. Accessed: 2019-05-05.
- [24] GSMarena. 2023. Xiaomi Redmi Go. https://www.gsmarena.com/xiaomi_redmi_go-9534.php.
- [25] ITU. 2021. Facts and Figures 2021: 2.9 billion people still offline. <https://www.itu.int/hub/2021/11/facts-and-figures-2021-2-9-billion-people-still-offline/>.
- [26] Jameson Rader. 2023. AudioQR. <https://github.com/jamesonrader/AudioQR>.
- [27] Soonwon Ka, Tae Hyun Kim, Jae Yeol Ha, Sun Hong Lim, Su Cheol Shin, Jun Won Choi, Chulyoung Kwak, and Sunghyun Choi. 2016. Near-ultrasound communication for tv's 2nd screen services. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. 42–54.
- [28] Jesutofunmi Kupoluyi, Moumena Chaqfeh, Matteo Varvello, Russell Coke, Waleed Hashmi, Lakshmi Subramanian, and Yasir Zaki. 2022. Muzel: assessing the impact of JavaScript dead code elimination on mobile web performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 335–348.
- [29] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhooob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS 2019)*. <https://doi.org/10.14722/ndss.2019.23386>
- [30] Hyewon Lee, Tae Hyun Kim, Jun Won Choi, and Sunghyun Choi. 2015. Chirp signal-based aerial acoustic communication for smart devices. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2407–2415.
- [31] Ye Geoffrey Li and Gordon L Stuber. 2006. *Orthogonal frequency division multiplexing for wireless communications*. Springer Science & Business Media.
- [32] Ivano Malavolta, Kishan Nirghin, Gian Luca Scoccia, Simone Romano, Salvatore Lombardi, Giuseppe Scanniello, and Patricia Lago. 2023. JavaScript Dead Code Identification, Elimination, and Empirical Assessment. *IEEE Transactions on Software Engineering* (2023).
- [33] Marcin Kondej. 2023. Use the Raspberry Pi as an FM transmitter. Works on every Raspberry Pi board. https://github.com/markondej/fm_transmitter.
- [34] NextRadio. 2023. NextRadio Free Live FM Radio. https://play.google.com/store/apps/details?id=com.nextradioapp.nextradio&hl=en_US&gl=US.
- [35] Zhen Qin, Qingliang Zeng, Yixin Zong, and Fan Xu. 2021. Image inpainting based on deep learning: A review. *Displays* 69 (2021), 102028. <https://doi.org/10.1016/j.displa.2021.102028>
- [36] Quiet Modem Project. 2023. Android app built with Quiet which allows o pass data through the speakers of an Android device. <https://github.com/quiet/org.quietmodem/quiet>.
- [37] Quiet Modem Project. 2023. Quiet modem project. <https://github.com/quiet/quiet>.
- [38] Retekeess. 2024. TR508 FM Transmitter Low Power Broadcast Station for Drive-in Church School. <https://www.retekeess.com/tr508-fm-transmitter/>.
- [39] Nirupam Roy, Haiham Hassanieh, and Romit Roy Choudhury. 2017. Backdoor: Making microphones hear inaudible sounds. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 2–14.
- [40] Olivier Rukundo and Hanqiang Cao. 2012. Nearest Neighbor Value Interpolation. *International Journal of Advanced Computer Science and Applications* 3, 4 (2012). <https://doi.org/10.14569/IJACSA.2012.030405>
- [41] G Enrico Santagati and Tommaso Melodia. 2016. A software-defined ultrasonic networking framework for wearable devices. *IEEE/ACM Transactions on Networking* 25, 2 (2016), 960–973.
- [42] Aaron Schulman, Dave Levin, and Neil Spring. 2014. RevCast: Fast, private certificate revocation over FM radio. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 799–810.
- [43] Sonicwall. 2019. Wireless: SNR, RSSI and Noise basics of wireless troubleshooting. <https://www.sonicwall.com/support/knowledge-base/wireless-snr-rssi-and-noise-basics-of-wireless-troubleshooting/180314090744170/>.
- [44] Isidora Stanković. 2014. Recovery of Images with Missing Pixels using a Gradient Compressive Sensing Algorithm. *ArXiv abs/1407.3695* (2014). <https://api.semanticscholar.org/CorpusID:15997634>
- [45] Don Torrieri. 2005. *Principles of spread-spectrum communication systems*. Vol. 1. Springer.
- [46] United Nations 2023. Radio: The universal medium that leaves no one behind. <https://news.un.org/en/story/2022/02/1111882>.
- [47] Vladislav Veluga. 2019. Real FM Radio. <https://github.com/vladislav805/RFM-Radio>.
- [48] Bob Watson. 1980. FSK: signals and demodulation. *Watkins-Johnson Company Tech-notes* 7, 5 (1980).
- [49] William T Webb and Lajos Hanzo. 1994. *Modern Quadrature Amplitude Modulation: Principles and applications for fixed and wireless channels: one*. IEEE Press-John Wiley.
- [50] Wikipedia. 2023. MSN Direct. https://en.wikipedia.org/wiki/MSN_Direct.