

Design of a custom 5G and 4G Signal of Opportunity Navigation Radio Receiver

Aiden Morrison^{1,*†}, Nadezda Sokolova^{1,†}

¹SINTEF Digital, SCT Dept., Strindveien 4, Trondheim, Norway, 7034

Abstract

While off-the-shelf radios exist which allow exploitation of cellular communications signals, their size, weight power and cost metrics are often prohibitive when four or more independently tunable channels are needed. To avoid these limitations as well as those imposed by monolithic driver architectures which require the controlling software to dictate the behavior of all channels simultaneously the authors undertook to design a smaller, lighter and lower cost SDR platform with an independently abstractable driver interface. This paper describes the motivations behind, design process used to arrive at and resulting implementation of a multi-channel Signal of Opportunity (SoOP) radio optimized for exploitation of 5G and 4G signals between 600 and 5800 MHz for synchronization and navigation purposes and its associated driver model which allows independent channel control while maintaining synchronization between channels, while physically fitting into a minimal size, weight, power and cost (SWAPC) envelope.

Keywords

5G, SoOP, GNSS, Navigation, Synchronization

1. Introduction

1.1. Motivation

The motivations of the authors for implementing this design can be divided into two categories. The first of these is the accumulation of significant experience with Global Navigation Satellite System (GNSS) based navigation systems and their well known limitations regarding indoor reception [1], but also the growing family of problems caused by malicious actors jamming, spoofing, or otherwise degrading access to Medium Earth Orbit (MEO) satellite signals [2, 3]. The second category of motivations stems from the observation that as the 3GPP standards body drives improvements to the performance of cellular communications standards from LTE to 5G to 6G and beyond, that the synchronization and synchronization requirements to the frequency and time references used by the transmitters will necessarily improve resulting in ever more implicit utility for PNT exploitation of these signals even for passive listeners. Based on these motivations the many researchers globally [4, 5, 6] have been investigating the suitability of the existing transmissions for exploitation as timing references, and navigation references. While commercial off the shelf (COTS) radios can serve this purpose, their designs are often tuned to be extremely flexible in multiple dimensions such as sampling rate, center frequency, filtering options and while they may have reasonable costs for single channel receivers, the expense tends to rapidly escalate when higher capabilities or more channels are required. While the authors had experience with custom L-band receivers capable of operating in an RFI measurement context [2], it was quickly realized that the challenges associated with receiving cellular signals for SoOP purposes were significantly different from those of GNSS. Furthermore, in the context of lightweight drone applications, many of the COTS multi-channel solutions are far too heavy, resulting in a desire and a need for a lightweight, compact, moderate cost solution to allow many-channel cellular reception on a drone.

WIPHAL'26: Work-in-Progress in Hardware and Software for Location Computation October 5–8, 2026, Rome, Italy

*Corresponding author.

† These authors contributed equally.

✉ aiden.morrison@sintef.no (A. Morrison); nadia.sokolova@sintef.no (N. Sokolova)

id 0000-0002-0617-9959 (A. Morrison); 0009-0005-9079-1380 (N. Sokolova)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

1.2. Signal of Opportunity Design Challenges

One of the core differences between MEO satellite signals and those of cellular towers as received by a moving user in an urban or indoor environment is the extreme dynamic range of the received signal strength. While the portions of the L-band used by GNSS are typically dominated by relatively static thermal noise levels (when not impacted by RFI), the cellular signals vary over several orders of magnitude of signal strength even under normal use. A user operating close to a cellular tower could measure a level of -60 dBm Reference Signal Received Power (RSRP), while towards the coverage boundary of a given cell a level below -110dBm could be encountered. Simultaneously, depending on geographic location and local regulatory decisions an urban environment could contain downlink signals on several non-adjacent frequency bands either interspersed with uplink signals from nearby devices in intermingled bands in the case of Frequency Division Duplexing (FDD), or even within the same band as the downlink signals in the case of the 5G Time Division Duplexing (TDD). In order for a SoOP radio receiver to be viable in these conditions it must have a combination of very wide dynamic power range, a wide and precise tuning range, robust band isolation, a low noise floor for weak signal reception, a stable local oscillator, and ideally the ability to access several bands of interest simultaneously. While off-the-shelf solutions such as Universal Software Peripheral Radios can meet all or most of these requirements with appropriate add-ons, the SWAPC demands can become prohibitive. When considering a pedestrian mobile or compact drone borne platform there are no known off the shelf solutions which are feasible, particularly with regard to weight and cost when considering the likelihood of loss of the platform during testing. Further, it is often the case that both USRPs and other Software Defined Radio (SDR) front-ends can have unforeseen choices within their designs that limit their utility in the general or cellular SoOP use cases. For example, with some radio implementations the sampling rate, center frequency or both may have residuals that are either not communicated through the software stack to the user [7] or are not representable due to quantization errors in floating point arithmetic. In other cases, radios which are advertised as multiple channel devices may have constraints that prevent their channels from tuning to different frequencies, such as the Ettus B210 radio which does have two channels useful for 2x2 Multiple Input Multiple Output (MIMO) applications, the two receive channels are driven from the same mixer reference thus binding them to the same operating frequency and preventing their use to track independent signals over different bands. Further complications can arise from software stack considerations such as the concept of many multi-channel radios presenting through drivers as a monolithic device which may be controlled by only a single user process at a time. This restriction in turn requires that SoOP software be adapted to support the number of channels presented by a given device or that a translation layer be implemented. This is feasible, but becomes a hindrance when it is desirable to dynamically retune individual channels to other frequency bands and transmitter identities in response to higher level navigation logic as is the case for users that wish to operate dynamically beyond the range of a single set of transmitters. In response to these constraints and lacking the size weight and budget to support an X410 radio, the authors set out to rapidly design and implement our own SoOP research radio implementation and driver stack tuned for the specific challenge of cellular SoOP operation on a drone platform. As the target application was to provide non-GNSS navigation constraints to a drone named 'Icarus' the front-end was named 'Waxwing' to match both the mythology behind the name of the drone as well as the hubris of trying to design a custom radio to fit the application. The remainder of this paper is dedicated to the explicit design rationale and process and the achieved results.

1.3. The Target Signals

One of the factors which complicates radios used for cellular signal reception is the wide range of frequencies allocated to carry these signals and the fact that different network operators may elect to divide their frequency allocations differently in different regions and over different time periods. The band allocations targeted during the design of Waxwing are shown in Table 1 and Table 2 for the technologies of Frequency Division Duplexing (FDD) and Time Division Duplexing (TDD), respectively.

A further factor demanding flexibility in the context of TDD signals is the potential of the network operator to dynamically reallocate where within the TDD band allocation the exploitable synchronization symbols are transmitted, in turn requiring front-ends targeting these signals to be wideband and flexible while also providing high isolation between the bands of interest and adjacent signals. Regardless of active management attempts the TDD bands are expected to require higher dynamic range and nearband isolation in receivers compared to what is needed for use of the FDD bands.

Within each of these bands the Long Term Evolution (LTE aka 4G) and 5G New Radio (NR) signals may take on different characteristics depending on the potential subdivision of the band among local operators and consequent available bandwidth. While both standards families utilize variations of Orthogonal Frequency Division Multiplexing (OFDM), the Sub-Carrier Spacing (SCS) of 4G is typically limited to 15 kHz while 5G NR signals may use 15 kHz, 30 kHz, or 60 kHz depending on which of the bands in Table 1 and Table 2 are considered per [8, 9]. The importance of the SCS options in use is that it informs the amount of bandwidth needed by radios attempting to passively utilize the signals for SoOP purposes. In the case of 4G the Primary Synchronization Symbols (PSS) are transmitted in the central 64 adjacent carriers of the downlink band, while in the case of 5G NR the analogous synchronization symbols are transmitted on 128 adjacent carriers on a defined raster of possible positions within TDD or FDD downlink band. Multiplying the maximum of the possible SCS values (60 kHz) and applying this to a 5G NR signal with 128 adjacent carriers used for synchronization we arrive at a minimum required bandwidth of 7.68 MHz per channel. By supporting center frequencies from 700 to 6000 MHz we ensure the radio is applicable to any current known global implementation of 5G or LTE in Frequency Range 1 (FR 1).

Table 1

Targeted Current FDD Bands.

Band Name*	Uplink (MHz)	Downlink (MHz)
N28	703-733	758-788
N20	832-862	791-821
N8	876-915	921-960
N3	1715-1785	1805-1880
N1	1940-1980	2130-2170
N7	2500-2570	2620-2690

* We have intentionally omitted the allocated but underutilized N31 band near 465 MHz.

Table 2

Targeted Current TDD Bands.

Band Name	Link Frequency (MHz)
N38	2570-2620
N78	3400-3800

2. RF Signal Handling

Because of the expected wide variability in the input power level of the signals of interest, along with the expectation that undesirable (from the point of view of SoOP use) signals will be present adjacent to the desired synchronization components, the design of the RF signal handling sections of the front-end had to have both flexible gain per band, high selectivity, and ideally also a low noise figure. To attain the low noise figure while also maintaining the ability to survive operation at high input power levels a Low Noise Amplifier (LNA) was selected that sacrificed frequency coverage below the selected limit of 700 MHz in exchange for a low noise figure and high input power rating. The specific model selected offered a Noise Figure (NF) of 0.35 dB near the mid-band of the signals of interest at 1.9 GHz, while also purporting a maximum input power of +19 dBm. Comparing these parameters with those of an X410 USRP, they represent a 4.5 dB advantage in 1 dB compression point (P1dB) compared to the -15 dBm rating of a USRP X410 (potentially 34 dB higher survivable input power if the maximum specified is treated as an absolute maximum), and a greater than 7 dB advantage in noise figure compared to the 8 dB NF of the X410. In short, for the portion of the spectrum containing the cellular signals indicated, the device is tuned to provide a much more sensitive yet much more power tolerant radio than is available from the very flexible USRP devices.

Per Figure 1, subsequent gain stages were interleaved with RF power dividers and ceramic High-Pass Filter (HPF) anti-aliasing stages to minimize degradation of the NF relative to the first stage while avoiding the saturation issues posed by directly cascading the gain stages prior to the passive signal handling components. In the prototype of Waxwing one of the four branches of the input gain stage tree was terminated prior to the Ceramic HPF to support the six populated channels.




Figure 1: The input gain stage of the Waxwing front-end.

Per Figure 2, each of the populated channels used an independently synthesized reference tone and low-side mixing to translate the input signal of interest to 284.16 MHz where a Surface Acoustic Wave (SAW) Band Pass Filter (BPF) with approximately 8 MHz of pass-band bandwidth was used to provide adjacent band rejection. Each of the populated channels then passed the generated signal through an independently controlled programmable attenuator prior to a last stage fixed amplifier. The combination of 31.25 dB of digitally programmable attenuation within each channel with approximately 9 dB of digital amplitude control of the synthesized tone sent to the individual channel mixers provided approximately 40 dB of digital gain feedback authority.




Figure 2: RF signal handling within each of the six implemented channels.

The remainder of the dynamic range of the Waxwing is attributable to the use of a 12 bit Analog to Digital Converter (ADC) whose rated 11.3 Effective Number of Bits (ENOB) converts to a Spurious Free Dynamic Range (SFDR) of 69.8 dB using Equation 1. This SFDR combined with the 31.25 dB of adjustable attenuation range and 9 dB of mixer reference tone power range provide a total of just over 110 dB of dynamic range to the system.

$$\text{SFDR [dB]} = 6.02 \cdot \text{ENOB} + 1.76 \quad (1)$$

The selection of the 284.16 MHz Intermediate Frequency (IF) operating target was made to enable the digitization of the signal to exploit sub-Nyquist sampling techniques whereby a sampling clock of 30.72 MHz would alias the 9th Nyquist zone around the SAW filter center frequency of 284.16 MHz directly to the band center of 7.68 MHz of the real valued baseband sampled signal spanning DC to 15.36 MHz. The primary motivation behind this choice was to map the IF to a portion of the spectrum where undesired artifacts from the synthesizer, the mixer, and other locally generated frequencies would fall outside of the pass-band of the SAW filter and thereby be nearly eliminated prior to the final gain stage as is illustrated in Figure 3.

In Figure 3 the nuisance components include the DC offset of the ADC, the reference tone fed to each of the synthesizers at six times the sampling frequency or 184.32 MHz, the orange polygon representing the pass band of the SAW filter, and a higher frequency mixing tone used to bring the input signal of interest to the center of the filter pass band is shown in yellow. For the defined 700 MHz minimum tuning frequency the minimum mixer frequency corresponds to 415.84 MHz resulting in a wide separation and therefore high attenuation of the f_{mix} leakage component from the desired signal in the SAW filter pass-band. The use of a reference tone to the synthesizers of 184.32 MHz ensures that any leakage passing through the synthesizers, and mixers are intentionally mapped to a point in the




Figure 3: RF signal handling within each of the six implemented channels showing illustrative signal amplitudes versus relevant frequencies. The orange polygon indicates the pass band of the installed SAW filter, with locally generated tones of concern shown in purple or yellow and DC offset error in red.

spectrum where they are both significantly eliminated by the SAW filter or can be eliminated through basic Digital Signal Processing (DSP). The apparent content of this signal assuming that the mixing tone is completely eliminated but that the intrinsic error of the ADC and a small leakage component of the 184.32 MHz tone remain is illustrated in Figure 4.




Figure 4: The real valued baseband sample stream from the ADC contains primarily the desired aliased signal from the SAW filter at the band center and both a bias error component intrinsic to the ADC and a residual leakage of the mixer reference tone superimposed at DC.

A final motivation for the selection of these frequencies is that in combination with the selection of an appropriate synthesizer for mixing tone generation the 184.32 MHz reference used to drive the synthesizer allows it to produce a raster of selectable frequencies of exactly 1 Hz over the entirety of the target frequency range from under 700 MHz to over 5.75 GHz thereby fulfilling the objective of precise tunability over the operating range. The DSP in this front-end realization is limited to the combined operation of conversion of the signal from a real valued stream to a complex baseband stream which has been frequency translated to map the band center at $\frac{1}{4}$ of the sampling frequency to complex baseband. This is illustrated in Figure 5 where the complex baseband representation of the signal after conversion and frequency translation but prior to digital filtering and decimation are shown.




Figure 5: Spectrum contents and aliasing after the real to complex conversion and $\frac{1}{4}$ rate mixing translation but before digital filtering. Beneficially the subsequent application of the quarter band filter indicated by the green outline can substantially eliminate both the aliased components as well as the ADC DC bias and residual reference tone leakage.

This mixing and filtering architecture has the benefit of being implementable with very limited resources within the Field Programmable Gate Array (FPGA) as the utilized quarter-band filter structure benefits from both symmetric filter coefficients and alternating zero values of those coefficients allowing a much higher order filtering operation to be implemented with a small amount of DSP and logic resources.

2.1. Local Clocking

As with any Time of Flight (ToF) or Doppler shift measurement system, instability in the local clock reference is undesirable as it manifests in the generated observables as a bias term that must be compensated for. Ignoring all other constraints we would select the most stable oscillator available to drive the signal handling and capture processes. Unfortunately, due to atomic clocks driving SWAPC by themselves costing more than the remainder of the system hardware, the Waxwing system compromises through the use of a thermally stabilized or ‘Ovenized’ crystal oscillator (OCXO). A further parameter of concern when implementing frequency synthesis for frequency translation or sampling clock generation is the period to period instability or ‘jitter’ of the generated clock. In concept any jitter introduced into the mixing clock will be superimposed on the down-converted signal of interest, while jitter in the sampling clock will serve to degrade the SNR of the signal per [10]. The relationship between an achievable SNR versus the highest frequency component f_{in} of the input signal and the timing jitter T_{jitter} of the sampling aperture is given in Equation 2.

$$SNR_{Jitter} [dBc] = -20 \log (2 \cdot \pi \cdot f_{in} \cdot T_{jitter}) \quad (2)$$

In clock synthesizers one of the primary sources of jitter derives from the need to support a wide range of input reference frequencies through the implementation of a high-bandwidth clock input, effectively trading the flexibility of input frequency range for instability in the output. To work around this source of jitter, the clocking system on the Waxwing uses a ‘jitter attenuator’ device as indicated in Figure 6. where the 10 MHz OCXO signal is effectively used as a low frequency or equivalently a long-period reference signal as interrogated by a tracking loop operating at a known frequency (here 54 MHz) via a narrow tracking bandwidth.




Figure 6: Generation of low-jitter sampling and reference clocks via a dedicated Jitter attenuator and synthesizer chip.

This configuration allows the produced synthesizer reference tones and the sampling clock to benefit from a narrow bandwidth low jitter tracking loop driven by the 54 MHz reference with very low cycle to cycle jitter of 72 femtoseconds over 12 kHz to 20 MHz while the underlying OCXO itself has an instability of 350 femtoseconds on this same bandwidth at the 10 MHz input.

3. Digital Signal Processing and Data Handling

Once digitized, the captured samples are streamed from the ADC to the system FPGA as real valued quantities with 14-bits per sample. While the ADC is expected to produce only 11.3 ENOB, the wider representation is retained and extended through the initial mixing and real to complex value conversion process as well as the quarter-band filtering to avoid loss of precision. For RF systems with wide dynamic power ranges, a feedback mechanism to best use the available quantization depth is mandatory and described as an Automatic Gain Control or AGC. In the operating context of cellular communications signals, the implementation of the AGC measurement and feedback system must take into account the intrinsically intermittent nature of the signal of interest, as well as the potential for rapidly varying power levels. In the case of the TDD multiplexed signals for example, the same portion of the band

which contains or is adjacent to the SoOP exploitable synchronization symbols may contain uplink transmissions from User Equipment (UE) that could be within meters of the SoOP receiver. While the peak transmission power of UEs is typically limited to 2 Watts for mobile phones, the very short distances can cause these signals to appear at much higher power levels than the tower downlink. While protocol awareness and dynamic interpretation of traffic in the band would allow fine-tuned adjustment of gain and gating of reception this adds significant systemic complexity. As a compromise the AGC was implemented to take into account the bit utilization maximums over 100 ms windows of time. By observing over this interval of time the intent is to ensure that at least one set of 5G synchronization symbols from any towers that may be present are observed and considered, which is important when bands are idle. Simultaneously, the AGC adjusts the gain of the receiver to attempt to maintain the bit utilization at -12 dB Full Scale (dBFS) power or stated otherwise to use 10 of the available 12 bits of dynamic range to leave a margin for sudden increases in received signal power without significantly degrading the captured signal. The feedback process is independent per reception channel and designed to quickly reduce gain when above of the target range but to slowly increase it when below the target amplitude range.

As the filtered samples are generated they are packed via a Intellectual Property (IP) block within the FPGA to a custom format which embeds a header into the sample stream at the top of every 1 Megabyte (MiB) of total data. This arrangement is shown in Figure 7. Where the 256 Byte 'total' header contains a synchronization pattern used to aid in aligning the host driver to the packet frame structure and all relevant information about the configuration and operating state of the front-end and each of its channels. While overall data can include aspects such as the ordinal sample number at the start of each packet, flags for events, and system status blocks the channel specific data defines critical parameters such as the operating frequency, AGC feedback state and bit utilization in addition to mixer lock flags critical to determine when a commanded frequency change operation has completed when changing band settings of a single channel while under operation. The packet structure is a trade off in that it first truncates the width of samples to a 12 bit representation for each of the I and Q branches (3 Bytes per IQ pair) in order to fit within the transport bandwidth limitations through the selected USB3 interface FIFO device to the host machine and commensurate with the expected ENOB provided by the front-end.




Figure 7: Packet format of front-end to host machine over USB3 link.

Once the packetized data has passed over the constrained USB3 to host channel the lower level driver software unpacks and reformats the data into independent streams of complex int16 samples for each operational device hardware channel as illustrated in Figure 8.




Figure 8: Packet format on ØMQ transport layer between low level driver and SoapySDR driver showing per-channel repacking of data.

4. Software Interface Design

The typical approach of writing a driver for an SDR is to treat the radio as a monolithic device with each individual channel exposed within this entity. This approach was considered undesirable for

the use cases of Waxwing as in the SoOP context we want to be able to command different channels independently of each other such that different SDR implementations can treat their assigned channels as independently tunable and controllable radios yet benefit from the coherent sampling and synchronized nature of the multi-channel radio. To achieve this, the software interface was split into three sub-components. The interface between the Waxwing FPGA and the host machine is via a USB3 First In First Out (FIFO) IC and its dedicated low level driver. A multi-threaded driver layer was implemented to separate the latency sensitive streaming process from the device from the sample decoding/unpacking process, the streaming workers that handle the further transmission of each individual channel, and the passage of returned configuration update requests to any channel of the device. To respect latency limitations and prevent loss of samples, return messages from the individual channels are coalesced by the high priority interface thread and only sent to the radio in the activity gap immediately following the reception of a packet from the radio. The second layer of the interface utilizes ZeroMQ(ØMQ) Inter-Process Communication (IPC) bind points to present each of the eight logical channels of the radio within the host system and provides a transition between the low level Waxwing device specific driver to a more abstract and flexible interface in the third layer. This is currently implemented for local operation on a single machine but is feasibly adapted to network operation if sufficient network bandwidth is available to support the sample streams. The third layer of the Waxwing software interface implements a SoapySDR driver. Per the website of the creators of this software “*The SoapySDR project is an open-source generalized API and runtime library for interfacing with SDR devices*” [11] which is a solution to the problem of having a proliferation of SDR devices and driver frameworks which can require SDR software to understand and implement the interface layers for each and every radio type they wish to support. By providing this compatibility layer the creators of SoapySDR simplify the SDR ecosystem by allowing SDR authors to implement a single driver model yet allow control over e.g. Lime SDR, HackRF, BladeRF, SDRplay, iio, airspy, rtlsdr, uhd (USRPs), and more models of SDR front-ends. Upon the connection to and initialization of any individual channel of the Waxwing via the SoapySDR Application Programming Interface (API) functions the radio can be configured to set channel parameters for the connected channel and to start streaming of data. If the radio is not yet running, this triggers the low level driver to commence the streaming process from all channels into a local circular buffer. Since other channels may or may not connect at any later time, the behavior of this buffer is to reset newly connecting channels to the current ‘top’ of the buffer to ensure that only fresh data is provided to each channel. In the event of an overflow caused by e.g. an incorrect disconnection, only the pointer of the impacted channel is reset, allowing arbitrary connection and disconnection on channels to not impact the operation of other channels. Further, API calls to alter gain levels, to change center frequency or to otherwise modify the state of any of the individual channels are executed transparently and without impacting the operation of the other hardware channels, allowing the SDR software instances connecting to each individual channel to operate without concern for the state of the other SDRs that may be using the Waxwing yet to benefit from the synchronized nature of the sample streams sent to each channel. Because the Waxwing prototype has a relatively stable oscillator onboard but not an absolute time scale, the process used for establishing a sample epoch on an absolute time scale is to note the current time at the start of streaming on whichever channel is ‘first’ to request data from the low level driver. As long as at least one channel remains active, the driver will remember this start epoch of the host machine time, and provides downstream propagated time stamps by referencing the ordinal sample number of the start of the current packet being sent to a channel to that start epoch by the nominal sampling rate. This process allows the Waxwing driver to both make use of the coarse synchronization of the host machine’s local clock to provide an absolute time scale, but also to ensure that the coherent sampling across channels is maintained such that every channel regardless of the number of connections and disconnections will continue to refer all collected data to a common time scale. A natural consequence of this approach is that the sample times of the active channels will stay locked to one another and benefit from the relatively low frequency error of the onboard OCXO while the host machine time may otherwise drift.

5. Hardware Realization

To limit risk and cost during initial development the system was dimensioned both physically and electrically for eight independent channels, but only six were populated on the prototype, with the unused area allocated to debug and test facilities such as secondary reference frequency outputs and probe test points. The partially complete prototype is shown in Figure 9.



Figure 9: Installation of the USB3 FIFO chip from a donor board during the initial hexachannel prototype implementation of Waxwing in August 2025. Dimensions of the system excluding the RF connector are 131 by 103 mm while mass was less than 300 g with an RF shielding frame installed.

6. Achieved Functionality

6.1. Raw Signals

Once implemented, the Waxwing front-end prototype successfully produced a continuous stream of six independently controllable sample streams tailored for capture and exploitation of 4G and 5G NR signals for SoOP purposes. An example of one of these signals captured in October 2025 in an indoor residential environment in Trondheim Norway is shown in Figure 10 where an LTE cell and data traffic were captured by channel ‘B’ of the front-end.




Figure 10: Waxwing received LTE Cell ID 294 broadcasting its FDD downlink on 806 MHz over 512 FFT bins. The central frequency bins contain the periodic primary and secondary synchronization symbols while OFDM downlink data of various Resource Block configurations is also apparent.

Simultaneously during this data capture, each of the other five Waxwing prototype channels independently captured five other FDD downlink transmissions on 796 MHz, 1830 MHz, 1850 MHz, 2120 MHz and 2160 MHz.

6.2. SDR Use

Through careful design of the Waxwing system drivers the system is capable of providing synchronized yet independent signal streams to six virtual Soapy SDR devices fed by the underlying Waxwing hardware per Figure 11 where the bidirectional exchange of data and configuration settings are shown.




Figure 11: The Waxwing system design discussed in this paper comprises the blue elements of the Hardware, the low level driver, and a SoapySDR driver abstraction layer along with their interfaces. The elements in green are the external off the shelf passive antenna and external SDR software packages.

This implementation allows individual and potentially heterogeneous SDR components to use each of the channels of the Waxwing to operate independently on their assigned channel without disruption to other instances of the same or other SDR utilizing other channels of the physical hardware. This in turn allows for a higher level navigation ‘supervisor’ program to control the individual instances of the SDR processes to drive acquisition and tracking processes as will be needed in highly dynamic platforms that cannot be told which sources to track a-priori.

6.3. Tracking Results

For the purposes of tracking cellular 4G and 5G signals the authors have utilized the STARE Software Defined Radio available from the European Space Software Repository [12, 13]. With the Waxwing front-end prototype connected to a ‘rubber duck’ style passive antenna indoors in a residential building, it was possible to track a different cellular downlink signal in each of the six hardware channels. When converted to Doppler measurements at 1 kHz rate a number of interesting properties of the data can be seen in Figure 12.




Figure 12: Tracking Doppler results generated via six parallel instances of STARE SDR fed by data from the Waxwing receiver.

The first interesting result is that the cell tracked on 2120 MHz contains a notable amount of high-frequency noise compared to the tracking results of the other transmitters, which the authors believe arises from different makes, models and vendors of base station equipment in use along with differing clock and frequency control systems and algorithms in play. A second observation is that the Doppler curves for the four highest frequency cells on 2120, 1850 and 1830 MHz are not in the order expected if the frequency tracking error was driven only by the Waxwing local OCXO. To isolate this factor we

next convert the results from frequency to meters per second in Figure 13 to expose the nature of the frequency errors of the transmitters.




Figure 13: Range rate results in meters per second after low pass filtering to 50 Hz bandwidth.

When converted to units of meters per second it suddenly becomes clear that while the Waxwing internal OCXO is a dominant source of range rate bias of approximately 25 meters per second, there is a residual shared by two distinct groupings of transmitters. The towers at 1830 MHz and 2140 MHz appear to share a common clock rate error while those on the other four transmitters share a different common clock error. This grouping indicates that either network operators or physical cell sites have common mode clock steering resulting in a situation analogous to constellation dependent clock offsets in GNSS though here in the clock rate errors. Beneficially this indicates that in environments where this holds true that error states may be shared between observables if it is known a-priori or determined online which reference clock a given tower is steered to or by. Past analysis using static receivers with GNSS disciplined oscillators indicates that the instability in locally available transmitters tends to be commensurate with centimeter per second line of sight velocity uncertainties over periods of milliseconds to hundreds of seconds [14] making the onboard OCXO of the Waxwing the primary source of Doppler residual that must be estimated, though with the ability to track several simultaneous signal sources with centimetric range rate uncertainty the online estimation and removal of the receiver clock error state should be feasible.

7. Implemented Evolutions

As the Waxwing prototype proved the implemented concepts were viable, the lessons learned were incorporated into a 'V1' full implementation which builds on the prototype shown in this paper. In addition to the population of all 8 channels a number of additional adjustments and improvements were made. The first of these desired features was provision for external event input such as an event marker from an external system which will allow an external event to be time-stamped in terms of the ordinal sample epoch of its arrival relative to the free-running sampling clock of the Waxwing. Further, the jitter attenuator's ability to prioritize and autonomously select between multiple input clock sources was exploited to add support for an optional 10 MHz input signal to allow the system to be driven by a precise external frequency reference. The overall intent behind the external event reference and optional 10 MHz input is to allow the Waxwing sample stream to be precisely synchronized to an external time scale such as a UTC realization driven by a GNSS disciplined oscillator. The authors now plan to use the designed system to engage in further activities related to GNSS denied and indoor SoOP navigation research and development applications for platforms such as drones where the size and weight of multi-band USRPs would otherwise be prohibitive.

8. Acknowledgments

This work was partially supported by the Norwegian Council of Research (project number 344275). The authors also acknowledge in-kind support from SINTEF Digital, including access to lab facilities and second HW revision development.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

- [1] P. W. Ward, J. W. Betz, C. J. Hegarty, Satellite signal acquisition, tracking, and data demodulation, in: E. D. Kaplan, C. J. Hegarty (Eds.), *Understanding GPS: Principles and Applications*, 2 ed., Artech House, 2006, pp. 237–238. Section 5.13.
- [2] N. Sokolova, A. Morrison, A. Diez, Characterization of the GNSS RFI Threat to DFMC GBAS Signal Bands, *Sensors* 22 (2022) 8587. doi:10.3390/s22228587.
- [3] A. Morrison, N. Sokolova, J. E. Håkegård, T. H. Bryne, L. Ruotsalainen, A multi-site quad-band radio frequency interference monitoring alerting and reporting system, in: *Proceedings of the European Navigation Conference (ENC)*, 2020. doi:10.23919/ENC48637.2020.9317522.
- [4] Z. M. Kassas, V. Ghadiok, T. E. Humphreys, Adaptive estimation of signals of opportunity, in: *Proceedings of the 27th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2014)*, 2014, pp. 1679–1689.
- [5] Z. M. Komodromos, T. E. Humphreys, Signal parameter estimation and demodulation of the OneWeb ku-band downlink, *npj Wireless Technology* 2 (2026) 7. doi:10.1038/s44459-025-00013-y.
- [6] W. Barrett, S. Kozhaya, Z. Kassas, D. Marsh, Navigating the arctic circle with Starlink and OneWeb LEO satellites, in: *Proceedings of MILCOM 2025 IEEE Military Communications Conference*, 2025, pp. 1–6. doi:10.1109/MILCOM64451.2025.11310048.
- [7] C. O'Driscoll, J. T. Curran, Carrier phase tracking considerations for commodity SDR hardware, in: *Proceedings of the 31st International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2018)*, 2018, pp. 4182–4196. doi:10.33012/2018.16117.
- [8] 3rd Generation Partnership Project (3GPP), NR; User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone, Technical Specification TS 38.101-1 V19.4.0, 3GPP, 2025. Table 7.3.2-1a.
- [9] J. Flores de Valgas, J. F. Monserrat, A. Hüseyin, Flexible numerology in 5G NR: interference quantification and proper selection depending on the scenario, *Mobile Information Systems* 2021 (2021) 1–9. doi:10.1155/2021/6651326.
- [10] T. Neu, Jitter vs SNR for ADCs, Texas Instruments Precision Labs Report 4704, Texas Instruments, 2017. URL: <https://www.ti.com/video/series/precision-labs.html>.
- [11] Pothosware, SoapySDR documentation and application programming interface description, n.d. URL: <https://github.com/pothosware/SoapySDR/wiki>.
- [12] I. Lapin, G. Seco-Granados, O. Renaudin, F. Zanier, L. Ries, Joint delay and phase discriminator based on ESPRIT for 5G NR positioning, *IEEE Access* 9 (2021) 126550–126563. doi:10.1109/ACCESS.2021.3111759.
- [13] European Space Agency, Stare v. 1.4.0, n.d. URL: <https://essr.esa.int/project/stare>, accessed version 3.2, 2025.
- [14] A. Winter, A. Morrison, N. Sokolova, Analysis of 5G and LTE signals for opportunistic navigation and time holdover, *Sensors* 24 (2024) 213. doi:10.3390/s24010213.