MORE DOLLARS, MORE SWaP – THE RIGHT TRADE-OFF?
Picking the Right SDR Architecture for the Task

Software defined radio (SDR) design is a constant tradeoff between form and function, or the desired performance attributes against what we have to pay for them in terms of size, weight, power and cost (SWaP-C). The typical list of radio desired attributes can include:

- RF performance such as noise figure (NF) or linearity (IP3)
- Number of receive (Rx) or transmit (Tx) channels
- Instantaneous bandwidth (IBW) and frequency coverage
- Integrated processing capability
- High speed interfaces supported

Optimizing any or all of these characteristics is almost always at the detriment of SWaP-C. Beyond obvious trade-offs, for example adding channels increasing power consumption, there are other more subtle decisions to be made depending upon the target application. To get an idea of some, we will use two Epiq products that seem very similar on the surface, but underneath are designed very differently and offer very different performance and SWaP-C operating points.

Comparing the NDR318 and NV800 SDRs

For this exercise we will use two 8-channel products with banner specifications that appear almost identical, shown in **Table 1**.

Specification	NDR318 	NV800 
# of Receivers	8	8
Max RF Frequency	6 GHz	6 GHz
Maximum BW	40 MHz	50 MHz
Typical Rx NF	10 dB	7 dB
Typical Rx IP3	0 dBm	0 dBm
High Speed Data IO	VITA 49 Over 2x 10 GbE	VITA 49 Over 2x 10 GbE

Table 1: *RF specification comparison*

However, they also diverge in some areas, as shown in **Table 2**.

Specification	NDR318 	NV800 
# of Transmitters	0	1
Dimensions	12″ x 8″ x 1.9″	9.92″ x 7.24″ x 2″
Volume	182.4 in^3	143.6 in^3
Weight	6.5 lbs	3.7 lbs
Typical Power Consumption	47 W	25 W
Cost	$$$$	$$

Table 2: *SWaP-C comparison*

What is so different between the two SDRs that would drive the NDR318 to be at first glance so similar to the NV800 for performance, but be ***nearly 2x higher across the board for SWaP-C***? The answer is architecture.

Sidekiq NV800 – Zero IF Architecture & Amplifier-First Design

The NV800 uses a radio architecture called Zero IF, or ZIF. This architecture has seen incredible advances over the last two decades, driven by cellular and Wi-Fi commercial technology investment. Specifically, the NV800 uses four of the Analog Devices ADRV9004 ZIF devices in its architecture. These ZIF devices are fully integrated Rx and Tx subsystems, complete with mixers, filters, amplifiers, data converters, and digital signal processing. That high level of integration significantly reduces SWaP-C, and thanks to a decade plus of investment and the maturation from 3G to 4G to 5G, the ADRV9004 has surprisingly high performance.

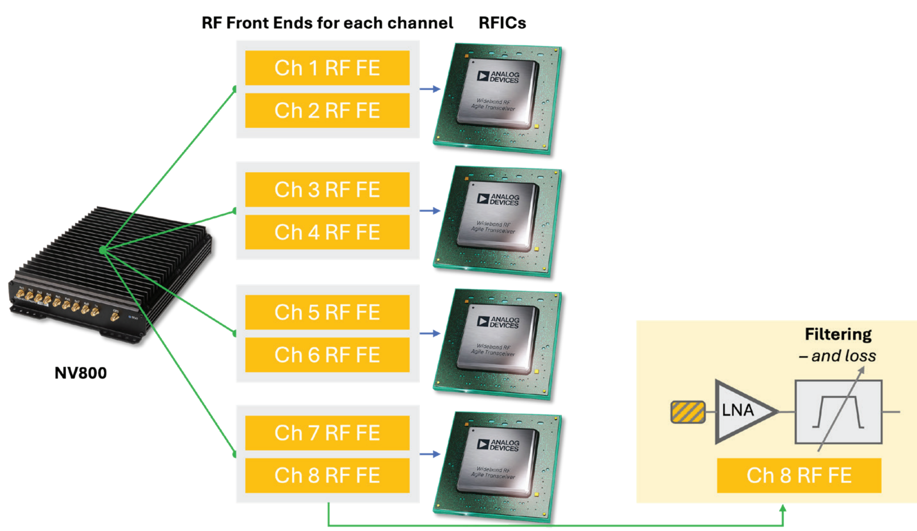

Figure 1: *A representation of the NV800 RF Front End (RFFE)*

Additionally, the NV800 RF front end design uses an "amplifier-first" design approach (**Figure 1**), meaning that the first major signal chain component that the received signal is exposed to is an amplifier – specifically a low noise amplifier. This architecture can minimize noise figure, and the NV800 has a slightly better NF than the NDR318, as shown in the table above.

Amplifier-First Design:

- The low-noise amplifier (LNA) boosts the signal before it undergoes the loss of the filter section. This is the main reason why the NV800 has better noise figure than the NDR318.
- It also enables better power consumption – if signal is lost early in the signal chain, designers have to spend power later to get it back.

However, the news is not all good - LNAs are non-linear, and big signals can create subtle challenges as we'll see later.

NDR318 – Super Heterodyne Architecture & Filter-First Design

The NDR318 does not use ZIF devices, it is a discrete super heterodyne (superhet) architecture (**Figure 2**), meaning that the entire signal chain is composed of individual components that have been selected to optimize performance. A superhet uses a multi-stage IF architecture, which inherently demands more components and thus more cost and power than the zero IFs that are in a ZIF design. The end result is that by selecting this architecture, the NDR318 designer has selected a path that requires more components, which is bad for SWaP-C, but allows for very detailed performance optimization across the entire signal path.

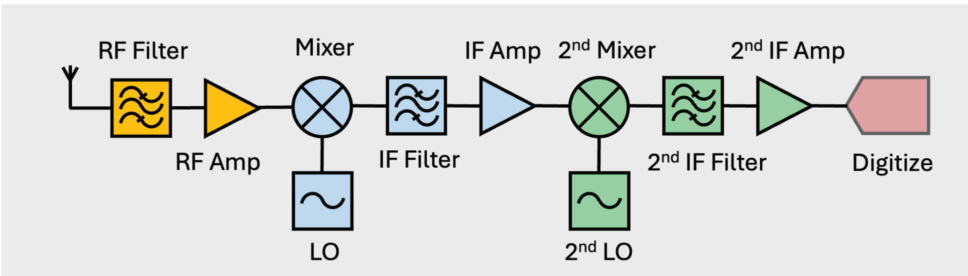

***Figure 2**: A generic super-het architecture such as is used in the NDR318*

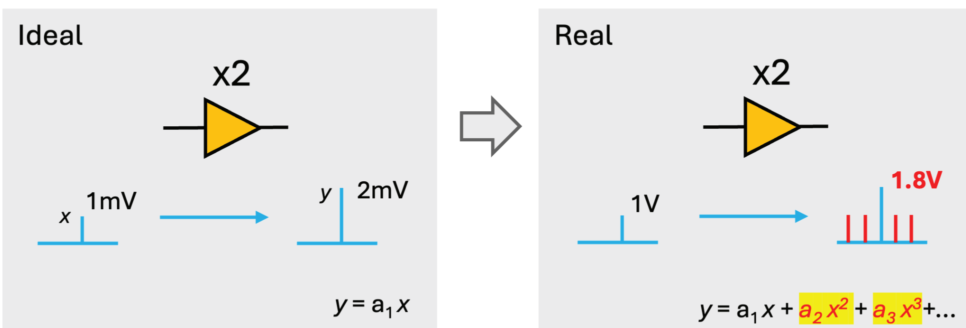

***Figure 3**: An ideal amplifier only provides gain (a). However, real amplifiers introduce unwanted spurious signals, which get worse as signals get large.*

The NDR318 also uses a filter-first approach, meaning that instead of the signal first hitting an amplifier in the signal chain, it goes through a filter, specifically a sub-octave pre-selection filter. There are many benefits to this approach; the largest relates to the fact that amplifiers are nonlinear devices (unlike filters) – they create harmonics and other spurious as the signal passes through them (**Figure 3**).

By having a filter-first approach the NDR318 designer has been able to remove all undesired signals from the signal chain before the signal is exposed to that first non-linear device. This will allow for much greater immunity in the presence of interference and jamming signals, which will be explored more in the following sections.

Performance Nuances

One challenge of looking at a specification table for complex products like the NDR318 and NV800 is that it is hard to fully capture all the ways the radio will behave in the real world when relying upon laboratory-based measurements. There are many ways this could be explored, but for the comparison of these two products there are two key nuances to examine: **quadrature correction** and **out of band signal rejection**.

Quadrature Correction

All modern SDRs do their digital signal processing on quadrature I and Q data, which represents the real RF signal received or transmitted at the antenna (**Figure 4**).

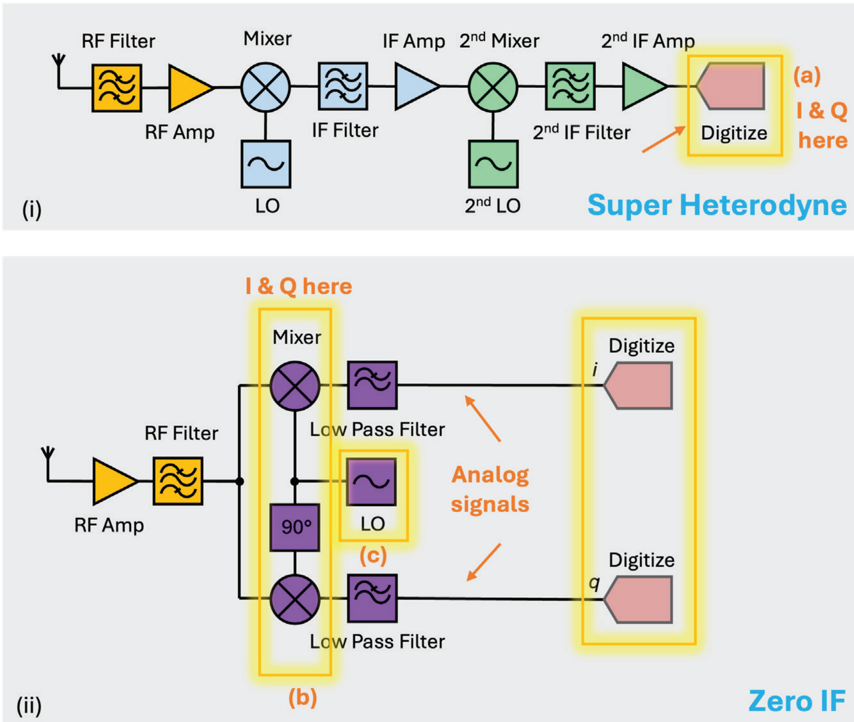

Figure 4: Comparing superhet (i) and ZIF architectures (ii) for quadrature error and DC offset

Superhet architectures split the signal into I & Q components only after it is digitized using a digital down converter (DDC) (a). As digital signals don't suffer from impairments in the same way as analog microwave signals, quadrature correction tends to be a non-issue for this architecture. However, Zero IF architectures split the analog RF signal in the down conversion mixer, then digitize I and Q independently (b). In this case, the two microwave signals must be perfectly in quadrature until digitized – despite any slight differences in components in each path. For devices like the ADRV9004, clever calibrations and algorithms inside the chip do an impressive, but not perfect, job of correcting them, leading to small spurious signals.

DC Offset

Zero IF architectures can also suffer from LO leakage (c), with small amounts of unwanted energy mixing down to 0 Hz, which puts these undesirable artifacts down where the wanted signals are and making them very difficult to filter out. Superhets don't have signals at 0 Hz, so this tends to be a non-issue.

Two-Tone IP3 Testing

A classic method of exploring spurious-free dynamic range is to perform two-tone third order intercept point (IIP3) testing, as shown in **Figure 5**. For our example, two clean microwave signal generators were set to output sine waves around 1.6 GHz and 4 MHz apart from each other before being fed into each radio in turn.

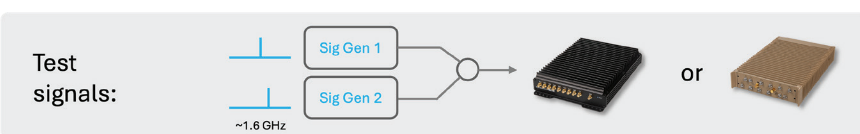

Figure 5: Two-tone IIP3 testing setup

Results are shown in **Figure 6**, with a key showing color-coded identification of spurious products. Overall, the NDR318 exhibits cleaner spur performance floor. The NV800 shows the expected ZIF defects discussed above, but at impressively low levels; older ZIF designs might not always perform quite so well.

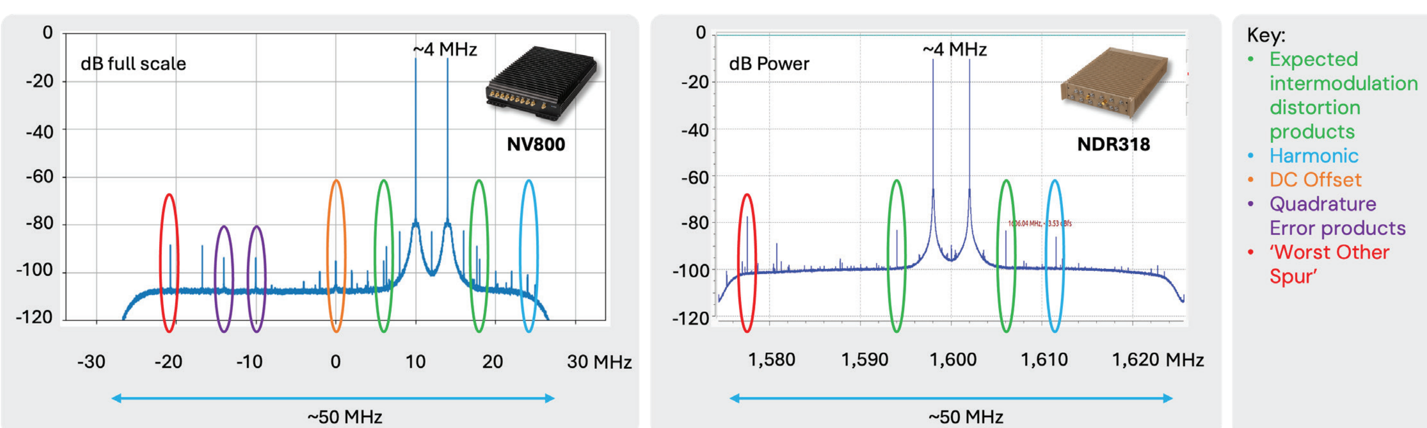

*Figure 6: Test results for the NV800 (**left**) and NDR318 (**right**) with input signals set to -10 dB_{FS}*

Out-of-Band Rejection

In a datasheet the RF specifications are typically measured within the desired frequency band of interest. For example, the two plots shown for the NDR318 and NV800 in the prior section placed two -10 dB_{FS} tones inside the receivers instantaneous bandwidth. Outside of those two tones the only other signals are those which the radio generates itself through nonlinearity and quadrature error.

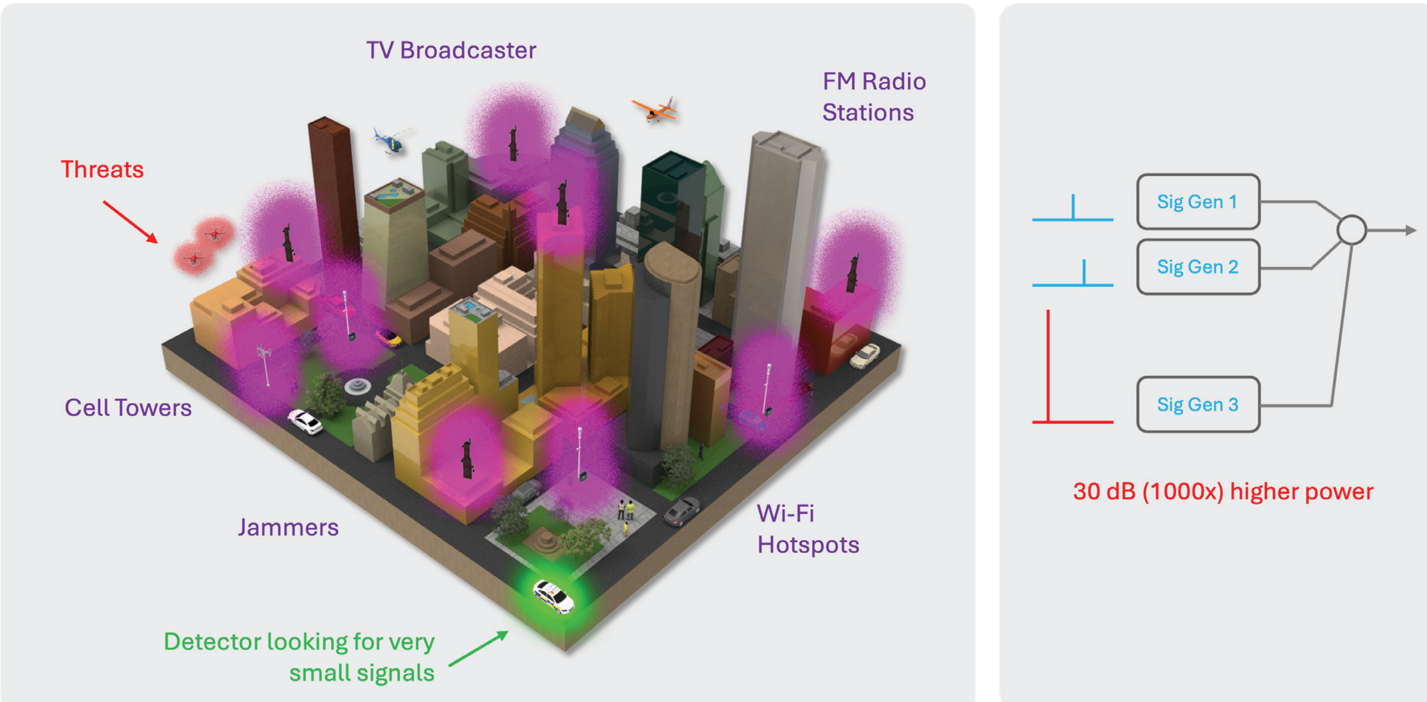

***Figure 7**: An example of using an SDR to detect drone threats within a congested urban environment (**left**) and an updated test setup (**right**)*

However, in the real world it is very rare that only the desired signal of interest is present at the antenna. Often there are dozens or hundreds of undesired signals in addition to the signal of interest (an example of UAS threat detection is shown in **Figure 7**). Those signals can be very large such as TV or FM radio broadcast or can be intentional to disrupt radio performance such as jamming signals. The IP3 number in a datasheet gives the user an indication of how well the receiver will hold up to these strong out-of-band signals.

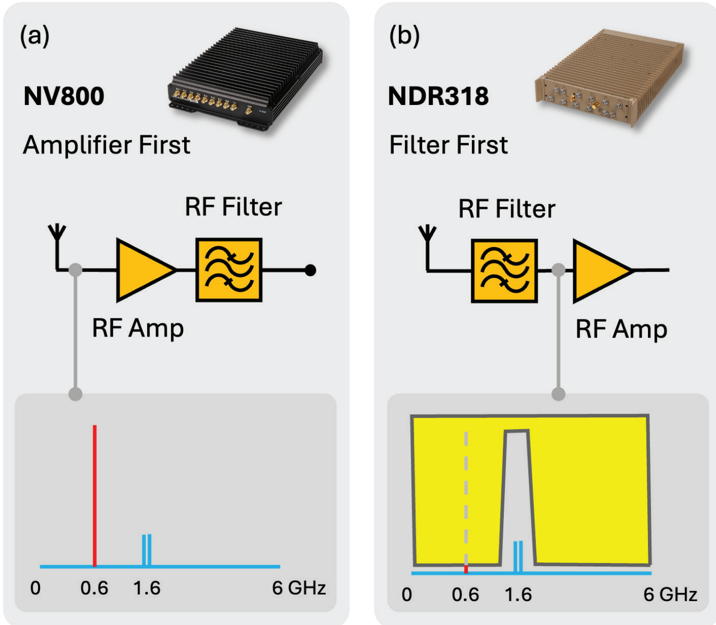

***Figure 8**: A strong interfering signal reaches the input amplifier without attenuation in the amplifier-first design (**a**) but at a much lower level in the filter-first design (**b**), leading to better performance in congested environments*

A simple step we can take to explore this is to modify our test setup and add another signal generator supplying a much larger out-of-band interferer.

For this example, the large signal is 1 GHz below our two original signals, and 30 dB (1,000x) higher power.

The two architectures respond very differently to this situation, as illustrated in **Figure 8.** For the amplifier-first design (***a***), any signal that falls within the entire band will reach the LNA, including our test interfering signal at full strength, contributing to overload of the receiver. In the filter-first design (***b***), a significant portion of the interferer will already have been removed, greatly aiding receiver performance.

The results of our simple test are shown in **Figure 9**, which also includes the results from our previous test for comparison. Looking at the amplifier-first design on the left, performance is starting to suffer. There are increased inter-mod products & spurious signals as the LNA is under compression. More strong interferers would make things worse. The situation for the filter-first design is very different; the only obvious difference between the upper and lower plots is a slight increase in noise floor which is coming from the signal generator – otherwise the receiver is untroubled by the extra test signal.

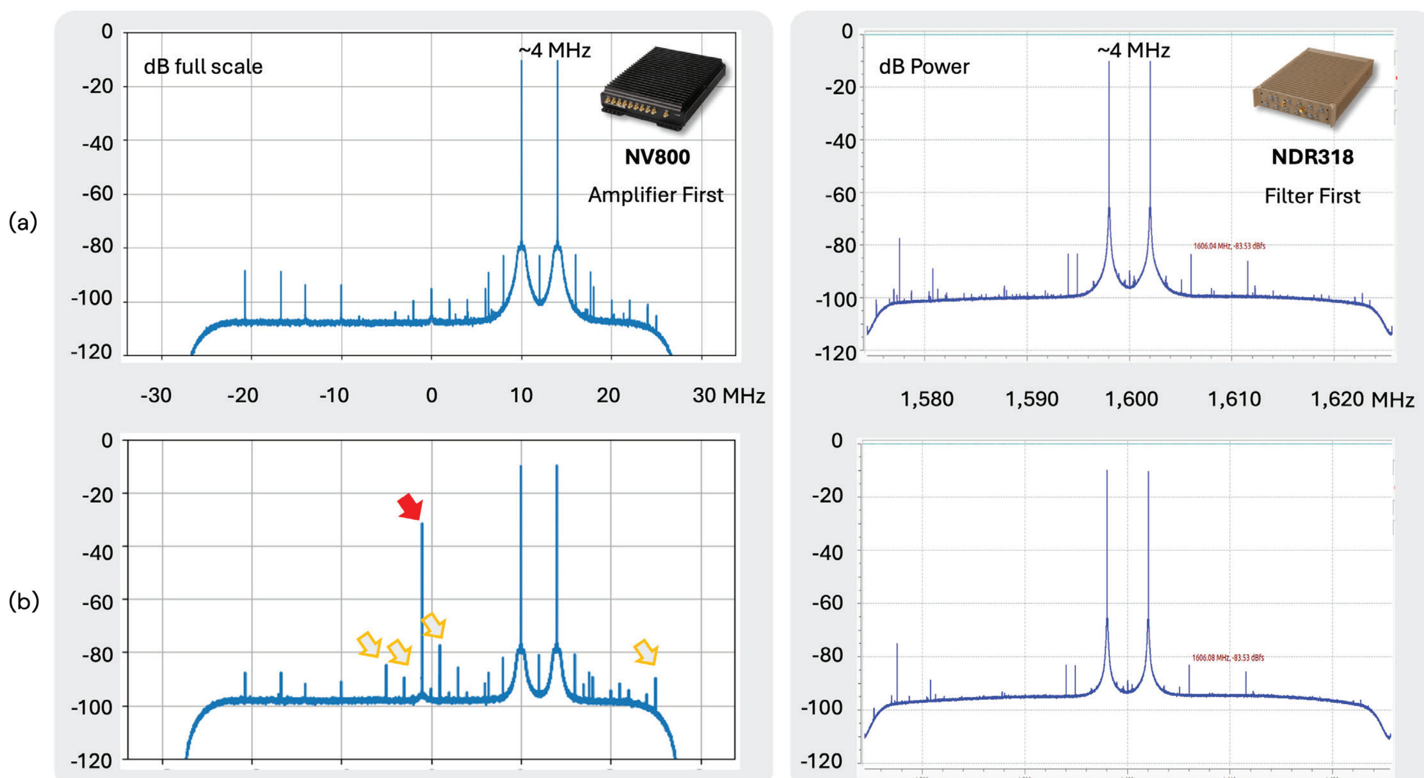

***Figure 9**: A comparison of the amplifier-first design (**left**) and filter-first (**right**) showing the results from our original two-tone test (**a**), above, and the impact of adding the strong interfering signal (**b**), below.*

In **Figure 10** we take this a step further as a thought experiment. As noted, our simple interference setup (***a***) is not representative of a very congested environment such as (***b***). A traditional method of simulating this in a test environment is noise power ratio, or NPR, (***c***) where a broadband noise source is applied to the receiver input having passed through a very high-quality notch filter with steep skirts. An ideal receiver would output a similar spectrum in the pass band (***d***), while lower performing radios would suffer from interfering energy encroaching from the sides and an increased apparent noise floor (***e***). The filter-first NDR318 will perform more closely to the ideal case than the NV800.

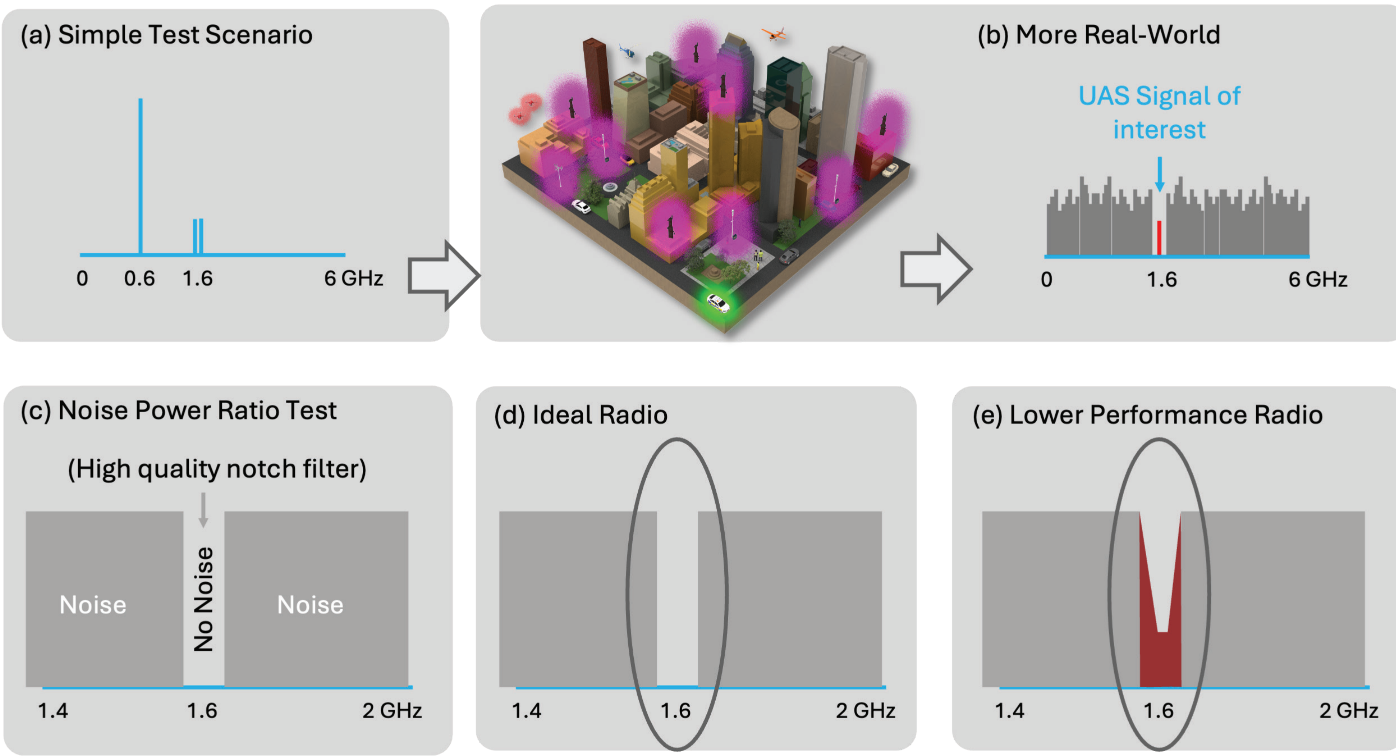

***Figure 10**: Going beyond the simple test setup to illustrate how the radios might perform in a crowded spectrum, using a noise power ratio (NPR) test as a thought experiment*

Phase Coherence/ Frequency Hopping

It turns out that the extra size, weight, power and cost of the NDR318 does enable other advantages, also. These relate to phase coherence and frequency hopping, as summarized in **Table 3**.

	NDR318 	**NV800** 
Phase Coherence	Sub-degree accuracy Determined by RF designer	Few degrees of accuracy Determined by RFIC
Tuning Speed/ Frequency Hopping	100 to 100's µs under all conditions	100 to 100's of µs, but varies depending upon what else the radio is doing, what the specific frequency is, if coherency is required, etc.
Frequency Hopping While Maintaining Phase Coherence	Hopping with or without coherency is transparent to the user – yielding a better user experience	Support varies widely by product and expected hopping rates For NV800, hopping slows to ms+ timescale

***Table 3**: Phase coherence between channels and frequency hopping characteristics for the two radios*

Conclusion

So which radio is better? It very much depends upon the situation, space and power constraints as well as the budget.

- The amplifier-first NV800 has impressive performance, particularly given its low SWaP and cost (and it has a transmit channel plus an option to extend one receive channel to 18 GHz)
- Seeing further under the worst signal conditions, doing the most accurate DF or tracking the fastest hopping signals, requires more SWaP-C but the real-world performance of the filter-first NDR318 is worth it for those critical applications

To answer the question posed in the title of this paper – more dollars, and more SWaP can be exactly the right choice depending upon the problem being solved - it's a question of choosing the right tool for the job.

References

1. Explanation of noise power ratio (NPR): https://www.analog.com/media/en/training-seminars/tutorials/MT-005.pdf
2. More about SDR architectures: https://blog.epiqsolutions.com/software-defined-radios-which-rf-architecture-should-i-choose

ABOUT EPIQ

Epiq Solutions develops cutting edge tools for engineering teams and government-focused organizations requiring situational awareness and detailed insight into their RF environments in order to identify and act against wireless threats.

www.epiqsolutions.com
sales@epiqSolutions.com
+1 847 598 0218
3740 Industrial Ave,
Rolling Meadows, IL 60008, USA