The Civilian Battlefield

Protecting GNSS Receivers from Interference and Jamming

MICHAEL JONES
ROKE MANOR RESEARCH, LTD.

Concern about GPS vulnerability has received a great deal of attention recently — and for good reason. GNSS receivers are highly susceptible to jamming and spoofing. Historically, signal jamming and the design of equipment to protect against it has been considered primarily a military problem. Now, signal interference threatens all GNSS receivers, military and civilian. But antenna and receiver design options can mitigate or eliminate the problem.

Growing dependence on GNSS for positioning, navigation, and timing (PNT) has raised a parallel concern about the potential risks of signal interference. The popular press has recently highlighted accounts of car thieves using GPS jammers, solar flares pumping out L-band radiation, and faulty television sets causing havoc to GPS receivers across an entire harbor.

Speculation has suggested such dire scenarios as the collapse of telecom, power, and banking networks, ships colliding, and planes falling out the sky. Responses to these stories can be equally extreme, with some arguing that "GPS is unreliable," or "We need an alternative PNT system."

Well, *perhaps*. But isn't it better to protect what we have, rather than cast it aside as unsuitable?

When computers were first threatened by viruses and hackers, we didn't toss them aside complaining that "computers are too vulnerable — we need an alternative system." No, we didn't resort to pen and paper for all our work; we simply installed firewalls and virus checkers.

So it is with GNSS — or should be. Instead of simply criticizing the technology's weaknesses, we need to explore solutions to the interference and jamming problem. And by "solutions," I'm talking about protecting what we have, rather than simply abandoning GNSS and resorting to less mature alternatives such as enhanced Loran (eLoran).

With this approach in mind, the following discussion will present a few of the ways in which we can make our GNSS receivers more resilient to interference, with a particular focus on the role of receiver antennas in mitigating its effects. But let's start by briefly considering why GNSS is so vulnerable in the first place.

Why is GNSS Vulnerable?

As most people are aware, the fundamental cause of GNSS vulnerability is the power of transmissions from satellites. GNSS signals are so "quiet" that they are easily swamped by the smallest of interfering signals.

Taking GPS as an example, the signals reaching the Earth's surface are around -163 dBW (50 x 10^{-18} watts). Pretty much any other radio frequency phenomenon is going to be larger than that. In fact, even the GPS signal itself is well below the noise floor of the receiver.

The reason that GPS signals survive at all is due to the spread-spectrum nature of the transmission, allowing receivers to correlate the satellite signal out from below the background noise. But each receiver exhibits a limitation to the amount of non-GNSS interference it can cope with, whilst still acquiring or tracking the desired signal. That is, each receiver has a maximum jammer-to-signal ratio (J/S) that it is able to tolerate.

Figure 1 illustrates the problem: the diagonal lines represent interfering sources of various powers, and the horizontal dashed lines show some typical receiver thresholds. Theoretically, at least, a 10-milliwatt jammer will prevent a receiver from acquiring the C/A code at a distance of 10 kilometers, and a receiver already tracking the C/A code will lose lock about a kilometer from the jammer.

This is pretty scary stuff, considering that 10 milliwatts is a very tiny jammer indeed. Even a P(Y)-code receiver will stop tracking when a few hundred metres from the jammer. Once the jammers get larger, it's pretty much "game over" for unprotected GNSS signals. In any case, higher power jammers aren't really necessary at ground level anyway, because once a receiver gets past 10 kilometers or so from the jammer, it's likely to be below the horizon and no longer in the jammer's line-of-sight.

Of course, jamming is not the only form of GNSS interference. Spoofing is a more recent, and perhaps more frightening, threat.

GNSS Spoofing

While jamming is concerned with denying the availability of a service, spoofing fools the receiver into thinking it is still happily tracking satellites, when in reality it is processing fake GNSS signals. Broadly speaking, we face two main types of spoofing: simulation and record-replay.

Simulation involves equipment that generates and transmits valid GNSS signals into which the spoofer can inject any information desired. Probably the easiest way to do this is with a GPS simulator as shown in **Figure 2**. Stories have been reported where such a setup has been used to hijack a truck in an experimental demonstration (See, for example, the article by S. Davidoff listed in the Additional Resource section near the end of this article.)

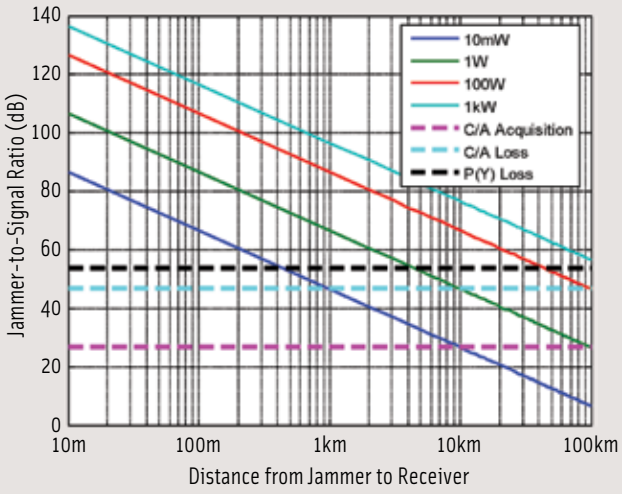

FIGURE 1 The effect of various jammers on GPS receivers

The second spoofing method involves simply recording a real GNSS signal, and replaying it at a later time. Although this method cannot be used to impose a user-defined scenario on a receiver, it is enough to wreak havoc in unwary receivers by making real satellite signals appear at a different time and from different locations, and this form of attack can also be used against encrypted GNSS services.

Spoofing can be considered a more dangerous form of interference than jamming, because it is not always obvious that you are being spoofed, while jamming constitutes a denial-of-service and is easily detected. Having said that, performing a spoofing attack is not altogether straightforward: in order to spoof a receiver that is already tracking, it would typically need to first be jammed and then re-acquire tracking on the spoofed satellites.

We can use various mechanisms to detect spoofing, such as monitoring absolute and relative signal strengths, monitoring satellite ID codes, performing time comparisons, and generally keeping an eye out for unusual or unlikely signal scenarios. But the problem remains, because even if you know you are being spoofed, you still have to deal with the spoofing!

Now that we've looked at the fundamental issue, let's take a whirlwind tour through some of the possible solutions to the problem.

Receiver versus Antenna Solutions

Broadly speaking, two general elements comprise a GNSS receiver, the antenna and the receiver itself (RFIC, downconverter, tracking channels, digital signal processor, and so forth). We can choose to enhance either or both of these elements in order to improve resilience to interference.

FIGURE 2 A GPS simulator "spoofs" a receiver into believing it is at a different location

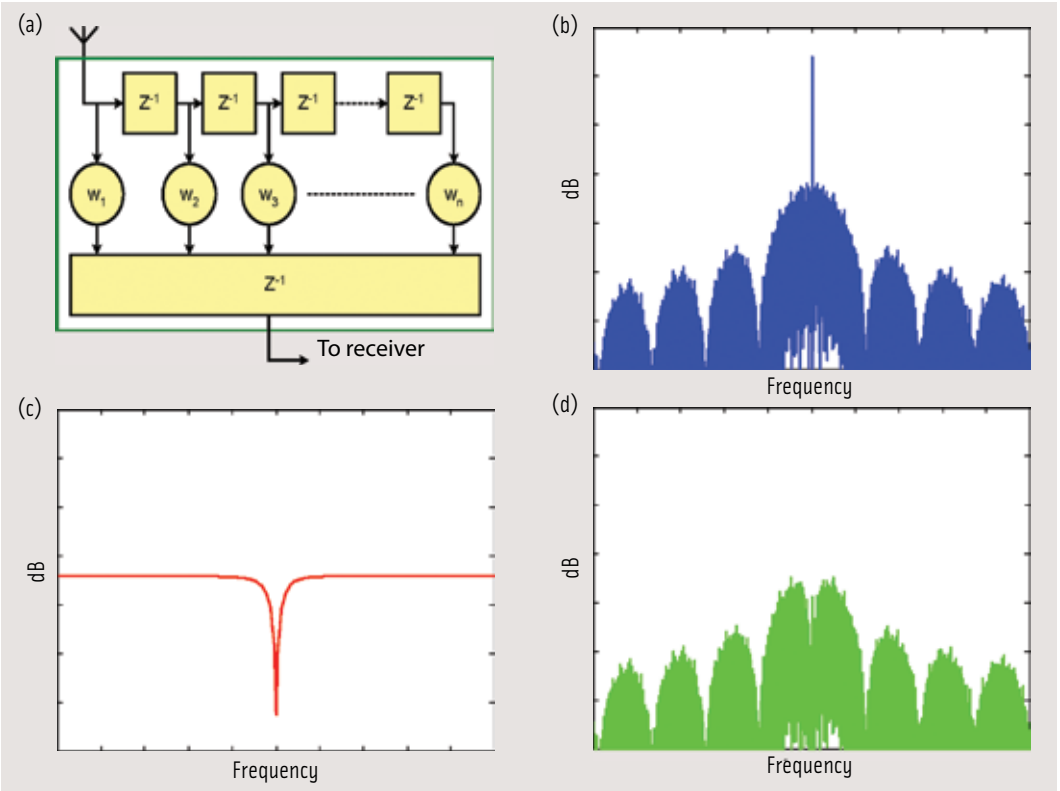

FIGURE 3 (a) Digital filter for adaptive notch filtering. (b) C/A code with tone jammer at center frequency. (c) Notch response of filter. (d) Filtered signal passed to receiver DSP

An antenna enhancement would typically involve replacing the antenna with a more sophisticated version, along with some associated electronics. Antenna enhancements are attractive because they do not affect existing receiver systems. As an appliqué solution, we are simply swapping an old component for a better one, with a cost range reflecting the sophistication and capabilities of the new antenna.

A receiver enhancement modifies the signal processing and tracking capabilities to render the equipment more robust against threats. These enhancements also add a range of expense, depending on the level of change required.

Receiver Solutions

In our arsenal of receiver enhancements is *adaptive notch filtering*. Here we are simply trying to filter out the jammers from the received signal and is a straightforward approach requiring minimal modifications to an existing system.

Figure 3a shows a simple digital filter structure that could be placed before the GNSS receiver functions. Simple notch filtering might be achieved with an analog filter following the LNA but, for maximum control, the signal should be downconverted and filtered in the digital domain. This could be done either in the receiver itself, or as part of an appliqué front-end.

The figure illustrates the use of a basic finite impulse response (FIR) filter, where the weights adapt to filter out any input power. Consider the spectrum in Figure 3b where a C/A-code signal is being jammed with a continuous wave (tone) jammer at center frequency. The filter adapts itself to give the frequency response shown in Figure 3c that, when applied to the input, produces the signal of Figure 3d to the receiver's digital signal processor (DSP).

With this approach, we can see that the jammer interference has indeed been removed, but at the expense of degrading the C/A-code signal itself. Consequently, we should consider adaptive notch filtering as generally a destructive process that must be used with caution.

Many jammers, rather than being simple tones, are wideband in nature and jam across the full GNSS frequency. Attempting to use an adaptive notch filter would then result in the GNSS signal itself being removed. Clearly, this technique is not ideal, and is certainly not suitable against broadband jamming.

A further issue with such a filtering technique is that raised time sidelobes can give rise to tracking errors, unless the receiver is intelligent enough to handle the situation and lock onto the main time lobe. In summary, adaptive notch filtering is rarely an acceptable solution , but can be effective for countering specific known types of narrowband interference.

Switching Frequencies

The second obvious thing to attempt when encountering a jammer in a GNSS frequency is to switch to a different GNSS band. Military users of GPS have the luxury of being able to choose between L1 and L2, and in the near future aviation receivers will have access to L5. Most mass-market civilian C/A users, however, are stuck with only L1, although GPS/GLONASS-capable cell phones and other mobile devices are beginning to reach the market.

The introduction of Galileo and other GNSS systems can only be a good thing here. Modifying receivers to be multi-GNSS–compatible allows for frequency diversity; so, in a few years when the L1 signals are being jammed, we might switch to the Galileo OS on E5a.

Although this sounds good, it might not always provide a satisfactory solution. Making multi-GNSS receivers adds costs and demands on power and the CPU; moreover, multiple frequencies and signals can still be jammed. If someone is willing to spend $10 on a

jammer to knock out L1, why wouldn't they spend another $10 to knock out L2 or E5 as well?

Integrating GNSS with INS

Combining GNSS with an inertial navigation system (INS) offers well-known benefits. The two systems are highly complementary. Although inertial technology typically provides very good short-term accuracy, it drifts and loses accuracy over time as errors accumulate. Conversely, GNSS may exhibit relatively poor short-term accuracy, but it makes up for it with exceptional long-term accuracy because a receiver generates independent, absolute position fixes with rapid update rates.

Furthermore, an INS cannot be jammed. So, by using inertial navigation, which is periodically corrected by input from a GNSS, we get the best of both worlds. Such an architecture may be a *loosely coupled* system as illustrated in **Figure 4a**. The green boxes are the GPS receiver functions, with additional INS functions shown in yellow.

In such INS designs, the GPS position and velocity information is passed to the navigation Kalman filter to provide periodic corrections to the data from the inertial measurement unit (IMU). This is often referred to as *reinitializing* the INS. Compared to GNSS on its own, the jamming immunity of the overall INS system has improved.

The level of integration between the two systems can be taken further, to give the architecture shown in **Figure 4b**, typically referred to as a *tightly coupled* system. Here the Kalman filters of the GNSS and INS are combined into a single filter. This avoids errors that are introduced by the cascading of separate filters, and gives a further immunity to jamming, though at the expense of significant modifications to the existing receiver.

Modifying receivers to be multi-GNSS–compatible allows for frequency diversity; so, in a few years when the L1 signals are being jammed, we might switch to the Galileo OS on E5a.

Integration can be taken yet further to give an *ultra tightly coupled* or *deeply integrated* system. In this architecture, the GNSS tracking loops themselves are removed and become absorbed into the single navigation Kalman filter.

Variations exist for each of these architectures as well as some argument as to what really constitutes a "tightly coupled" or "deeply integrated" system.

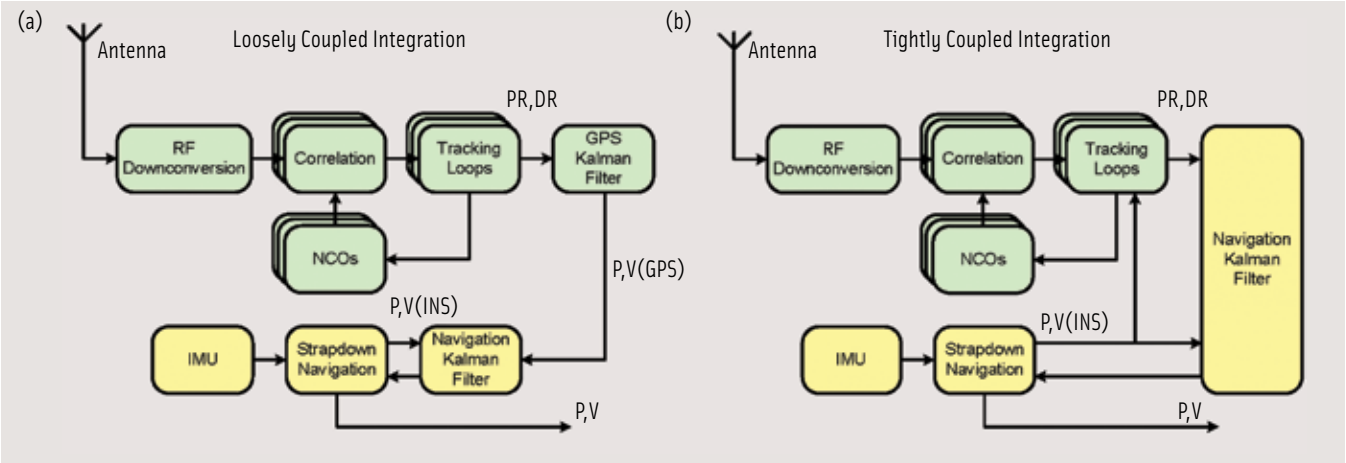

FIGURE 4 (a) Loosely coupled integration. (b) Tightly coupled integration

But, fundamentally, they are all trying to achieve the same thing: to continue to provide position and velocity in the presence of jamming.

A careful examination of the preceding sentence reveals two issues with this type of protection. Firstly, if jamming is present for a long period, then the GNSS will be unable to operate and provide data with which to reinitialize the INS. So, the position/velocity solution will tend to rely more and more on the INS data as time goes on. Consequently, the GNSS/INS will start to accumulate errors as the jamming continues.

Secondly, and perhaps more importantly, these GNSS/INS architectures only provide position and velocity, and not time. Applications and infrastructures — such as phone, utility and banking networks — that rely on a precise time reference from GNSS will not see any benefit here.

Moreover, inertial technology inevitably brings additional costs for sensors and their integration, ranging from microelectromechanical systems (MEMS) on the low end to ring laser gyros at the high end. Integrated GNSS/INS architectures typically find use in high-value moving platforms.

Antenna Solutions

Arguably the single most powerful antenna-based technique for mitigating interference is to exploit spatial diversity — that is, to make use of the fact that satellite signals and jamming signals usually come from different directions.

Figure 5a illustrates how a typical GNSS antenna pattern might look, with satellite signals S1 and S2, and jamming signal J, all received equally. If we can choose an antenna such that the gain pattern is more like **Figure 5b**, then the receiver does not "see" the jammer.

Can this be done? Of course it can, and this is where adaptive antennas shine.

An adaptive antenna, also referred to as a controlled radiation pattern antenna (CRPA), consists of a number of smaller antenna elements spaced apart from one another, with their outputs all summed together (**Figure 5c**).

A signal arriving from a particular direction hits each antenna element at a slightly different time. Depending on the difference in arrival times, the antenna element signals might add together *constructively*, giving a large output signal, or *destructively*, giving a small output signal.

Notice in Figure 5c that each element also has a controllable phase shift, or

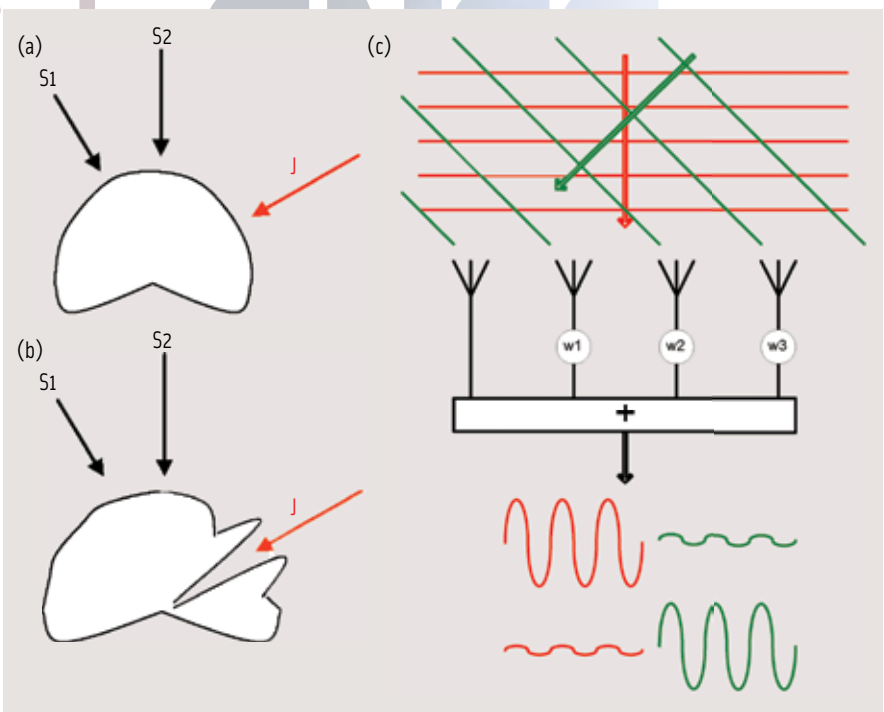

FIGURE 5 (a) Illustration of a typical GNSS antenna pattern. (b) The adaptive antenna concept. (c) The principle of sidelobe cancelling in a controlled radiation pattern antenna

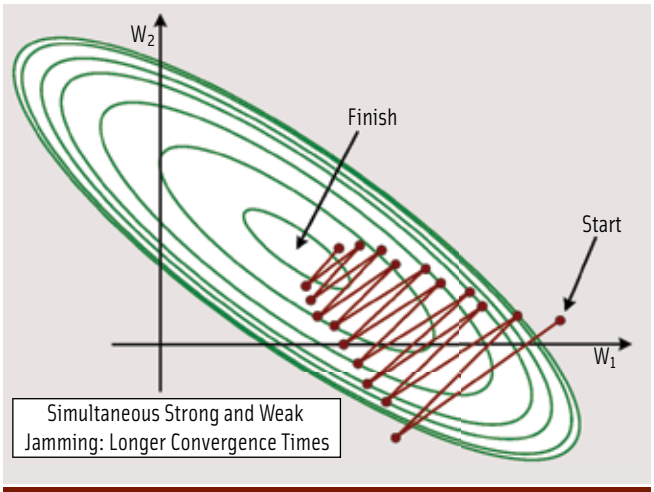

FIGURE 6 Convergence time issues with gradient-descent algorithms

weight, *w*. By altering these weights, or delaying the signal from each element, we can get the antenna to cause destructive interference in a direction, or several directions, that we choose.

In the case of N antenna elements, it is possible to achieve a satellite signal gain of 10log(N) dB, whilst greatly attenuating any interfering signals from other directions. The level to which jamming signals are attenuated depends on various factors, including the number of separate jamming sources, the nature of the jamming waveforms, and the numerical accuracy of the signal processing.

So, we can choose to pass the green signal in Figure 5c, whilst attenuating the red signal, or vice versa. In other words we can make the antenna "blind" to jammers and spoofers, so that the receiver never has to worry about them. This technique is known as *null steering*, or *sidelobe cancellation*.

The next question, then, is how do we make the antenna steer nulls in the directions we want?

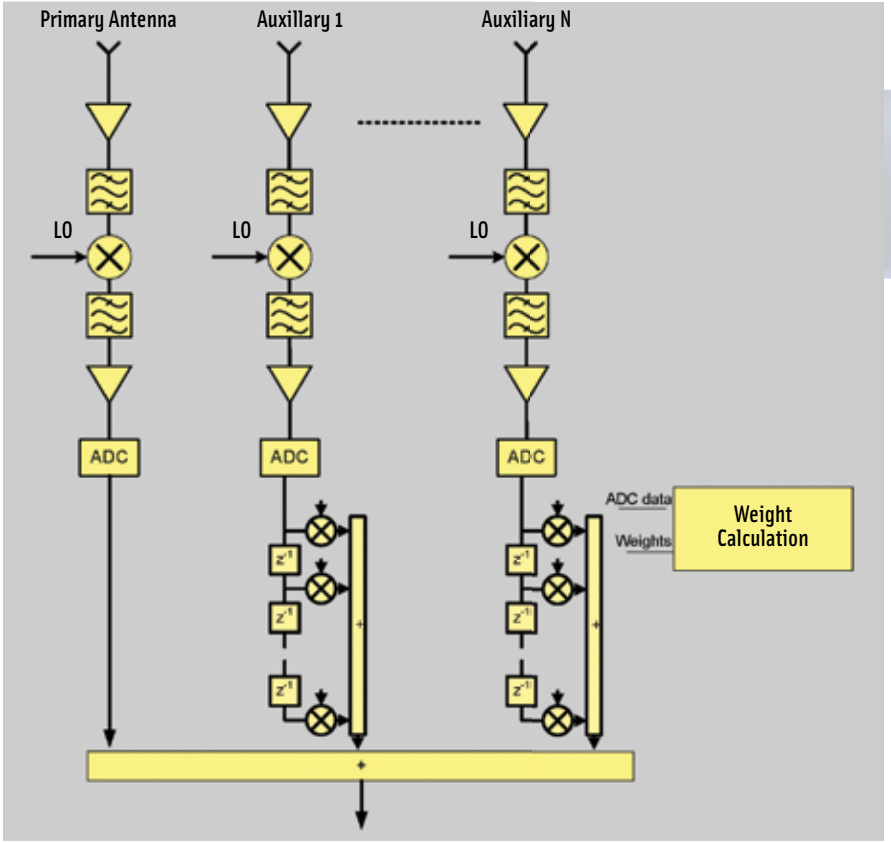

FIGURE 7 Digital jammer cancellation with space-time adaptive processing

An Optimal Solution

Continuing our reference to the antenna of Figure 5c, we can denote the signal received by the antenna as the vector **x(t)**, and the summed antenna output as **e(t)**. The leftmost element is designated as the primary element, *p*, while the others have variable weights, ***w***. So, the antenna output is given by

$$e(t) = p(t) - \mathbf{w}^T \mathbf{x}(t) \quad (1)$$

Given that GNSS signals are below the noise floor, the objective is to minimize the output power of the antenna, which will have the effect of cancelling the interfering signals. The average output power is given by

$$E\left\{\left|e^2(t)\right|\right\} = \quad (2)$$
$$= E\left\{\left|p^2(t)\right|\right\} - \mathbf{w}^H \mathbf{r}_{xp} - \mathbf{r}_{xp}^H \mathbf{w} + \mathbf{w}^H \mathbf{R}_{xx} \mathbf{w}$$

Manipulating Equation (2) to find the optimum weights that minimize the output power leads to the well-known Wiener equation

$$\mathbf{w}_{opt} = \mathbf{R}_{xx}^{-1} \mathbf{r}_{xp} \quad (3)$$

The matrix **R** is the data covariance matrix, and **r** is the vector of cross-correlations between the primary and auxiliary antenna elements. We can solve this equation in many ways, which require various levels of computational power.

The first observation of interest is that the output power is a quadratic function of the weights, **w**, which indicates that a single global optimum set of weights exists. Given this knowledge, the optimum weights could be easily computed using a steepest-descent method, such as least mean squares (LMS),

$$\mathbf{w}(k+1) = \mathbf{w}(k) + \mu \mathbf{x}^*(k) e(k) \quad (4)$$

This is, in fact, a very viable and low-cost way to adapt the antenna such that it cancels out any jammers.

However, use of such a gradient-descent algorithm has a well-known drawback. The convergence time, that is, the time taken for the algorithm to cancel the jamming sources, is highly dependent on the nature of the jamming signals. Mathematically, the algorithm is sensitive to *eigenvalue spread*.

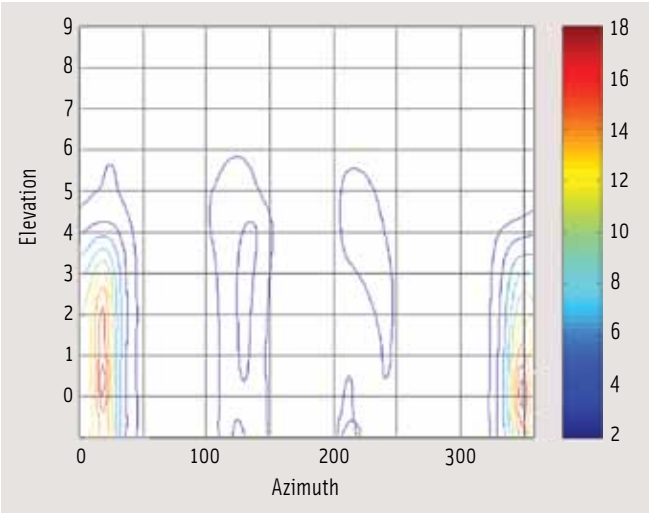

FIGURE 8 Direction finding of jamming sources

Figure 6 illustrates this for a case in which both a strong jammer and a weak jammer present are present, and the search space has become an elliptic paraboloid. Consequently, the strong jammer is cancelled quickly, but the time taken to cancel the smaller jammer is much longer.

This may not be an issue, particularly if operating in a fairly static jamming environment. However, in cases where the jamming sources are rapidly changing, moving, or blinking, then the gradient descent solution may not be ideal.

Digital Jammer Cancellation

The best anti-jam performance is achieved when digital techniques are employed to directly solve the Wiener equation, $w_{opt} = \mathbf{R}^{-1}{}_{xx}\,\mathbf{r}_{xp}$, where $\boldsymbol{R}$ is an n by n data covariance matrix and $\boldsymbol{r}$ is a vector of cross-correlations between the primary and auxilliary antenna elements. **Figure 7** provides an example of a receiver architecture designed to accomplish this.

In this approach, each antenna input is down-converted to an intermediate frequency (IF) using a superheterodyne receiver, before sampling with an analog to digital converter (ADC). Notice that the weights, which were previously used to apply a phase shift to each channel, have been replaced by FIR filters. This technique, known as *space-time adaptive processing* or STAP, is one of the most powerful techniques available.

The STAP filters allow gain and phase corrections to be applied across the GNSS band, which gives superior performance against broadband jamming sources. The Wiener equation is solved directly by a process that is both fast and non-dependent on the jamming environment. Moreover, other signal processing functions can be optionally performed such as direction finding on the jamming signals.

Figure 8 shows how jammers can be located. In this example, two jammers were present at low elevations and with azimuths of 15 and 350 degrees. After digital processing has been performed in the antenna electronics, the signal is upconverted back to RF, for connection to the GNSS receiver.

Once again, the beauty of such a solution is that no changes are required to the existing system, other than swapping the old antenna for the adaptive one.

Anti-Spoofing

The digital anti-jam architecture, just described, is also capable of protecting against spoofing attacks. However, anti-spoofing requires a somewhat different approach. Because spoofer signals look just like GNSS signals, we can't get rid of them by power minimization techniques alone.

Through the addition of a GNSS code correlator, we can determine the direction-of-arrival of any GNSS signal and, by application of a constraint to the weight calculation, the antenna will then additionally cancel the spoofing source.

Adaptive Beamforming

Although jammer and spoofer cancellation techniques are extremely powerful, we achieve the ultimate protection with a further step: beamforming the receiver antenna toward GPS satellites and their signals.

In addition to enabling an adaptive antenna array to steer minimal antenna gain towards jammers, beamforming simultaneously ensures that the antenna steers *maximum* gain towards satellites. Put another way, instead of minimizing the interference-to-noise ratio (INR), the aim is to maximize the signal-to-interference-plus-noise ratio (SINR).

Figure 9 illustrates the SINR maximization architecture, where adaptive beamforming is applied. Because multiple desired signals (one from each satellite) are available, multiple beamformers are required. The diagram shows 8 beamformers, although this number will be a trade-off between performance and complexity. For example, an all-in-view receiver that can track 12 satellites on both L1 and L2 could use up to 24 beamformers, although high performance can be achieved with far fewer.

A beam can be formed in a particular direction by constraining the adaptive array gain to be unity in that direction, such that

$$\mathbf{g}(\mathrm{az},\mathrm{el})^T\,\mathbf{w} = 1 \qquad (5)$$

Using Lagrange multipliers or another method, we can determine the optimum weights by minimizing the power subject to this constraint:

$$\mathbf{w}_{opt} = \mathbf{R}^{-1}_{xx}\,\mathbf{g}^*\left(\mathbf{g}^T\mathbf{R}^{-1}_{xx}\,\mathbf{g}^*\right)^{-1} = \mu\mathbf{R}^{-1}_{xx}\,\mathbf{g}^* \qquad (6)$$

As with the standard Wiener equation, there are various methods to solve this. If there is sufficient processing power, the whole equation can be solved directly. If a simpler approach is desired then, given that μ is a scaler, the approximation $\mathbf{w}_{opt} \approx \mathbf{R}^{-1}_{xx}\,\mathbf{g}^*$ could be used. Alternatively, we can apply some mathematical tricks in order to solve the whole equation with minimal computational effort.

Of course, in order to steer beams in the direction of satellites, it's important to know where the satellites are! As the satellite almanac is generally known, and the orientation of the antenna could be known, the adaptive beamformers can know the direction of arrival of

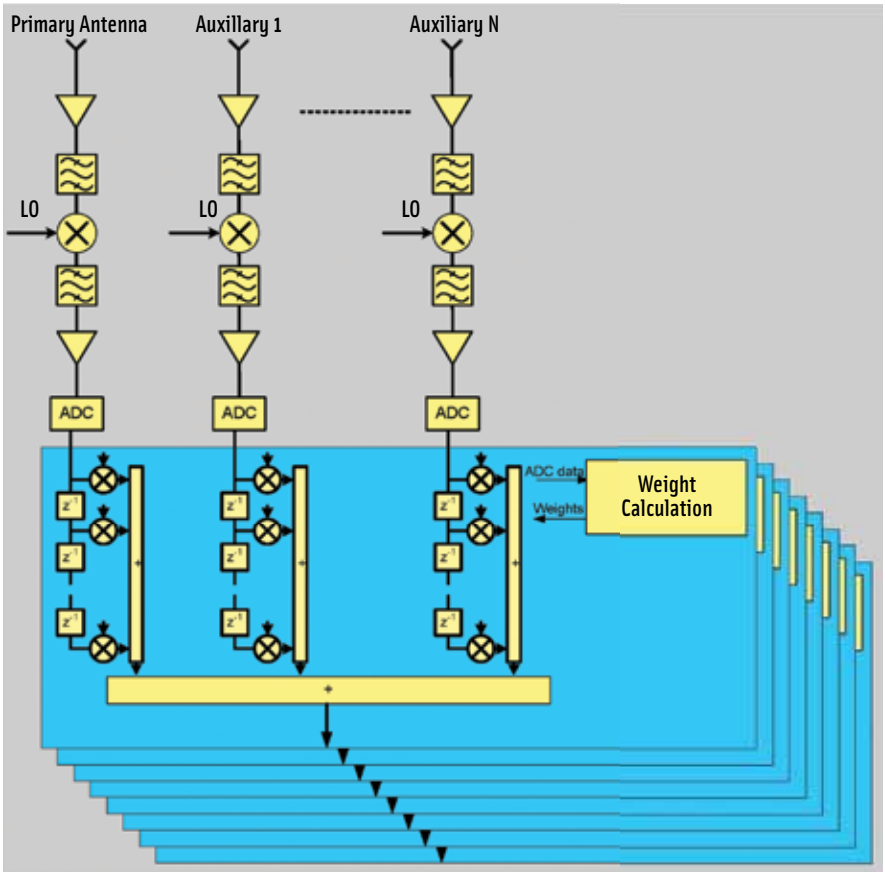

FIGURE 9 GNSS protection using adaptive beamforming with STAP

each satellite signal. If a solution of this kind is used on a moving platform, such as an aircraft or other vehicle, then some form of inertial reference is required in order to know the orientation of the antenna.

Finally, although a beamforming architecture offers the highest performance in terms of anti-jam capability, it does have a potential drawback. Because we have an optimized signal output for each satellite of interest, the beamformer outputs cannot be combined back into a single RF input to a GNSS receiver. To achieve this, we require a digital interface directly into each tracking channel of the receiver.

If a receiver already has the option of separate digital inputs, then no modifications to it will be necessary.

Conclusion

Threats to GNSS receivers are real, and the wider community is gradually awakening to the idea that something needs to be done. GPS jammers are widely available, and the threat from spoofing is growing. Despite our enormous and ever-growing dependence on GNSS, little has been done to truly address the problem, and yet there *are* answers to the problem.

This brief article has attempted to highlight some of the many possible solutions, each of which has its own advantages and disadvantages. Adaptive antennas offer excellent protection from both jamming and spoofing sources, and with modern technology they can be also be made inexpensively.

Risk assessment is fundamentally important, however. Conducting a threat analysis for one's own GNSS application is crucial in order to decide what method of protection — if any is needed — is best.

Just as any computer can be compromised if a hacker tries hard enough, so can any GNSS receiver. Accordingly, a hacker will try harder to interfere with an application where the consequences are more severe, in which case suitable countermeasures are called for.

As our dependence on GNSS grows, protection technology that was once only available to the military is now becoming available to the commercial world. If we embrace the options available, to protect our receivers, we can continue to rely on the wonderful systems that we call GNSS.

Acknowledgment

This article was adapted from a presentation given by the author at the GNSS-10 conference on April 29, 2010, hosted by the Institute of Engineering and Technology at Savoy Place, London.

Additional Resources

1. Davidoff, S. "GPS Spoofing," *Philosecurity*, <http://philosecurity.org/2008/09/07/gps-spoofing>, 2008

2. Last, D., "GPS: The Present Imperfect," *Inside GNSS*, vol. 5, no. 3, May 2010

3. Schmidt, G., "INS/GPS Technology Trends", The Draper Technology Digest, vol. 2, pp. 143-154, 1998

Author

Michael Jones is a consultant engineer at Roke Manor Research, specializing in advanced GNSS adaptive antenna systems and state-of-the-art radar. He gained a First-Class Masters degree in electronic engineering from the University of York and was also awarded the Racal prize for best avionics degree. Since joining Roke Manor Research, he has worked on a multitude of advanced navigation and radar systems contracted by the U.K. and U.S. governments. He has expertise in adaptive antenna GPS systems and was jointly responsible for a number of navigation protection systems using interference cancellation, adaptive beamforming and direction finding. Jones specializes in the simulation and hardware implementation of advanced signal-processing algorithms, and has worked on a number of FPGA and ASIC designs for radar and GNSS systems. IG