16. Interference

Todd Humphreys

Global navigation satellite system (GNSS) signals are so weak near the Earth's surface that they can be easily squelched by natural or man-made interference. Moreover, the most popular GNSS signals – those offered with unrestricted access – are unencrypted and unauthenticated, which means they can be counterfeited, or spoofed. Strict international laws protect the radio frequency bands allocated to GNSS, but mother nature does not respect these laws, and man-made interference – whether accidental or intentional – is a growing concern.

This chapter examines sources of GNSS signal interference and the interference effects on GNSS signal tracking. It offers a systematic treatment of natural, unintentional, and intentional interference, with emphasis on intentional jamming and spoofing. Theoretical performance bounds are developed for the simplest cases of narrowband and wideband interferences. The chapter finishes with a review of the state of the art in antenna-oriented and signal-processing-oriented interference detection and mitigation techniques.

All GNSS waveforms are spread-spectrum signals, which are uniquely resilient to interference. Indeed, robustness in the face of jamming was one of the primary features, along with low probability of intercept and good multiple access properties, which motivated the original development of spread-spectrum techniques for military systems. Nonetheless, GNSS signals are extremely vulnerable to jamming because, near the surface of Earth, they have no more flux density than light received from a 50 W bulb at a distance of 2000 km. To blandly remark that GNSS signals are weak is to understate their fragility: They are so weak that most modern electronics jam GNSS receivers at close range, requiring special precautions be taken to isolate receivers embedded in computers, mobile phones, vehicles, and other modern GNSS-dependent systems.

Table 16.1 ITU space-to-Earth radio navigation satellite service (RNSS) frequency allocations (after [16.2, 3]). ARNS refers to the Aeronautical Radionavigation Service. Bands that are designated as both RNSS and ARNS enjoy, in principle, no greater International Telecommunication Union (ITU) protection from harmful interference than RNSS bands, but in practice they are granted more conservative safety margins (see, e.g., ITU-R M.1903) and they are likely to be monitored more assiduously by ITU member nations

Frequency interval (MHz)	Bandwidth (MHz)	GNSS bands	Notes
1164–1215	51	L5/E5a/E5b/L3/B2	ARNS band; pulsed DME/TACAN interference present [16.1]
1215–1240	25	L2	Legacy GPS L2 band
1240–1260	20	L2	Legacy GLONASS L2 band
1260–1300	40	E6/B3/LEX	
1559–1610	51	L1/E1/B1	ARNS band; legacy GPS and GLONASS L1 band
5010–5030	20	C1	

Unintentional and intentional GNSS interferences are distinguished from each other more by motive than by effect. Both can be narrowband or wideband (relative to the bandwidth of the desired GNSS signal), structured or random. The user of a GNSS receiver suffering from interference may care little about the jammer's intent: What is important is a clean spectrum. Indeed, the recent emergence of so-called personal privacy devices (PPDs) – low-cost GNSS jammers used to ward off GNSS tracking – blurrs the lines between unintentional and intentional interference: The privacy device user only intends to jam GNSS receivers in an imaginary bubble around himself; he may never intend to disrupt the GNSS-dependent timing system at the bank down the street.

Interference that mimics GNSS signal structure and content is a special threat to GNSS receivers. Instead of simply degrading the accuracy of the position, velocity, and time (PVT) solution, transmission of such structured interference, referred to as spoofing, can fool a receiver into producing a precise but erroneous solution. Worse yet, the induced solution can be entirely dictated by the spoofer operator, who may have malevolent intentions. All GNSS signals are spoofable to one degree or another – at the very least, they can all be recorded and replayed into a target receiver, as is routinely done for receiver testing. But the most popular GNSS signals, the so-called open signals, are especially vulnerable because they are (so far) almost entirely predictable, lacking encryption or authentication of any form. For radionavigation as for communication, predictability is the enemy of security.

From the origins of GNSS, national and international policy has afforded special protection to the GNSS radio bands, and now that GNSS receivers have become pervasively embedded in the infrastructure that supports the global economy, such protection is of special importance. The International Telecommunication Union (ITU) forbids any interference *which endangers the functioning of a radionavigation service* [16.2] in the GNSS bands, which are designated as radionavigation satellite service (RNSS) bands by the ITU. Table 16.1 summarizes the ITU's current frequency allocations for GNSS signals.

In some regions, the penalty for emitting unauthorized signals in the GNSS bands is severe: In response to a rising number of so-called PPDs, the United States Federal Communications Commission (FCC) levies costly fines on intentional violators [16.4], and the penalty for intentional transmission in Australia can include a 2 yr prison term [16.5]. But despite government protections of the GNSS bands, they remain cluttered with interference, and there is every indication that such interference will worsen in the decades to come as more GNSS constellations begin broadcasting [16.6], as people respond to pervasive GNSS tracking by employing PPDs [16.7], and as communications signals ineluctably encroach on the enormously valuable GNSS spectral bands [16.8].

This chapter examines the effects of interference on GNSS receivers. The chapter begins with a presentation of the general analysis technique that will be used to evaluate the effect of interference that is statistically independent of the GNSS signals. The technique will then be applied to study the effects of canonical narrowband, wideband, and multiaccess interference. Following this, other specific interference waveforms such as pulsed interference will be discussed. Thereafter, GNSS spoofing, a particular type of interference that cannot be considered statistically independent of the GNSS signals, will be given a focused treatment. The chapter finishes with an examination of interference detection and mitigation strategies. Note that GNSS multipath, while a genuine type of interference, is treated separately in Chap. 15.

16.1 Analysis Technique for Statistically Independent Interference

Beyond the statement that GNSS interference always degrades PVT accuracy, one can say little in general about interference effects on late-stage signal processing products because these effects are highly receiver-dependent: A vector-tracking low-tracking-bandwidth receiver will, for example, produce a much more robust PVT solution than a scalar-tracking wide-bandwidth receiver. At earlier processing stages, however, interference effects are substantially common across receiver types and thus a general treatment becomes possible. Accordingly, this section presents an analysis of interference effects on the primitive correlation-and-accumulation products that form the basis of signal tracking in all GNSS receivers.

16.1.1 Received Signal Model

Consider the following generic representation of a received GNSS signal exiting a receiver's radio frequency (RF) front-end downconversion chain. For notational compactness, the signal is expressed by its complex baseband representation as

$$r_S(t) = \sqrt{P_S} D(t-\tau(t)) C(t-\tau(t)) \exp(j\theta(t)), \quad (16.1)$$

where P_S is the received signal power in watts, $D(t)$ is the binary navigation data modulation, $C(t)$ is the binary spreading (ranging) code, $\tau(t)$ is the code phase, and $\exp(j\theta(t))$ is the carrier with phase $\theta(t)$. The code phase $\tau(t)$ varies slowly and, for purposes of interference modeling and analysis, can be modeled as constant; thus, it will be denoted τ hereafter.

Let $r_I(t)$ represent a complex-valued interference signal, and let $n(t) = n_I(t) + jn_Q(t)$ be a zero-mean complex-valued Gaussian process that models thermal noise. Then, the full received signal-plus-interference-and-noise is given by

$$r(t) = r_S(t) + r_I(t) + n(t) .$$

The received components $r_S(t)$, $r_I(t)$, and $n(t)$ are assumed to be limited by a bandpass filter in the RF front end having a noise-equivalent bandwidth of W_{FE} Hz. The quadrature processes $n_I(t)$ and $n_Q(t)$ are modeled as spectrally flat on the range, $|f| < W_{FE}/2$ with two-sided density $N_0/2$, where N_0 has units of W/Hz. Consequently, on this range the full complex thermal noise process $n(t)$ has a two-sided density of N_0. The data $D(t)$ and spreading code $C(t)$ are assumed to be normalized to unity power so that

$$P_S = \lim_{T\to\infty} \frac{1}{T} \int_{-T/2}^{T/2} |r_S(t)|^2 dt .$$

If $r_S(t)$, $r_I(t)$, and $n(t)$ are statistically independent, then the total received power in the bandwidth W_{FE}, denoted by P_T, is

$$P_T = P_S + P_I + P_n , \quad (16.2)$$

where P_I is the total power in $r_I(t)$, and $P_n = W_{FE}N_0$. The carrier power to thermal-noise density ratio is $C/N_0 = P_S/N_0$, and the signal-to-thermal-noise ratio is $SNR_{FE} = P_S/P_n$. Similarly, the signal-to-interference-and-thermal-noise ratio is $SINR_{FE} = P_S/(P_n + P_I)$. Figure 16.1 offers an example illustration of the relationship between the power spectra of $r_S(t)$, $r_I(t)$, and $n(t)$.

16.1.2 Thermal-Noise-Equivalent Approximation

A key insight greatly simplifies GNSS interference analysis: The effect of interference on almost all GNSS receiver functions can be accurately modeled as if it were caused by spectrally flat thermal noise of a certain density. This subsection explains when this thermal-noise-equivalent approximation is valid and notes its limitations.

GNSS signal processing is founded on correlation of the received signal $r(t)$ with a local replica

$$l(t) = C_l(t-\hat{\tau}) \exp(j\hat{\theta}(t)),$$

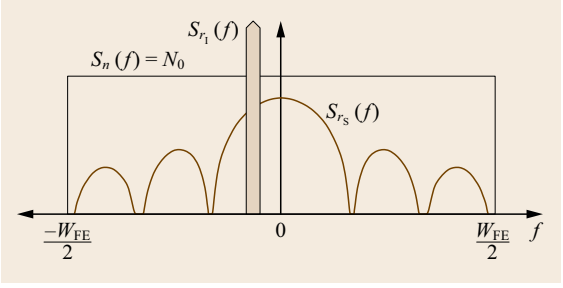

Fig. 16.1 Stylized depiction of the power spectra $S_{r_S}(f)$, $S_{r_I}(f)$, and $S_n(f)$ that correspond, respectively, to the received components $r_S(t)$, $r_I(t)$, and $n(t)$. The spectra are assumed to be significant only within the interval $|f| \le W_{FE}/2$, where W_{FE} is the bandwidth of the RF front end's narrowest bandpass filter. The total power in $S_{r_S}(f)$, $S_{r_I}(f)$, and $S_n(f)$ within this interval is, respectively, P_S, P_I, and P_n

where, ignoring the effects of band-limiting, $C_l(t)$ is often taken to be equal to $C(t)$, though it may differ from $C(t)$ when modeling early-minus-late correlation or when a specialized code replica is generated to reduce multipath. Suppose that a GNSS receiver is tracking the carrier phase of $r_S(t)$ so that $\hat{\theta}(t) \approx \theta(t)$. Then, the complex correlator output

$$Y(t) \equiv r^*(t)l(t) = S(t) + I(t) + N(t) \tag{16.3}$$

is composed of the desired component

$$S(t) \approx \sqrt{P_S} D(t-\tau)C(t-\tau)C_l(t-\hat{\tau}),$$

an interference component

$$I(t) = r_I^*(t)C_l(t-\hat{\tau})\exp(\mathrm{j}\hat{\theta}(t)),$$

and a random noise component $N(t) = n^*(t)l(t)$.

If the components $r_I^*(t)$, $C_l(t-\hat{\tau})$, and $\exp(\mathrm{j}\hat{\theta}(t))$ are wide-sense stationary and mutually statistically independent, as is a reasonable approximation for non-spoofing interference, then the autocorrelation function of $I(t)$ can be expressed as

$$\begin{aligned} R_I(\tilde{\tau}) &\equiv E[I^*(t)I(t-\tilde{\tau})] \\ &= E[r_I(t)r_I^*(t-\tilde{\tau})] \\ &\quad \times E[C_l(t-\hat{\tau})C_l(t-\tilde{\tau}-\hat{\tau})] \\ &\quad \times E(\exp(\mathrm{j}\hat{\theta}(t))\exp(\mathrm{j}\hat{\theta}(t-\tilde{\tau}))) . \end{aligned} \tag{16.4}$$

In other words, $R_I(\tilde{\tau})$ is the product of the autocorrelation functions corresponding to each of the three components of $I(t)$. Consequently, the power spectral density of $I(t)$, $S_I(f) = \mathcal{F}[R_I(\tilde{\tau})]$, where $\mathcal{F}$ denotes the Fourier transform, can be found by convolving the power spectra of the three components. Let $S_{C_l}(f)$, $S_{r_I}(f)$, and $\delta(f+\hat{f}_D)$ be the respective power spectra of $C_l(t)$, $r_I(t)$, and $\exp(\mathrm{j}\hat{\theta}(t))$, where

$$\hat{f}_D = -\frac{1}{2\pi}\frac{\mathrm{d}\hat{\theta}}{\mathrm{d}t}$$

is the receiver's estimate of the desired signal's apparent Doppler frequency, in Hz, and $\delta(f)$ is the Dirac delta function. It follows that

$$\begin{aligned} S_I(f) &= S_{C_l}(f) * S_{r_I}(f) * \delta(f+\hat{f}_D) \\ &= S_{C_l}(f) * S_{r_I}(f+\hat{f}_D) , \end{aligned}$$

where $*$ denotes convolution.

The values of $S_I(f)$ within a narrow neighborhood about $f = 0$ are a useful starting point for predicting GNSS interference effects. To understand why, consider the block diagram in Fig. 16.2, which illustrates correlation of the received signal $r(t)$ with the local signal replica $l(t)$ followed by an accumulate-and-dump operation that produces the discrete complex accumulation products $Y_k = I_k + \mathrm{j}Q_k, k = 1, 2, \ldots$. The accumulate-and-dump operation acts as a low-pass filter having a squared frequency response

$$|H_a(f)|^2 = \mathrm{sinc}^2(fT_a) ,$$

where $\mathrm{sinc}(x) \equiv \sin(\pi x)/\pi x$ and T_a is the accumulation interval in seconds. The interference power that passes through the accumulate-and-dump filter into the complex accumulation products – and thereafter into the code and carrier tracking loops – is given by

$$P_{aI} = \int_{-\infty}^{\infty} |H_a(f)|^2 S_I(f)\mathrm{d}f .$$

Let the noise-equivalent bandwidth of the accumulate-and-dump filter be defined as

$$W_a \equiv \int_{-\infty}^{\infty} \mathrm{sinc}^2(fT_a)\mathrm{d}f = \frac{1}{T_a}$$

and let $I_0 \equiv S_I(0)$. Then, so long as $S_I(f)$ is nearly constant (flat) over a few multiples of W_a, P_{aI} can be approximated as

$$P_{aI} \approx \tilde{P}_{aI} \equiv I_0 W_a .$$

For typical values of T_a, and for typical spreading code replicas $C_l(t)$, the quasi-constant condition on $S_I(f)$ is easily satisfied. To understand why, consider Fig. 16.3 in connection with the following argument. Assume that $S_{C_l}(f)$ and $S_I(f)$ are smooth (no spectral lines) with respective frequency derivatives $S'_{C_l}(f)$ and $S'_I(f)$. The error in the approximating P_{aI} by $\tilde{P}_{aI}$ can be expressed in dB as

$$\Delta P_{aI} \equiv 10\log_{10}\left(\left|\frac{\tilde{P}_{aI}}{P_{aI}}\right|\right) ,$$

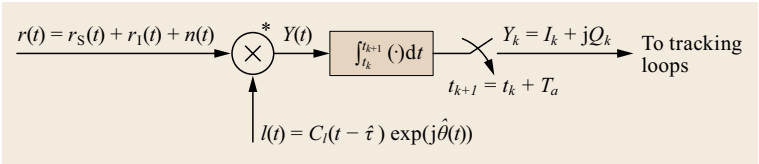

Fig. 16.2 Block diagram of the standard correlation and accumulation process in a GNSS receiver. The complex product of the incoming signal $r(t)$ and the local replica $l(t)$ is accumulated over T_a seconds to produce the discrete complex-valued accumulation product Y_k

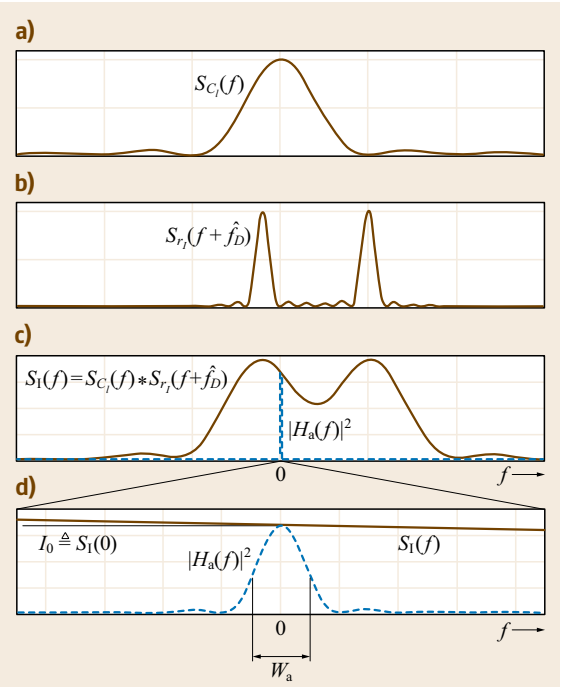

Fig. 16.3a–d Example power spectra and filtering involved in interference analysis: **(a)** $S_{C_l}(f)$, the spectrum of the GNSS replica code; **(b)** $S_{r_l}(f+\hat{f}_D)$, the spectrum of the received interference convolved with $\delta(f+\hat{f}_D)$; **(c)** $S_I(f)$, the spectrum of $I(t)$, together with $|H_a(f)|^2$, the squared frequency response of the accumulate-and-dump filter; **(d)** zoomed view of $S_I(f)$ and $|H_a(f)|^2$ near $f=0$ showing that, despite the interference being fairly narrowband, $S_I(f)$ is approximately flat over the noise-equivalent bandwidth W_a

which for practical $C_l(t)$ satisfies

$$\Delta P_{aI} < 10\log_{10}\left(1+\frac{|S_I'(0)|}{S_I(0)}W_a\right).$$

But, due to the properties of convolution,

$$\frac{|S_I'(0)|}{S_I(0)} \le \max_f \frac{|S_{C_l}'(f)|}{S_{C_l}(f)}.$$

And note that when performing the maximization, one need only consider f values within

$$\mathcal{U}_\epsilon = \{f | \epsilon < S_{C_l}(f)\}$$

for some $\epsilon > 0$ because for $f \notin \mathcal{U}_\epsilon$ the possible contribution of $|S_{C_l}'(f)|/S_{C_l}(f)$ to P_{aI} is small, making large values of $|S_{C_l}'(f)|/S_{C_l}(f)$ immaterial. Putting these pieces together, ΔP_{aI} can be upper bounded as

$$\Delta P_{aI} < \max_{f\in\mathcal{U}_\epsilon}\left[10\log_{10}\left(1+\frac{|S_{C_l}'(f)|}{S_{C_l}(f)}W_a\right)\right].$$

Consider an example designed for large ΔP_{aI}. Let $C_l(t)$ be matched to the relatively narrowband GPS L1 C/A code (ignoring spectral lines), for which

$$S_{C_l}(f) = T_C\,\mathrm{sinc}^2(fT_C)$$
$$S_{C_l}'(f) = \frac{2T_C}{f}\left[\mathrm{sinc}\,(2fT_C) - \mathrm{sinc}^2(fT_C)\right],$$

where $T_C \approx 1\,\mu s$ is the spreading code chip interval. Choosing $\epsilon = S_{C_l}(0)/100$, it can be shown that $|S_{C_l}'(f)/S_{C_l}(f)|$ achieves a maximum of approximately $25T_C$ so that, even assuming $W_a = 1\,\mathrm{kHz}$ – the widest typical accumulate-and-dump bandwidth for the GPS L1 C/A signal – the ratio ΔP_{aI}, and thus the error in approximating P_{aI} by $\tilde{P}_{aI}$, remains less than 0.105 dB, which can be considered insignificant for most applications.

The thermal-noise-equivalent approximation to interference effects can be summarized as follows. At the input to the low-pass accumulate-and-dump filter that produces the complex accumulations $Y_k = I_k + jQ_k$, the carrier-power to thermal-noise density ratio is $C/N_0 = P_S/N_0$; at the output of the filter, the signal-to-thermal-noise ratio is $\mathrm{SNR} = P_S/N_0W_a$. When, in addition to thermal noise, interference is present, then at the filter input the carrier-power to interference-and-thermal-noise ratio (CINR) can be approximated as

$$\mathrm{CINR} = \frac{C}{N_{0,\mathrm{eff}}} = \frac{P_S}{N_0 + I_0},$$

where $N_{0,\mathrm{eff}} \equiv N_0 + I_0$ is the effective thermal noise density, which accounts for both thermal noise and interference. At the filter output, the signal-to-interference-and-thermal-noise ratio can be approximated as

$$\mathrm{SINR} = \frac{P_S}{N_{0,\mathrm{eff}}W_a}.$$

Thus, apart from the limitations described below, analysis of GNSS receiver behavior in the presence of interference can proceed just as analysis of receiver behavior in the presence of thermal noise, which is well understood [16.9–11], by substituting CINR (or $C/N_{0,\mathrm{eff}}$) for C/N_0, and SINR for SNR.

16.1.3 Limits of Applicability

Approximating interference that is statistically independent of the code and carrier replicas as if it were

thermal noise with spectral density I_0 at the input of the accumulate-and-dump filter yields excellent agreement with the full theoretical error statistics for acquisition, carrier tracking, and data demodulation [16.12]. The approximation is also accurate for predicting the statistics of any coherent correlation with code replica $C_l(t)$. For example, it accurately predicts the statistics of the coherent early-minus-late code phase error so long as data bits are estimated correctly, and $C_l(t)$ is taken to be the difference between early and late code replicas [16.12]. But the thermal-noise-equivalent approximation is known to produce biased code phase error statistics for noncoherent code phase discriminators [16.13, 14]. In this case, narrowband interference maximizes code tracking error not when the interference is centered at $f = 0$ Hz (i. e., when aligned with the desired signal's carrier frequency), as one would expect, but rather when it is centered at $f \approx 1/T_C$ Hz. However, if one properly accounts for squaring loss, then even the noncoherent phase error statistics can be reduced to an accurate thermal-noise-equivalent representation [16.12]. In short, the thermal-noise-equivalent approximation has wide applicability for analysis of interference effects.

It is worth noting that if the received interference $r_I(t)$ is not statistically independent of $C_l(t-\hat{\tau})$ and $\exp(j\hat{\theta}(t))$, then factorization of $R_I(\tilde{\tau})$ as in (16.4) is not possible and the thermal-noise-equivalent approximation is not valid. This case arises, for example, when the interference is structurally similar to the desired signal $r_S(t)$ and is approximately code-phase aligned with $r_S(t)$ – in other words, when the interference is a spoofing signal. For this reason, spoofing-type interference will be treated separately later in this chapter; meanwhile, all $r_I(t)$ will be assumed to be independent of $C_l(t-\hat{\tau})$ and $\exp(j\hat{\theta}(t))$. Furthermore, all code and carrier-phase measurements will be assumed to be produced by coherent phase discriminators. Under these conditions, the thermal-noise-equivalent approximation whereby CINR is substituted for C/N_0 can be expected to accurately predict receiver effects.

16.1.4 Overview of Interference Effects on Carrier Phase Tracking

Assuming the thermal-noise-equivalent approximation to be valid, this subsection gives an overview of interference effects on carrier-phase tracking. Attention is focused on phase tracking because the phase-tracking loop, or phase lock loop (PLL), is the weakest link in the signal tracking chain. Typically, if the PLL can maintain lock, then a frequency-tracking loop and a code-phase-tracking loop can as well.

Phase Error Variance

Consider a standard (nonsquaring) PLL with true phase input $\theta(t)$ and phase estimate $\hat{\theta}(t)$. When the phase error $\varphi(t) = \theta(t) - \hat{\theta}(t)$ is small enough that the PLL's phase detector can be regarded as linear, then, for zero-mean white driving noise, the PLL's phase error variance $\sigma_\varphi^2 = E[\varphi^2(t)]$ (in rad^2) is accurately approximated by [16.15]

$$\sigma_\varphi^2 = \frac{B_n N_0}{C} \equiv \frac{1}{\rho_L} , \tag{16.5}$$

where B_n is the PLL's single-sided noise bandwidth and ρ_L is the loop SNR. GNSS carrier-phase tracking of data-modulated signals requires a squaring (e.g., Costas) PLL, which is insensitive to the half-cycle phase changes induced by the data modulation. In a squaring PLL, the actual phase error tracked is 2φ, with the corresponding variance denoted by $\sigma_{2\varphi}^2$. Furthermore, ρ_L is reduced by a squaring loss factor approximately equal to [16.16]

$$S_L = \left(1 + \frac{N_0}{2T_a C}\right)^{-1} ,$$

where $1/T_a$ is the predetection bandwidth. Thus, for the squaring loop,

$$\sigma_\varphi^2 = \frac{\sigma_{2\varphi}^2}{4} = \frac{1}{\rho_L S_L}$$

is a useful approximation for σ_φ^2 in the linear regime. For analysis of the squaring loop, an equivalent loop SNR is defined as [16.17, p. 206]

$$\rho_{eq} \equiv \frac{\rho_L S_L}{4} , \tag{16.6}$$

which leads to $\rho_{eq} \approx 1/\sigma_{2\varphi}^2$ for small φ.

At large values of φ, the assumption of PLL linearity breaks down and analysis becomes more difficult. An exact expression for σ_φ^2 for a first-order nonsquaring PLL driven by white Gaussian noise is found in [16.18, Chap. 4]. Precise phase error statistics for all but this standard first-order loop are typically obtained via simulation. Fortunately, one can show that the exact phase error variance for the standard first-order loop is a reasonable proxy for that of higher-order loops. Thus, one can identify the region of approximate linear PLL operation by noting that, for the standard first-order loop, the linear model in (16.5) is reasonably accurate (within 20%) for $\rho_L > 4$, or $\sigma_\varphi < 28.6°$ [16.18, Chap. 4]. Likewise, a squaring loop behaves approximately linearly for $\rho_{eq} > 4$, or $\sigma_\varphi < 14.3°$.

Cycle Slipping

A PLL's phase detector is periodic, meaning that it cannot distinguish between the phase errors φ and $\varphi + 2n\pi$ (nonsquaring loop) or φ and $\varphi + n\pi$ (squaring loop), where n is an integer. As a result, an infinite set of stable attractors exists for the nonlinear difference equations that describe the PLL error dynamics. At low loop SNR, the phase error can slip from one stable attractor to another, leading to infinite σ_φ^2 in the steady state. This is the familiar cycle slip phenomenon associated with PLLs [16.19, 20], [16.15, Chap. 6].

The mean time to first cycle slip T_s is defined as the average time required for the loop phase error to reach $\pm 2\pi$ ($\pm\pi$ for the squaring loop) for the first time, starting from an initial condition of zero phase error. For first-order loops, and in other cases where cycle slips occur as isolated events, T_s is the same as the mean time between cycle slips; if cycle slips occur in bursts – as may happen for $\rho_L, \rho_{eq} < 5$ in second- or higher-order loops – then T_s and the mean time between cycle slips are not related simply [16.20].

As with the calculation of σ_φ^2, an analytical solution for T_s has only been possible for the simple case of a first-order unstressed (zero static phase error) PLL driven by white Gaussian noise, in which case [16.18, p. 101]

$$T_s = \frac{\pi^2 \rho_L I_0^2(\rho_L)}{2B_n} \tag{16.7}$$

is the time to first slip/mean time between slips for a nonsquaring loop, $I_0(\cdot)$ being a modified Bessel function of the first kind. An approximate T_s for first-order squaring loops is obtained by substituting ρ_{eq} for ρ_L. Unstressed second- and higher-order loops have lower T_s than unstressed first-order loops, and stressed loops are more prone to cycle slipping than unstressed loops; nonetheless, (16.7) remains a useful upper bound. For GNSS applications, a second- or third-order loop is required to accurately track carrier-phase in the presence of Doppler-induced quadratic phase growth. In fact, even the second-order loop experiences significant loop stress ($\approx 1^\circ$ static phase error) during the largest GNSS line-of-sight accelerations. Only the third-order loop maintains near-zero static phase error for all GNSS geometries.

Frequency Unlock

The general term *phase unlock* refers to single or successive cycle slips. At very low loop SNR, a PLL may never recover phase lock after a long succession of cycle slips. This phenomenon, called *drop lock* in the PLL literature, is related to the PLL's frequency pull-in range. For reasons that will become clear, the term *frequency unlock* is a more precise descriptor than drop lock for the phenomenon as it relates to the discrete-time PLLs used in modern GNSS receivers.

A PLL's frequency pull-in range is the maximum frequency step input that a PLL is able to *pull in* and eventually achieve phase lock. For example, a continuous-time first-order nonsquaring PLL has a pull-in range equal to the loop gain K [16.19]. For higher-order PLLs, the frequency pull-in range can be thought of as the maximum tolerable mismatch $\Delta\omega = |\omega_c - v|$ between the carrier frequency ω_c and the PLL's internal estimate of carrier frequency v, assuming that higher-order loop filter states (e.g., the estimate of carrier frequency rate) are relaxed, where applicable.

Continuous-time PLLs whose loop filters contain one or more perfect integrators have an infinite frequency pull-in range [16.15, Chap. 8]. On the other hand, the frequency pull-in range of second- and higher-order discrete-time PLLs is limited by the loop update (accumulation) interval T_a. When the frequency mismatch $\Delta\omega$ exceeds a certain threshold $\Delta\omega_m$, then v is attracted toward a stable equilibrium value that satisfies $T_a\Delta\omega = n\pi$ (nonsquaring loop) or $T_a\Delta\omega = n\pi/2$ (squaring loop), $n = 1, 2, 3, \ldots$. Intuitively, these equilibrium values exist because the loop cannot detect a phase error change of $2n\pi$ (nonsquaring loop) or $n\pi$ (squaring loop) between loop updates. The value of $\Delta\omega_m$ is a function of the particular loop configuration. It can be surprisingly small for PLLs common in GNSS receivers: for a third-order Costas loop with $T_a = 10\,\text{ms}$ and $B_n = 10\,\text{Hz}$, $\Delta\omega_m = 81\,\text{rad/s} \approx 13\,\text{Hz}$. At very low loop SNR, cycle slips can occur in bursts as noise and phase dynamics force v momentarily away from ω_c [16.20]. If, due to such forcing, $\Delta\omega$ exceeds $\Delta\omega_m$, then there is a high probability that v will become trapped at one of the incorrect stable equilibrium values. Thus, the PLL experiences frequency unlock.

Frequency unlock and momentary phase unlock have rather different practical consequences. Unlike momentary phase unlock (i. e., cycle slipping), frequency unlock often leads to complete loss of the GNSS signal link – a result of signal attenuation due to frequency detuning. If v settles on an equilibrium value such that $n \geq 2$ (nonsquaring loop) or $n \geq 4$ (squaring loop), then the baseband signal power drops by more than 13 dB, making it likely that the PLL will experience further frequency detuning and eventually lose the signal entirely. Worse yet, re-acquisition may not be possible at low SNR.

16.2 Canonical Interference Models

16.2.1 Wideband Interference

The simplest variants of $r_I(t)$ are the extreme cases of wideband and narrowband interferences. Consider first wideband interference. Suppose that $r_I(t)$ is spectrally flat with power density $S_I(f) = P_I/W_{FE}$ over a two-sided front-end bandwidth $W_{FE} >> 1/T_C$, where T_C is the chip interval of $C(t)$ (e.g., $1/T_C = 1.023\,\text{MHz}$ for the GPS L1 C/A code). In this case, $S_I(f) = S_{C_l}(f) * S_{r_I}(f + \hat{f}_D) \approx S_{r_I}(f) = P_I/W_{FE}$, which implies that $I_0 \equiv S_I(0) = P_I/W_{FE}$. Hence, post-correlation error analysis can proceed by approximating the carrier-to-noise ratio as

$$\text{CINR} = \frac{C}{N_{0,\text{eff}}} = \frac{P_S}{N_0 + P_I/W_{FE}} \,. \tag{16.8}$$

Continuous Gaussian wideband interference is interesting because it is dense in both frequency and time and its amplitude distribution is shaped like that of receiver thermal noise. Thus, from the perspective of an adversarial jammer, wideband Gaussian interference is a conservative strategy: Although it demands significant power, it affords receivers in the target area no more effective interference mitigation techniques than those commonly applied for weak GNSS signal tracking.

16.2.2 Narrowband Interference

Suppose $r_I(t)$ is a narrowband interference signal offset by f_I Hz from the GNSS carrier frequency. As an extreme case, consider perfect tone interference

$$\begin{aligned} r_I(t) &= \sqrt{P_I}\exp(\mathrm{j}2\pi f_I t) \\ S_{r_I}(f) &= P_I\delta(f - f_I) \,. \end{aligned}$$

In this case, the power spectrum $S_I(f)$ is simply a scaled and frequency-shifted version of $S_{C_l}(f)$

$$\begin{aligned} S_I(f) &= S_{C_l}(f) * S_{r_I}(f + \hat{f}_D) \\ &= P_I S_{C_l}(f) * \delta(f + \hat{f}_D - f_I) \\ &= P_I S_{C_l}(f + \hat{f}_D - f_I) \,. \end{aligned}$$

Smooth Spectrum Approximation

As a first approximation, let $S_{C_l}(f)$ be any smooth (no spectral lines) function with an equivalent rectangular bandwidth of $W_C > 2|f_I|$. Then, interference power P_I/L_C passes into the correlation products, where $L_C = W_C/W_a$ is termed the spread-spectrum processing gain. In this approximation, $I_0 = P_I/W_C$, so that

$$\text{CINR} = \frac{P_S}{N_0 + P_I/W_C} \,.$$

For a large jamming-to-signal power ratio P_I/P_S, N_0 becomes negligible compared with P_I/W_C, in which case CINR can be approximated as

$$\text{CINR} = 10\log_{10}(W_C) - 10\log_{10}\left(\frac{P_I}{P_S}\right) \quad \text{dB Hz} \,.$$

For example, if $W_C = 1\,\text{MHz}$, then a tone interference source with a jamming-to-signal power ratio of $P_I/P_S = 25\,\text{dB}$ would result in a CINR of approximately $60 - 25 = 35\,\text{dB Hz}$.

Moving toward a more accurate analysis of tone interference, consider now the actual shape of $S_{C_l}(f)$ while retaining the assumption of smoothness (no spectral lines). In particular, suppose that $S_{C_l}(f) = T_C\text{sinc}^2(fT_C)$, which would be the case for a local replica matched to a random binary spreading code $C(t)$ with chip interval T_C. Then, for tone interference with power P_I it follows that

$$\begin{aligned} S_I(f) &= P_I S_{C_l}(f) * \delta(f + \hat{f}_D - f_I) \\ &= P_I T_C \text{sinc}^2[(f + \hat{f}_D - f_I)T_C] \,. \end{aligned}$$

From this expression, it is clear that the tone interference will minimize CINR (by maximizing $I_0 \equiv S_I(0)$) when $f_I = \hat{f}_D$. In other words, under the smooth spectrum approximation with $S_{C_l}(f) = T_C\text{sinc}^2(fT_C)$, the greatest degradation to CINR occurs when the tone is aligned with the Doppler-shifted carrier frequency of the desired signal.

One can apply a similar analysis to modern GNSS signals with binary offset carrier (BOC) spreading code modulation. In this case, the worst-case tone interference occurs when f_I coincides with the Doppler-shifted peak of one of the offset side lobes. However, due to the additional spreading afforded by BOC-type signals, the resulting interference is, in general, less severe than for a sinc^2-type waveform with equivalent T_C [16.21].

Effect of Spectral Lines

The smooth-spectrum approximation is appropriate for pseudorandom spreading codes $C(t)$ with a long code repetition period, such as the encrypted legacy military GPS spreading codes, for which the period is not publicly known but surely exceeds one week [16.22], and for the GPS L2CL code, which has a period of 1.5 s [16.23]. For short-period pseudorandom codes, however, the approximation is not appropriate because interference can be narrower than the spacing between spectral lines. Assume that $C(t)$ is a repeating code with period $T_p = T_C N_p$, where $N_p \in \mathbb{N}$ is the number of chips per code period. As a periodic function, $C(t)$ can be

decomposed as a Fourier series, which means that its power spectrum $S_C(f)$ is expressable as a weighted sum of Dirac delta functions

$$S_C(f) = \sum_{i=-\infty}^{\infty} c_i \delta\left(f - i\Delta f_\mathrm{p}\right), \quad i \in \mathbb{Z} \tag{16.9}$$

with constraint

$$\sum_{i=-\infty}^{\infty} c_i = 1$$

and spectral line spacing $\Delta f_\mathrm{p} = 1/T_\mathrm{p}$. Assuming a matched local code replica $[C_l(t) = C(t)]$, Fig. 16.4 shows the spectral line structure of $S_{C_l}(f)$ for an example GPS L1 C/A code.

For tone interference $S_{r_\mathrm{I}}(f) = P_\mathrm{I}\delta(f - f_\mathrm{I})$, $S_\mathrm{I}(f)$ is simply a scaled and shifted version of $S_{C_l}(f)$

$$\begin{aligned} S_\mathrm{I}(f) &= S_{C_l}(f) * S_{r_\mathrm{I}}\left(f + \hat{f}_D\right) \\ &= P_\mathrm{I} S_{C_l}\left(f - f_\mathrm{I} + \hat{f}_D\right). \end{aligned} \tag{16.10}$$

Interestingly, if none of the tines in the comb of spectral lines that constitute $S_\mathrm{I}(f)$ falls within the passband of the accumulate-and-dump filter $H_\mathrm{a}(f)$, then the tone interference will have a negligible effect on the accumulation products. This can be quantified probabilistically as follows. If the frequency offset f_I is modeled as a random variable uniformly distributed over a range wider than Δf_p, then the probability that one of the spectral lines in $S_\mathrm{I}(f)$ will fall within the noise-equivalent bandwidth W_a of the accumulate-and-dump filter is

$$P_X = \left[\mathrm{mod}(|f_\mathrm{I}|, \Delta f_\mathrm{p}) \le W_\mathrm{a}\right] = \frac{W_\mathrm{a}}{\Delta f_\mathrm{p}}.$$

For N_s signals tracked, each with independent random f_D, the probability of significant interference in any tracking channel rises to

$$P_{X_\mathrm{T}} = 1 - (1 - P_X)^{N_\mathrm{s}}.$$

By way of example, for GPS L1 C/A-code tracking with $T_\mathrm{a} = 20\,\mathrm{ms}$ and $N_\mathrm{s} = 10$, $P_X = 0.05$ for each tracking channel and $P_{X_\mathrm{T}} = 0.4$ for the ensemble.

From (16.9) and (16.10), it is evident that tone interference is most damaging when f_I is aligned with the Doppler-shifted spectral line having the largest weighting coefficient c_i. For example, for the spectrum shown in Fig. 16.4, the largest c_i, located at $\pm 72\,\mathrm{kHz}$, is 23 dB below the total power in $S_{C_l}(f)$. Therefore, when targeting this signal, a tone interferer with power P_I would be attenuated by at least 23 dB before passing into the accumulate-and-dump filter. (Interestingly, tone interference targeting a C/A signal at exactly the Doppler-shifted L1 carrier frequency is ineffective because the balanced C/A Gold codes, which have only one more 1 than 0, produce a nearly insignificant $-60.2\,\mathrm{dB}$ line component at zero offset.) In general, the largest spectral line components among all GPS L1 C/A Gold codes attenuate tone interference by only 18.3 dB [16.24]. By way of comparison, a perfectly random code sequence with the same chip interval ($T_\mathrm{C} \approx 1\,\mu\mathrm{s}$) would attenuate the interferer by at least 60 dB.

In general, one can say that spectral lines in $S_{C_l}(f)$ have two contrary effects on tone interference: (1) line sparsity reduces the probability that interference will have a significant effect – most likely the interference will fall harmlessly between the lines, but (2) in the event that tone interference does coincide with a powerful line component, the interference effect is severe.

Of course, pure tone interference is only a convenient fiction; all interference encountered in practice will have a nonzero spectral width. Convolving an arbitrary $S_{r_\mathrm{I}}(f)$ with an $S_{C_l}(f)$ of the form in (16.9) results in an interference spectrum of the form

$$\begin{aligned} S_\mathrm{I}(f) &= S_{C_l}(f) * S_{r_\mathrm{I}}\left(f + \hat{f}_D\right) \\ &= P_\mathrm{I} \sum_{i=-\infty}^{\infty} c_i S_{r_\mathrm{I}}\left(f - \Delta f_\mathrm{p} + \hat{f}_D\right). \end{aligned} \tag{16.11}$$

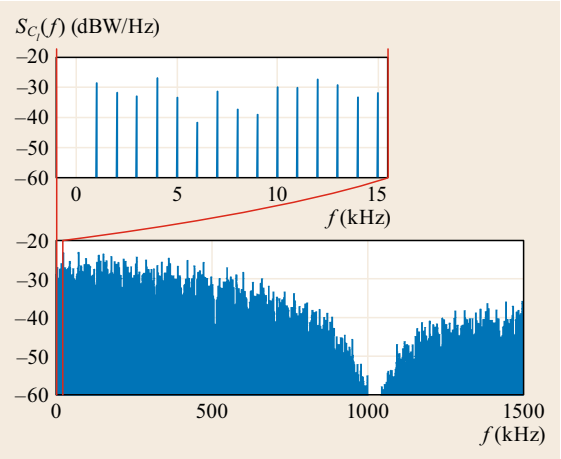

Fig. 16.4 Power spectrum $S_{C_l}(f)$ corresponding to the GPS L1 C/A code replica for pseudo-random number sequence (PRN) 31. The units of $S_{C_l}(f)$ assume that the power of $C_l(t)$ is normalized to 1 W. Because $S_{C_l}(f) = S_{C_l}(-f)$, only positive frequencies are shown. Bottom panel: The interval $0 \le f \le 1500\,\mathrm{kHz}$ showing the code's approximate $T_C \mathrm{sinc}^2(fT_C)$ spectral envelope. Top panel: Expanded view of the first 15 kHz, showing distinct spectral lines with irregular weighting spaced at $\Delta f_p = 1/T_p = 1\,\mathrm{kHz}$

Thus, each tine in the comb now assumes the shape of $S_{r_I}(f)$. For interference that is narrow with respect to Δf_p, each tine remains distinct from its neighbors and is weighted according to the corresponding c_i; as the interference widens, the tines blend together and the spectrum flattens.

16.2.3 Matched-Spectrum Interference

An inescapable property of multiaccess spread-spectrum systems such as GNSS is that, from the perspective of a receiver channel tracking a particular GNSS signal (a unique combination of spreading code and center frequency), all other signals at the same frequency act as interference. Moreover, many of these interfering signals will have a power spectrum that is closely matched with that of the desired signal. This matched-spectrum interference is a particularly potent nuisance because it allocates power, as a function of frequency, in exact proportion to the weighting that the receiver applies with its local replica in attempting to track the desired signal. Thus, the most powerful spectral lines – the most important contributors to the total received GNSS signal power – are affected by the greatest amount of noise. In recognition of this, adversarial interferers often adopt matched-spectrum interference as their waveform of choice. In the case of nonmalicious intrasystem (e.g., within GPS) or intersystem (e.g., between GPS and Galileo) interference, the competing waveforms are by design weak and approximately power-matched so that the interference is small compared to the ever-present thermal noise, though not entirely insignificant – especially with the proliferation GNSS satellites.

When matched-spectrum interference originates from GNSS satellites, it is termed multiaccess interference. As an illustration of the effects of such interference, consider a pseudorandom binary spreading code whose power density under a smooth-spectrum approximation is

$$S_C(f) = P_C T_C \mathrm{sinc}^2(fT_C) \,,$$

where P_C is the received signal power and T_C is the spreading code chip interval. This model applies, for example, to the spreading codes of GPS L1 C/A and P(Y), L2 C and P(Y), and L5 I and Q. Assume, for simplicity, that the receiver's power-normalized code replica is perfectly matched to the incoming code so that $S_C(f) = P_C S_{C_l}(f)$ (i. e., band-limiting effects in the RF front end are ignored).

Treating $S_C(f)$ as an interference spectrum and assuming $\hat{f}_D$ is negligible compared to the bandwidth of $S_C(f)$, we have

$$\begin{aligned} S_I(f) &= S_{C_l}(f) * S_{r_I}(f) \\ &= P_C S_C(f) * S_C(f) \\ &= P_C \int_{-\infty}^{\infty} S_C(f-\nu) S_C(\nu) \mathrm{d}\nu \\ &= P_C \int_{-\infty}^{\infty} S_C(\nu-f) S_C(\nu) \mathrm{d}\nu \,, \end{aligned}$$

where the last equality follows from $S_C(f) = S_C(-f)$. Hence,

$$\begin{aligned} I_0 \equiv S_I(0) &= P_C \int_{-\infty}^{\infty} S_C^2(\nu) \mathrm{d}\nu \\ &= P_C \int_{-\infty}^{\infty} \left[T_C \mathrm{sinc}^2(\nu T_C)\right]^2 \mathrm{d}\nu \end{aligned}$$

which, by the change of variables $q = \nu T_C$, becomes

$$I_0 = P_C T_C \int_{-\infty}^{\infty} \mathrm{sinc}^4(q) \mathrm{d}q = \left(\frac{2}{3}\right) P_C T_C \,.$$

Thus, the effect of a single multiaccess interference signal with received power P_C is to raise the effective thermal noise density from N_0 to

$$N_{0,\mathrm{eff}} = N_0 + \left(\frac{2}{3}\right) P_C T_C \,.$$

The significance of multiaccess interference is measured with respect to N_0. Suppose there are M multiaccess signals whose average received power is $\bar{P}_C$. Then, from the perspective of a single desired signal, the multiaccess power density becomes equivalent to N_0 when

$$\left(\frac{2}{3}\right) \bar{P}_C T_C (M-1) = N_0 \,.$$

Thus, to ensure that multiaccess density does not exceed N_0 requires

$$M \le 1 + \frac{3/2}{\left(\bar{P}_C/N_0\right) T_C} \,.$$

Figure 16.5 shows this bound for $T_C = 1\,\mu\mathrm{s}$, which applies to GPS L1 and L2C, and for $T_C = 0.1\,\mu\mathrm{s}$, which

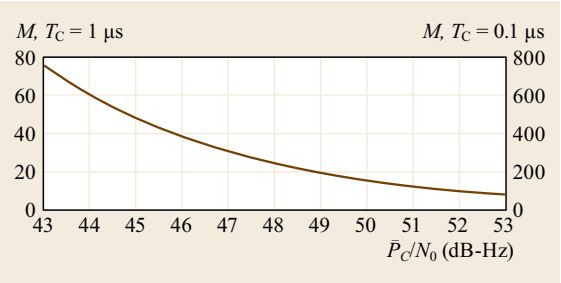

Fig. 16.5 Maximum number of simultaneously received multiaccess GNSS signals with power spectrum $S_C(f) = P_C T_C \mathrm{sinc}^2(fT_C)$ such that $I_0 \le N_0$, as a function of $\bar{P}_C/N_0$, where $\bar{P}_C$ is the average power of the $M-1$ multiaccess interferers. The left- and right-hand scales correspond, respectively, to $T_C = 1\,\mu s$ and $T_C = 0.1\,\mu s$

applies to GPS L5 I and Q. Assuming that, for the average user, the number of received signals M is approximately one-fourth of the total number of orbiting GNSS satellites and that $\bar{P}_C = 47\,\mathrm{dB\,Hz}$, and assuming all satellites broadcast only the GPS L1 C/A signal, the multiaccess interference density exceeds N_0 when the constellation size grows beyond 124 satellites.

It is worth noting that, although a 3 dB rise in the effective thermal noise floor (from N_0 to $N_0 + I_0 = 2N_0$) is significant, most GNSS users would gladly trade this degradation for the vastly improved dilution of precision and reduced convergence times for carrier-phase differential GNSS (CDGNSS) positioning and precise point positioning (PPP) that a larger multi-GNSS constellation would afford.

Finally, observe that, from the perspective of an adversarial interferer, matched-spectrum interference is the most efficient use of transmit power among all interference waveforms. For example, in the case of a local replica with density $S_{C_l}(f) = T_C \mathrm{sinc}^2(fT_C)$, it can be shown that for a fixed interference power P_I, the interference density assumes its maximum value $I_0 = (2/3)P_I T_C$ when $S_{I_1}(f) = P_I T_C \mathrm{sinc}^2(fT_C)$.

16.3 Quantization Effects

The effect of signal quantization on interference depends less on the bandwidth of the interference – whether wideband or narrowband – than on its amplitude distribution. The salient result in this regard is as follows: For white, Gaussian-distributed interference, the quantizer's output SNR is always degraded relative to its input SNR, whereas for constant-amplitude interference (e.g., a swept tone), the quantizer output SNR can actually exceed its input SNR. In any case, an optimal quantization strategy seeks to minimize the SNR degradation through the quantizer.

16.3.1 One-Bit Quantization

If the discrete samples entering a one-bit (two-level) quantizer are Gaussian distributed and uncorrelated, then the SNR is degraded by a factor $2/\pi$ or $-1.96\,\mathrm{dB}$ [16.25]. Designers of low-cost GNSS receivers often view this modest loss as a small price to pay for a one-bit quantizer's economy of implementation and low power consumption, which explains the popularity of one-bit quantization in consumer devices.

However, one-bit quantization performs poorly in the presence of strong tone interference [16.24]. To understand why, consider a simple case in which thermal noise is absent and a pure tone interference signal is received phase coherently (in-phase) with the carrier of a desired biphase-modulated GNSS signal. In this case, it is clear that, if the interference amplitude α is greater than the GNSS signal amplitude, then the interference completely suppresses the GNSS signal in one-bit quantization because the signal's noise-free biphase transitions are dominated at every sampling instant by the coherent interference.

In the presence of thermal noise, the desired GNSS signal is no longer completely suppressed by coherent tone interference, but the quantizer SNR degradation remains severe whenever $\alpha > \sigma$, where σ is the thermal noise standard deviation. Note that if the tone interference is out of phase by some angle θ, then its effective amplitude becomes $\alpha \cos\theta$. Thus, if θ is slowly varying and $\alpha > \sigma$, then the GNSS signal is periodically suppressed. When θ varies rapidly compared to the reciprocal integration time $1/T_a$, as with tone interference significantly offset from the desired GNSS signal carrier frequency – or, more generally, with any constant-amplitude interference – SNR degradation is less severe than in the case of coherent tone interference but still increases rapidly with increasing $\alpha > \sigma$.

It follows from these observations that one-bit quantization is a serious design flaw for receivers meant to operate in the presence of strong constant-amplitude interference.

16.3.2 Multibit Quantization

Multibit quantization is preferable to one-bit quantization when constant-amplitude interference may be

present. Not only can multibit quantization prevent total suppression of the desired GNSS signal, but, with properly chosen quantization levels, it can substantially suppress constant-amplitude interference.

Two-bit (four-level) quantization is an especially attractive option for GNSS receivers because it is simple to implement and amenable to low-power processing yet yields significantly less SNR degradation than one-bit quantization in wideband Gaussian noise (0.55 dB versus 1.96 dB [16.24, 26, 27]). The two-bit quantization function $q_2(x)$ is graphically shown in Fig. 16.6. For uncorrelated zero-mean Gaussian noise with standard deviation σ, both the minimum mean-square-error distortion criterion [16.28] and the minimum SNR degradation criterion [16.26] (in the limit of low SNR) are optimized when the magnitude threshold is chosen as $L = 0.98\sigma$ and the ratio of the quantization levels is approximately $a_2/a_1 = 3.3$. This remains true whether the noise is thermal in origin (i. e., proportional to the receiver system temperature) or is a combination of thermal noise and ambient interference, so long as the combined noise-plus-interference amplitude distribution remains Gaussian and sample-wise uncorrelated. Implementation of this quantization strategy within a GNSS receiver is typically realized by setting $a_1 = 1$, $a_2 = 3$ and adjusting the automatic gain control (AGC) so that $|q_2(x)| = a_2$ with probability 0.33.

When significant non-Gaussian interference is present in the received analog signal, the probability distribution $p(x)$ of the input to the quantizer is no longer approximately Gaussian and the above values for a_1, a_2, and L become suboptimal. If $p(x)$ is known, then new mean-square-distortion-minimizing values can be calculated numerically as described in [16.28]. For the special case of unity-amplitude tone interference with a phase that varies rapidly relative to $1/T_a$, and in the limit of low SNR, $p(x)$ assumes the shape shown in Fig. 16.7. In this case, it can be shown numerically that the mean-square distortion is minimized when $L = 0.573$ and $a_2/a_1 = 2.89$. But, importantly, and in contrast to the Gaussian noise-plus-interference case, these distortion-minimizing values do not also minimize SNR degradation. Instead, for spread-spectrum signals with large processing gain (such as GNSS signals), SNR degradation is minimized as L approaches the upper limit of $p(x)$ [16.26]. The key insight is that, for this choice of L, the quantizer maximizes the number of captured code transitions, as illustrated in Fig. 16.8.

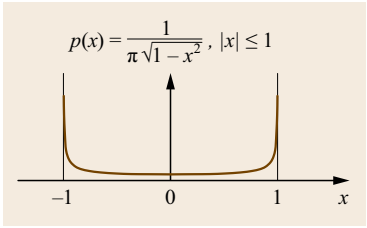

Fig. 16.7 Probability distribution of the quantizer input x for unity amplitude tone interference in the limit of low SNR

More generally, a properly configured multibit quantizer exhibits *negative* SNR degradation (i. e., there is a positive *conversion gain*) when the incoming interference has a fixed amplitude (e.g., a swept tone). This result holds even when the interference is a combination of fixed-amplitude and Gaussian interference, so long as the fixed-amplitude interference dominates [16.29]. This contrasts with Gaussian interference, for which a two-bit quantizer's output SNR is always degraded by at least 0.55 dB relative to its input SNR.

Within a GNSS receiver, adaptive two-bit quantization for suppression of constant amplitude interference can be implemented as follows. When significant constant-amplitude interference is detected, the adaptive quantizer raises the threshold L from the Gaussian-noise-optimized value for L (approximately

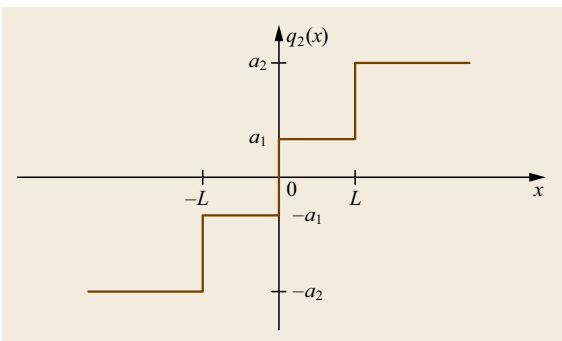

Fig. 16.6 Quantization function $q(x)$ for two-bit (four-level) quantization, showing the magnitude threshold L and the quantization levels $\{-a_2, -a_1, a_1, a_2\}$

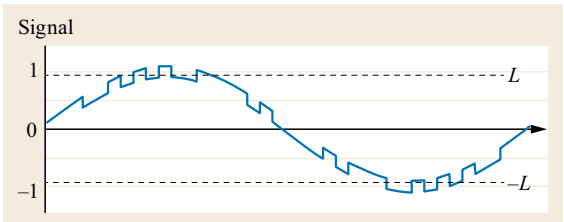

Fig. 16.8 Example threshold value L for two-bit quantization of a binary spread-spectrum signal in the presence of strong unity-amplitude tone interference. As the signal-to-interference power ratio decreases from the −20 dB ratio shown, the curve's distribution approaches that of Fig. 16.7, and the optimal value of L approaches 1

$L = \sigma$) to a new value that places L near the edge of the $p(x)$ distribution (equivalently, the AGC can lower its gain until this condition is reached). The optimal value of L depends on the relative strengths of the GNSS signal, the constant-amplitude interference, and the Gaussian noise and interference. Figure 16.9 shows the quantizer conversion gain for several example scenarios with different relative signal, noise, and interference strengths. A simple suboptimal approach sets L so that $|q(x)| = a_2$ with a predetermined probability (e.g., 10%); in an alternative, higher-performance approach, a feedback signal from the GNSS receiver's baseband processor adjusts L to maximize the average C/N_0 of the tracked GNSS signals. Note that as the constant-amplitude interference power increases relative to the Gaussian interference, the quantizer can more effectively suppress the former, but its performance becomes more sensitive to choice of L. For best performance, the ratio a_2/a_1 should also be adjusted upward from its Gaussian-adapted setpoint (approximately $a_2/a_1 = 3$), but this is less important than adjusting L. An example of adaptive multibit quantization implementation can be found in [16.14, Fig. 6.1].

Three-bit (8-level) and higher quantization bring further reduction of SNR degradation for all interference and noise types, but the marginal improvement above two-bit quantization is modest and decreases rapidly with additional bits. In uncorrelated Gaussian noise and interference, the SNR degradation through a three-bit quantizer is 0.272 dB (versus 0.55 dB for a two-bit quantizer) [16.27]. Details on three-bit quantizer performance can be found in [16.24].

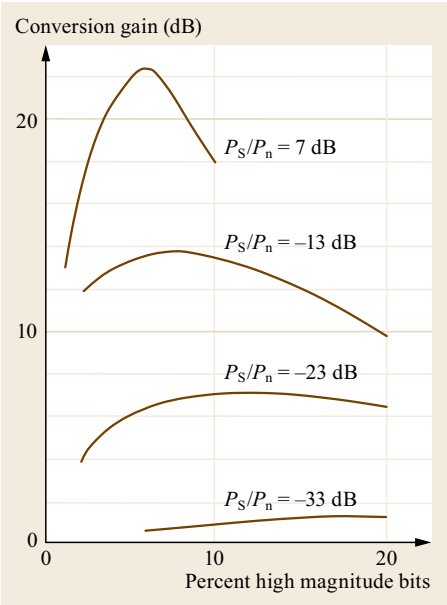

Fig. 16.9 Two-bit quantizer conversion gain (ratio of quantizer output SNR to input SNR) for a scenario in which the incoming spread-spectrum signal is corrupted by both Gaussian noise (or interference) and constant-amplitude interference, as a function of the percentage of high magnitude bits (percentage of samples for which $|q(x)| = a_2$). The different curves correspond to different values of the signal power to Gaussian noise (or interference) ratio P_S/P_n. For all curves, the ratio of the signal power to the constant-amplitude interference is $P_S/P_{ca} = -40$ dB, and $a_2/a_1 = 8$ (after [16.29], courtesy of the Institute of Electrical and Electronics Engineers (IEEE))

16.4 Specific Interference Waveforms and Sources

16.4.1 Solar Radio Bursts

Solar radio bursts (SRBs) are intense outbursts of radio emissions from the Sun, with spectral power ranging from HF to above the L band. They are typically associated with solar flares, which are caused by the acceleration of electrons in the solar atmosphere and whose rate of occurrence follows the 11 yr sunspot cycle [16.30, 31]. SRBs' jamming effect on radio equipment was first noted during World War II when strong SRBs jammed British anti-aircraft radar on many occasions [16.32]. SRBs can cause greater than 10 dB fades in a GNSS signal's C/N_0 [16.33, 34].

Given their broad-spectrum power distribution, SRBs are typically modeled as contributing to a receiver's thermal noise $n(t)$. In particular, they raise a GNSS receiver's antenna temperature T_A, which is related to the receiver's noise density N_0 by

$$N_0 = k_B(T_R + T_A) ,$$

where k_B is Boltzmann's constant and T_R and T_A are respectively the receiver and antenna noises in degrees Kelvin. T_R is the equivalent temperature of noise sources internal to the receiver, primarily those in the first-stage low-noise amplifier (LNA). T_A is the temperature equivalent of noise impinging on the antenna, including radiation from the warm Earth, cosmic noise, and solar radio noise. T_A varies with antenna motion (as more or less warm Earth radiation is visible), antenna blockage (e.g., an increase in T_A due snow accumulation [16.35]), and variable solar radiation. Note that

these are difficult or impossible for a stand-alone (non-networked) GNSS receiver to predict. Of these, solar radiation is least site-specific: All GNSS receivers in view of the Sun are similarly affected.

To judge the impact of SRBs on GNSS receivers, it is instructive to examine the rate of occurrence of those SRBs that significantly increase a receiver's P_T. Such events not only reduce C/N_0 but also lead to false alarms in received power monitoring, a technique whereby intentional interference is detected based solely on P_T (discussed further in Sect. 16.6.2). Table 16.2 shows the SRB occurrence rate for three different levels of increased P_T. Let $P_T/P_{T,nom}$ be the ratio of received power in the presence of a SRB to nominal received power. Assume that non-SRB interference is negligible so that $P_I = 0$, leaving $P_T = P_S + P_n$, where

$$P_n = W_{FE}N_0 = W_{FE}k_B(T_R + T_A) \, .$$

Let the antenna temperature be $T_A = T_{A0} + T_{As}$, where T_{A0} is a nominal value for T_A and T_{As} is the increase in T_A due to solar radiation.

Table 16.2 is interpreted as follows. Each value of $P_T/P_{T,nom}$ can be related to a value of T_{As} by

$$\frac{P_T}{P_{T,nom}} = \frac{P_S + k_B B(T_R + T_{A0} + T_{As})}{P_S + k_B B(T_R + T_{A0})}$$

assuming the following reasonable parameter values: $P_S = -146\,\text{dBW}$, $W_{FE} = 2\,\text{MHz}$, $T_R = 188\,\text{K}$, and $T_{A0} = 100\,\text{K}$. Each T_{As}, in turn, is related to a change in C/N_0 by

$$\Delta C/N_0 = \frac{T_R + T_{A0}}{T_R + T_{A0} + T_{As}}$$

and to a solar flux density S_1 by

$$S_1(\text{SFU}) = \frac{2k_B T_{As}}{A_e 10^{-22}} \, ,$$

where the effective antenna area is taken to be $A_e = 7.23 \cdot 10^{-3}\,\text{m}^2$, which is a good approximation for a single-element GNSS antenna, and where the additional factor of 2 in the numerator reflects the assumption that only half the total-polarization solar radiation contributes to T_{As} through a GNSS antenna, which is designed to received right-hand circularly polarized signals [16.34]. The factor 10^{-22} converts $\text{W/m}^2/\text{Hz}$ to solar flux units (SFU). The resulting S_1 values listed in Table 16.2 are those above which P_T would increase by the amount shown. As a final step, the model $N(S > S_1, \nu_1, \nu_2)$ from [16.36] is invoked (with the correction factor C_{geo}) to approximate the total number of bursts exceeding S_1 in the frequency range ($\nu_1 = 1\,\text{GHz}$, $\nu_2 = 1.7\,\text{GHz}$) over a 40 yr historical period. This is used to estimate T_e, the time between triggering events, for solar maximum years and for all years.

Table 16.2 Time between threshold-exceeding solar radio burst events for various values of the ratio $P_T/P_{T,nom}$

Threshold values				T_e (days)	
$P_T/P_{T,nom}$ (dB)	T_{As} (K)	$\Delta C/N_0$ (dB)	S_1 (SFU)	Solar max.	All years
0.44	40.9	−0.6	1560	9.2	22.0
0.93	91.3	−1.2	3488	17.3	42.9
1.5	157.7	−1.9	6022	26.5	67.4

Table 16.2 reveals that solar radio bursts causing a degradation in C/N_0 of 1.9 dB or greater are rare, occurring approximately once per month during solar maximum. Truly intense SRBs causing 10 dB or more of degradation and interrupting signal tracking, as in the 2006 storm [16.33], are extremely rare. Nonetheless, SRBs can be problematic for signal authentication techniques based solely on P_T, as will be discussed in Sect. 16.6.2.

16.4.2 Scintillation

A transionospheric radio wave can exhibit temporal fluctuations in phase and intensity caused by electron density irregularities along its propagation path, a phenomenon called scintillation, or fading. At GNSS frequencies (L band), strong scintillation is manifest in deep power fades (> 15 dB) that are often associated with rapid phase changes. Such vigorous signal dynamics stress a receiver's carrier tracking loop and, as their severity increases, lead to navigation bit errors, cycle slipping, and complete loss of carrier lock [16.37, 38].

Signal refraction, caused by large-scale irregularities, results in low-frequency variations in group delay (measured by the code phase, or pseudorange, observable) and carrier phase. Signal diffraction, caused by smaller-scale (approximately 400 m) irregularities, scatters L-band signals so that the radio waves reach terrestrial receivers through multiple paths. Interaction between signals from multiple directions occurs at the carrier-phase level, yielding constructive and destructive interference patterns that produce variations in both the phase and amplitude of received signals.

It may at first seem out of place to treat ionospheric scintillation as interference, but the mutual interference caused by diffraction can challenge signal tracking as much as intermittent jamming, and diffractive interference shares characteristics with structured interference such as GNSS spoofing. The same argument can be made for nonionospheric multipath effects – those due to signal reflections – but these are treated separately in Chap. 15. Chapter 39 also treats scintillation, but

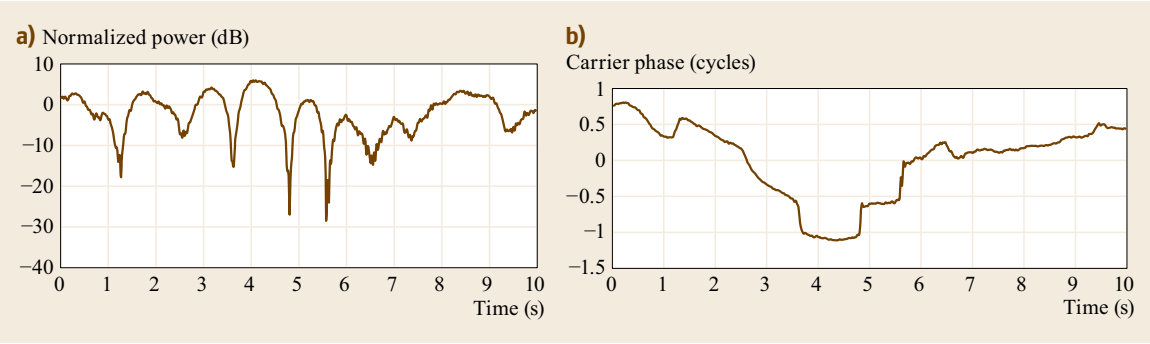

Fig. 16.10a,b Normalized signal power **(a)** and carrier phase **(b)** time histories from a record of GPS L_1 data with $S_4 \approx 0.9$ (after [16.37], courtesy of IEEE)

with an eye to phenomenology rather than receiver effects.

Severe L-band scintillation is both infrequent and geographically confined. The type known as equatorial scintillation, or equatorial spread F, generally occurs between local sunset and 2400 local time in the region extending $\pm 15°$ about the magnetic equator [16.39]. Another common type of scintillation occurs at high latitudes [16.40]. Significant effects have also been noted in the mid-latitude region, but they occur infrequently [16.41]. This section concentrates on equatorial scintillation because it is the most interference-like, making signals particularly difficult to track.

The severity of scintillation can be succinctly characterized by two parameters, the scintillation index, S_4, and the decorrelation time τ_0 [16.42]. S_4 measures the intensity of scintillation, and is defined by

$$S_4^2 = \frac{\langle I^2 \rangle - \langle I \rangle^2}{\langle I \rangle^2} \,,$$

where $I = \alpha^2$ is signal intensity, α being the signal amplitude, and $\langle \cdot \rangle$ denotes time average. The scintillation decorrelation time $\tau_0 > 0$ is a measure of the rapidity of scintillation. A small τ_0 (e.g., < 0.5 s) implies a scintillating channel that changes rapidly with time.

A short sample from the scintillation library introduced in [16.37] is presented in Fig. 16.10. The sample manifests strong scintillation, with $S_4 \approx 0.9$. The most striking features of the plot are the deep power fades that occur simultaneously with abrupt, approximately half-cycle phase changes whose sense (downgoing or upgoing) appears random. Such fades appear to be a universal feature of strong equatorial scintillation, and they are the primary cause of phase unlock for PLLs tracking strongly scintillating signals.

PLLs are affected by scintillation in two related ways: (1) increased phase error variance and (2) phase unlock.

Phase Error Variance

The phase error variance models given in Sect. 16.1.4 assume that all phase errors are due to constant-intensity white measurement noise. Furthermore, (16.5) and (16.6) assume PLL linearity. These assumptions are violated during severe scintillation: Amplitude fading causes variations in the loop SNR, phase changes are time correlated, and, when attempting to track through the large, rapid phase changes associated with deep fading, the PLL cannot be expected to operate in its linear regime. For these reasons, calculating the phase error variance for a PLL tracking through strong scintillation is not straightforward [16.38]. Figure 16.11 shows how

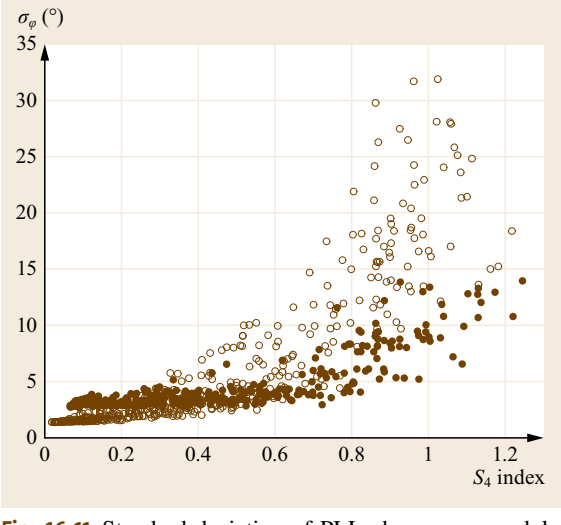

Fig. 16.11 Standard deviation of PLL phase error modulo π for a decision-directed arctangent phase discriminator over 30 s test records versus S_4 for ultra-high frequency (UHF) signals at $C/N_0 = 43$ dB Hz (*open circles*) and for GPS L_1 signals within $40 < C/N_0 < 44$ dB Hz with mean $C/N_0 = 43$ dB Hz (*filled circles*) (after [16.38], courtesy of IEEE)

σ_φ, the standard deviation of the phase measurement error modulo π, increases with increasing S_4, a dependence that is both due to the fade-induced reductions in loop SNR and to phase scintillation with frequency components that exceed the PLL's bandwidth. The large values of σ_φ at high S_4 contribute to the degradation of carrier-phase-dependent GNSS systems during strong scintillation.

Phase Unlock

The general term *phase unlock* refers to single or successive cycle slips. Phase and amplitude scintillation cause cycle slipping by either deep rapid fading or prolonged fading. In the limit as the fade depth increases, the accompanying abrupt, nearly π-rad phase transition looks like bi-phase data modulation, to which a squaring-loop PLL is insensitive by design. Hence, the PLL detects no phase shift and a half-cycle slip occurs. In marginal cases, where the PLL might be capable of distinguishing a scintillation-induced phase transition from a data-bit-induced phase transition, the sudden drop in loop SNR increases the likelihood of a cycle slip. In short, simultaneous power fades and abrupt phase changes are a particularly challenging combination.

Prolonged amplitude fading is the second mechanism by which scintillation causes cycle slipping. This phenomenon may be considered a special case of fading in which the fading time scale is elongated so that the amplitude fade is accompanied by phase dynamics that are slow compared to a typical 10 Hz PLL noise bandwidth. In this case, broadband measurement noise dominates and (16.7) applies. Cycle slips occur rarely by this mechanism.

Figure 16.12 presents results in terms of cycle slip rate on the left vertical axis, and, for convenience, in terms of the mean time between slips, T_s, on the right vertical axis. As would be expected, a general increase in the rate of cycle slips accompanies increasing S_4. The lack of cycle slips below $S_4 \approx 0.4$ suggests that, whatever its other characteristics (e.g., τ_0), scintillation with $S_4 \lesssim 0.4$ can be considered benign.

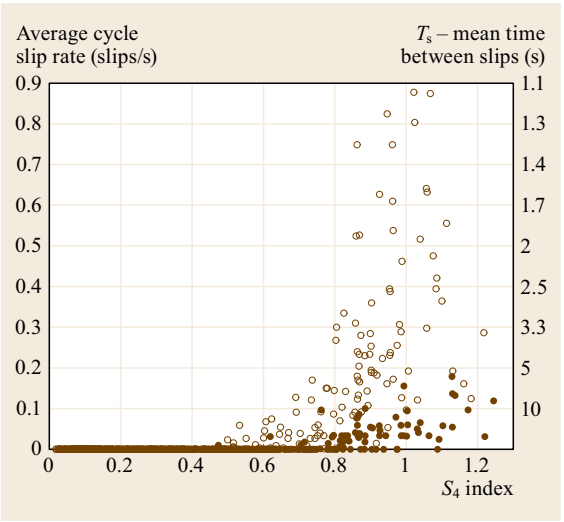

Fig. 16.12 Average cycle slip rate for the decision-directed arctangent phase discriminator over 30 s test records versus S_4 for UHF signals at $C/N_0 = 43$ dB Hz (*open circles*) and for GPS L_1 signals within $40 < C/N_0 < 44$ dB Hz with mean $C/N_0 = 43$ dB Hz (*filled circles*). The *right vertical axis* expresses the cycle slip rate in terms of T_s (after [16.38], courtesy of IEEE)

16.4.3 Unintentional Interference

Spectral surveys of the GNSS bands reveal that in rural areas the bands are largely free of interference, but in urban areas they are often corrupted by intermittent interference sources [16.43]. Most of these interference events are unintentional. Similarly, radio frequency interference (RFI) can disturb signal tracking when a GNSS receiver's antenna is packaged closely to other electronic equipment, as on a small satellite. Following are some examples of unintentional interference sources.

Harmonics

Nonlinearity in any one of several stages involved in RF transmission generates power not only at the intended transmission frequency but also at integer multiples, or harmonics, of that frequency. For example, UHF television signals with carrier frequencies near 525 MHz are notorious for injecting third-harmonic power into the GNSS L1 band [16.44, 45].

When broadcast transmitters are powerful, as with television transmitters, a harmonic near the GNSS bands can substantially degrade GNSS tracking performance. If a harmonic lies within a GNSS band of interest, then it cannot be attenuated by standard RF filters designed to isolate the GNSS signals. If powerful enough, the interfering harmonic will drive a GNSS receiver's dominant LNA into its nonlinear regime, causing a loss of sensitivity and leaving spurious tones across the target GNSS band [16.45].

DME/TACAN

The GPS L5 band and the Galileo E5a and E5b bands are situated in an ARNS band also allocated to distance measuring equipment (DME) and Tactical Air Navigation (TACAN) systems whose strong pulsed emissions act to significantly degrade GNSS

tracking [16.1]. DME/TACAN systems, which operate between 960 and 1215 MHz, produce emissions that are sparse in both the time and frequency domains. Pulses are transmitted in pairs 12 μs apart, with each pulse lasting 3.5 μs. The maximum practical transmission rate is 2700 pulse pairs per second, which means that interference from a single DME/TACAN transmitter is limited to less than 2% of a 1 s time interval. In the frequency domain, a single DME/TACAN signal occupies only 100 kHz, with channels spaced by 1 MHz. Thus, the total time-frequency occupancy of a single DME/TACAN transmitter in a 10 MHz band is only 0.02%. Such sparsity permits mitigation techniques that render DME/TACAN interference harmless even when GNSS receivers are airborne over so-called hot spots having a high density of DME transmitters [16.1].

Powerful Near-Band Transmissions

The radio spectrum between 700 MHz and 2 GHz, which includes all current GNSS bands, is particularly attractive for the provision of data to mobile units such as smartphones because the wavelengths of signals in this band are short enough that small antennas can be effective yet long enough to penetrate indoors. These desirable properties, coupled with the intense and rising demand for mobile data, portend the eventual placement of powerful transmissions in the radio bands adjacent to GNSS bands.

The 2010–2012 debate over whether to allow powerful terrestrial long term evolution (LTE) signals to be broadcast in the mobile satellite service (MSS) band just below the GNSS L1 band brought to the fore the susceptibility of contemporary GNSS receivers, especially high-precision receivers, to powerful near-band transmissions [16.46]. It was shown, for example, that typical GPS and Galileo receivers tracking signals centered at 1575.42 MHz suffered C/N_0 degradation greater than 3 dB when exposed to communications signals with received power exceeding −80 dBm in the 1545.2–1555.2 band even when the latter were filtered with a high-quality bandpass filter [16.8].

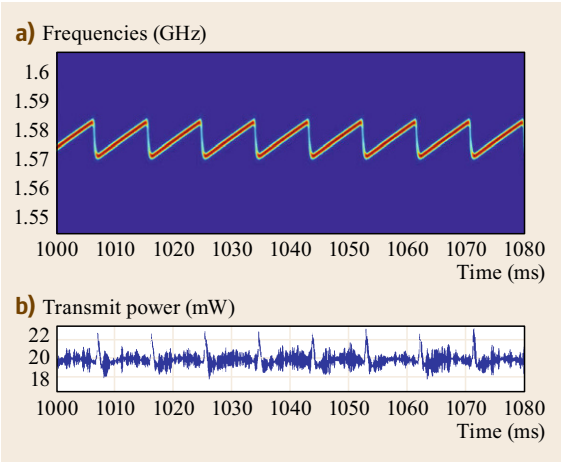

Fig. 16.13a,b Time histories of frequency spectrum **(a)** and transmit power **(b)** for a typical chirp-style PPD (after [16.47])

16.4.4 Intentional Interference

Intentional interference, or jamming, has been a staple of navigation warfare since World War II [16.32]. With the emergence of PPDs [16.7] and incidents of nation-scale intentional disruption of civil GNSS [16.48], intentional interference is now also a civil concern.

PPDs are by far the most common source of intentional interference. The PPD user may intend only to jam GNSS tracking devices in his near vicinity (e.g., on his person or vehicle), but in fact such devices can disrupt GNSS signal tracking out to an effective radius of from 100 m to several kilometers [16.47].

Virtually all PPDs transmit a swept tone waveform (chirp) similar to that shown in Fig. 16.13. This waveform can be generated from inexpensive components and is quite effective in rendering GNSS receivers inoperable unless these have been especially designed for jam resistance [16.49]. The frequency sweep period of the 18 units tested in [16.47] ranged from 1 to 27 μs, with total transmit power in a 20 MHz band centered at L1 ranging from −14 to 28 dBm.

16.5 Spoofing

A GNSS spoofing signal is a type of structured interference that adheres closely enough to a GNSS signal specification so as to appear authentic to an unsuspecting GNSS receiver. Whether intentional, as in a deliberate attempt to manipulate the PVT readout of a target GNSS receiver [16.50, 51], or unintentional, as in an errant GNSS simulator or repeater signal, spoofing signals similarly affect a GNSS receiver. For convenience of presentation, the following discussion will treat all spoofing as intentional, with the term spoofer referring both to the spoofing device and its operator.

Spoofing was once only a threat to military GNSS receivers and applications, but has now become a more

general concern as civil GNSS spoofing becomes easier and its consequences are more serious. The emergence of low-cost off-the-shelf software-defined radio hardware has significantly reduced the cost and complexity of spoofing. With such hardware, a competent programmer sufficiently familiar with the openly documented GNSS protocols [16.23, 52] can generate realistic civil GNSS signals despite having minimal knowledge of RF electronics. Easier still, low-cost GNSS signal simulators and record-and-replay devices enable even GNSS neophytes to conduct a limited but potent form of spoofing. Against a backdrop of increasing economic dependence on civil GNSS for transportation, communication, finance, and power distribution, the increased accessibility of civil GNSS spoofing raises the risk of attack and the urgency of finding effective antispoofing measures.

Spoofing is different from unstructured interference in two primary respects. First, it can be surreptitious: Neither the target GNSS receiver nor its operator may detect that an attack is underway because the spoofer can seamlessly supplant counterfeit signals for their authentic counterparts. Second, in a spoofing attack, the received interference $r_I(t)$ is statistically correlated with the received authentic signal $r_S(t)$; consequently, the total received power P_T is neither the sum of P_S, P_I, and P_n, as in (16.2), nor does the autocorrelation function of the interference component $I(t)$ decompose, as in (16.4), because the cross-terms do not average to zero. As a result, the analysis of spoofing effects is, in general, more challenging than the analysis of statistically independent interference. To be sure, spoofing effects bear a strong resemblance to multipath effects, but multipath-induced structured interference is accidental, whereas spoofing may involve a strategic attacker who can arbitrarily adjust signal power, code phase, carrier phase, and signal structure for maximum effect.

To generalize the treatment of spoofing in what follows, the authentic signal model will allow for digital modulation that is unpredictable to a would-be spoofer. A modulation sequence that is entirely unpredictable or has unpredictable segments will be termed a security code, and a security-code-bearing GNSS signal will be termed security enhanced [16.53–57]. A nonsecurity-enhanced GNSS signal can be represented by a special case of this model in which the security code is replaced by a sequence of ones.

16.5.1 Generalized Model for Security-Enhanced GNSS Signals

From the perspective of a GNSS receiver, current and proposed security-enhanced GNSS signals can be represented by a simple adaptation of the baseband received signal model introduced in (16.1):

$$\begin{aligned} r_S(t) &= \sqrt{P_S}W(t-\tau)D(t-\tau)C(t-\tau)\exp(j\theta(t)) \\ &= \sqrt{P_S}W(t-\tau)X[\tau,\theta(t)]\,. \end{aligned} \tag{16.12}$$

Compared to (16.1), the novel component here is $W(t)$, which represents a ± 1 valued security code with chip length T_W. For notational simplicity, the product of the authentic signal's navigation data stream $D(t-\tau)$, spreading (ranging) code $C(t-\tau)$ and baseband phasor $\exp(j\theta)$ is abbreviated as $X(\tau,\theta)$ for code phase τ and carrier phase θ. The chip length of the spreading code $C(t)$ is denoted as T_C. For convenience, receiver time t is assumed to be equivalent to true time (e.g., GPS system time).

The security code $W(t)$ is either fully encrypted or contains periodic authentication codes. The defining feature of $W(t)$ is that some or all of its symbols are unpredictable to a would-be spoofer prior to broadcast from a legitimate GNSS source. The unpredictable symbols in $W(t)$ serve two related functions: (1) they enable verification of $W(t)$ as originating from a GNSS Control Segment (standard message authentication), and (2) they increase the complexity of a spoofing attack by forcing the spoofer to either replay a received $W(t)$ or attempt to estimate $W(t)$ on-the-fly. Note that if a GNSS signal is not security enhanced (has no unpredictable modulation), the model in (16.12) still applies, with $W(t) = 1$.

16.5.2 Attacks Against Security-Enhanced GNSS Signals

The unpredictability of the security code $W(t)$ is an obstacle for a would-be spoofer. A simple spoofing technique, such as discussed in [16.58], relies on the known signal structure of the GPS L1 C/A signal and the near-perfect predictability of its navigation data stream. However, if a GNSS signal is security enhanced, then the spoofer of [16.58] cannot perfectly match its counterfeit signals chip-for-chip to the authentic signals.

A spoofer could, of course, ignore the broadcast security codes altogether, filling in dummy values for $W(t)$, but such a scheme is easily detected. In an attack against a GNSS signal modulated by a low-rate security code ($T_W \gg T_C$) (e.g., navigation message authentication (NMA), as proposed in [16.55–57, 59]), the dummy $W(t)$ values would fail the cryptographic validation test. Against a high-rate security code ($T_W \approx T_C$), the dummy $W(t)$ values would yield zero av-

erage power when correlated with the true $W(t)$ sequence [16.53, 59].

Therefore, to be effective while evading detection, a spoofer must attempt to match both the structure and content of the authentic signal. It can do this via one of the following specialized spoofing attacks.

Meaconing

A meaconing, or replay, attack is a specialized spoofing attack in which an entire segment of RF spectrum is captured and replayed [16.60]. If the meaconer employs a single receiving antenna element, then no individual signal is isolated in a meaconing attack. Thus, in this case, a GNSS meaconer cannot arbitrarily manipulate the PVT of a target receiver. Rather, the target receiver will display the position and velocity of the meaconer's receive antenna and a time in arrears of true time. If this antenna is on a dynamic platform, then the meaconer can adjust the position and velocity implied by its signals for greater effect in the attack.

If the meaconer employs multiple antenna elements whose RF signals are individually digitized, then it can isolate individual GNSS signals by pointing a gain enhancement toward each overhead GNSS satellite. For example, a 16 element antenna array could be used to direct a narrow $\approx 12\,\mathrm{dB}$ enhancement toward each satellite. By combining the separate digital streams while manipulating the phasing of each stream within the ensemble, a meaconer can dictate the ensemble's implied PVT within a wide range about the true PVT (with the implied timing always in arrears of true time).

For a single GNSS signal corresponding to a particular satellite, the combined meaconed and authentic received signals can be modeled as (16.1) but with $r_S(t)$ as in (16.12) and

$$r_I(t) = \alpha\sqrt{P_S}W(t-\tau_c)X[\tau_c, \theta_c(t)] + n_c(t)\ .$$

Here, $\tau_c > \tau$ and θ_c are the code phase and carrier phase of the counterfeit meaconing signal, respectively, and $n_c(t)$ is the noise introduced by the meaconer's RF front end. The meaconed signal arrives at the target receiver's antenna with a delay $d = \tau_c - \tau > 0$ seconds relative to the authentic signal, an unavoidable consequence of the triangle inequality and the processing delay through the meaconing device. The coefficient α is the meaconed signal's amplitude advantage factor relative to the authentic signal.

High-performance digital signal processing hardware permits a meaconer located close to its intended target to drive the delay d to under a few tens of nanoseconds. In the limit as d approaches zero, the attack becomes a zero-delay meaconing attack in which the meaconed signals are code-phase-aligned with their authentic counterparts. Such alignment enables a seamless liftoff of the target receiver's tracking loops, following which a meaconer can increase d at a rate that is consistent with the target receiver's clock drift and gradually impose a significant timing delay.

Note that, unless $d \approx 0$, a meaconer with $\alpha \approx 1$ will cause significant variations in the target receiver's PVT estimate: the meaconing signals will act as severe multipath. Thus, if the meaconer cannot ensure $d \approx 0$, it is better off transmitting with an overwhelming amplitude advantage ($\alpha \gg 1$) to quickly stabilize the target's perceived PVT at the meaconer's intended value. Therefore, a meaconer with d a significant fraction of T_C is detectable at $\alpha \approx 1$ due to multipath-like PVT variations and at $\alpha \gg 1$ due to anomalous high received power. Furthermore, if $d > 2T_W$, then the meaconer will be unable to capture a code tracking loop that is locked to an authentic signal for any value of α: The meaconing signal will not be close enough in time to the authentic signal to dislodge the receiver's code tracking loop. Instead, the meaconer will be forced to jam the target receiver to force re-acquisition, which will alert the target to the attack. In any case, GNSS system designers have an incentive to make T_W as small as possible to increase the difficulty of a meaconing attack.

Security Code Estimation and Replay Attack

A Security Code Estimation and Replay (SCER) attack allows greater flexibility than a meaconing attack in manipulating the target receiver's PVT solution. In a SCER attack, a spoofer receives and tracks individual authentic signals and attempts to estimate the values of each signal's security code on-the-fly. It then reconstitutes a consistent ensemble of GNSS signals, with the security code estimates taking the place of the authentic security codes, and transmits the ensemble toward the target receiver. For a single GNSS signal corresponding to a particular satellite, the combined SCER-spoofed and authentic received signals can be modeled as (16.1) but with $r_S(t)$ as in (16.12) and

$$r_I(t) = \alpha\sqrt{P_S}\hat{W}(t-\tau_c)X[\tau_c, \theta_c(t)] + n_c(t)\ ,$$

where $\hat{W}(t-\tau_c)$ represents the security code estimate arriving with a delay of $d = \tau_c - \tau > 0$ seconds relative to the authentic security code $W(t-\tau)$, $n_c(t)$ is noise introduced by the spoofer (e.g., due to quantization effects in the signal generation), and other quantities are as introduced previously. The delay d can be modeled as the sum $d = p + e$ of a processing and transmission delay $p > 0$ and an estimation and control delay $e > 0$. The delay p represents the combined minimum signal processing delay and additional propagation time and

does not contribute to better estimates of the security code chips. The delay e represents an additional delay imposed by the spoofer to improve its estimate of the security code chip values and to control the relative phasing of the spoofing signals so as to impose spoofer-defined position and timing offsets on the defender.

Mounting a stealthy SCER attack is challenging if the target receiver has been designed to detect SCER spoofing. The attacker must keep $d = p + e$ small enough to remain within the target receiver's clock uncertainty but must extend e enough to reliably estimate the security code chip values. The following two SCER attack strategies serve to illustrate this tradeoff.

Zero-Delay Attack. Consider a spoofer that is colocated with the target GNSS receiver's antenna and has negligible processing delay so that $p \approx 0$. Assume that $e = 0$, meaning that the spoofer adds no estimation and control delay. Thus, $d = p + e \approx 0$. In this zero-delay attack, $\tau_c \approx \tau$, which implies that each spoofing signal is approximately code-phase-aligned with its authentic counterpart as received by the target receiver.

Despite such code phase alignment, a zero-delay attack can still alter the target receiver's position and time by injecting false messages through $D(t)$ (e.g., erroneous satellite ephemeris or clock model parameters or an erroneous leap second). However, with $e = 0$, the spoofer's security code estimate $\hat{W}(t)$ will be highly erratic for the first few microseconds following an unpredictable chip transition in $W(t)$. This is illustrated in Fig. 16.14, which shows simulated time histories of $\hat{W}(t)$ for two different chip value estimation strategies over the first 20 μs after the beginning of a security code chip with $T_W > 20$ μs. In this scenario, for which the spoofer C/N_0 is an unusually high 54 dB Hz, the spoofer's chip estimates become reliable after about 8 μs. For each 3 dB drop in spoofer C/N_0, the interval required for reliable chip estimates doubles.

The key to zero-delay SCER attack detection, as explained in [16.54], is to develop a detection statistic that is sensitive to the increased error variance in $\hat{W}(t)$ in the crucial early moments immediately following unpredictable transitions in $W(t)$.

Nonzero-Delay Attack. In a nonzero-delay SCER attack, the spoofer rebroadcasts a counterfeit signal that arrives at the defender's RF front end with a delay $d > 0$ relative to the authentic signal. Any significant delay d (e.g., greater than about 20 ns) in the spoofer's counterfeit signal at the beginning of an attack would be immediately obvious to a target receiver that has been continuously tracking authentic signals since before the beginning of the attack. Therefore, the spoofer's strategy in the nonzero-latency SCER attack is typically to break the target receiver's tracking continuity by jamming or blocking the authentic signals for an interval of time before initiating the spoofing attack, thus, widening the target receiver's timing uncertainty, or *window of acceptance* [16.53, 55, 61]. The required duration of the signal-denial interval depends on the desired delay d and on the assumed stability of the target receiver's clock (for stationary receivers) or clock and inertial measurement unit (for moving receivers). For the low-cost temperature-compensated crystal oscillators (TCXOs) typical in commercial GNSS equipment, in-the-field stability is approximately 10^{-7}. Ovenized crystal oscillators (OCXOs), common in more demanding timing applications, have stability of approximately 10^{-10}. Thus, widening a TCXO-driven static target receiver's time uncertainty by 8 μs would require approximately 80 s of jamming or blockage, and widening an OCXO-driven static receiver's time uncertainty by the same amount would require approximately one day of jamming or blockage.

After the jamming-or-blockage prelude, the nonzero-delay SCER attacker initiates a spoofing attack in which d can be as large as the target receiver's tim-

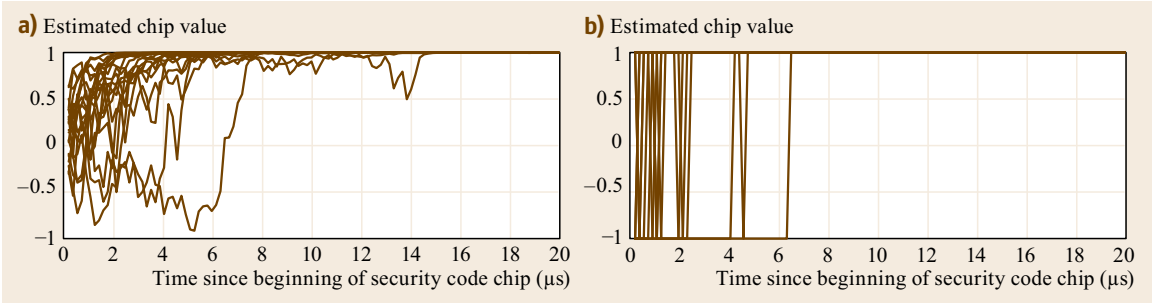

Fig. 16.14a,b Simulated time histories of security code chip estimates $\hat{W}(t)$ for a minimum mean square error (MMSE) estimator **(a)** and for a maximum a posteriori (MAP) estimator **(b)** over the first 20 μs after the beginning of a unity-valued security code chip for a spoofer with received $C/N_0 = 54$ dB Hz (after [16.54], courtesy of IEEE)

ing uncertainty. The attacker exploits the component e of this delay to more accurately estimate the value of each unpredictable chip in $W(t)$ so that $\hat{W}(t)$ appears accurate to the target receiver. Long security code chips (e.g., $T_W = 40\,\text{ms}$ as suggested for civil navigation message (CNAV) NMA in [16.54, 56]) allow the spoofer to significantly increase e and thereby generate highly accurate chip estimates. However, a large delay $d = p + e$ is itself a liability for the spoofer because of the long jamming-or-blockage interval required. Thus, the spoofer finds itself vulnerable to detection at low d due to poor security code chip estimates and at high d due to a noticeable timing delay.

Note that, with a SCER attack, the attacker can eventually specify an arbitrary position and an arbitrary delayed time as the spoofer slowly pulls each signal's code phase to the desired offset. Note also that if $W(t) = 1$ (i. e., the GNSS signal is not security enhanced), then the attacker need not delay at all: He can exploit the near-perfect predictability of $D(t)$ to anticipate the next navigation data symbol value and ensure that it arrives at the target receiver's antenna just on time – perfectly aligned with the true $D(t)$ [16.58]. Thus, the unpredictability of the security code – even a low-rate code such as in NMA – forces a SCER spoofer to expose himself with a jamming-or-blockage attack prelude. Finally, note that signal jamming or blockage for any significant interval of time (relative to the receiver clock stability) must be viewed not only as a temporary nuisance but also as a security threat that persists even after the interference apparently subsides. This is because, in the absence of some other means of verifying the authenticity of GNSS signals, a SCER attack detector's probability of detection is irrecoverably reduced by a loss of signal continuity [16.55].

Effect of Coherence

In a spoofing attack, the complex correlator output modeled in (16.3) contains a desired component $S(t) \equiv r_S^*(t)l(t)$ and an interference component $I(t) \equiv r_I^*(t)l(t)$, both of which are dependent on the local replica's code phase $\hat{\tau}$ and carrier phase $\hat{\theta}$. Denote these as $S(t, \hat{\tau}, \hat{\theta})$ and $I(t, \hat{\tau}, \hat{\theta})$. Also, for a given authentic and spoofing signal pair $r_S(t)$ and $r_I(t)$, let $\varphi(t) \equiv \theta_c(t) - \theta(t)$ be the relative carrier phase.

If a spoofing attack is code-phase aligned so that $|\tau_c - \tau| < T_C$, and Doppler matched so that

$$\frac{1}{2\pi}\left|\frac{d\varphi}{dt}\right| < \frac{1}{T_a}$$

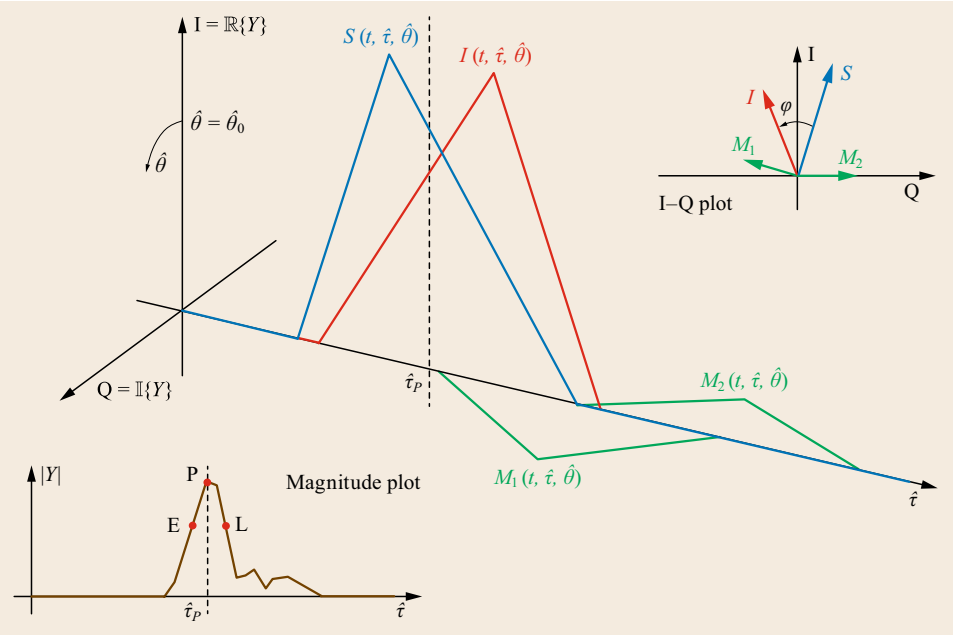

Fig. 16.15 Stylized complex correlation functions depicting a spoofing attack in which $|\tau_c - \tau| < T_C$ and $d\varphi/dt \approx 0$. The *blue trace* marked $S(t, \hat{\tau}, \hat{\theta})$ represents the desired signal correlation function, the *red trace* marked $I(t, \hat{\tau}, \hat{\theta})$ represents the interference (spoofing) signal correlation function, and the *green traces* marked $M_i(t, \hat{\tau}, \hat{\theta})$, $i = \{1, 2\}$, represent two multipath correlation functions. The receiver's code and carrier tracking loops track the composite correlation function, $Y(t, \hat{\tau}, \hat{\theta})$, whose magnitude is shown in the *lower inset plot* along with the early, prompt, and late correlation taps

with T_a is the accumulation interval from Fig. 16.2, then $r_S(t)$ and $r_I(t)$ are substantially frequency coherent and thus cannot be considered statistically independent. As a consequence, the combined signal power P_T is not simply the sum $P_T = P_S + P_I + P_n$, as in (16.2), but depends on $\tau_c - \tau$, φ, and the relative spoofing amplitude α. Figure 16.15 shows the relationship between $S(t, \hat{\tau}, \hat{\theta})$ and $I(t, \hat{\tau}, \hat{\theta})$ in this regime.

The interference power P_I can be decomposed as $P_I = \alpha^2 P_S + P_{nc}$, where P_{nc} is the power in the noise component $n_c(t)$. If code-phase alignment and Doppler matching are approximately achieved in a spoofing attack ($|\tau_c - \tau| \approx 0$ and $d\varphi/dt \approx 0$), the possibility of which was demonstrated in [16.50] against a nonsecurity-enhanced GNSS signal, then P_T can be expressed as

$$P_T = \left[\sqrt{P_S} + \sqrt{\alpha^2 P_S}\cos(\varphi)\right]^2 + \alpha^2 P_S \sin^2(\varphi) + P_{nc} + P_n \,. \quad (16.13)$$

This expression indicates that the noise components P_{nc} and P_n, which are noncoherent with the authentic signal, add directly to P_T, as does $\alpha^2 P_S \sin^2(\varphi)$, which is the power in the spoofing signal's frequency-coherent component that lies in phase quadrature to the authentic signal. By contrast, $\alpha^2 P_S \cos^2(\varphi)$, which is the spoofing power component that is phase aligned with the authentic signal, does not add directly to P_T but instead interacts with the authentic signal as shown. For $k \in \mathbb{Z}$, the spoofing signal contributes maximally to P_T when $\varphi = k2\pi$ (phase alignment), minimally when $\varphi = (1+2k)\pi$ (antiphase alignment), and power-additively – as if it were a purely noncoherent signal – when $\varphi = (1/2+k)\pi$ (orthogonal alignment).

It is interesting to note that if φ is treated as a random variable uniformly distributed on $[0, 2\pi]$, then the expected value of P_T is equivalent to the P_T that arises in the case of purely noncoherent interference signals; that is, $E[P_T] = P_S + P_I + P_n$. Hence, for an ensemble of statistically independent spoofer-and-authentic signal pairs, (16.2) remains a useful approximation for the power contributed by each pair even when the spoofer can achieve Doppler frequency alignment ($d\varphi/dt = 0$) but has no finer control over the carrier phase. By distinction, if the spoofer has knowledge of the target receiver's antenna position to within a small fraction of a carrier wavelength, then it can arbitrarily adjust α and φ to exercise full control over P_T according to (16.13). Figure 16.16 demonstrates that frequency-coherent spoofing signals affect P_T as expected.

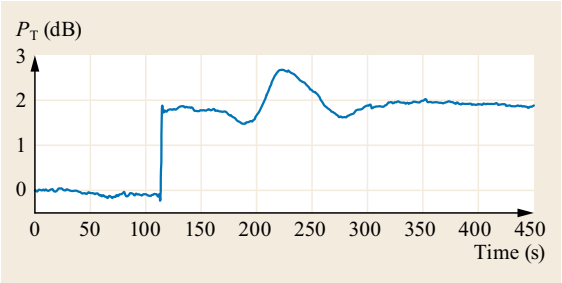

Fig. 16.16 Total received power P_T in a 2 MHz band centered at the GPS L1 frequency showing the onset of a spoofing attack using the testbed described in [16.62], normalized by the average value of P_T prior to the attack. The attack begins with a sudden increase in P_T just before 100 s. Thereafter, the total authentic signal power and total spoofing power were maintained constant; thus, the oscillations in P_T are due to the frequency coherence between the spoofing and authentic signals, with each pair of spoofing-and-authentic signals having similar values of φ

An important consequence of a spoofer's having arbitrary control over α and φ is that, by choosing $\alpha = 1$ and $\varphi = \pi$ for each spoofing and authentic signal pair, a spoofer can effectively annihilate the authentic signals at the location of the target antenna. Such a nulling attack has the effect of jamming the target receiver while *reducing* the total received power P_T in the GNSS band of interest. Moreover, the nulling signals could be paired with an independent ensemble of spoofing signals to simultaneously eliminate the authentic signals while presenting clean counterfeit signals to the target receiver. The attacker could thus evade tests, such as the received power test proposed in [16.35] and the pincer defense proposed in [16.63], designed to detect anomalies in the total received power or distortion in the correlation function caused by interaction of the authentic and spoofing signals. GNSS antennas that are clearly visible to the public from close range and those whose coordinates are publicly posted to subdecimeter accuracy are at greatest risk of such nulling attacks.

16.6 Interference Detection

Many schemes for detecting and mitigating GNSS interference have been proposed since the early days of GPS. These schemes apply at one or more of three application points in the GNSS signal processing chain, as shown in Fig. 16.17: (1) the analog stage, (2) the post-digitization but precorrelation stage, and (3) the correlation and post-correlation stage. Several effective interference detection schemes are detailed in this section; the following section treats interference mitigation.

16.6.1 C/N_0 Monitoring

A drop in a receiver's measured C/N_0 on any channel that cannot be explained by signal shadowing indicates interference of some type. C/N_0 is related to the SNR of the complex accumulations Y_k (Fig. 16.2) on which code and carrier tracking are based by SNR $= CT_a/N_0$. As C/N_0 measurements are generated post-correlation, C/N_0 monitoring applies at point (3) in Fig. 16.17.

Given measured C/N_0, one can be assured that code and carrier tracking will perform no better than what would be expected for SNR $= CT_a/N_0$. Nominal C/N_0 values across all tracking channels do not, however, guarantee the absence of interference, since spoofing interference, whether intentional or not, can cause the affected receiver to report perfectly normal C/N_0 values. For example, the spoofer described in [16.62] can dictate the received C/N_0 for each signal by adjusting the relative magnitudes of its output signals and adding artificial noise to the signal ensemble.

Given that C/N_0 loss is often caused by signal shadowing, and that nominal C/N_0 values are no guarantee of the absence of interference, a C/N_0 monitor such as proposed in [16.64] is best applied in combination with other complementary techniques for GNSS interference.

16.6.2 Received Power Monitoring

Monitoring the total received power P_T in a GNSS band of interest, known as received power monitoring (RPM), is one of the simplest and most effective strategies for detecting interference [16.35, 65, 66]. For systems with multibit-quantized sampling and automatic gain control (AGC) in the RF front end, estimating P_T is as easy as measuring the voltage applied by the AGC unit to adjust the signal amplitude before quantization. In a constant-gain system with sufficient dynamic range to prevent quantization saturation, P_T can be estimated directly from the precorrelation samples. In any case, RPM can be thought of as applying at point (2) in Fig. 16.17.

Figure 16.18 shows the nominal power spectrum about the GPS L1 frequency as measured at the output of a high-quality GNSS antenna and front-end system. Despite their statistical independence and low power, the received GPS L1 C/A signals combine to yield an obvious enhanced density in the familiar $\mathrm{sinc}^2(fT_C)$ pattern near L1 that rises above the noise floor.

For interference detection with a suitably low false alarm rate, one must examine the size and predictability of variations in P_T that can be considered natural or otherwise innocuous. Figure 16.19 shows a two-day record of P_T for the setup in Fig. 16.18 in the 2 MHz band centered at L1. The time history reveals marked diurnal variations, the result of diurnal patterns in temperature, solar radiation, and the overhead satellite constellation. Even though the record's diurnal repeatability is evidently only good to approximately 0.3 dB, its predictability given knowledge of local temperature and satellite orbital ephemerides is actually better than this.

Figure 16.20 offers an expanded view of a 7.5 min interval using the same setup and showing both the 2 and 10 MHz traces. The different size of the variations in the two traces at time scales less than about 150 s indicates that the variations do not originate in broadband noise; they are likely due to multipath effects at the carrier-phase level caused by reflections off nearby surfaces and by atmospheric diffraction and refraction. Close examination of multi-day records of P_T reveals that these short-time-scale variations do not repeat appreciably at the solar or sidereal day. In sum-

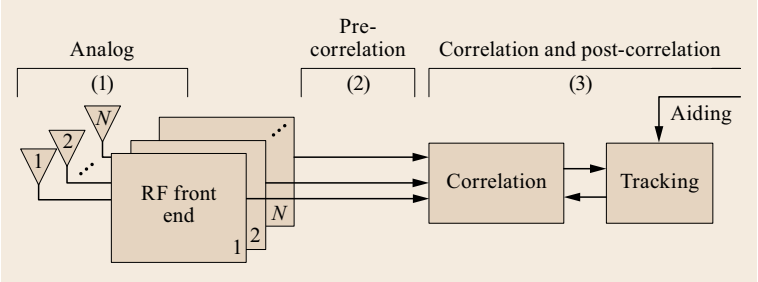

Fig. 16.17 Application points for interference detection and mitigation: (1) in the analog stage prior to digitization, (2) after digitization but before correlation, and (3) in correlation and in post-correlation tracking and PVT estimation

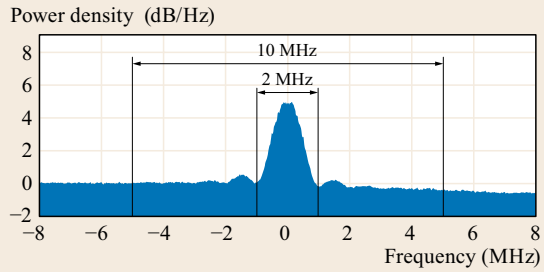

Fig. 16.18 Power spectrum centered at the GPS L1 frequency as estimated from a 1 s interval of data captured via a high-quality static antenna and RF front-end combination in a moderately quiet outdoor RF environment. Bands for 2 and 10 MHz power measurements are shown. The power density scale has been centered near the noise floor for ease of viewing. In absolute units, the noise floor sits at approximately −204 dBW/Hz

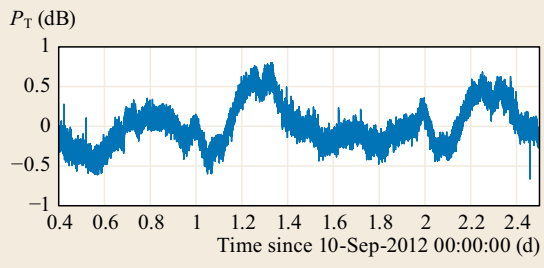

Fig. 16.19 A two-day record of received power P_T in the 2 MHz band shown in Fig. 16.18, normalized by the average received value over the interval

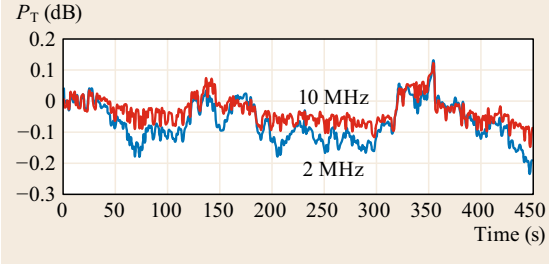

Fig. 16.20 A 7.5 min record of received power in the 2 and 10 MHz bands shown in Fig. 16.18, normalized by the initial values of P_T in each band

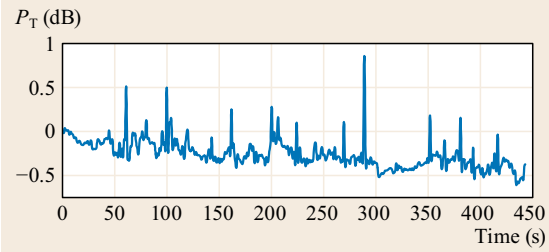

Fig. 16.21 Received power P_T in a 2 MHz band centered at the GPS L1 frequency averaged over 1 s intervals for a receiver mounted on a vehicle driving through the streets of downtown Austin, Texas. The data correspond to the *clean dynamic* data record from [16.67]

mary, it appears that for a static antenna, the practically unpredictable variations in P_T about L1 have root-mean-squared deviations of at least 0.1 dB for a 2 MHz band and 0.05 dB for a 10 MHz band.

For a dynamic antenna, P_T can be much more variable. Figure 16.21 shows a time history of P_T for a receiver mounted on a vehicle driving through the streets of downtown Austin, Texas. The P_T excursions, the largest of which exceeds 1 dB, would be unpredictable to a GNSS user without an up-to-date RF interference map of the area.

Against background variations that are unpredictable at the 0.1 dB level, or even the 1 dB level, deliberate jamming from close range remains obvious, as revealed by the effect on P_T of highway motorists using PPDs shown in Fig. 16.22. Naive spoofing also has an obvious effect: consider the sudden 2 dB uptick of P_T in Fig. 16.16. However, contrary to the claims in [16.35], RPM is not a generally effective means of detecting spoofing. This is because the increase in P_T during a spoofing attack may be smaller, or not significantly larger, than unpredictable variations in P_T due to causes other than spoofing. As mentioned in Sect. 16.5.2, a spoofer able to arbitrarily control the relative amplitude α and phase φ of each spoofing signal can annihilate the authentic signals and supplant them with counterfeit signals of equal power, thereby, maintaining P_T constant.

A spoofer lacking precise control over φ cannot prevent an increase in P_T while successfully capturing the target receiver's tracking loops, but the increase in P_T can be small: For a commercial-grade GNSS receiver, the uptick in P_T may be as small as 0.56 dB [16.62]. If unpredictable natural variations in P_T are modeled as a Gaussian process with a 0.1 dB standard deviation and a 150 s decorrelation time, then a detection threshold equal to $\gamma = 0.44$ dB would be sufficient to detect such an uptick with high probability while maintaining a once-per-year false alarm rate. However, the natural variations in P_T have a much thicker high-side probability distribution tail than a Gaussian process. For example, as detailed in Table 16.2, solar radio bursts would cause P_T to exceed $\gamma = 0.44$ dB every 9.2 days on average during solar maximum. Note that

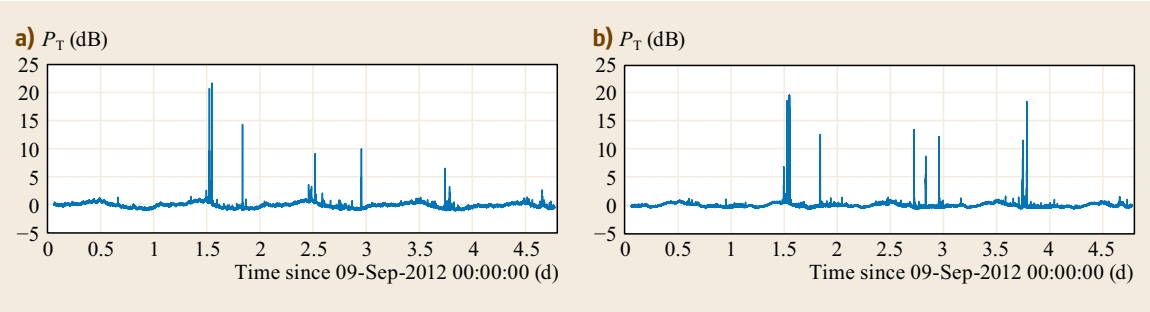

Fig. 16.22a,b Received power in the 10 MHz band centered at GPS L1 at two sites 1 km apart that straddle State Highway 1, west of Austin, TX. **(a)** Data from site located at the Center for Space Research. **(b)** Data from site located at Applied Research Laboratories. Both traces are normalized by the average value of P_T over the interval. The large excursions in P_T are due to motorists using PPDs as they travel along the highway

although spoofing alarms could be dismissed during known solar radio burst events, which can be independently monitored – even predicted [16.68], this offers little protection, for a clever attacker could time his attack to coincide with the arrival of a sizable burst.

Besides solar radio bursts, nonspoofing interference endemic in urban environments and near major thoroughfares can often cause an increase in P_T exceeding $\gamma = 0.44\,\text{dB}$, as shown in Figs. 16.21 and 16.22. One might argue that it is perfectly appropriate for a spoofing detector to alarm in the presence of a solar radio burst or an intentional jammer, but the consequences of spoofing can be much more malign than those of natural interference or jamming, and so it behooves a defender to distinguish between these.

16.6.3 Augmented Received Power Monitoring

When acting alone, RPM is effective at detecting strong interference but cannot be considered a reliable detector of weak interference such as low-power spoofing. It can, however, be paired with other tests that are sensitive to GNSS-like structure in the received signal to yield a powerful joint detection test for spoofing, provided the spoofer cannot arbitrarily manipulate α and φ. Three RPM augmentation strategies are discussed in the following sections.

Augmentation with C/N_0 Monitoring

A simple C/N_0 monitor will not detect spoofing signals whose C/N_0 values are matched to those of the authentic signals. But when paired with RPM, C/N_0 monitoring becomes a reasonably reliable detection strategy because it is challenging for a spoofer to ensure nominal received C/N_0 values without significantly increasing P_T. Only with a nulling attack, such as described in Sect. 16.5.2, can a spoofer ensure that C/N_0 matching does not increase P_T. Without nulling, C/N_0 matching (with no unusual variations) requires overwhelming spoofing power, which manifests as increased P_T.

Augmentation with Precorrelation Structural Power Content Analysis

The precorrelation structural power content analysis method advanced in [16.69] detects the presence of spoofing based on the excessive power content of GNSS-like signals in the received raw samples. In the absence of RPM, a spoofer can evade this detector by transmitting with overwhelming power, thus, driving the received authentic signals into the noise floor as the receiver's AGC compensates for the high received total power. The method of [16.69] will then only measure precorrelation structural power content commensurate with a single signal for each expected received waveform, and will thus fail to alarm. However, when combined with RPM, a structural power detector becomes powerful for spoofing detection. As for C/N_0 monitoring, augmentation with RPM forces the spoofer to either mount a nulling attack or be exposed with high likelihood in the joint test statistic.

Augmentation with Distortion Monitoring

The pincer defense advanced in [16.63] thoroughly embraces the concept of augmenting RPM for improved spoofing detection. Its name is meant to evoke a pincering, or trapping, of the spoofing signals between an RPM and a signal distortion monitor. As with C/N_0 and precorrelation structural power monitoring, distortion monitoring acting on its own cannot detect a spoofing attack executed with overwhelming power because the interaction between the authentic and false signals, which is the source of the signal distortion sought, is eliminated by action of the AGC as the spoofing-to-authentic power ratio increases.

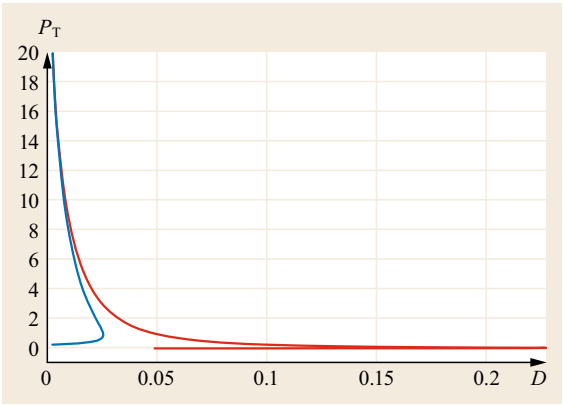

Fig. 16.23 Distortion (in the same units as accumulation), as a function of P_T for in-phase (*blue*) and antiphase (*red*) multipath or spoofing interference at a fixed delay of 0.15 chips. For the same delay, all other relative phases yield distortion profiles that lie within this envelope (after [16.63]; reprinted with permission)

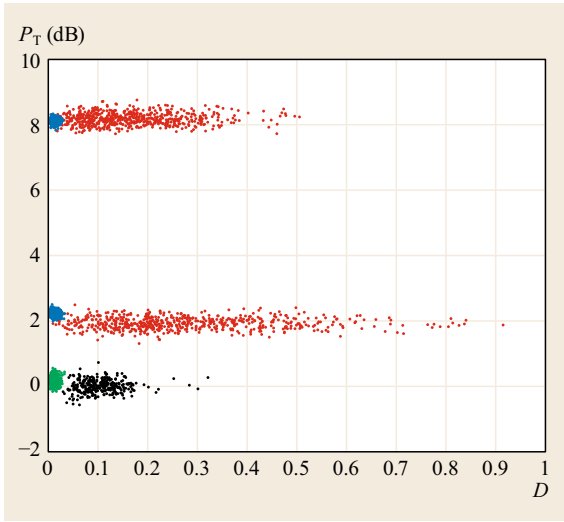

Fig. 16.24 Scatter plot showing simulated D and P_T for clean (only thermal noise; *green*), multipath (*black*), spoofing (*red*), and jamming (*blue*) scenarios. The spoofing and jamming scenarios are simulated at two different power levels. The simulated accumulation amplitudes were chosen so that D was allowed to range from 0 to 1 (after [16.63]; reprinted with permission)

The GNSS signal quality monitoring literature has proposed several metrics for signal distortion [16.70]. These metrics are all calculated based on correlation products and so apply at point (3) in Fig. 16.17. The pincer defense adopts the so-called symmetric difference D. Let Y_E and Y_L be the early and late complex accumulations with a predetermined early late spacing, respectively. Then, D is defined as the magnitude of the complex early-late difference: $D \equiv |Y_E - Y_L|$. Thus, D is sensitive to early-late asymmetry in both magnitude and phase.

Unless a spoofer is capable of a nulling attack, then distortion caused by the interaction between authentic and spoofing signals of comparable amplitude will be evident as $D > 0$. Figure 16.23 shows that D approaches zero in the limit of both weak and powerful spoofing. But weak spoofing affects a GNSS receiver no more than multipath, and powerful spoofing can be detected by a significant increase in P_T. Such is the basic premise of the pincer defense.

The pincer defense seeks to classify interference as either spoofing, jamming, or multipath, and to distinguish these categories from normal thermal noise, all on the basis of D and P_T. The challenge can be appreciated in reference to Fig. 16.24, which shows a scatter plot of D and P_T values under simulated spoofing (red), jamming (blue), multipath (black), and clean (only thermal noise; green). Clearly, there is overlap between the categories, especially between low-power spoofing and severe multipath.

The pincer defense detection and identification problem can be stated as follows. Given a time history of measurements $\mathbf{z}_k \equiv [D_k, P_{T,k}]^\top, k \in \mathcal{K} \equiv \{1, 2, \ldots, N\}$, determine whether the receiver experienced no interference (the null hypothesis, H_0), or, whether for $k \in \mathcal{K}_I \equiv \{k \in \mathcal{K} | k \geq k_o\}$, the receiver experienced multipath (H_1), jamming (H_2), or spoofing (H_3), where k_o is the interference onset index. The problem reduces to a set of generalized likelihood ratio tests conditioned on estimates of k_o, on the interference amplitude α, and, for H_2 and H_3, on an estimate of the code delay τ_c.

Figure 16.25 shows an example observation space for a single measurement $\mathbf{z}_k$, partitioned into decision regions for the four hypotheses. The region boundaries depend on the estimates of α and τ_c, on the cost of deciding H_i when H_j is true, $i, j \in \{0, 1, 2, 3\}$, and on the prior probabilities of the four hypotheses.

The problem formulation introduced above is not unique to the pincer defense; indeed, the detection and identification problem for all interference detection techniques can be formulated in terms of H_0, H_1, H_2, and H_3. Joint detection and classification offer the dual benefit of increased detection power and actionable information about the nature of the interference; these benefits, however, come at the cost of additional computational complexity [16.71].

16.6.4 Spectral Analysis

If the discrete-time quantized samples produced by a receiver's RF front end are accessible to a module capable

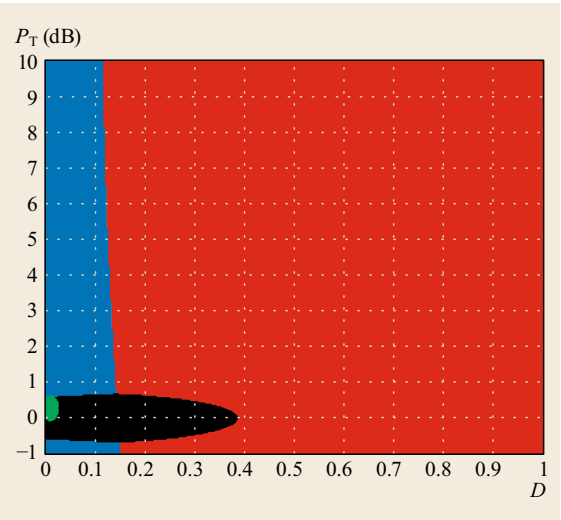

Fig. 16.25 Example observation space for a single measurement $\mathbf{z}_k$ divided into decision regions for clean (only thermal noise; *green*), multipath (*black*), spoofing (*red*), and jamming (*blue*) (after [16.63]; reprinted with permission)

of performing a discrete Fourier transform (DFT), then the received signal power spectrum can be periodically estimated and analyzed. On multifrequency receivers, this may entail analysis of six or more individual GNSS bands. The computational burden of such analysis can be reduced by use of an efficient DFT implementation and by extending the interval between production of power spectra.

Power spectrum analysis is both a simple and powerful interference diagnostic technique, indicating not only the presence but also the nature of interference, whether wideband or narrowband, constant or fleeting. Figure 16.18 shows the power spectrum centered at L1 produced by a 1 s interval of data from a high-quality static receiver in a quiet RF environment. The spectrum shown is an estimate based on the usual periodogram technique of averaging the spectra produced by overlapping sections of the original data, with each time segment weighted by a windowing function.

The key challenge of interference detection and identification via power spectral analysis is distinguishing actual interference from spectral variability due to signal shadowing, multipath, temperature variation, and the changing overhead GNSS signal constellation. As shown in the example data set in Fig. 16.19, the aggregate power in the 2 MHz band centered at L1 can vary by more than 1 dB even when no interference is present. Much of this variation is periodic and therefore predictable. Sophisticated spectral analysis techniques could apply models or machine learning to distinguish

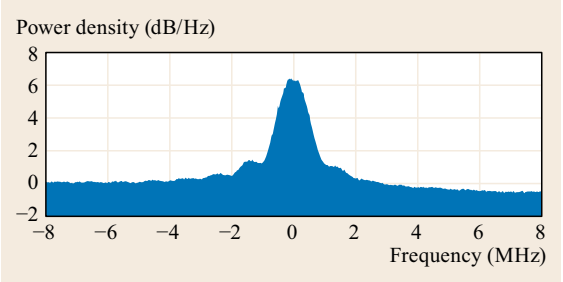

Fig. 16.26 Power spectrum under the same conditions as Fig. 16.18 except that the receiver is now subject to a GPS spoofing attack using the testbed described in [16.62]

novel interference from background variability. Naturally, the problem is much less challenging for static receivers than for mobile ones.

Spectral analysis, even acting alone, can be effective at discovering spoofing. Figure 16.26 shows the same 16 MHz wide power spectrum as in Fig. 16.18 and for the same receiver but for data captured during a spoofing attack in which a false signal was generated for each authentic signal. The profile in Fig. 16.26 thus represents the power spectrum of an admixture of spoofing and authentic signal ensembles. The attack was designed to be stealthy, achieving approximate authentic signal nulling (as described in Sect. 16.5.2) during the interval of data from which the spectrum was computed. Even so, obvious differences are evident between Figs. 16.26 and 16.18. Besides the approximately 2 dB increase in power in the 2 MHz band centered at L1, the side lobes on both sides of the main lobe are more prominent in the spoofed spectrum. Such differences offer hope that a useful degree of spoofing detection could be provided based solely on power spectral measurements.

16.6.5 Cryptographic Spoofing Detection

A GNSS signal modulated with an unpredictable but verifiable security code $W(t)$, as in (16.12), is much more resistant to spoofing than a GNSS signal with no purposeful unpredictability. The security code $W(t)$ is best implemented as a cryptographic sequence. In NMA, $W(t)$ is a low-rate (e.g., 50–250 Hz) binary sequence containing periodic digital signatures that are unpredictable at transmission but can be verified upon receipt to certify the origin of the complete data sequence $D(t)$ [16.55–57]. Alternatively, $W(t)$ can be implemented as a high-rate (e.g., 500–10 000 kHz) binary sequence whose chip interval can be as short as that of the underlying spreading code $C(t)$, as is the case for the GPS Y and M signals, the Galileo PRS signal, and spread-spectrum security codes proposed for civil applications [16.53].

The security of the military GPS Y and M codes is based on symmetric-key cryptography. The GPS control segment generates a pseudorandom binary spreading code sequence based on a combination of secret keys. A military receiver generates a local replica of the same sequence based on a functionally equivalent set of secret keys, enabling despreading and signal tracking. Unauthorized agents are presumably denied access to the secret keys, so, in theory, they can neither generate nor predict the spreading sequence, which means they can neither track nor anticipate the military GPS signals for purposes of spoofing.

It is neither practical nor prudent to base civil security codes on symmetric-key cryptography. Instead, all proposed civil schemes are based on public-key cryptography or on delayed disclosure of secret keys. Even the technique proposed in [16.72], which leverages the military Y code to secure civil GPS receivers, assumes that the Y code is revealed to the receiver some time after receipt.

Detection

Spoofing of a security-code-enhanced GNSS signal is easily detected if the counterfeit signal's security code fails digital signature verification (for low-rate security codes) or fails to generate significant power when correlated against a replica security code (for high-rate security codes). Only meaconing and SCER attacks are capable of generating counterfeit signals that could satisfy these preliminary tests.

For both meaconing and SCER attacks, the detection techniques discussed previously can be quite effective, particularly augmented received power monitoring and spectral analysis. For SCER attacks, another powerful tailored detection test can be formulated [16.54, 55]. The test's decision statistic is based on received power P_T and on a specialized correlation statistic L. Given its dependence on P_T, SCER attack detection can be thought of as another type of received power monitoring augmentation, much like C/N_0 monitoring or the pincer defense.

The SCER attack detector's specialized correlation statistic L is designed to be sensitive to the high error variance of the spoofer's security code estimate $\hat{W}(t)$ in the moments immediately following each unpredictable chip transition. Reference [16.54] develops the statistic and describes its distribution under H_0 (no attack) and H_1 (SCER attack). What follows briefly describes how the statistic is generated within a receiver and offers an example test result.

Let W_k be the value of the security code $W(t)$ during the k-th chip. For convenience, assume that the receiver's accumulation interval is equivalent to the length of W_k, as for NMA. Then, the correlation statistic L can be generated as shown in Fig. 16.27. The lower signal path is the standard matched-filter-type correlation operation previously depicted in continuous time in Fig. 16.2. The product of the incoming samples r_i and a complex local signal replica $l_i = W_k C_l(t_i - \hat{\tau}_i)\exp(-\mathrm{j}(2\pi f_{\mathrm{IF}} t_i + \hat{\theta}(t_i)))$ is accumulated over the interval spanned by W_k to produce the prompt complex correlation products $I_k + \mathrm{j}Q_k$ that get fed to code and carrier tracking loops. The code tracking loop also ingests correlation products from identical paths – not shown – involving early and late versions of $C_l(t_i - \hat{\tau}_i)$.

The upper path in Fig. 16.27 produces the SCER attack detection statistic L. The real part of the product $r_i l_i$ is multiplied by a smooth weighting function $\beta(n_{ki})$, defined in [16.54], that gives full weight to the i_k-th sample but decays rapidly toward zero for subsequent samples. This weighting has the effect of suppressing those samples over which the error variance in the spoofer's security code chip estimate $\hat{W}_k$ has become small because the spoofer has had sufficient time to obtain an accurate estimate of W_k; as illustrated in Fig. 16.14, only the early high-variance samples are useful in distinguishing H_1 from H_0. The weighted product $\beta(n_{ki})\mathcal{R}(r_i l_i)$ is accumulated over the interval spanned by W_k to produce the single-chip detection statistic S_k, N of which are biased, squared, and accumulated as shown to produce the final statistic L. The constants a and b are related to the theoretical mean

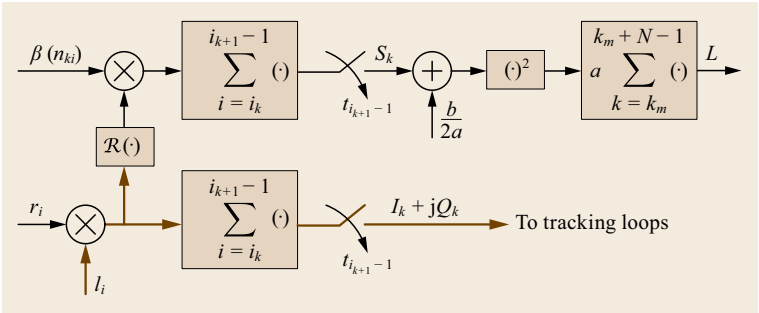

Fig. 16.27 Block diagram illustrating how generation of the SCER attack statistic L relates to standard GNSS signal correlation. *Thick brown lines* denote complex signals, whereas *thin black lines* denote real-valued signals

μ_p and variance σ_p^2 of S_k under $H_p, p = 0, 1$ by

$$a = \frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}, \qquad b = 2\left(\frac{\mu_1}{\sigma_1^2} - \frac{\mu_0}{\sigma_0^2}\right).$$

SCER Attack Detection Example

The test results shown in Fig. 16.28 are expressed in terms of the empirical distribution of L at various stages of an example SCER attack performed in the testbed of [16.62]. The top panel shows the attack prelude during which only the authentic signal is present. At this stage, the histogram of L values exhibits good correspondence with the theoretical null-hypothesis probability distribution $p_{L|H_0}(\xi|H_0)$, where ξ is the value at which the probability density of the detection statistic L is evaluated. The center panel shows the situation during the initial stage of the attack when the authentic and spoofing signals are aligned to within a small fraction of the $\approx 1\,\mu$s spreading code chip interval. Because the counterfeit and authentic signals in this test are so nearly matched in power, this stage manifests strong interaction between the two in the defender's complex-valued prompt correlator. Such interaction violates the either/or assumption of the SCER detection test. The detection statistic does exceed the threshold more than half the time, but instead of clustering within $p_{L|H_1}(\xi|H_1)$, it exhibits spreading driven by variations in the relative carrier phase of the interacting authentic and spoofing signals.

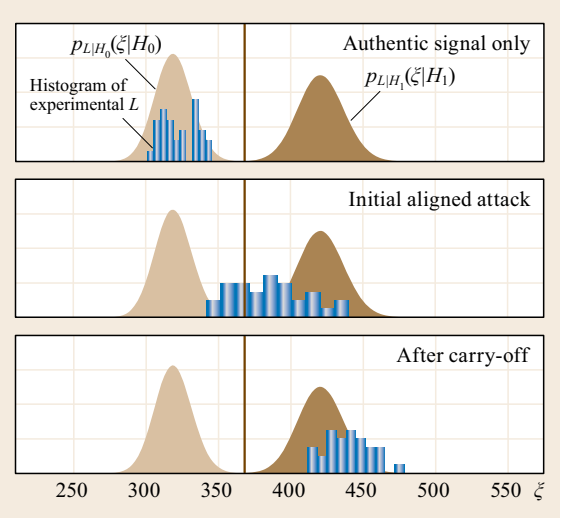

Fig. 16.28 Histograms of experimentally generated detection statistics L (*bar plots*) compared with the detection threshold (*thick vertical line*) and the theoretical distributions $p_{L|H_j}(\xi|H_j), j = 0, 1$ at various stages of a zero-delay SCER attack

After the spoofer has successfully carried off the defender's tracking points and the authentic and spoofed correlation peaks are separated by more than two spreading code chips, the SCER detector's attack model again becomes valid. The bottom panel of Fig. 16.28 shows that at this stage, the detection statistic clearly clusters beyond the detection threshold and roughly within the theoretical $p_{L|H_1}(\xi|H_1)$ distribution.

16.6.6 Antenna-Based Techniques

A GNSS receiver employing only a single, static antenna cannot measure the arrival direction of incoming signals, but a receiver with a moving antenna or multiple antenna elements can discern arrival direction and can use this information to detect interference. Antenna-based techniques are powerful for interference detection because an interference source commonly transmits from a single antenna whereas GNSS signals come from a spatially diverse set of overhead satellites. A spoofing detector based on a single moving antenna is developed in [16.73], and one based on a pair of static antennas is developed in [16.74]. The latter demonstrates nearly immediate spoofing detection with a low-cost system in a live spoofing attack.

16.6.7 Innovations-Based Techniques

A final opportunity for detecting spoofing interference arises in the PVT estimation algorithm that draws in the GNSS pseudorange and carrier-phase observables produced by the tracking loops, or, in the case of a vector tracking architecture, in the consolidated tracking and PVT estimation algorithm. The tracking block in Fig. 16.17 is intended as a generic reference to such tracking and estimation functions, and would be the application point for innovations-based spoofing detection techniques.

PVT estimation algorithms typically employ a model of the receiver dynamics – including clock dynamics – and may have access to non-GNSS aiding data such as from an inertial measurement unit (IMU), barometer, magnetometer, etc. Sequential estimators such as the Kalman filter are commonly used for this purpose, processing a regular cadence of observables and generating a regular output of PVT estimates.

Significant inconsistency between the estimator's predictions and GNSS observables can be detected by standard hypothesis testing applied to the estimator residuals, or innovations (Chap. 24). Reference [16.51] offers a framework for innovations analysis that is optimized for sensor deception, including GNSS spoofing. The framework applies an integrity risk performance

index to account for the fact that a sensor attack only causes harm when the target system exceeds its alert limit – when a ship leaves its assumed transit corridor or a timing system exceeds its required timing accuracy specification, for example. The framework adopts a minimax detection strategy for robustness to unknown spoofer actions. It is shown that an attacker can cause the target system to exceed its protection limits without detection whenever the attack-induced dynamics lie comfortably within the drift envelope of the PVT estimator's model-based propagation process. For example, PVT estimation based on pseudorange and Doppler observables and inertial sensors, a common combination, can be led astray by a spoofer whose induced error trajectory gradually departs from the true trajectory as if driven by the drift processes in the inertial sensors [16.50].

In response to this vulnerability, [16.75] proposes a powerful detection test for GNSS-guided vehicles that exploits high-frequency platform dynamics caused by environmental disturbances (e.g., wind gusts buffeting an aircraft). These dynamics are practically unpredictable to a would-be spoofer yet easily measured by both the inertial sensors and high-rate (e.g., 20 Hz) carrier-phase observables. An innovations test on the GNSS carrier-phase measurements that exploits such natural dithering, or even purposeful dithering if natural disturbances offer inadequate excitation, poses great difficulty for a spoofer unless the spoofer is physically attached to the target platform.

16.7 Interference Mitigation

GNSS interference detection is the key to avoiding hazardously misleading information in a GNSS-based PVT solution: Once interference has been detected, the user or larger system can make decisions with full knowledge that the trustworthiness of the PVT solution may be compromised. But mere detection does not ensure *continuity* of reliable PVT information, which is a requirement for many systems and users. PVT continuity may be achieved by human intervention: A ship's crew can fall back to visual, radar, or even celestial navigation once alerted to GNSS interference. But, increasingly, navigation and timing systems are expected to maintain PVT continuity *automatically* in the face of GNSS interference.

One design philosophy gaining traction in recent years views GNSS as so vulnerable to interference that it must be backstopped with an entirely GNSS-independent PVT source. According to this philosophy, the sensible response to detection of threatening GNSS interference is to abandon GNSS, at least temporarily, by failing over to a non-GNSS backup PVT system. But despite impressive advances in IMU and clock stability, in the use of non-GNSS signals of opportunity for PVT, in non-GNSS time distribution, in electro-optical navigation, and in dedicated terrestrial PVT systems, this approach has only proven useful for short intervals of time (a few minutes) or restricted areas of operation (a radius of a few tens of kilometers). So far, GNSS remains irreplaceable because no combination of non-GNSS PVT systems has yet to rival the essential suite of GNSS benefits: (1) global coverage, (2) high PVT accuracy over indefinitely long time intervals, and (3) low cost to users. Accordingly, this section focuses on GNSS interference mitigation techniques that ensure PVT resilience not by abandoning GNSS but by toughening and augmenting it.

16.7.1 Spectrally or Temporally Sparse Interference

Effective techniques exist for mitigating interference that is sparse in frequency (narrowband) or time (pulsed). Mitigation of spatially sparse interference, that is, interference with a small number of narrow directions of arrival, will be treated in Sect. 16.7.3.

Sparse interference mitigation techniques exploit time correlation in an interference signal's phase or amplitude to estimate and excise the interference signal, thereby, increasing the desired signal power to noise ratio. The more highly time correlated an interference signal's amplitude or phase, the more accurately it can be reconstructed and excised, sparing the downstream acquisition and tracking routines from harmful interference effects.

Filtering

Without proper early stage RF filtering, even interference far from GNSS frequency bands of interest can be problematic for a GNSS receiver when the interference is sufficiently strong: The out-of-band signal rejection of the receiving antenna and the first-stage LNA may not be sufficient to prevent a strong out-of-band signal from saturating the LNA. Thus, in mobile handsets and at cellular base stations, one finds GNSS receivers with stringent RF filtering before first-stage amplification despite the direct C/N_0 reduction (equivalent to the filter impedance loss) that such filtering entails.

Narrowband interference within the GNSS band is more challenging to mitigate than out-of-band interference. Selective (high quality factor) analog filtering within a GNSS band of interest requires large and expensive analog filters. Likewise, LNAs with a linear range wide enough to prevent saturation in the face of strong interference are expensive, as are antenna arrays capable of pointing a null toward the interference source. Thus, attenuation of the received signal before low-noise amplification may in some cases be the only economical recourse to prevent LNA saturation. Unfortunately, one pays the full measure of such attenuation in reduced C/N_0.

Assuming LNA saturation is avoided, properly configured multibit quantization can be a first defense against narrowband interference. As mentioned in Sect. 16.3.2, multibit quantization can yield a conversion gain (an increase in C/N_0 relative to the unquantized discrete-time samples) when the amplitude of the incoming interference is approximately constant. However, for the one-bit (two-level) quantization employed in many low-cost GNSS receivers, quantizer SNR is severely and irrecoverably degraded by the presence of strong narrowband interference. Even two-bit (four-level) quantization may be insufficient to prevent capture of the quantization process by a strong narrowband interferer, if the interference amplitude varies rapidly or if there are multiple narrowband interferers present.

Assuming sufficient quantization resolution, adaptive digital filtering in the precorrelation stage (point (2) in Fig. 16.17) is a low-cost and highly effective way to mitigate in-band narrowband interference. This technique, commonly referred to as adaptive notch filtering, exploits the time correlation of narrowband interference signals to distinguish them from thermal noise and from the desired spread-spectrum signal, both of which look uncorrelated at chip-length sampling intervals.

Adaptive notch filtering can be implemented either as a transversal filter in the time domain or as shaping in the frequency domain. In the time-domain approach, the weights of a transversal filter are adjusted to minimize the filter's output power [16.76]. Solution of the optimal tap weight vector has complexity $\mathcal{O}(n^2)$, where n is the number of samples in the block used to determine the optimal weights. One may trade off performance for reduced computational demand by extending the interval between subsequent computation of the optimal weight vector. Straightforward implementation can yield highly effective interference suppression even for multiple narrowband interferers: *Dimos* et al. [16.77] show that three pure tone interference sources with a combined interference-to-thermal-noise power of 30 dB in the GPS L1 C/A band can be suppressed by 28 dB. For the same interference power and number of interferers, but with bandwidths of 25, 50, and 100 kHz, suppression performance reduces to 24.25, 20.75, and 16 dB, respectively, showing that time-domain notch filtering performance degrades as the interference bandwidth increases.

The frequency domain approach entails Fourier transformation of a block of n precorrelation samples (possibly weighted by a windowing function), multiplication of the transform by some appropriate filter, and inverse Fourier transformation of the product. The interference suppression filter applied in the transform domain can be generated automatically to whiten the transformed samples. In the simplest approach, regions containing interference peaks exceeding a predefined threshold can be simply blanked out. The transform approach has complexity $\mathcal{O}(n\log(n))$ and so is less computationally burdensome than time-domain notch filtering with continuous updating of the filter tap weighting. Another benefit of the transform approach is that successive transforms can be averaged to produce a power spectrum estimate, which, as mentioned earlier, is a useful tool for general situational awareness of the interference environment.

The distinctive swept tone interference of PPDs can also be considered sparse given its high regularity [16.47]. A model-based technique is developed in [16.49] that effectively estimates the frequency sweep parameters of PPD signals, allowing the interference to be excised. Such model-based filtering is the logical extension of notch filtering for interference signals that are highly predictable and easily distinguished from the desired GNSS signals.

Blanking

Interference signals that are sparse in time, for example, pulsed interference, can be substantially suppressed by so-called pulse blanking [16.1]. Blanking degrades C/N_0 in proportion to the fraction of RF front-end samples that are discarded. A combined adaptive notch filtering and blanking technique is explored in [16.1] to mitigate DME/TACAN interference, which is sparse in both time and frequency.

16.7.2 Spectrally and Temporally Dense Interference

Interference that is both wideband and continuous is spectrally and temporally dense, unlike narrowband or pulsed interference. It may yet be spatially sparse, but a GNSS receiver with a single, static antenna is unable to exploit such sparseness for mitigation. In this section, dense interference will refer to interference

which is both spectrally and temporally dense regardless of its spatial characteristics. The focus will be on signal-processing-based interference mitigation techniques that do not rely on multiple or moving antennas. The next section treats mitigation of spatially sparse interference using multiple or moving antennas.

Dense interference has substantially time-uncorrelated amplitude and phase at the RF front-end sampling rate, making it appear as thermal noise or as a spread-spectrum GNSS signal to the receiver. Spoofing interference (including meaconing) is an example of interference that is especially difficult to mitigate, because by construction it is intended to masquerade as a legitimate GNSS signal. Faced with multiple identically shaped and sized autocorrelation peaks for the same pseudorandom number code, a receiver can easily recognize that a spoofing attack is underway but cannot mitigate the attack – that is, cannot identify and track only the authentic signal – unless the receiver's combined timing and positioning uncertainty is well within the inter-peak separation. For this reason, post-detection mitigation of a subtle spoofing attack is often only possible by exploiting multiple or moving antennas and will therefore be left to the next section.

It is convenient to treat dense nonspoofing interference such as continuous wideband Gaussian interference as if it were thermal noise for purposes of mitigation. Thus, the dense interference mitigation problem becomes identical to the problem of acquiring and tracking weak GNSS signals in an indoor environment except that the multipath effects in the indoor environment are likely to be more severe than in an outdoor interference environment. Mitigation is applied at the correlation and post-correlation stage, or point (3) in Fig. 16.17. Given a front-end bandwidth of W_{FE} Hz and an in-band interference-to-signal power ratio of P_I/P_S, the resulting effective C/N_0 will be as in (16.8), which for strong interference becomes $C/N_{0,\text{eff}} = P_S W_{FE}/P_I$. Thus, to withstand interference exceeding $P_I/P_S = 50\,\text{dB}$ in a $W_{FE} = 10\,\text{MHz}$ bandwidth, a receiver would need to acquire and track GNSS signals below $C/N_{0,\text{eff}} = 10\log_{10}(10^7) - 50 = 20\,\text{dB Hz}$.

Consumer-grade GNSS receivers offer surprisingly good protection against dense interference despite their low cost, because they have been designed for operation at low C/N_0. Even without network aiding, a consumer-grade GNSS receiver can acquire signals from a cold start at −148 dBm, which corresponds to $C/N_0 = 26\,\text{dB Hz}$ for a typical $N_0 = -174\,\text{dBm/Hz}$. This amounts to resilience against P_I/P_S up to 37 dB in a 2 MHz bandwidth. Tracking and performance can be substantially better than cold-start acquisition, achieving remarkable thresholds as low as −167 dBm, or $C/N_0 = 7\,\text{dB Hz}$ assuming $N_0 = -174\,\text{dBm/Hz}$ [16.78].

The receiver presented in [16.11] can be considered a benchmark for what is possible with a stand-alone scalar-tracking architecture when computational limitations are ignored. Its algorithms can acquire and maintain lock on signals down to $C/N_0 = 18\,\text{dB Hz}$ by assuming a low-cost TCXO and moderate acceleration uncertainty. Clearly, the superior tracking performance of the consumer-grade receiver in [16.78] implies a vectorized tracking architecture.

The current state-of-the art in low-C/N_0 acquisition and tracking is embodied in the DINGPOS high-sensitivity GNSS platform for deep indoor scenarios [16.79]. The platform records synchronized data from a micro-electromechanical system (MEMS) IMU, a barometer, a magnetometer, and a GNSS RF front-end driven by an OCXO-quality reference clock. The data are combined with known navigation data symbols in a software-defined GNSS receiver employing a vector tracking architecture to achieve coherent integration over 2 s intervals under pedestrian dynamics. In dynamic simulation scenarios, DINGPOS acquires down to $C/N_0 = 6\,\text{dB Hz}$ and tracks down to $C/N_0 = -1\,\text{dB Hz}$. This represents remarkable interference immunity: up to $P_I/P_S = 71\,\text{dB}$ in a $W_{FE} = 10\,\text{MHz}$ bandwidth for tracking. Even higher P_I/P_S immunity can be achieved by combining DINGPOS-style signal processing with antenna array processing, the subject of the next section.

16.7.3 Antenna-Based Techniques

Though currently expensive, multielement antenna arrays are perhaps the most effective general tool for interference mitigation. Antenna array interference mitigation exploits spatial sparseness in the direction of arrival of interference sources and spatial diversity in the direction of arrival of desired GNSS signals from overhead satellites. Early array processing methods passed the RF signal from each array element through a variable phase shifter. The phase-shifted RF signals were then combined into a single RF stream that was directed to the RF front end for conditioning and digitization. In this approach, the GNSS receiver saw only a single antenna gain pattern (e.g., a pattern with a null directed toward an interference source) at any given instant.

The modern approach to array processing is much more flexible. The RF feed from each antenna is independently digitized, as shown in Fig. 16.17. A complex weight vector is applied across the individual digitized streams to achieve a desired gain pattern. Importantly, any number of weighted combinations of the digital

streams can be created simultaneously, with the unique combinations fed to a bank of separate GNSS processing channels. In this way, each channel sees an alternative antenna array gain pattern, which permits a beam to be steered toward the satellite whose signal the channel is intended to track, for example.

Continuously calculating the set of optimal weighting vectors is the primary computational challenge of array processing, with the primary practical challenge being the need to periodically calibrate the array as temperature and other environmental variations cause minute but significant changes in the phase shift through each antenna element.

A computationally efficient approach to weighting vector calculation is offered in [16.80], but this approach requires the direction of arrival of the desired signal to be known, which entails knowledge of the antenna array's attitude in global coordinates. Preferable are blind adaptive techniques such as the one proposed in [16.81], which automatically maximizes the ratio of power in the desired signal to power in the interference signal plus thermal noise in the correlation products. Better still, though more computationally demanding, are joint space–time interference mitigation techniques that exploit interference time correlation or spatial correlation, or both, in a joint space–time mitigation framework [16.82]. A single interferer is detected in this framework based on estimates of the spatial correlation matrix. A narrowband interferer is detected based on estimates of the time correlation matrix (or based on time correlation evident in the Fourier domain). Such space–time array processing thus combines the virtues of adaptive notch filtering with adaptive beam forming. The beamforming aspect of the approach works equally well whatever the nature of the interference source – intentional or not, GNSS-like or not – so long as the source presents a compact direction of arrival.

References

16.1 G.X. Gao, L. Heng, A. Hornbostel, H. Denks, M. Meurer, T. Walter, P. Enge: DME/TACAN interference mitigation for GNSS: Algorithms and flight test results, GPS Solutions **17**(4), 561–573 (2013)

16.2 Radio Regulations (ITU-R International Telecommunication Union, Radiocommunication Sector, Geneva 2012) http://www.itu.int/pub/R-REG-RR

16.3 FCC Online Table of Frequency Allocations (Federal Communications Commission, 2013) http://transition.fcc.gov/oet/spectrum/table/fcctable.pdf

16.4 FCC Enforcement Advisory No. 2011-03, DA 11-249 (Federal Communications Commission, 2011) http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-11-249A1.pdf

16.5 Mobile phone and GPS jamming devices FAQ (Australian Communications and Media authority, 2014) http://www.acma.gov.au/theACMA/faqs-mobile-phone-and-gps-jamming-devices-acma

16.6 G.X. Gao, P. Enge: How many GNSS satellites are too many?, IEEE Trans. Aerosp. Electron. Syst. **48**(4), 2865–2874 (2012)

16.7 T.E. Humphreys: The GPS dot and its discontents: Privacy vs. GNSS integrity, Inside GNSS **7**(2), 44–48 (2012)

16.8 M. Rao, C. O'Driscoll, D. Borio, J. Fortuny: Light-squared effects on estimated C/N_0, pseudoranges and positions, GPS Solutions **18**(1), 1–13 (2014)

16.9 P. Misra, P. Enge: *Global Positioning System: Signals, Measurements, and Performance*, 2nd edn. (Ganga-Jumana, Lincoln 2012)

16.10 A.J. van Dierendonck: GPS receivers. In: *Global Positioning System: Theory and Applications*, Vol. 1, ed. by B.W. Parkinson, J.J. Spilker (AIAA, Washington DC 1996) pp. 329–407

16.11 M.L. Psiaki, H. Jung: Extended Kalman filter methods for tracking weak GPS signals, Proc. ION GPS 2002, Portland (ION, Virginia 2002) pp. 2539–2553, 24-27 Sep. 2002

16.12 C. Hegarty, M. Tran, Y. Lee: Simplified techniques for analyzing the effects of non-white interference on GPS receivers, Proc. ION GPS 2002, Portland (ION, Virginia 2002) pp. 620–629

16.13 J.W. Betz: Effect of narrowband interference on GPS code tracking accuracy, Proc. ION NTM 2000, Anaheim (ION, Virginia 2000) pp. 16–27

16.14 P.W. Ward, J.W. Betz, C.J. Hegarty: Interference, multipath, and scintillation. In: *Understanding GPS: Principles and Applications*, ed. by E.D. Kaplan, C.J. Hegarty (Artech House, Boston 2005) pp. 243–299

16.15 F.M. Gardner: *Phaselock Techniques*, 3rd edn. (Wiley, Hoboken 2005)

16.16 T.E. Humphreys, M.L. Psiaki, B.M. Ledvina, P.M. Kintner Jr: GPS carrier tracking loop performance in the presence of ionospheric scintillations, Proc. ION GNSS 2005, Long Beach (ION, Virginia 2005) pp. 156–167

16.17 M.K. Simon, M. Alouini: *Digital Communications over Fading Channels* (Wiley, New York 2000)

16.18 A.J. Viterbi: *Principles of Coherent Communication* (McGraw-Hill, New York 1966)

16.19 S.C. Gupta: Phase-locked loops, Proc. IEEE **63**(2), 291–306 (1975)

16.20 G. Ascheid, H. Meyr: Cycle slips in phase-locked loops: A tutorial survey, IEEE Trans. Comm. **COM-30**(10), 2228–2241 (1982)

16.21 B. Motella, S. Savasta, D. Margaria, F. Dovis: Method for assessing the interference impact on GNSS receivers, IEEE Trans. Aerosp. Electron. Syst. **47**(2),

1416–1432 (2011)
16.22 J.J. Spilker Jr.: GPS signal structure and theoretical performance. In: *Global Positioning System: Theory and Applications*, Vol. 1, ed. by B.W. Parkinson, J.J. Spilker (AIAA, Washington DC 1996) pp. 57–119
16.23 Navstar GPS Space Segment / Navigation User Segment Interfaces, Interface Specification, IS-GPS-200, Rev. H (Global Positioning Systems Directorate, Los Angeles Air Force Base, El Segundo 2013)
16.24 J.J. Spilker Jr.: Interference effects and mitigation techniques. In: *Global Positioning System: Theory and Applications*, Vol. 1, ed. by B.W. Parkinson, J.J. Spilker (AIAA, Washington DC 1996) pp. 717–771
16.25 J.H. van Vleck, D. Middleton: The spectrum of clipped noise, Proc. IEEE **54**(1), 2–19 (1966)
16.26 F. Amoroso: Adaptive A/D converter to suppress CW interference in DSPN spread-spectrum communications, IEEE Trans. Comm. **31**, 1117–1123 (1983)
16.27 C.J. Hegarty: Analytical model for GNSS receiver implementation losses, Navigation **58**(1), 29–44 (2011)
16.28 J. Max: Quantizing for minimum distortion, IRE Trans. Inf. Theory **6**(1), 7–12 (1960)
16.29 F. Amoroso, J.L. Bricker: Performance of the adaptive A/D converter in combined CW and Gaussian interference, IEEE Trans. Comm. **34**(3), 209–213 (1986)
16.30 D.J. McLean, N.R. Labrum: *Solar Radiophysics: Studies of Emission from the Sun at Metre Wavelengths* (Cambridge Univ. Press, New York 1985)
16.31 P.M. Kintner Jr., T.E. Humphreys, J. Hinks: GNSS and ionospheric scintillation: How to survive the NextSolar maximum, Inside GNSS **4**(4), 22–30 (2009)
16.32 R.V. Jones: *Most Secret War* (Penguin UK, London 2009)
16.33 A.P. Cerruti, P.M. Kintner, D.E. Gary, A.J. Mannucci, R.F. Meyer, P. Doherty, A.J. Coster: Effect of intense December 2006 solar radio bursts on GPS receivers, Space Weather **6**(10), 1–10 (2008)
16.34 A.P. Cerruti, P.M. Kintner, D.E. Gary, L.J. Lanzerotti, E.R. de Paula, H.B. Vo: Observed solar radio burst effects on GPS/wide area augmentation system carrier-to-noise ratio, Space Weather **4**(10), 1–9 (2006)
16.35 D.M. Akos: Who's afraid of the spoofer? GPS/GNSS spoofing detection via automatic gain control (AGC), Navigation **59**(4), 281–290 (2012)
16.36 G.M. Nita, D.E. Gary, L.J. Lanzerotti, D.J. Thomson: The peak flux distribution of solar radio bursts, Astrophys. J. **570**, 423–438 (2002)
16.37 T.E. Humphreys, M.L. Psiaki, B.M. Ledvina, A.P. Cerruti, P.M. Kintner: A data-driven testbed for evaluating GPS carrier tracking loops in ionospheric scintillation, IEEE Trans. Aerosp. Electron. Syst. **46**(4), 1609–1623 (2010)
16.38 T.E. Humphreys, M.L. Psiaki, P.M. Kintner: Modeling the effects of ionospheric scintillation on GPS carrier phase tracking, IEEE Trans. Aerosp. Electron. Syst. **46**(4), 1624–1637 (2010)
16.39 J. Aarons: Global morphology of ionospheric scintillations, Proc. IEEE **70**(4), 360–378 (1982)
16.40 J. Aarons: Global positioning system phase fluctuations at auroral latitudes, J. Geophys. Res. **102**, 17219–17231 (1997)
16.41 B.M. Ledvina, J.J. Makela, P.M. Kintner: First observations of intense GPS L1 amplitude scintillations at midlatitude, Geophys. Res. Lett. **29**(14), 4-1–4-4 (2002)
16.42 T.E. Humphreys, M.L. Psiaki, J.C. Hinks, B. O'Hanlon, P.M. Kintner Jr.: Simulating ionosphere-induced scintillation for testing GPS receiver phase tracking loops, IEEE J. Sel. Top. Signal Process. **3**(4), 707–715 (2009)
16.43 J. Do, D.M. Akos, P.K. Enge: L and S bands spectrum survey in the San Francisco Bay area, Proc. IEEE PLANS 2004, Monterey (IEEE, New York 2004) pp. 566–572
16.44 R. Johannessen, S.J. Gale, M.J.A. Asbury: Potential interference sources to GPS and solutions appropriate for applications to civil aviation, IEEE Aerosp. Electron. Syst. Mag. **5**(1), 3–9 (1990)
16.45 A.T. Balaei, A.G. Dempster: A stasistical inference technique for GPS interference detection, IEEE Trans. Aerosp. Electron. Syst. **45**(4), 1499–1511 (2009)
16.46 C. Kurby, R. Lee, L. Cygan, E. Derbez: Maintaining precision receiver performance while rejecting adjacent band interference, Proc. ION ITM 2012, Newport Beach (ION, Virginia 2012) pp. 574–597
16.47 R.H. Mitch, R.C. Dougherty, M.L. Psiaki, S.P. Powell, B.W. O'Hanlon, J.A. Bhatti, T.E. Humphreys: Signal characteristics of civil GPS jammers, Proc. ION GNSS 2011, Portland (ION, Virginia 2011) pp. 1907–1919
16.48 K.D. Wesson, T.E. Humphreys: Hacking drones, Sci. Am. **309**(5), 54–59 (2013)
16.49 R.H. Mitch, M.L. Psiaki, S.P. Powell, B.W. O'Hanlon: Signal acquisition and tracking of chirp-style GPS jammers, Proc. ION GNSS 2013, Nashville (ION, Virginia 2013) pp. 2893–2909
16.50 A.J. Kerns, D.P. Shepard, J.A. Bhatti, T.E. Humphreys: Unmanned aircraft capture and control via GPS spoofing, J. Field Robotics **31**(4), 617–636 (2014)
16.51 J. Bhatti: Sensor Deception Detection and Radio-Frequency Emitter Localization, Ph.D. Thesis (University of Texas, Austin 2015)
16.52 European GNSS (Galileo) Open Service Signal In Space Interface Control Document, OS SIS ICD, Iss. 1.2 (European Union 2015)
16.53 L. Scott: Anti-spoofing and authenticated signal architectures for civil navigation systems, Proc. ION GNSS 2003, Portland (ION, Virginia 2003) pp. 1542–1552
16.54 T.E. Humphreys: Detection strategy for cryptographic GNSS anti-spoofing, IEEE Trans. Aerosp. Electron. Syst. **49**(2), 1073–1090 (2013)
16.55 K.D. Wesson, M.P. Rothlisberger, T.E. Humphreys: Practical cryptographic civil GPS signal authentication, Navigation **59**(3), 177–193 (2012)
16.56 A.J. Kerns, K.D. Wesson, T.E. Humphreys: A blueprint for civil GPS navigation message authentication, Proc. IEEE/ION PLANS 2014, Monterey (ION, Virginia 2014) pp. 262–269
16.57 I. Fernandez Hernandez, V. Rijmen, G. Seco Granados, J. Simon, I. Rodriguez, J.D. Calle: Design drivers, solutions and robustness assessment of navigation message authentication for the Galileo

open service, Proc. ION GNSS 2014, Tampla (ION, Virginia 2014) pp. 2810–2827
16.58 T.E. Humphreys, B.M. Ledvina, M.L. Psiaki, B.W. O'Hanlon, P.M. Kintner Jr.: Assessing the spoofing threat: Development of a portable GPS civilian spoofer, Proc. ION GNSS 2008, Savannah (ION, Virginia 2008) pp. 2314–2325
16.59 G. Hein, F. Kneissl, J.-A. Avila-Rodriguez, S. Wallner: Authenticating GNSS: Proofs against spoofs, Part 2, Inside GNSS **2**(5), 71–78 (2007)
16.60 Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System (John A. Volpe National Transportation Systems Center 2001)
16.61 O. Pozzobon, C. Wullems, M. Detratti: Security considerations in the design of tamper resistant GNSS receivers, Proc. NAVITEC 2010, Noordwijk (IEEE, New York 2010)
16.62 T.E. Humphreys, D.P. Shepard, J.A. Bhatti, K.D. Wesson: A testbed for developing and evaluating GNSS signal authentication techniques, Proc. Int. Symp. Certif. GNSS Syst. Serv. (CERGAL), Dresden (DGON, Düsseldorf 2014)
16.63 K. Wesson: Secure Navigation and Timing Without Local Storage of Secret Keys, Ph.D. Thesis (University of Texas, Austin 2015)
16.64 V. Dehghanian, J. Nielsen, G. Lachapelle: GNSS spoofing detection based on receiver C/N_0 estimates, Proc. ION GNSS 2012, Nashville (ION, Virginia 2012) pp. 2878–2884
16.65 P.W. Ward: GPS receiver RF interference monitoring, mitigation, and analysis techniques, Navigation **41**(4), 367–391 (1994)
16.66 S. Lo, D. Akos, F.M. Eklof, O. Isoz, H. Borowski: Detecting false signals with automatic gain control, GPS World **23**(4), 38–43 (2012)
16.67 T.E. Humphreys, J.A. Bhatti, D.P. Shepard, K.D. Wesson: The texas spoofing test battery: Toward a standard for evaluating GNSS signal authentication techniques, Proc. ION GNSS 2012, Nashville (ION, Virginia 2012) pp. 3569–3583
16.68 National Oceanic and Atmospheric Administration: Space Weather Alerts Description and Criteria, http://legacy-www.swpc.noaa.gov/alerts/description.html#electron
16.69 A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, G. Lachapelle: Pre-despreading authenticity verification for GPS L1 C/A signals, Navigation **61**(1), 1–11 (2014)
16.70 S. Gunawardena, Z. Zhu, M.U. de Haag, F. van Graas: Remote-controlled, continuously operating GPS anomalous event monitor, Navigation **56**(2), 97–113 (2009)
16.71 A.S. Willsky: A survey of design methods for failure detection in dynamic systems, Automatica **12**(6), 601–611 (1976)
16.72 M.L. Psiaki, B.W. O'Hanlon, J.A. Bhatti, D.P. Shepard, T.E. Humphreys: GPS spoofing detection via dual-receiver correlation of military signals, IEEE Trans. Aerosp. Electron. Syst. **49**(4), 2250–2267 (2013)
16.73 M.L. Psiaki, S.P. Powell, B.W. O'Hanlon: GNSS spoofing detection using high-frequency antenna motion and carrier-phase data, Proc. ION GNSS 2013, Nashville (ION, Virginia 2013) pp. 2949–2991
16.74 M.L. Psiaki, B.W. O'Hanlon, S.P. Powell, J.A. Bhatti, K.D. Wesson, T.E. Humphreys, A. Schofield: GNSS spoofing detection using two-antenna differential carrier phase, Proc. ION GNSS 2014, Tampa (ION, Virginia 2014) pp. 2776–2800
16.75 S. Khanafseh, N. Roshan, S. Langel, F.-C. Chan, M. Joerger, B. Pervan: GPS spoofing detection using RAIM with INS coupling, Proc. IEEE/ION PLANS 2014, Monterey (ION, Virginia 2014) pp. 1232–1239
16.76 L.B. Milstein: Interference rejection techniques in spread spectrum communications, Proc. IEEE **76**(6), 657–671 (1988)
16.77 G. Dimos, T.N. Upadhyay, T. Jenkins: Low-cost solution to narrowband GPS interference problem, Proc. Nat. Aerosp. Electron. Conf. NAECON 1995, Dayton (IEEE, New York 1995) pp. 145–153
16.78 MAX-M8 GNSS Module datasheet (u-Blox), UBX-15031506, https://www.u-blox.com/sites/default/files/MAX-M8-FW3_DataSheet_(UBX-15031506).pdf
16.79 H. Niedermeier, B. Eissfeller, J. Winkel, T. Pany, B. Riedl, T. Wörz, R. Schweikert, S. Lagrasta, G. Lopez-Risueno, D. Jiminez-Banos: DINGPOS: High sensitivity GNSS platform for deep indoor scenarios, Proc. Int. Conf. Indoor Position. Indoor Navig. (IPIN), Zurich (IEEE, New York 2010)
16.80 G. Seco-Granados, J.A. Fernández-Rubio, C. Fernández-Prades: ML estimator and hybrid beamformer for multipath and interference mitigation in GNSS receivers, IEEE Trans. Signal Process. **53**(3), 1194–1208 (2005)
16.81 M. Sgammini, F. Antreich, L. Kurz, M. Meurer, T.G. Noll: Blind adaptive beamformer based on orthogonal projections for GNSS, Proc. ION GNSS 2012, Nashville (ION, Virginia 2012) pp. 926–935
16.82 M.H. Castañeda, M. Stein, F. Antreich, E. Tasdemir, L. Kurz, T.G. Noll, J.A. Nossek: Joint space-time interference mitigation for embedded multi-antenna GNSS receivers, Proc. ION GNSS 2013, Nashville (ION, Virginia 2013) pp. 3399–3408