REPORT ITU-R M.2076

Factors that mitigate interference from radiolocation and Earth exploration-satellite service/space research service (active) radars to maritime and aeronautical radionavigation radars in the 9.0-9.2 and 9.3-9.5 GHz bands and between Earth exploration-satellite service/ space research service (active) radars and radiolocation radars in the 9.3-9.5 and 9.8-10.0 GHz bands

(2006)

1 Introduction

Question ITU-R 234/8 calls for study of technical characteristics, performance criteria, and other factors of radiolocation and radionavigation systems in the bands 9 000-9 200 MHz and 9 300-9 500 MHz and of the interference criteria for those systems. In addition, Resolution 747 (WRC-03) has established Agenda item 1.3 for WRC-07 to consider upgrading the allocations to the radiolocation service in the 9 000-9 200 and 9 300-9 500 MHz bands to co-primary and to consider extending the primary allocation to the Earth exploration-satellite (EES) (active) service and space research (SR) (active) service in the band 9 500-9 800 MHz band contiguously by 200 MHz. Characteristics of representative terrestrial radars in the 8 500 MHz-10.5 GHz band are contained in Recommendation ITU-R M.1796. This Report is a further contribution to the studies required by Question ITU-R 234/8 and Resolution 747 (WRC-03).

Recommendation ITU-R M.1372-1 – Efficient use of the radio spectrum by radar stations in the radiodetermination service, describes some of the most important interference suppression techniques that are used in radars generally. The emphasis in that Recommendation is on post-detection processing, although one of the techniques described there can be implemented prior to detection. The factors discussed herein include some of those covered in Recommendation ITU-R M.1372 as well as some that complement those.

1.1 Summary of findings

The main form of interference degradation that pulsed interference is likely to cause is an increase of the rate of false alarms. This is naturally mitigated by some common characteristics of radars, including low antenna sidelobes and asynchronous pulsing. Responses to individual pulses, including fast time constant, matched filtering effects, and other pulse-shortening effects, are beneficial. The form of coupling of most concern is sidelobe-to-main-beam coupling.

Prudent radar design can mitigate pulsed interference in numerous ways. These include:

– multiple-pulse techniques, including *M*-out-of-*N* processing;
– deliberate removal of individual asynchronous pulses;
– sensing of asynchronous-pulse effects in post-processing review of Doppler-filter outputs;
– nonlinear and time-varying processes such as limiting and sensitivity time control;
– scan-to-scan correlation.

2 Types of radars in the bands

Several types of radionavigation radars operate in the 9 000-9 200 and 9 300-9 500 MHz bands. Ground-based aeronautical radionavigation radars operate in the 9 000-9 200 MHz band; they include precision-approach radars (PARs) and airport surface detection equipment (ASDE) radars. These are discrete-target surveillance radars. The 9 300-9 500 MHz band is used by a large number of maritime radionavigation radars, the great majority of them being aboard ships, and by airborne weather-avoidance radars. The maritime systems are discrete-target radars while the airborne systems are distributed-target radars.

The radiolocation service operates on a secondary allocation basis in the 9 000-9 200 and 9 300-9 500 MHz bands. Land-based weather radiolocation radars operating in the 9 300-9 500 MHz band are privileged with respect to other radiolocation radars (Radio Regulations (RR) No. 5.475). Radiolocation radars also operate in the 9 500-9 800 MHz and 9 800 MHz-10.0 GHz bands on a primary-allocation basis.

Spaceborne synthetic-aperture radars (SARs) in the EES/SR (active) services currently operate in the 9.5-9.8 GHz band on a co-primary allocation basis. The proposal to extend that allocation by 200 MHz is driven by a desire to enhance the range resolution of the SARs.

3 Types of potential interference effects

The two most prominent types of performance degradation that radiolocation or EES/SR (active) radars could inflict on discrete-target surveillance radars such as PARs, ASDEs, or maritime navigation radars fall into the categories of:

– missed target detections;

– generation of false target detections or “false alarms” and false target tracks.

These two effects can be thought of as a decrease in probability of detection and an increase in probability of false alarm, respectively.

Although radiolocation or EES/SR (active) radars could conceivably inflict some degree of desensitization (missed target detections, etc.), that effect is expected to be minor, as has been demonstrated in several measurement programs, so attention will focus on the generation of false targets.

Pulsed signals from other radars create a potential for generation of false target detections even when a well-designed “constant-false-alarm-rate” (CFAR) operation is provided in the terrestrial radar. However, the remainder of this Report shows that these effects can largely be avoided by good design. Discrete-target radars, including target-dedicated tracking radars, are also subject to aggravation of position-estimation errors and target-classification errors due to unwanted signals. However, these effects are more likely to be inflicted by continuous, noise-like interference than by pulsed interference from other radars.

Performance degradation that radiolocation and EESS radars could inflict on distributed-target radars, including weather-avoidance radars or weather surveillance radars, consists of discrete (e.g. single-pixel) false alarms (referred to in the weather-radar community as speckle) and introduction of inaccuracy into derived measures of weather phenomena. The degradation that interference of any kind can inflict on synthetic-aperture imaging radars is being expressed by the space-science community as an increase of the variance of processor output power in any pixel[1].

[1] Recommendation ITU-R RS.1166 – Performance and interference criteria for spaceborne active sensors.

These effects are in contrast to the effect of continuous noise-like interference on discrete-target radar that has effective control over its false-alarm rate. In that case, the probability of false alarm tends to remain unchanged, but the curve of probability of detection as a function of target range or radar cross section (RCS) inexorably suffers a shift to shorter range or higher RCS as the undesired signal becomes stronger. This is generalized desensitization, predominately affecting targets that are small, distant, or poorly illuminated due to adverse propagation conditions such as multipath propagation or adverse ducting. It also degrades other functions such as tracking precision. However, continuous noise-like interference is outside the purposes of this Report.

4 Interference-mitigating characteristics commonly found in radars

Interference can be mitigated by weak or transient power coupling, certain receiver nonlinearities, time-varying gain, signal processing, post-processing, and separation in carrier frequency. In radar-to-radar interactions, separation in frequency is not always necessary for compatible operation because high degrees of isolation in power coupling and in time either occur naturally or can be achieved by good design. Isolation through polarization mismatch occurs in some combinations of radiolocation and spaceborne radars and navigation radars, but cannot be relied upon in the general case because radars of a given allocated service often use horizontal, vertical, and/or circular polarization.

Specific mechanisms that contribute to such mitigating factors are identified in the following sections. Many of them apply to pulses coupled from radiolocation or spaceborne-sensor radars to maritime, airborne, and air-traffic-control radars, while some apply mainly to radars in just one or another of those categories.

4.1 Isolation in power coupling (antenna-mediated effects)

Interactions between two radars of different types almost always involve asynchronism between the scanning of the two antenna beams. This is virtually assured when one of the radars is a radiolocation radar and the other is a radionavigation radar, because differences between their missions lead to differences between their system characteristics. Asynchronous scanning is enhanced further in interactions involving radiolocation radars that are “3-dimensional”; those radars use pencil beams scanned in elevation as well as azimuth, whereas navigation radars for surface use (maritime and air-traffic-control) are usually “2-dimensional”; i.e. they scan only in azimuth. Eight of approximately 14 radiolocation radars described in Recommendation ITU-R M.1796 have pencil beams that scan in elevation as well as azimuth. Thus, the pencil beams of these radiolocation radars normally spend much of the timesearching regions either above the horizon, where they cannot couple strongly to the surface-based radionavigation radars or, in the case of airborne radars, at varying depression angles, so they illuminate a particular surface-based or airborne navigation radar only occasionally. The most powerful radiolocation radars are surface-based and have radiation nulls on the horizon, so they couple poorly with surface-based radionavigation radars. Further, radiolocation radars often use electronic steering and scan in patterns that are deliberately pseudo-random or in patterns that are quasi-random because they adapt to the target environment. In such cases, the main beam of the radiolocation radar revisits the direction of the navigation radar only at irregular intervals instead of periodically. This makes it unlikely that discrete-target radionavigation radars will interpret main-beam-to-main-beam interfering radar signals as a valid target. In any event, the fact that main beams of all radars are narrow causes the fraction of time during which main-beam-to-main-beam conjunctions prevail to be extremely small. Consequently, the situations that are normally of concern are limited to:

– radiolocation radar sidelobes to radionavigation radar sidelobes;

– radiolocation radar main beam to radionavigation radar sidelobes; and

– radiolocation radar sidelobes to radionavigation radar main beam.

4.1.1 Sidelobe-to-sidelobe coupling

The bulk of the sidelobes of both radiolocation and radionavigation radars have gains that are at least 30 dB below the main-beam gains. In fact, median sidelobe levels of such high-gain antennas tend to be approximately –10 dBi, so that median sidelobe suppression factors are typically about 40 dB. Maritime navigation radars operating around 10 GHz normally use slotted waveguide array antennas. Consequently, they have rather good sidelobe suppression. In addition, they have relatively narrow beams in the azimuth plane. An example of an azimuth-plane antenna pattern measured on a commercial maritime navigation radar operating in the 9.3-9.5 GHz band is presented in Fig. 1. As that figure shows, the strongest sidelobe is suppressed by about 25 dB and the median sidelobe level is at least 47 dB weaker than the main beam gain.

FIGURE 1

Azimuth-plane antenna-gain pattern of 10 GHz band maritime navigation radar

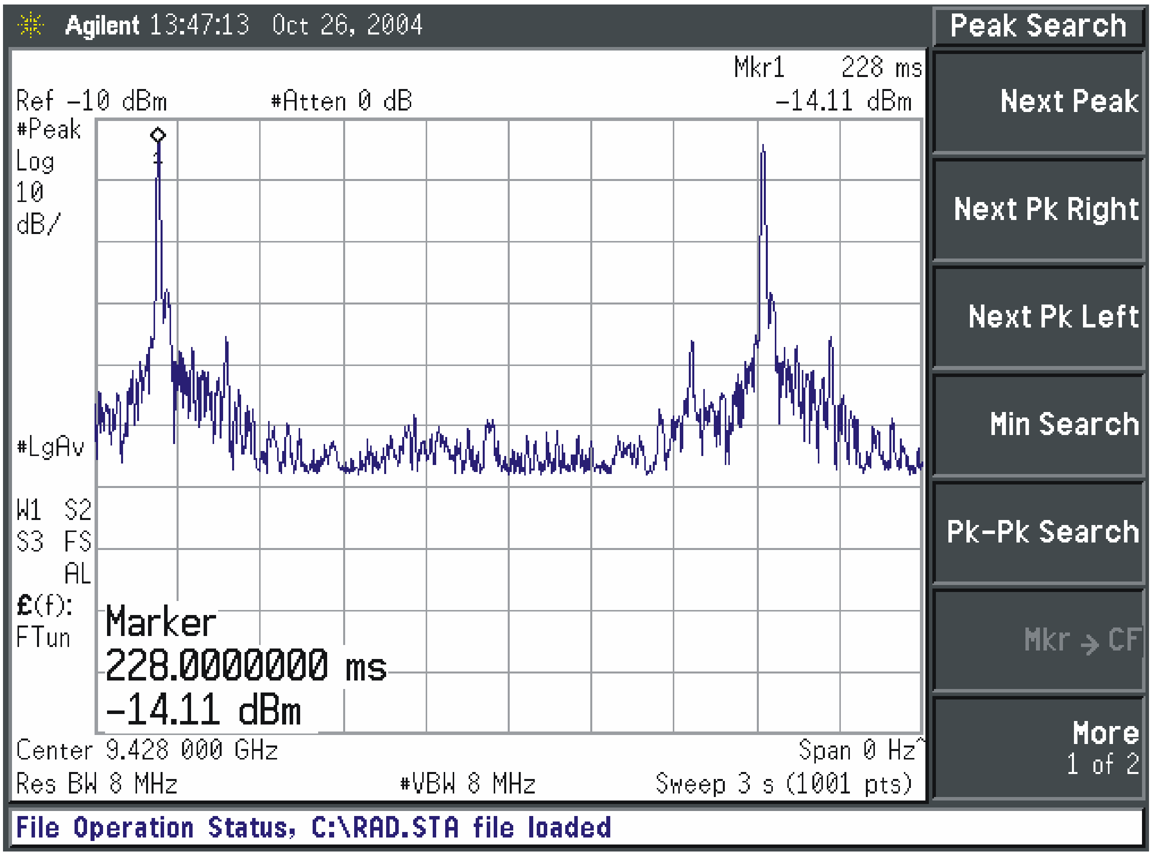

Rap 2076-01

This kind of performance is not reflected in most published sidelobe gain values, including those presented in Recommendation ITU-R M.1796, because specifications and standards usually state only the levels of the highest, close-in sidelobes. But it is readily understandable. Since an antenna can only concentrate energy and not amplify it, any gain in its main beam can only be achieved by lowering the directive gain in most other directions below the average of directive gain over all directions 4π(sr), which is necessarily 0 dBi. The stated values of main-beam gains are power gains, which account for ohmic losses; i.e. dissipation of the energy that the antenna fails to radiate. They are therefore usually several dB lower than the associated directive gains. The power gain of the entire antenna pattern over all 4π(sr) of angle is lower than the corresponding directive gain by the same factor, so the average power gain in the sidelobe region cannot possibly exceed about –3 dBi.

Good design concentrates more of the radiated energy in the main-beam region and suppresses most of the sidelobes further. Consequently, most sidelobe-to-sidelobe coupling is typically 66 to 80 dB weaker than main-beam-to-main-beam coupling.

Except when separation distances are quite short, therefore, sidelobe-to-sidelobe-coupled pulses are usually too weak to evoke false alarms.

It does happen that antennas having rectangular or quasi-rectangular apertures concentrate their sidelobe gain into ridges lying in planes that contain the longitudinal and transverse axes of the aperture, in which sidelobe gains can average higher than –10 dBi, but in those cases the sidelobes in all other planes are suppressed to values averaging less than –10 dBi. In addition, any false alarms that are evoked by sidelobe-to-sidelobe coupling will be spread randomly over a wide range of azimuth values, so they tend not to appear as targets.

4.1.2 Main-beam-to-sidelobe coupling

Apart from low-powered beacon transponders, the radiolocation radars in these bands, described in Recommendation ITU-R M.1796, typically have antenna gains ranging from about 28 to 42 dBi, with weather radars having gains as high as 46 dBi. The primary radars have narrow azimuth beamwidths, ranging from 1.5° to 5.75° at 3 dB down, with weather radars having beams as narrow as 0.9°. If their azimuth coverage is uniform over 360°, as is typically the case, their main beams will illuminate other radars no more often than 1.5/360 * 100 = 0.42% to 5.75/360 * 100 = 1.6% of the time, and as little as 0.9/360 * 100 = 0.25% for weather radars, and the many radars that scan in elevation will illuminate them via the radiolocation radar main beams much less often than that. The low values of these percentages do not assure compatibility by themselves, but they are important by virtue of several facts:

– the occasional illuminations occur at intervals differing from the radionavigation radars' scan period;
– the interference is pulse-like and asynchronous;
– any interference effect tends to take the form of false alarms.

Thus, false alarms, including apparent weather blips, inflicted via the main beam of a rotating-beam radiolocation radar will normally migrate through the apparent azimuths of the radionavigation radar, typically falling along spiral loci on the plan-position indicator (PPI) display. Unless they are extremely dense, those can be discarded visually or in track-while-scan processing algorithms. Radiolocation radars that have electronic beam steering often scan pseudo randomly since they are not constrained to scan at a uniform angular rate; when such nonuniform scanning is done in the same plane as the (usually uniform) scanning of a radionavigation radar, false alarms that might be inflicted via the radiolocation main beam will be spread randomly over a wide azimuthal sector of the radionavigation radar, thus not correlating to form false-target tracks.

4.1.3 Sidelobe blanking

An optional feature that is sometimes incorporated into radars is that of sidelobe blanking [Skolnik, 1990; Maisel, 1968]. That arrangement supplements the high-gain antenna that is typical of radars with an auxiliary low-gain antenna feeding a separate receiver having the same gain as the main receiver. Logarithmic amplifiers are provided in both channels so that the ratio of signal powers in the two channels can be conveniently derived by a subtraction network. The purpose of such blankers is to prevent detection or other processing of strong target return pulses and interference pulses via radar antenna sidelobes. That is accomplished by using an appropriate ratio of auxiliary-antenna gain to main-antenna gain and appropriate values of blanking threshold for the ratio of signals received via the two antennas. This technique cannot protect against continuous interference since if such interference were strong enough to blank the receiver it would do so most of the time, drastically degrading system effectiveness. The value of sidelobe blanking for interference suppression accrues only on low-duty-ratio interference. If this technique were used in navigation radars, it would further restrict the interactions of interest to those in which interference is received via the navigation radar’s main beam.

4.1.4 Sidelobe-to-main-beam coupling

Let us postulate that a radiolocation radar impinges unwanted energy on a radionavigation radar, creating a potential for evocation of false alarms. The most troublesome false alarms are those that are detected at nearly the same azimuth and range on successive scans of the radionavigation radar's antenna beam, since they could then be correlated, either by manual observation or automatically, to appear like a valid target. One of the necessary conditions for that to occur is that false alarms must occur consistently when the radionavigation radar’s main beam is directed toward a given bearing. This focuses attention on the case of coupling from radiolocation radar sidelobes to the radionavigation radar's main beam. That coupling can be quite strong on occasion, since the radionavigation radars in these bands typically have fairly high gains. Maritime navigation radars in this band are required to have at least 20 dB of sidelobe suppression outside of the 10° sector centred on the main beam, which provides at least 20 dB of interference suppression for 97% of the time in any direction. (Actual sidelobe suppression is often much better than that, as illustrated in Fig. 1). By itself, this infrequence will not prevent false alarms from correlating to appear as targets, since it will tend to confine them to a single narrow sector. But asynchronous pulsing will cause any false alarms to appear at essentially random ranges, often changing non-monotonically from scan to scan, which reduces the chance that they will be correlated by an automatic tracking algorithm or by visual observation.

4.2 Processor-mediated effects

Within the signal processor, the effect of unwanted signals from other radars can be influenced by processes that operate in the time scale of individual pulses (termed “fast time” in the SAR literature) and by processes that operate in the time scale of several pulses (termed “slow time” in the SAR literature). Any two radars of differing types, especially if they serve different missions such as radiolocation and radionavigation, almost invariably use different pulse repetition intervals, particularly at a given point in time. This provides opportunities for use of powerful techniques for mitigating radar-to-radar interference. Because this suppression results from the absence of synchronism between the pulses generated by the victim radar and those received from interfering radar, it amounts to isolation in time rather than in space or radio frequency. These techniques comprise several processing methods to be described below.

4.2.1 Individual-pulse processes

4.2.1.1 Fast time constant

Many maritime navigation radars provide a fast time constant (FTC), or differentiation, feature for reducing the obscuring effect of precipitation clutter. The FTC technique is applied in the video, or post-detection, circuitry but it precedes the “pulse-to-pulse correlation” or noncoherent integration processes. Because of that, it can enhance the effectiveness of those processes by constraining the undesired pulses to narrow widths and low duty ratios as they are operated on by the pulse-to-pulse correlation processes. It is standard practice to provide some form of FTC in the maritime navigation radars operating around 10 GHz as well as those operating around 2 GHz; the FTC, in conjunction with the logarithmic IF amplifier/detectors used in those radars, facilitates suppression of sea clutter. It also has the serendipitous effect of shortening longer pulses that might be received from other radars. In fact, FTC is closely related to the technique of pulse width discrimination, which is used in some radars for the express purpose of defeating active interference.

4.2.1.2 Off-tuned effects on duty ratio

Independent of FTC use, long pulses from radiolocation radars that are off-tuned from the radionavigation radars will evoke responses in the latters' IF sections that are much narrower than the pulses transmitted by the radiolocation radars. Transitions at the beginning and end of the transmitted pulse will evoke responses that resemble the radionavigation radar's impulse response, with a width of only about 1 ms or less (depending on the navigation radar's pulse width mode). During dwells between transmitted-pulse transitions (rise times, fall times, and some subpulse transitions), response levels will be low, approximating those that an unmodulated off-tuned carrier would evoke. Like the use of FTC, this effect can lower the effective duty ratio of alien pulses substantially and thereby greatly potentiate the effectiveness of "pulse-to-pulse correlation" (to be described) in lowering the likelihood of false-target detections.

4.2.2 Multiple-pulse integration techniques

In general, these techniques include both "pre-detection" or coherent integration and "post-detection" or noncoherent integration. As used in this context, "detection" refers only to the process that extracts the waveform envelope and discards its carrier, not to the process of comparing a signal level against a threshold to determine whether a target is present. Coherent integration is normally used to implement Doppler processing and to maximize overall radar sensitivity. Precision-approach radars may use coherent integration. In contrast to coherent integration, noncoherent integration operates on only the magnitudes of the received pulses after their phase information has been discarded by an "envelope detector". Whether they perform coherent integration or not, most navigation radars perform some kind of noncoherent integration. A variety of post-detection integration techniques exist; they are surveyed in Skolnik's Radar Handbook [Trunk, 1990].

4.2.2.1 Linear integration

Besides the distinction between coherent and noncoherent integration, an important distinction exists between linear integration and nonlinear, or quantized, integration. Linear integration weights received pulse trains not only by the number of received pulses in individual range/angle or range/angle/Doppler cells but also by the amplitude of each pulse. Because linear integration retains the amplitude weighting of each pulse, a strong pulse contributes proportionately more to the integrator output than a weak pulse does, so it permits isolated interference pulses to produce relatively strong outputs if they are sufficiently powerful. However, the weighting by the number of pulses in a given range/angle/Doppler cell discriminates against asynchronous pulsed interference to a degree, which is especially valuable when the interference pulses are relatively weak.

4.2.2.1.1 Multiple-pulse interference mitigating characteristics peculiar to synthetic-aperture radars

In addition to the processing-gain advantage that accrues on individual pulses, an additional processing gain occurs in SARs due to integration of the many pulses that form the synthetic aperture. The desired signal power is raised by a factor equal to the square of the number of pulses, N, integrated during the synthetic aperture time, which is typically very long. In response to asynchronous pulsed interference, however, the azimuth processing gain will be close to unity.

Typically, the overall (range and azimuth) processing gain for low-duty-ratio asynchronously-pulsed interference will be no more than a few dB. In contrast, continuous noise (or noise like-interference) will experience an azimuth processing gain equal to N.

4.2.2.2 Rejection of asynchronous pulse interference by means of Binary integration

In contrast to linear integration, Binary (nonlinear), also discussed in Recommendation ITU-R M.1372, integration discards varying amounts of pulse amplitude information; in the extreme case, and every pulse is weighted equally. Nonlinear integration thus tends to equalize the weighting of individual valid target-return pulses and strong interference pulses, so it discriminates against isolated asynchronous interference pulses even if they are very strong. Coherent integration is normally linear, while noncoherent integration can be either linear or nonlinear.

“Binary integration”, “sequential-detection” or “double-threshold detection” is a noncoherent and nonlinear process. In this case, “detection” refers to the output of a threshold comparator that is itself downstream from an “envelope detector”. Sequential detection combines threshold-comparator outputs or “first detections” in each range/angle cell during individual pulse repetition intervals (PRIs) or “sweeps”. These processes are often referred to as either integrators or correlators, although they are seldom true integrators or correlators in the strict mathematical sense. The individual detections are either limited or quantized to simple binary quantities (zero or one). The various designs make differing tradeoffs among target-detection or tracking sensitivity, accuracy of target azimuth estimation (“centroiding”), and suppression of detections evoked by asynchronous pulses. In some cases, the operator has some latitude to adjust the tradeoffs by adjusting an operating setting. The asynchronous-pulse-rejection characteristics of such processes has been summarized in Recommendation ITU-R M.1372.

Double-threshold integrator/detectors of the binary type are of particular interest because they are especially powerful in discriminating against target declarations caused by asynchronous pulses. Two types of binary double-threshold integrator/detectors can be distinguished:

– sliding window, *M*-out-of-*N*;

– up-down counter with arbitrary counting rules. This is sometimes called an accumulator, binary integrator, or exponential integrator.

Both of these techniques are sometimes referred to as Markov processes. They are addressed in Recommendation ITU-R M.1372.

4.2.3 Asynchronous pulse removal and replacement techniques

The possibility of removing isolated asynchronous pulses has been recognized for many decades, preceding the advent of digital signal processing. Early versions (as well as simpler versions used today) operated only on amplitude, without using phase information. The simplest form is basically a 2-out-of-2 binary integrator, as shown in Fig. 2. Such circuitry is effective on “normal” video; i.e. in the absence of MTI cancellers.

FIGURE 2

Simple video asynchronous-pulse remover

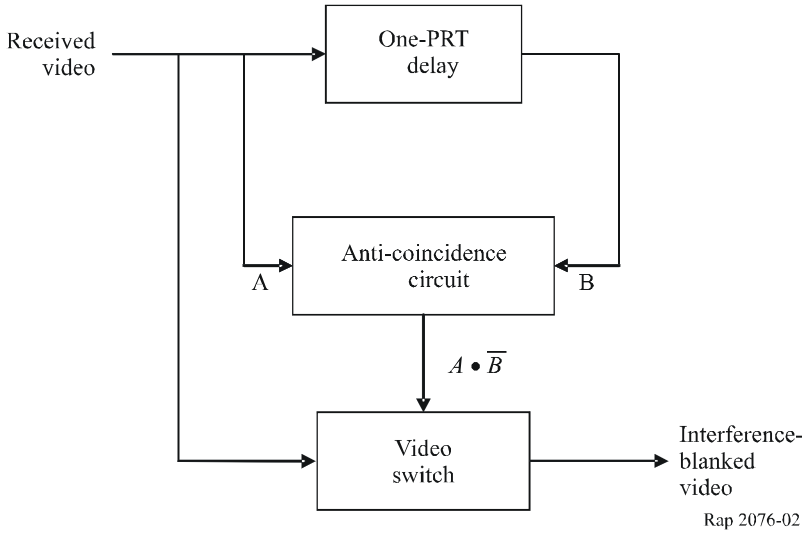

When MTI processing is performed, the problem arises that each isolated asynchronous pulse generates several synchronous pulses, with more synchronous pulses being generated when recursive or feedback cancellers are used than when only feed-forward cancellers are used. Even in the absence of MTI processing, feedback integrators can create the same problem. Two remedies naturally suggest themselves:

– Remove the asynchronous pulses before the signal reaches the MTI canceller.

– Use only feed-forward, or finite-impulse-response cancellers and integrators.

With the advances that have been made in digital signal processing capabilities, it has become feasible for radars to identify individual pulse samples as interference by virtue of being inconsistent, in either magnitude (i.e. power) or phase angle, with the received pulses samples that precede and/or follow them. This can be done prior to coherent processing such as Doppler filtering or autocorrelation. These processes use a sliding window of successive samples in the same range bin but from different PRIs to provide an approximation of what the sample value should be in the PRI under test. That window might span 8 to 16 or so PRIs or it might span just a couple of PRIs, with the accuracy in estimating the valid return signal better with the longer background windows. The processes compare the magnitude or magnitude squared ($I^2 + Q^2$), or the individual I and Q sample values, for the PRI under test (the "current" PRI) with their counterparts in the background window. Thus they can operate on either the pre-detection (coherent) or post-detection (noncoherent) forms of the received pulses. When the differences are sufficiently large, the current sample value is replaced by a value derived from the background values. Since the processing is performed on I and Q "coherent video" samples, it is applicable without regard to the operating band. Techniques of this kind are used in some radars in both the military and civil fields. Since these processes are necessarily performed prior to multiple-pulse integration, the estimates formed from the background samples, as well as the tested sample, are subject to noise variations when they are near the threshold of detectability, so weak interference pulses tend to be missed. On the other hand, the efficiency of interference removal improves progressively with increase of the interference-to-noise ratio, so removal of strong interference becomes almost categorical. This property makes such processes excellent complements to the suppression of weak asynchronous interference that is contributed by multiple-pulse linear integration such as Doppler filtering and post-detection (noncoherent) integration.

One maritime navigation radar that can operate in the 9.3-9.5 GHz band contains a post-detection process similar to the one just described. For each range cell, it establishes a sliding window spanning only 3 PRIs. If one of the pulses in the window is disproportionately strong relative to the other two pulses, a weighted average of the other two is substituted for the outlier pulse. The radar operator can switch that interference-rejection process on or off. Tests performed on that radar included application of a train of 10 ms pulses with a 1% duty factor, without gating to simulate antenna scanning modulation. Figure 3 shows how the radar's PPI appeared with the interference-rejection feature turned off. Each interference pulse produced a long radial line corresponding to the 10 ms interference pulse width.

FIGURE 3

Pulsed interference on maritime navigation radar PPI with interference-rejection feature off

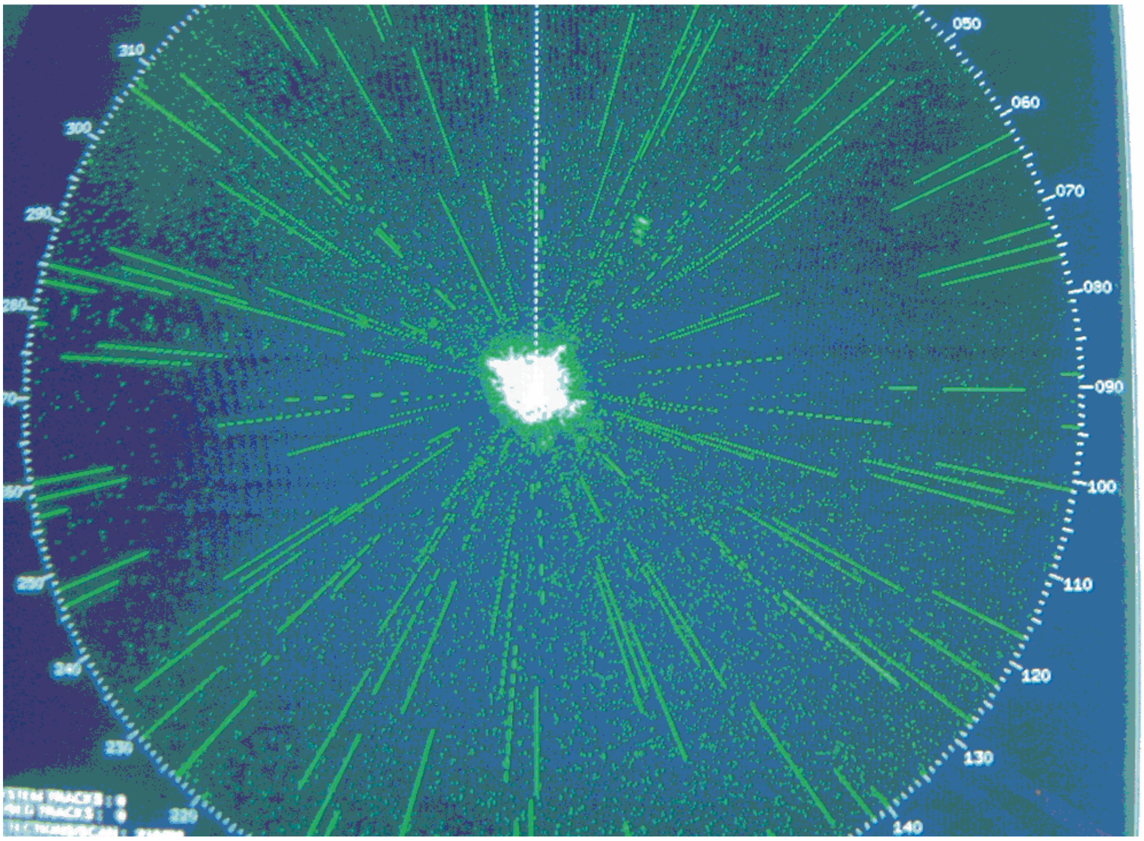

Rap 2076-03

Figure 4 shows how the radar's PPI appeared with the interference-rejection feature turned on. The interference blips are completely removed. Apart from some scattered speckle, the only blips on the PPI are return from fixed clutter near the test site plus a line of deliberately-applied simulated target returns at a bearing of about 280°.

FIGURE 4

Pulsed interference on maritime navigation radar PPI with interference-rejection feature on

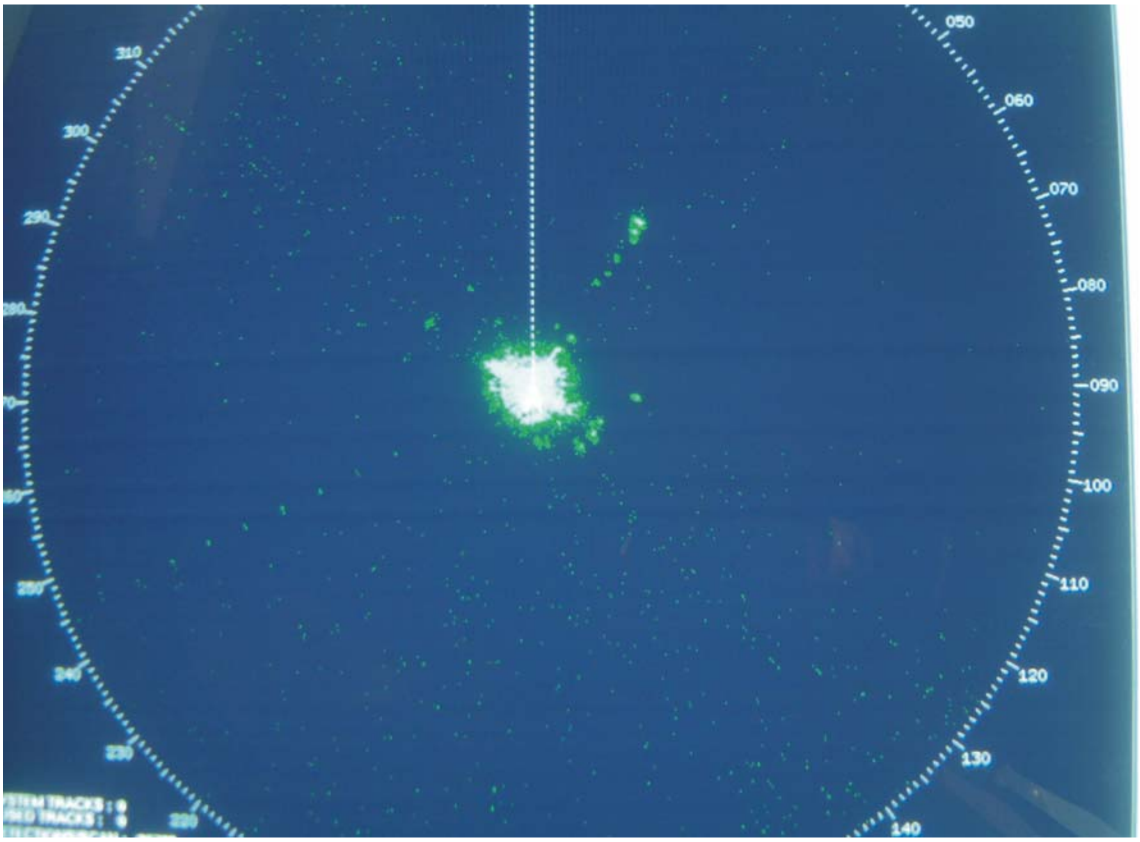

Rap 2076-04

4.2.4 Asynchronous pulse rejection requirements and practices

An IMO Resolution states that maritime navigation radars should provide means for the adequate reduction of interference from other radars (as well as unwanted echoes from various forms of clutter) (Section 5.3.2.1 of IMO Resolution MSC.192(79)). That exactly mirrors recommendations that appear in IEC standards for both SOLAS radars and non-SOLAS radars (IEC 60936 and IEC 62252). Although all the aforementioned factors contribute to such interference reduction even if only as a side effect, maritime navigation radars typically provide a feature that is explicitly intended to address this issue. They perform what is known in the community as “pulse-to-pulse correlation”, “sweep-to-sweep correlation”, or “line-to-line correlation”[2]. That feature normally uses a sliding window spanning *N* successive PRIs or “sweeps” and applies an *M*-out-of-*N* criterion for the 2nd detection.

A signal processor that is marketed by an industry leader for use in ground-based weather radars, some of which can operate in the 9.3-9.5 GHz band, contains a pre-detection process similar to the post-detection process used in the maritime navigation radar. For each range cell, it establishes a sliding window spanning only three PRIs. If certain conditions relating the powers of the three pulses are satisfied, the phasor value of the most recent pulse is replaced by that of the one immediately preceding it. The designer or user of the radar can select among several different conditions for pulse replacement, but the basic principle is that interference that visits a given range cell only once in three PRIs can be replaced by a reasonable value. Because the process spans only

---

[2] In Recommendation ITU-R M.1372, the 2-out-of-2 sliding-window process is referred to as a “PRF discriminator”.

three PRIs, it is capable of replacing interference having a duty ratio as high as $33^{1/3}$%, but that could be achieved only when a very special relationship prevails between the PRF of the interference pulses and that of the host radar. As the interference duty ratio becomes lower, the tolerance of interferer/victim PRF relationships becomes less particular and more robust.

Airborne weather-avoidance radars are also likely to incorporate asynchronous pulse-rejection processes. A Radio Technical Commission for Aeronautics (RTCA) standard contains a requirement for such radars to suppress radar-to-radar interference. It states: "Transmission from an identical-type radar, operating on an aircraft flying a parallel approach to an adjacent runway or following the equipped aircraft as closely as two nautical miles, shall not cause false alerts, missed detections or other observable interference"[3].

4.3 Nonlinear and time-varying gain effects

4.3.1 Limiting

Limiting can be performed at various points in the receiver/processor. It can be implemented by various means and can be either deliberate or inadvertent. Inadvertent limiting occurs if receiver RF and/or IF circuitry is driven beyond its linear range. In radars that use digital signal processing, this circuitry includes the A/D converter.

Regardless of its particular implementation, limiting clearly tends to equalize the amplitudes of undesired pulses and valid return pulses when either or both are strong. This enhances the capacity of the other processes mentioned herein to discriminate against asynchronous pulses that might be received from radiolocation radars.

Limiters are sometimes included in the receiver/processor chain in attempts to prevent detections on strong clutter. That was formerly common practice when moving-target-indicator (MTI) cancellers were used without adequately effective cell-averaging CFAR processes or clutter maps, the intent being for the operator to adjust the limit level so as to prevent almost all false alarms due to clutter residue. Such use of limiting is tending to decline as high dynamic range A/D converters and digital processing become more available.

As indicated above, processes in which amplitudes are retained tend to maximize the detection sensitivity, but those in which amplitudes are limited or quantized tend to suffer varying loss of sensitivity while providing strong discrimination against asynchronous pulses. When a limiter is provided, the operator can sometimes vary the tradeoffs from one that optimizes sensitivity to one that protects against interference-caused false alarms by lowering the limit level; i.e. by limiting more heavily. This is especially significant when a feedback integrator is used, since feedback processes breed synchronous pulses from isolated asynchronous pulses and therefore tend to respond poorly to asynchronous pulses.

Hard limiting in a wide bandwidth (Dickie Fix) prior to filtering to a narrower matched-filter bandwidth is a technique that has been used in the past as a cheap form of CFAR. It is also useful for attenuating the effect of narrow interference impulses and is sometimes used as an electronic counter-counter mesure technique against swept FM noise jamming. However, it is risky if it is not optional, since it permits strong and moderately strong interfering signals, even if off-tuned from the desired signal, to inflict small-signal suppression. Binary integration, discussed in § 4.2.2.2 can be regarded as a form of limiting, but it is implemented very far downstream in the receiver/processor and consequently causes no small-signal suppression.

---

3 Minimum operational performance standards for airborne weather radar with forward-looking windshear capability, Document No. RTCA/DO-220, September 21, 1993, RTCA, Inc.

4.3.2 Logarithmic amplifiers

Since the mid 1970s, almost all maritime navigation radars have used logarithmic amplifiers. Use of logarithmic amplifiers has a similar effect to that of hard limiting, though not as complete, in reducing the effect of high individual-pulse amplitude as a contributor to performance degradation. Certainly these amplifiers lessen the likelihood that strong interference pulses might saturate the receiver. Further, these amplifiers are inevitably linear or quasi-linear at low signal levels, but to achieve greatest effectiveness when used in conjunction with ac-coupling (or FTC; see § 4.2.1.1) in mitigating sea-clutter return, the transition between the linear region and the logarithmic region occurs at about 20 dB below the average noise level [Croney, 1956].

4.3.3 Sensitivity time control (STC)

STC is a form of time-varying gain. It is a form of deliberate desensitization that varies within each PRI or “sweep”. The receiver or processor is desensitized only at times corresponding to returns from short-range targets, since those targets produce such strong returns that full receiver sensitivity is not needed to detect them. STC often is such that, at ranges less than a selected value, the detection threshold for radar return, referred to the antenna port, varies at a rate that compensates approximately for the inverse-4th-power relationship between return power and target range, for a given target radar cross section. (In actuality, some receivers typically produce outputs proportional to the logarithm of the received signal's amplitude or power. The STC circuitry, implemented in the video section, weights those responses with a gain function derived from an exponential decay). STC helps to suppress clutter return, which is normally stronger for short-range clutter than for longer-range clutter and which might otherwise exceed the receiver/processor's linear dynamic range. It also tends to suppress detections due to “angels” (usually caused by reflections from birds) at short range.

Of course, STC helps to reduce the number and amplitudes of detections that might be evoked by radiolocation-radar pulses as well, to the extent that the detections would appear as short-range false targets. In many situations, it is more important to eliminate short-range false-targets than to eliminate false targets that appear at longer ranges, because short-range targets typically require more urgent action to avoid collisions than do long-range targets. Maritime navigation radars use STC, implemented in the video circuitry, as a means of suppressing sea clutter return. Air-traffic-control radars in this band have also had STC capability.

4.4 CFAR processing

Discrete-target surveillance radars need some mechanism to maintain the rate of detections within reasonable bounds. Since most detections are usually evoked by noise or clutter return and hence are false alarms, such mechanisms are referred to as CFAR processes. In the simplest system, this could take the form of a manual control of gain or detection threshold. Typically, however, the regulation is automated. That can be done on a rather global basis; i.e. affecting wide angular sectors and a large fraction of all ranges, or it can be done by means of locally adaptive thresholding. In the former category, techniques might include automatic gain control, which averages over all or most ranges, Dickie Fix, STC, and/or the combination of a logarithmic amplifier and FTC or pulse differentiation. Locally adaptive forms of CFAR include clutter mapping and especially local-average-and-threshold circuits. The clutter map is a matrix of signal levels, averaged over many antenna scans, for each of many small range/azimuth cells. Local-average-and-threshold circuits, or cell-averaging CFAR circuits, provide a detection threshold that adapts to the clutter (and interference) level in the immediate vicinity of each range/Doppler/azimuth cell that is being tested for target presence. Local-average-and-threshold CFAR processes operate by constructing a window that slides out in range (for each Doppler channel, in the case of Doppler radars) during each PRI. Each such window straddles the range cell for which a detection decision is to be made plus roughly 10 to 30 adjacent range cells (usually half of them at shorter range and half at longer range). Typically, the signal amplitudes in those adjacent cells are averaged and the average value is multiplied by a factor such as 4 or 8 to establish the local detection threshold. However, numerous variations on that basic design scheme are often used. For example, the cells prior to the tested cell might be averaged separately from those beyond the tested cell and the greater of the two average values might be used in setting the threshold.

In cell-averaging CFAR processes, individual cells that contain the strongest signals among the adjacent range cells are usually excluded from the averaging. This is probably done in part to prevent residue of return from point clutter scatterers from raising the detection threshold level unnecessarily. However, it also prevents isolated asynchronous pulses from contaminating the threshold value and producing inappropriately elevated threshold levels. This does not contribute to the reduction of false targets evoked by undesired pulses, but it does mitigate any tendency of undesired pulses to desensitize the victim radar.

4.5 Software-mediated effects (post-processing)

Range/azimuth clustering of target detections can also be examined automatically in various algorithms to decide whether a given cluster of detections represents a valid target or not. (Such processes also discriminate against “angels”, which are mostly reflections from birds.)

Any “track-while-scan” processing, which associates target reports from successive antenna-beam scans and estimates the targets' positions and vector velocities, also has the potential for censoring out false targets. Such post-processing is standard in air-traffic-control radars, and many maritime navigation radars also have a scan-to-scan-correlation feature for suppressing sea clutter. This feature tends to be associated with use of raster-scan displays, as distinguished from the traditional radial-scan displays. The original advantage of raster-scan displays was that they are much brighter than radial-scan displays because their screens are written numerous times, instead of only once, during each antenna scan. Their implementation requires that the data on target detections from all the range-azimuth cells in at least one complete antenna scan must be held in memory before being displayed. Great advances in digital memory circuitry in recent years have made use of raster-scan displays economically practical. The memory capability also permits storing several scans of radar-return data for all individual-scan detections and comparing the detection data in each range-azimuth cell on one antenna scan with the detection data in the same range-azimuth cell in one or two previous scans before displaying them. If returns are categorized in a binary fashion (present or absent) or in terms of a few levels, rules can be applied to derive any appropriate brightness level to be displayed at each range-azimuth pixel depending on the combination of return levels in the separate scans. Such a feature will intensify pixels for which returns appear consistently on two or more scans. More importantly for purposes of compatibility between radiolocation radars and radionavigation radars, it will dim or blank pixels in which an apparent return appears during one scan but not during the next one or two scans. Even in navigation radars that lack this feature, the operator can infer whether responses on the display represent valid target returns or not based on visual observation of the consistency of pixel illumination. When interference from another radar arrives via sidelobe-to-main-beam coupling, detections can be evoked recurrently on the same bearing. However, asynchronous pulsing of the two radars will tend to prevent those detections from recurring in the same range cell. Multiple-scan processing will therefore tend to prevent such detections from being displayed.

4.6 Spectrally-mediated effects

Among frequency-related phenomena, receiver selectivity and spurious-response suppression are factors in rejecting radar-to-radar interference just as they are in any other interference interaction. Receiver selectivity can be expected to suppress spectral components that are sufficiently outside the pass band of the victim receiver by at least 60 dB. The full benefit of the radionavigation-radar's stop band suppression accrues only if unwanted components of the emission spectra of radiolocation and spaceborne-sensor radars are comparably suppressed. In the 9 000-9 200 and 9 300-9 500 MHz bands, most of the radiolocation radars would use transmitters other than crossed-field devices, so the noise components of unwanted emissions would be relatively low.

Likewise, all of the spaceborne-sensor radars that are proposed to operate partly in the 9 300-9 500 MHz or 9 800-10 000 MHz bands would also use other than crossed-field transmitter devices. The levels of unwanted emissions from radiolocation and spaceborne–sensor radars in the bands of concern would be determined more by rise-fall-, and chip-transition ramps, with associated incidental angle modulation, than device noise. The spectral efficiency of several radar transmitter output devices is described in Recommendation ITU-R M.1314 – Reduction of unwanted emissions of radar systems operating above 400 MHz.

In addition, two-signal intermodulation tends to be less significant with regard to radar-to-radar interactions than it is with respect to interference between communications systems, because radar transmissions usually have comparatively low waveform duty ratios. Temporal conjunctions between pulses from two radars tend to be rare because each radar is unlikely to emit a pulse that arrives at a victim simultaneously with a pulse from the other radar.

Some radiolocation radars transmit long pulses with duty ratios that are low relative to those of communications systems but high relative to those of radionavigation radars. However, the longest pulses are usually modulated with swept frequency, or chirp, waveforms to support pulse compression in the radiolocation radar receiver, and the frequency sweep in such long pulses is usually much wider than the pass bands of radionavigation radars. Even if the frequency sweep were to span a radionavigation radar's pass band fully, substantial frequency-dependent rejection would prevail while the instantaneous frequency is being swept through points below and above that pass band. This effect has been confirmed empirically on a maritime navigation radar operating in the 9.3-9.5 GHz band, and the results are being reported on separately. This dilution of on-tuned energy can be viewed in several different ways. Clearly, the average power spectral density is lessened by the frequency spread. A more useful way to view the dilution effect is to recognize that the effective pulse width; i.e. the width of the pulses that emerge from the radionavigation radar’s IF section, is often much smaller than the radiolocation radar's transmitted pulse width. That pulsewidth reduction can be ensured by good design of radiolocation and EES/SR radars. In this way, truncation in frequency by the radionavigation-system's receiver translates into truncation in time. If the frequency sweep rate exceeds the square of the victim radar’s IF pass band width, the effective pulse width (at the output of the IF amplifier) is fixed by its impulse-response duration, but the amplitude of the received pulse is attenuated relative to the amplitude of the pulse captured by the victim's antenna. In that case, the peak response power becomes inversely proportional to the sweep rate. This of course also contributes mitigation.

Frequency diversity and frequency agility are used by many radars, especially aeronautical radionavigation radars. Use of frequency agility lessens the risk that emissions from one radar will consistently overlap the acceptance pass band of another radar.

5 Conclusions

Mutual compatibility between a radiolocation radar and a radionavigation radar is fostered first of all by the scanning of their antenna beams, so that undesired energy is seldom received via either main beam-to-sidelobe coupling and even more seldom via main-beam-to-main-beam coupling.

More importantly, differences between radiolocation-radar and radionavigation-radar scanning rates prevent recurrences of interference via the source’s main beam at the same victim bearing, mitigating any risk of scan-to-scan correlation among interference false alarms that might occur. Much additional mitigation is afforded by differences between the waveforms of the two types of radars and the associated rejection of undesired pulses via receiver filtering and signal processing. In some radars, the latter includes coherent (pre-detection) processing as well as noncoherent (post-detection) processing. In other navigation radars, coherent processing is largely or completely absent but noncoherent processing is potentially quite effective in enhancing compatibility between radiolocation and radionavigation radars.

Few if any of these mechanisms apply to interference from non-scanning, continuous-wave communications transmitters using noise-like waveforms to radars of any type.

References

CRONEY, J. [April 1956] Clutter on radar displays. *Wireless Engineer*, Vol. 33, p. 83-96.

MAISEL, L. [March 1968] Performance of sidelobe blanking systems. *IEEE Trans. on Aerospace and Electron. Systems*, Vol. AES-4, **2**, p. 174-180.

SKOLNIK, M.I. (Ed.) [1990] *Radar Handbook*, McGraw-Hill.

TRUNK, G.V. [1990] *Automatic detection, tracking and sensor integration*. Chapter 8 in M.I/ Skolnik (Ed.) *Radar Handbook*, 2nd edition, McGraw-Hill.