Recognising and Mitigating the Effects of Sidelobes

How to recognise if the radar video picture is being affected by sidelobe returns entering the radar video and what to do alleviate them or minimise their impact.

Summary

The ideal radar antenna would concentrate all of the RF energy into a single directional main beam. However, all radar antennas naturally have "sidelobes" – additional beams, away from the main beam, which radiate and receive some fraction of the radar's energy.

Normally, a target return will only show up in the radar video when the antenna is pointing at the target, to within about one width of the main lobe. However, if the target is a particularly strong reflector, its reflections may "leak" in via the sidelobes.

The presence of radar video returns due to sidelobes can result in a heavily-cluttered and confusing radar picture and also, potentially, false targets.

Introduction

Sidelobes returns are a natural consequence of the radar's antenna radiation pattern. The perfect radar antenna would have a single, narrow main beam/lobe, as shown below in Figure 1, with all of the radar's transmit/receive power concentrated within this one beam. Real radar antennas, however, have multiple, lower-level lobes around the full 360-degree coverage, for example as depicted in Figure 2.

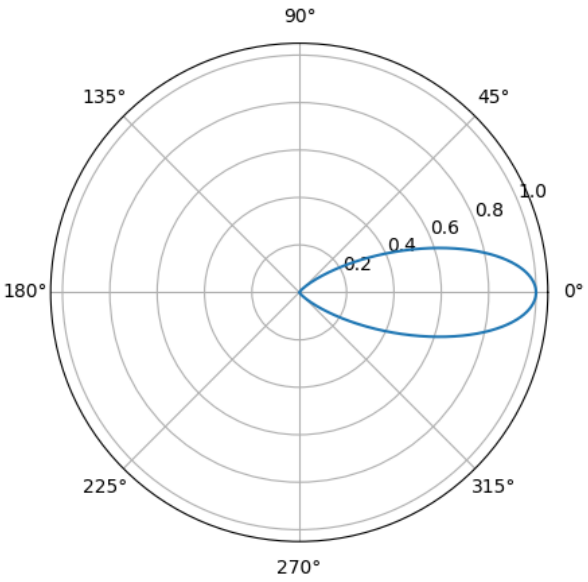

Figure 1: Sketch of "ideal" radar antenna pattern, showing single main lobe. Radial axis is plotted as power relative to boresight (0°).

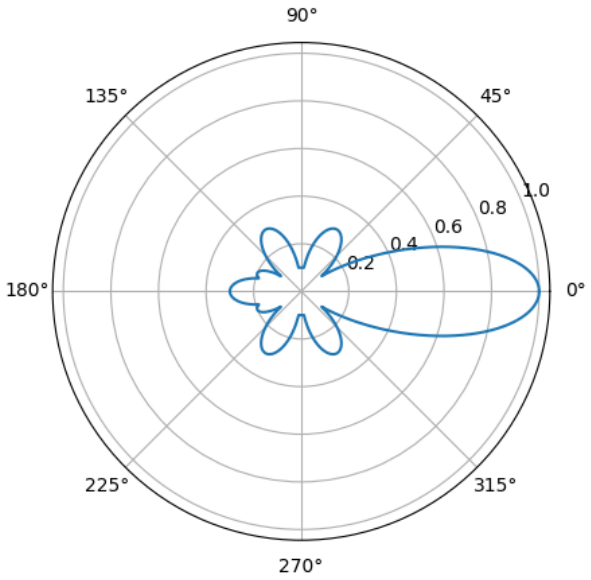

Figure 2: Sketch of real radar antenna pattern, showing (exaggerated) sidelobes.

Generally, the power transmitted/received in these sidelobes is far lower than the power within the main lobe, usually by several orders of magnitude. Therefore, unless the target is a particularly strong reflector and/or very close to the radar, any returns in the sidelobes will be extremely small and not visible within the radar video output. However, when returns in the sidelobes do make it through to the radar's video output they can cause problems.

Recognising Sidelobe Returns

The characteristic video from sidelobes looks like an extended arc of targets. The video returns are all at the same (or very similar) range as the originating target but offset in azimuth, in either direction from the real target position. In the most extreme case, the sidelobe returns may coalesce into a continuous ring of radar video.

The amplitude of the video will often be stronger as the main lobe sweeps past the target. Also, depending on the radar processing, the main lobe return may have a greater range extent in the video and therefore look bigger. However, in some cases the return from the main lobe may be entirely indistinguishable from the sidelobe returns.

It can be extremely difficult to discriminate between sidelobe video returns and genuine target returns from other physical targets that happen to be at the same range. Attempting to remove sidelobe video therefore carries the potential risk of removing genuine targets of interest.

The image below shows green radar video data which includes returns from two very large targets (marked). There are also a number of potential sidelobe returns shown. Only by knowing the "ground truth" or by observing the behaviour of the video over time can the presence of sidelobes be determined for sure.

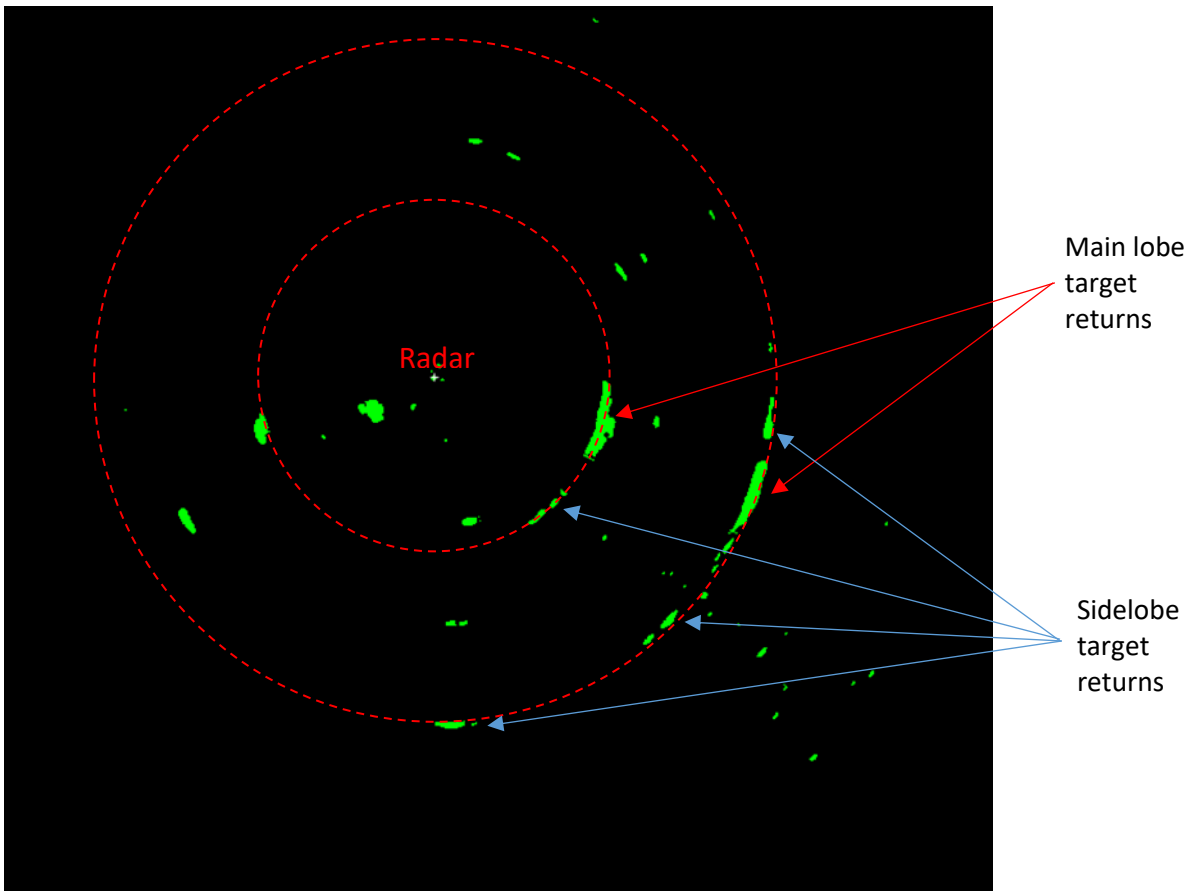

Figure 3: Example radar video picture showing suspected sidelobe returns (red dashed rings added for clarity)

Handling Sidelobe Effects

The single best way to mitigate sidelobes is for the radar itself to supress/handle them so that the effects never make it through to the video stage.

Some military and air surveillance radar systems have an additional, isotropic antenna which receives radar returns from all directions but at a power level between the main lobe and the sidelobes. If the signal amplitude received by the isotropic antenna exceeds that of the main radar antenna signal, then it is suppressed because it must be a sidelobe return.

Other radars, including small maritime models from Simrad, may provide active controls to suppress sidelobe returns using adaptive methods. Where such controls are provided, they should be enabled, at least at a low level, and certainly when operating close to buildings, shorelines or other large reflectors.

Decreasing transmission power (for example, by reducing the operating range of the radar) may also help, because fewer sidelobe returns will fall within the radar's sensitivity limits.

Where sidelobe returns are an unavoidable presence in the radar video, SPx Server includes a sidelobe suppression mode within the tracking stage. Full details are provided in the "SPx Server Tracker Configuration Guide", but essentially this processing acts to stop new tracks being initiated from plots if they fall within some configurable azimuthal offset of a larger one at similar range. This assumes that: a) the main lobe return is larger than the sidelobe ones and b) the returns are fragmented in azimuth, rather than being a single continuous arc.

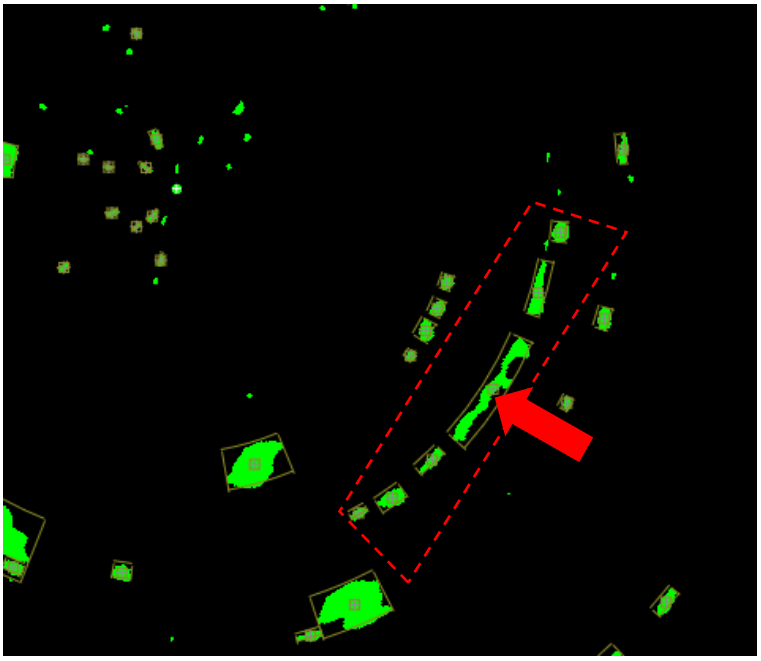

Figure 4: An example where SPx Server's sidelobe suppression mode would work well. The main target (pointed at) is clearly larger than the individual sidelobe targets either side, within the dashed red line.

In some cases, where the sidelobes do cause the radar video to appear as a single extended arc, it may not affect the tracking (at least in the short term). Assuming there is some symmetry in the video, such that the main lobe return is evenly flanked by sidelobe returns, then the centroid position value may fall close to the actual target position. In this case, the tracker will be able to maintain the track. However, this would require the plot extraction stage to allow larger azimuth sizes than may be desirable.

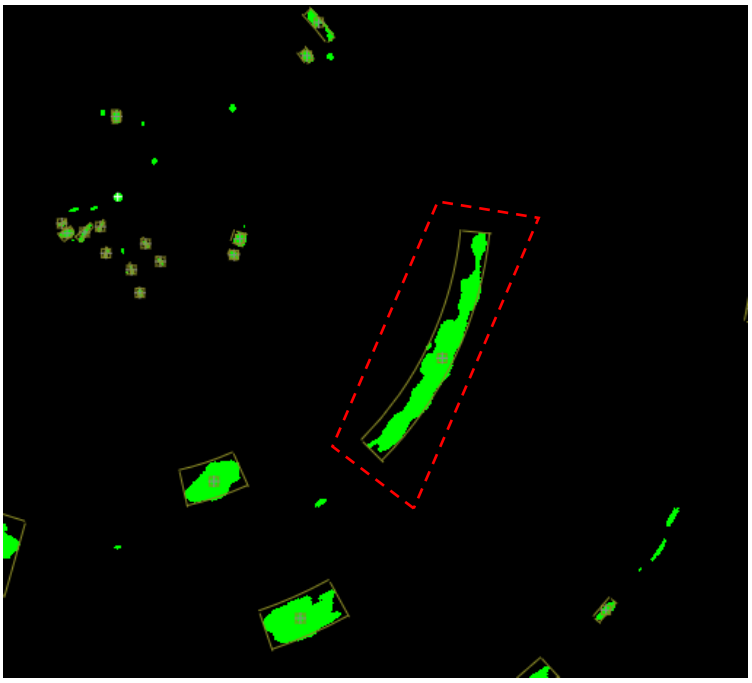

Figure 5: An example of sidelobe returns coalescing together with the main target return to form a single, extended target return.

Sidelobe returns may be less persistent than the main lobe ones. The inherently lower power means that they may be marginally detected by the radar and as the radar-to-target geometry changes slightly they fall below the radar's sensitivity limits. In which case, simply increasing the tracker's initiation time may be sufficient to mitigate their effects.

< End of Document >