computers are bad

https://computer.rip - me@computer.rip - fax: +1 (505) 926-5492

2022-12-04 over the horizon radar pt II

Previously on Deep Space Nine, we discussed the MUSIC and MADRE over-the-horizon radar (OTH) programs. MUSIC and especially MADRE validated the concept of an HF radar using ionospheric (often called "skywave" in the radio art) propagation, with a novel hybrid digital-analog computerized signal processing system. MADRE was a tremendous success, ultimately able to detect ICBM launches, aircraft, and ship traffic in the North Atlantic region. What was needed next seemed simple: a very similar radar, perhaps more powerful, aimed at the Soviet Union.

In 1964, final planning began on the full-scale OTH radar for missile defense. Code-named "Cobra Mist," the AN/FPS-95 [1] radar was initially slated for Turkey. Design proceeded for some time (a year or so) for a site in Turkey, but ultimately no site could be obtained. The documents I have are somewhat vague on the reason for this, but it seems likely that Turkey was hesitant to host a installation that would seem a direct affront to the USSR's nuclear power. Another host would have to be found, some time already into the design process. Finally, a site was selected that was not quite ideal but workable: Orford Ness, England.

Orford Ness (alternately, according to the typical complexity of English place names, Orfordness) is a sort of long, sandy semi-island, of a type prone to forming on the coast of Great Britain but not seen much here in the western United States. Orford Ness has been under the control of the Ministry of Defence since before WWII, and by the '60s was well in use as a research range for the Atomic Weapons Establishment (AWE). As is often the case with odd bits of land in military use, it contains a number of interesting historic sites which include a very early marine radionavigation system, a lighthouse notable for its involvement in the Rendlesham Forest UFO incident which occurred nearby, and several curious instrumentation bunkers left from AWE high-explosives research.

It also contains a lot of land, and that land would be needed for Cobra Mist. Construction proceeded from 1966 to 1972. Due to the experimental nature of OTH radar and the large scale of the system, during construction a set of experiments was designed under the name Design Verification System Testing (DVST). This was essentially field acceptance testing, but with the added complication that Cobra Mist was advancing the state of the art in OTH radar so much that the expected performance of the system was largely unknown. Cobra Mist fell into sort of an uncomfortable in-between: In part a production defense system, in part an experimental apparatus.

The Cobra Mist antenna, constructed by RCA (then a major defense contractor), consisted of 18 strings 620 meters long arranged in a circle, radiating from the center. A buried mesh ground plane extended forward from the antenna strings to provide vertical shaping of the beam, for a total built length on each string of about 900m.

A complex switching network connected six of these strings at a time to the transmitter and receiver, applying phase shifts to four of the strings (those not in the center of the active array) to maintain the phase of the emitted signal despite the varying distances of the strings from the target area (due to the arc shape of the array). This is an early version of the phased-array technology which is heavily used in military radar today... but a very early version indeed. The radar could be "steered" or aimed by selecting different sets of strings, but this was a fairly lengthy process and was done while offline. Nonetheless, it allowed the radar to cover a total azimuth of about 90 degrees with a smaller target area, about seven degrees, selected from within that range.

The altitude of the antenna was also somewhat steerable: each string contained two sets of radiating elements in phase relationships that would produce two different vertical angles. Further, varying the transmit frequency across its 6 to 40 MHz range resulted in different ranges as the ionospheric propagation shifted closer and further.

The minimum effective range was 500nmi, because the antenna could only emit radiation upwards (being located on the ground) and the maximum elevation angle produced a reflection from the ionosphere centered at about that distance. The maximum range is a somewhat more theoretical matter, but was placed at about 2000 nmi assuming simple single-hop propagation at the antenna's lowest achievable radiating angle. In practice, the vagaries of ionospheric propagation would make this range overoptimistic in some cases but an underestimate in favorable conditions where multi-hop or ducting patterns occur.

Behind the antenna, a large squat building housed a transmitter specified for 600 kW average power and 10 MW peak, although a report on the project from the Mitre Corporation (from which much of the information here is derived [2]) states that only 3.5 MW was achieved in practice. A specially designed receiver with a very large dynamic range (140dB specified) was connected to a set of RF filters and then an analog-to-digital converter system that provided the received signal to the computer.

Computers are ostensibly what I write about here, and this story has a good one. The Signal Analysis Subsystem (SAS) was a hybrid digital-analog computer driven by a Xerox Sigma V. The Sigma V was actually a fairly low-end computer, 32-bit but without the virtual memory support found on higher-end Sigma models. The inexpensive (relatively) computer was enabled by the fact that most actual signal processing was performed in the SAS. The SAS was constructed by Interstate Electric Corporation (IEC), under contract with the Army Security Agency which would later become the NSA. Like MADRE, the SAS functioned by converting the digitally stored signals back to analog---but at a much accelerated speed. This higher-frequency signal was then fed through a series of analog filters including a matched filter to discriminate reflections of the radar's distinctive pulse shape. The results of the analog filter process were then redigitized and returned to the computer, which allowed operators to display the results in graphical form. Cobra Mist's designers also provided a useful secondary mechanism: received signals could be written out to magnetic tape, allowing them to be analyzed off-line later by other computer systems. This would prove an important capability during DVST, when multiple other analysis systems were developed to interpret these tapes.

DVST began in 1972 and almost immediately hit a serious snag. Cobra Mist, an implementation of a technology reasonably well proven by MUSIC and MADRE, turned out to be pretty much useless. It struggled to detect targets in even the most favorable conditions, much less past Moscow.

Cobra Mist was never expected to be especially precise. Its huge antenna produced wide beams that reflected off of the ionosphere at an uncertain point. It would never be able to locate targets to more than general regions. But, like the eyes of insects, Cobra Mist was expected to make up for its blurry vision with a very fine sensitivity to motion. Its narrow Doppler filters should have been able to discriminate any object moving at more than 1.5 knots towards or away from the radar. The ability to discriminate target speeds very finely should have allowed Cobra Mist to not only differentiate moving targets from the ground, but to determine the number of moving targets based on their varying speeds.

The Mitre corporation had devised a series of twelve DVST experiments to demonstrate the radar's ability to detect and track aircraft. Despite its excellent sensitivity to motion, only three of these were even performed. The radar's poor performance on the initial three tests made the remainder hopeless, and efforts shifted to identifying the cause of the radar's unexpected blindness.

To make a rather long story short, DVST researchers identified a large, consistent problem with the radar: in every configuration, a great deal of radio noise was received at roughly the same range as the ground. This noise was of such great magnitude, not previously observed with OTH radar, that it drowned out the actual radar returns, preventing discrimination of the pulse reflections.

There were other problems as well. The 6MHz frequency floor proved a challenge, the computer display system was difficult to use to detect faint targets, and the antenna somewhat under-performed expectations. Various adaptations were made on most of these issues, including a significant overhaul of parts of the antenna system to improve gain, offline analysis of the recorded tapes, and later an overhaul of the computer system to provide a more flexible display capability. But none of these improvements could overcome the most basic problem, the surprising noise, which researchers labeled "range-related noise" due to its appearance at the same range and Doppler gates as the typical ground clutter.

I will spare a multi-page discussion of efforts to characterize and eliminate this noise, which you can read in the original paper linked at the bottom of this article. Suffice to say that numerous experiments showed that the noise appeared at around the same range as the ground in nearly all operating conditions, and that it was far greater in magnitude than would be explained by ground reflection (the usual clutter) or typical atmospheric, galactic, or man-made radio noise. The noise did not originate in the receiver or antenna system, it did not appear over signals transmitted from nearby the antenna, and the same noise could even be detected by a completely separate test antenna system installed for this purpose.

DVST engineers felt that the problem could likely be identified, but that for the time being Cobra Mist could not be effective in its mission. An Air Force panel reviewed the problem, and the result was a bit surprising. Defense money was clearly tighter in the '70s than it is in the era of the F-35, and perhaps the recently negotiated ABM treaty's restrictions on large radar overseas was a factor (Cobra Mist was an overseas radar, since although placed in the UK it was built and operated by the US). In any case, the Air Force pulled the plug. Cobra Mist shut down in 1973, without ever being operational and less than 18 months after the beginning of DVST. The project had cost about a billion dollars in today's money, and its short life left more questions than answers.

Or, perhaps more accurately, it left one very big question: what was the noise?

History has borne out the design of Cobra Mist. Multiple similar OTH radars have been constructed and provided good performance, including the Jindalee radar in Australia (designed with the participation of the same NRL researchers as Cobra Mist) and the Duga radar in the USSR, present-day Ukraine (which I will discuss more in part III). The OTH-B system, based on similar principles and first designed at around the same time, was greatly expanded over following decades and remained in service into the 21st century (this too will be discussed more in the next part). The point is that Cobra Mist *should* have worked, it was all correct in principle, but no one could explain that noise.

Given the passage of a great deal of time since then, the most likely explanation is probably some subtle defect in the receiver design. The Mitre report, written mostly to argue against the radar's premature shutdown, admits the possibility of a problem in the analog-to-digital conversion stage that possibly contributed to the noise but was not thoroughly investigated before the shutdown. The authors also considered the possibility that the noise was a result of one or more of several speculated effects: perhaps there are so many aircraft in the air as to create a broad set of Doppler-shifted reflections that cannot be discriminated from one another. Or perhaps there are some sort of objects on the ground that vibrate or rotate in such a way to produce a set of strong and broadly Doppler-shifted reflections. Subsequent OTH radars have not suffered from these problems, though, making it unlikely that the cause of the noise was some incidental property of Europe at the time.

The reason I find this story so fascinating is that the Mitre authors also suggest another possibility, one that they had intended to evaluate with additional tests that were never performed due to the 1973 shutdown. They thought it possible that Cobra Mist was being jammed.

The basic idea is this: if the USSR became aware of Cobra Mist (not at all a far stretch considering its large physical size) and managed to obtain some detailed information on its operation (via espionage presumably), it would have been quite possible to build an active countermeasure. A system of relatively small transmitters distributed across various locations in the USSR, the Mitre authors estimated 15 sites would do, could use computer technology to detect Cobra Mist's emitted pulses and then transmit the same pulse back at a wide variety of Doppler offsets. It would be difficult or (at least at the time) impossible to differentiate these transmitted "distractors" from actual returns, and they would appear to the operators as random noise centered around the target range. The slight latency the computer system would introduce even provides an explanation for one of the observed properties of the noise, that it occurred mostly at ranges slightly further than the peak ground reflection. Because of the extreme sensitivity of the radar, these active countermeasures would only require a few watts of RF power to overwhelm the radar.

The DVST team was able to perform one experiment that bolstered the theory that the noise was a result of *something* on the ground, whether an intentional countermeasure or an incidental property of something widely used in Europe. When targeting the radar at azimuths and ranges that consisted primarily of open sea, they did not observe the noise. The noise only seemed to occur on land, and more generally in their target region. The Mitre paper states near its end:

We are forced to conclude that the jamming technique is quite feasible,
and it is not clear that the experiments conducted at the AN/FPS-95
would have discovered the jamming had it occurred. If experiments
confirming or denying the possibility had been conducted, they would
have perhaps resolved the issue. They were not conducted.

With fifty years of hindsight it seems unlikely that there was any meaningful scale electronic countermeasures effort in the USSR that has still not come to light, but the possibility is certainly interesting. It is, at least in my mind, within the realm of reason: the USSR were working on their own OTH radar efforts at the time along much the same technical lines and so would have been familiar with the principles that Cobra Mist relied on. The USSR's own premiere OTH radar, Duga, suffered occasional jamming and the operators had to develop countermeasures. In other words, both the means and the motivation were in place. Considering that the scale of the jamming effort would have been relatively small (the Mitre authors mention 15 sites would have been sufficient and that Mitre had constructed similar antennas at the cost of about $25,000, low for high-end RF equipment at the time), perhaps there is something to it and the Soviet program simply faded into obscurity, little documented and known to few. This could be a minor front of early electronic warfare that has been lost to history.

The failure of Cobra Mist, frustrating as it was, did little to dissuade the United States or other countries from the broader concept of OTH radar. In fact, by the time Cobra Mist DVST was in progress, the US had already begin major work on an OTH radar system in the United States (and thus compliant with the ABM treaty), called OTH-B. By the end of the Cold War, OTH-B reached almost 10 MW in combined operating power across multiple sites, and was well on the way to complete 360 degree coverage from the US extending over a large portion of the planet.

As with many things late in the Cold War, it also suffered an ignominious fate: repeatedly replanned at the whims of politics, partially constructed, partially dismantled, repurposed for the war on drugs, and ultimately forgotten. We'll talk more about OTH-B, and its Soviet contemporary Duga made incidentally famous by the Chernobyl disaster, in Part III. As with most of my interests, it involves secret sites in barren remote areas near failed cities... not California City or the Rio Estates this time, but Christmas Valley, Oregon.

[1] I think I've used these part numbers several time without explaining them. Dating back to WWII, many military technology systems have been assigned article numbers according to this system. Initially, it was called the Army-Navy Nomenclature System, from which the "AN" prefix is derived (for Army-Navy). The part after the slash is a type code, for which we most often discuss FPS: fixed (F) radar (P) search (S). P is used for "radar" because R had already been claimed for "radio." When researching radar, you will also often see AN/TPS - transportable search radar. The number is just, well, a number, assigned to each project sequentially.

[2] This report, titled "The Enigma of the AN/FPS-95 OTH Radar," was helpfully retrieved via FOIA and preserved by the Computer UFO Network---not so much due to a latent interest in radar history but because Cobra Mist's nature as a secret facility in close proximity to the Rendlesham Forest has suggested to some a connection to the prominent UFO incident there. You can read it at CUFON