computers are bad

You are receiving this facsimile because you signed up for fax delivery of this newsletter. To stop delivery, contact Computers Are Bad by email or fax.

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

2022-12-11 over the horizon radar pt III

Programming notes:

1. While I am making few to no promises, I am once again attempting to fill the hole that social media left in my heart by using Cohost. If you are one of its ten other users please follow me: jbcrawford
2. After something like two years of saying I would, I finally got around to implementing my own software for the email newsletter. I have of course moved over existing subscribers but if you notice any issues with the email version please let me know. If you aren't an email subscriber, now is a good time to make my self esteem number go up and subscribe.

Most OTH radar have faded into obscurity since the end of their era of relevance (something we will get to in the chronology soon). One, though, looms above all the others. This is mostly a figurative expression, but it's also literally true to some degree, as it's a pretty tall one. I am referring of course to Duga, or as some might know it, the brain scorcher.

As I have previously mentioned, development of OTH radar proceeded in the USSR more or less in synchronization to developments on this side of the iron curtain. It's a bit difficult to nail down the Soviet history exactly due to a combination of lingering Cold War secrecy and the language barrier, but it appears that the USSR may have developed their first functional OTH radar somewhat before the US, at least as a pilot project. Soviet OTH radar did not become well known to the western world, though, until the commissioning of their first full-scale ballistic missile surveillance radar in 1972. This would become known as Duga-1.

To discuss Duga we first need to return briefly to the fundamentals of radar. A monostatic radar system transmits and receives at the same site, observing for reflections. A bistatic radar system transmits and receives from separate sites. The earliest bistatic radar systems, and many later on, actually looked less for reflections than variations in the signal received directly from the transmit site. We can understand this concept by considering the simplest possible type of bistatic radar: if you transmit a signal at one antenna, and receive it at another, than any reduction in the strength of the received signal might indicate the presence of an object in between.

This naive scheme is actually surprisingly practical and sees real use mostly in intrusion-detection applications, where microwave or millimeter wave radar acts as an RF analog to a laser or infrared beam gate---which has the advantage of being more tolerant of rain and fog. A notable example is the AN/TPS-39 microwave perimeter protection radar used to protect the entrance points to Atlas missile silos, but a similar system (made both more reliable and more compact by digital technology) is sometimes used today, particularly around the perimeters of prisons.

Bistatic radar can also be appreciably more complex. For example, the "Pinetree Line" radar system deployed across southern Canada for detection of Soviet bombers employed the simple bistatic radar principle enhanced with the use of Doppler detection. If you imagine the field between a radar transmitter pointed at a remote radar receiver, you might realize that objects not exactly between the antennas, but near the line between them, might reflect RF radiation at a glancing angle and cause a return to the receiver delayed by the angle distance between the object and the two sites. The Pinetree line exploited this effect as well, allowing it to detect aircraft before they reached the line proper.

At the extreme, a bistatic radar system detecting objects not directly between the sites appears more like a monostatic system, and the separation of the transmit and receive sites becomes more a detail of the antenna construction than the operating principle of the radar. This blurring of the line between monostatic and bistatic radar is apparent in OTH radar, very much as a design detail of the antennas.

There is a general challenge in radio: transmit and receive can be difficult functions to combine. There are several reasons for this. Considering the practicalities of antenna design, different antennas are sometimes more optimal for transit vs. receive. Moreover, it is very difficult to receive at the same time as transmitting, because the transmit power will easily overwhelm the front-end filtering of the receiver causing extremely poor reception, if not outright damage (this is the basic reason most conventional radio applications are "half duplex," where you cannot hear when you are talking).

For radar, this problem is typically addressed by transmitting a pulse and then switching to receive operation immediately afterwards. The same problem rears its head again when considering the relationship between range and pulse repetition frequency. When a radar detects an object 1500km away (slant range for simplicity), it takes about 0.01s for the RF pulse to make the round trip. This means that if the radar is avoiding duplex operation, it can transmit pulses a maximum of 100 times per second. In reality the problem is much worse, because of the need for buffer time past the longest desired range and more significantly the time required to switch between transmit and receive, which in the case of radar was typically being done by "RF contactors..." huge relays designed to tolerate RF power [1].

OTH radar systems are intended to operate at both extremely long ranges (3000km being moderate in this context) and desire high PRFs, since high PRFs are very useful in identifying and tracking fast-moving objects. For this reason, most practical OTH radars actually transmit and receive at the same time. There is only one practical way to achieve this: geographic separation, or in other words, bistatic radar. Most of the OTH radar systems we will discuss from this point use this method, with transmitter and receiver separated by at least 50km or so and usually placed on opposite sides of a mountain range or other ridge to provide additional geographic attenuation of the direct transmit power at the receive site.

That was all a long preamble to say where Duga is, which is of course the reason that it's so well known. The Duga-1 transmit site was located very near Chernobyl in present-day Ukraine, and the receive site about 50km to the northeast. The imposing Duga-1 antenna, which might be 150m tall but it's hard to find that number for sure, is visible from many parts of the Chernobyl exclusion zone and is a typical stop when visiting the area. The other antenna of Duga-1 no longer survives, having been dismantled decades ago.

It's hard to find much detailed information about Duga-1. Its transmit power is often cited as 10MW, but given the general history of OTH radar in the United States I find it likely that that was a target peak power that may never have been achieved. The practical average output power was probably more in the range of 1MW or lower, similar to US systems. There are actually quite clearly two separate antennas at the Chernobyl transmit site, something that is oddly rarely mentioned. I haven't been able to find details on why, but I will speculate: first, it is not at all unusual for radio transmit sites to have two antennas, primary and auxiliary. The auxiliary antenna serves as a backup and can be switched into use when the primary antenna is damaged or being serviced. The auxiliary antenna is often smaller, both to save money on the antenna and because it's often used with an auxiliary transmitter, smaller for the same reasons.

It's also possible that the two antennas have different radiation patterns, allowing basic steering functionality such as seen with Cobra Mist. It's most likely that these patterns would be two altitude angles, intended for longer and shorter range use. Considering the different sizes of the antennas it might make sense for the smaller antenna to have a higher altitude angle, since less transmit power would be required at shorter ranges (remember that the radar signals reflect of the upper atmosphere, so aiming higher moves the target area closer).

Duga-1 pointed north, towards the United States by great circle distance. It operated from 7-19 MHz (the lower minimum frequency than Cobra Mist would improve performance at night), and tended to operate on a given frequency for a short time of up to ten minutes before switching to a different frequency. This is a common operating pattern for HF radars because of the need for human operators to find clear frequencies and because different frequencies will tend to propagate better to certain regions. Cobra Mist tended to operate in a similar fashion, with operators selecting a likely-looking frequency and operating the radar on it for long enough to receive and analyze a fair amount of data before moving on.

Duga is best known in the western world for its transmissions. There is an aspect of OTH radar that I have not so far discussed, but is important: HF radio can propagate over extremely long distances, and OTH radars operate at very high power levels. The pulses transmitted by OTH radar can thus often be received over a large portion of the planet. Duga became famous for its widely heard chirps, earning it the nickname "Russian woodpecker." It was particularly well-known because it operated without regard for typical European and American band plans (but presumably according to USSR spectrum management conventions), so it was often heard on broadcast and amateur bands where it became irritating interference.

There was a famous fight back against the interference: some amateur radio operators began actively jamming Duga, transmitting recorded pulses back on the same frequency. These false returns apparently did cause frustration to the Duga operators, because shortly after jamming started Duga would switch to a different frequency. This casual electronic warfare went on for years.

Duga-1 apparently performed well enough in practice, as the USSR built another one. Duga-3, also located in Ukraine, faced east instead to observe for launches from Asia. It appears to have been a substantially similar design and was likely located in the same region to allow for shared engineering efforts and parts stockpiles.

Duga, at least its general design, is not unique to the USSR, and might be considered a latecomer depending on how exactly you place the dates. Despite the mysterious failure of Cobra Mist, the United States had nearly simultaneously began work on a full-scale OTH radar for defense against inbound ICBMs: OTH-B.

The history of OTH-B is fuzzier than Cobra Mist, although not quite as fuzzy as that of Duga. Although OTH-B is no longer operational, it formally left operation recently enough that it's likely some of the documentation is still classified. Moreover, OTH-B became a bureaucratic embarrassment to the Air Force for much the same reason as other late Cold War bureaucratic embarrassments: it took so long to finish that by the time it was ready, the war was over. There are still documents available from DTIC, though, enough that we can put together the general details.

Work on OTH-B began in 1970, before completion of Cobra Mist, but progress was slow for the kind of reasons that are later detailed in a GAO report. Cost and schedule issues delayed the first working prototype to 1977, years after the failure of Cobra Mist. OTH-B faced similar challenges: the prototype system apparently did not perform to specification, but it was functional enough that the Air Force made the decision to proceed with full-scale implementation.

Sometime in the late '70s, construction began on the first full OTH-B site. Like Duga, OTH-B employed a bistatic design. The transmitter was located at Moscow, Maine (a fitting town name), and the receiver near Columbia Falls, Maine. While the timeline of OTH-B is often confusing, multiple sources agree that the Moscow and Columbia Falls sites entered service in 1990---an astoundingly late date considering the program's 1970 start (with limited experimental validation of antenna and radio designs in the first year!). The reasons for the tremendously slow development of OTH-B seem lost to history, or at least deeply buried in Air Force files. We do know that over the course of its development the GAO issued multiple reports questioning not only the extremely long schedule and large budget of the program, but the necessity of the program itself. In 1983 the GAO estimated a final cost of about one billion dollars per site, and given that the program continued for another seven years before achieving its first commissioning, it seems more likely that this was an underestimate than an overestimate.

The same 1983 GAO report suggests that OTH-B befell much the same fate as many ambitious military acquisition programs: the full production system was contracted before experimental validation was complete, resulting in parallel construction and R&D activities. This situation nearly always results in changes being made to the design after relevant components were assembled, creating a work-rework cycle that can easily drag on for years and hundreds of millions.

For construction and operations purposes, OTH-B was divided into sectors. The word sector here can be understood in the same way as "sector antenna." A sector of OTH-B comprised a complete antenna and transmitter chain that provided coverage of an area of roughly 60 degrees. The East Coast system installed in Maine included three sectors to provide 180 degree coverage off of the coast, and the minimum practical takeoff angle resulted in a range of a bit over 3000km, placing a surprisingly large chunk of South America and a sliver of Europe within the detection area. Each sector operated at 1MW of power, although it's not clear to me if this was peak or average---more likely, peak.

OTH-B's late commissioning date provided the benefit of much more sophisticated computer technology. OTH-B target discrimination appears to have been entirely digital rather than a hybrid digital-analog system as seen in Cobra Mist. The system implemented beam steering, and used a basic digital beam steering design in which the transmit site emitted a 7.5 degree beam and the receive site processed it as three, separate 2.5 degree slices simultaneously based on a simple phased-array technique.

OTH-B's basic radio design varied appreciably from earlier systems like Cobra Mist or Duga. Rather than a pulsed radar, it was a continuous wave radar. The transmitter continuously emitted a signal that was frequency modulated with a "chirp" signal that could be used to "align" received reflections with the original transmitted signal. The CW design of the radar gave it significantly improved sensitivity, at the cost that the "duplex problem" (of the transmitter overwhelming the receiver) was much greater. As a result, OTH-B transmit and receive sites required an even greater separation of 100-200km.

OTH-B worked, apparently well although details on performance are hard to come by. It worked, though, in 1990: the year before the fall of the Soviet Union. By 1990 the Cold War was already effectively over, having been replaced by a period of both warming relations and Soviet economic collapse. OTH-B seemed a system without a purpose.

The DoD seldom lets lack of a purpose come between it and a large contract, and despite the clearly decreasing relevancy of the system construction went ahead on the second and third OTH-B sites. In 1988, the Air Force activated two West Coast stations for construction of three additional OTH-B sectors: Christmas Valley, Oregon (transmit) and Tulelake, California (receive). Christmas Valley is an unusual place, a failed real estate development in the pattern set by Horizon and Amrep in the Southwest. Platted to be a large city, it is now a tiny town centered on an artificial lake. A short distance from what might be called city center, the 300+ acre OTH-B site once loomed on the horizon. Its three sector antennas were metal frames reaching 135' feet high for the lowest bands, supporting insulated radiating elements projected in front of the "rack" at a jaunty angle. Each of the three antennas, nearly a mile wide, sat behind a 4000' buried ground screen. In total, each antenna sector antenna occupied nearly a square mile. Behind each antenna, a transmitter building of 14,000 square feet housed the radio equipment, shielded behind the antenna back screen.

Although construction of the Christmas Valley and Tulelake sites was substantially complete, they were turned down in 1991 a few months before the Soviet Union itself was deactivated. The radar never operated.

Meanwhile, construction had begun on the first sector of the north system, in Alaska. The transmitter would be located at Gakona, and the receiver at Tok. Little is known about this northern site. Radomes Inc. (an association for veterans of air defense radar operations) reports that construction on the first sector transmitter at Gakona had begun but was never completed, and the Tok site was never started.

As of 1991, the constructed and activated east installation of OTH-B formed sectors 1-3, the constructed but never activated west installation sectors 3-6, and the partially constructed Alaska system sectors 7-8. The final planned sector, 9, would have been located somewhere in the central United States facing south. It is not clear to me if a final site was ever selected for the south system, as it was canceled before construction was contracted.

The east system, fully functional when it was deactivated, won a second life. From 1991 to 1997 it was operated in partnership with the Border Patrol on a drug interdiction mission, detecting aircraft involved in smuggling. It was apparently useful enough in this role that its antenna was modified in 1992 to expand coverage over a portion of the US-Mexico border.

The western site, never activated, was kept in an idle state ready for potential reactivation until 1997. In 1997, the eastern system was deactivated and both east and west sites were transitioned to a storage status with only a few staff members to maintain site security and integrity against the elements. In 2002, east and west were reduced to "cold storage" and the radio equipment was removed. In 2007, at 17 years of disuse after their near-completion, the Christmas Valley and Tulelake sites were partially demolished. Antennas were removed and sold for scrap and some site remediation was performed. The equipment buildings and some infrastructure remain, their ownership having reverted to the respective states who have unclear plans for their future use. The Maine sites appear to be in a similar state, with antennas removed, although I haven't been able to find documentation on this action.

The partially constructed site at Gakona, Alaska met a more interesting fate. In 1993, it was reused by the Air Force in partnership with other organizations for the construction of HAARP. HAARP is a research instrument that uses RF radiation to ionize the upper atmosphere and observes the results. While HAARP is now operated by the University of Alaska for public research purposes, its original lifespan under Air Force management was secretive enough to make it the center of many conspiracy theories. While the Air Force is cagey about identifying its research function, it seems likely that it was used for at least two disciplines.

First, to develop an improved understanding of the radio reflection properties of the upper atmosphere. The lack of a clear model of this behavior was repeatedly identified as a limitation in the performance of OTH radar systems, and so this research would have directly contributed to the Air Force's interest in long-range radar. Second, by this time the NSA had a well-established interest in ionospheric radio propagation. Unusual propagation patterns sometimes allowed radio signals originating within the USSR and China to be received in nearby US-allied territories or even within the US proper. The NSA operated several research facilities that both performed this type of interception and basic research into ionospheric propagation with an eye towards making this signals intelligence method more effective and predictable. While not proven to my knowledge, it seems likely that HAARP served as a part of this effort, and that active ionization of the atmosphere may have been at least evaluated as a possibility to improve the chances of interception.

Amusingly, a 2015 paper from a HAARP researcher suggests that the University of Alaska is once again investigating use of the site for an OTH radar directed at observing vessel traffic through the arctic ocean. The arctic ocean north of Canada is becoming increasingly navigable due to climate change, and so is thought likely to become a strategically important area for naval combat and enforcement of Canadian sovereignty. Indeed, confirming a connection between HAARP's research goals and OTH radar, the paper notes that HAARP has successfully demonstrated the creation of an "artificial ionosphere" at lower altitudes which can be used to perform OTH radar at shorter ranges than previously practical---filling in the "beyond visual horizon but before atmospheric skip" range gap in existing radar methods.

What of OTH radar as a broader art? I have framed the decline of OTH radar systems mostly in terms of the fall of the Soviet Union. This was indeed the killing blow for both Duga and OTH-B, but by 1983 the GAO was recommending that OTH-B be canceled. OTH radar faced a challenge even greater than the end of the war: satellites. By the '90s, satellite remote sensing had advanced to the point that it could perform the functions of OTH radar, and usually better. The MIDAS (Missile Defense Alarm System) satellites, launched in the '60s, were the first working space-based ICBM launch detection system. MIDAS was quickly replaced by the more covertly named Defense Support Program or DSP. These satellite systems, based mostly on infrared imaging, remain the United States' primary missile defense warning system today. By the '80s, they provided both better sensitivity and wider coverage than HF OTH radar.

Today, HF OTH radar has faded nearly into obscurity. Some military OTH radar systems still operate. The Navy, for example, built a relatively small OTH radar system in 1987 directed south for surveillance of the Caribbean. It remains in operation today, mostly for counter-narcotics surveillance similar to the east OTH-B site's second chance. China constructed and operates a system very similar to OTH-B for airspace surveillance. China's OTH-B is thought to use interferometry methods for significantly improved resolution. The antenna design looks extremely similar to OTH-B, although I suspect this is more a case of common constraints producing a common design than China having based their design on the US one.

Interestingly, one of OTH-Bs most enduring contributions is not military but scientific: NOAA, through a partnership with the Air Force, discovered that the "ground clutter" received by OTH-B's east site from the ocean could be analyzed to collect information on wave heights and current directions. Today, a system of small coastal HF radars on both sides of the United States collect information useful for marine weather forecasting. Data from this system is also used for improved tidal predictions and by the Coast Guard to model surface currents, allowing improved predictions of where a drifting vessel in distress (or its lifeboats or survivors) might be found.

Ionospheric propagation remains an area of scientific interest. For some time, the Arecibo observatory operated an ionospheric modification observatory using transmitters repurposed from OTH-B's eastern sectors. HAARP operates to this day. The NSA is mum as always, but seems to still be busy at their propagation research facilities. But large OTH radar, on the scale of OTH-B, seems unlikely to return to the United States. It is costly, complex, and struggles to provide anything that satellites can't. As with terrestrial radio-based navigation systems, perhaps the specter of ASAT warfare will motivate renewed interest in OTH methods. But I wouldn't hold my breath. For now, all we have is a set of large trapezoids in the Oregon desert.

The end of OTH radar is, of course, far from the end of military radar. Improved radar methods, especially phased-array systems, made it far more practical to detect ICBMs at their extremely high, orbital approach altitudes. Terrestrial missile warning systems re-concentrated on this method, and many are in operation today. I'll get to them eventually.

[1] Some readers might be familiar with duplexers, filtering devices that employ dark magic RF trickery to effectively "split" transmit and receive power. There are two reasons that filtering-based approaches are not very practical for radar: first, radar must receive the same frequency it transmits, preventing the use of common cavity filters for separation. Second, RF filters become larger and dissipate more heat as they handle higher powers, and most radar operates at very high power levels. Even the simplest RF filters become impractical to build when considering megawatt power levels.