Ghost busting on the ocean floor

Some years ago, I worked for a start-up that was engaged in the manufacture and operation of side-scan sonars in seafloor-mapping operations. Like any other good start-up, my company performed its development work on a shoestring budget and was successful only by dint of young engineers like me doing long hours of overtime and late-night debugging sessions.

Although my primary responsibility was the development of digital-signal-processing hardware and software, during survey operations, I was responsible for operation and maintenance of any piece of electronic gear on the ship, including the stuff that we didn't build.

Side-scan sonars comprise two major assemblies: the "tow fish" and the surface-electronics package. The tow fish is a submersible assembly ranging in size from a torpedo to a large car, depending on the sonar operating frequency. A ship tows it at a depth that both isolates it from the effects of sea-surface-wave motion and places it below the acoustic-path-distorting thermocline. The tow fish houses a pair of port- and starboard-looking transducer arrays and their associated front-end electronics, as well as the telemetry system for sending digitized sonar returns to the surface and receiving command and control messages from the surface. The surface-electronics package consists of a real-time signal processor, which converts the sonar returns into both calibrated intensity and phase data—my specialty—and an enterprise-class computer system that produces the final backscatter images and bathymetric charts. Backscatter images are essentially monochrome maps of the reflectivity of the seafloor. Bathymetric charts are simply seafloor-contour maps.

One of the systems that we built had a chronic problem. It appeared as if data from the port and the starboard sides were somehow intermingling, causing ghost images of the port side to appear in the starboard data and vice versa. There were numerous theories about why this situation was occurring, but our development cycle never allowed us the time to get to the bottom of the problem.

I performed the tedious exercise of putting the system in a self-test mode. In this configuration, I could inject well-behaved test signals into the sonar front end and trace them through the entire system. Although we had built this capability into the system, no one until now had ever thought to use it to try to isolate the ghosting problem. Fortunately, the problem had become a persistent one, and, after a couple of hours, I isolated it to one board in the signal-processor rack. This board's function was to synchronize and demultiplex the digital-data stream that the tow fish had telemetered to the surface.

The board used mid-1980s technology and consisted mainly of dozens of 74LSxxx chips wired up as state machines, multiplexers, and the like. Oscilloscope probe in hand, suspect board on a card extender, schematics spread out on the improvised work bench, I followed the signal path from the board input into the belly of the beast. My efforts soon paid off. I discovered that an 8-bit gate that marshaled the incoming data to either the port or the starboard processing streams had a cold solder joint on its A/B control pin. So, data only sometimes made it to the right place. A quick touchup with the soldering iron fixed the problem.

Why didn't it occur sooner to us to try this test? I can't really say. In any case, our attempts to fix the problem during postprocessing amounted to sweeping the problem under the carpet. Returning to first principles saved the day. EDN

Hugh Shane is lead communications engineer at The MITRE Corp. Like Hugh, you can share your Tales from the Cube and receive $200. Contact edn.editor@reedbusiness.com.

www.edn.com/tales

DANIEL VASCONCELLOS