

The power of the hot seat



In the mid-1980s, I was designing visual systems for flight simulators for computer-graphics company Evans & Sutherland. These systems were huge—dozens of 19-in. racks full of cards working together to produce the out-the-window cockpit scene. The primary technologies were wire-wrap and 74LS logic. One of these big, multimillion-dollar systems was ready to ship to NASA, but, two days before the ship date, final testing turned up “sparkles” in several visual channels.

The only obvious thing was that the problem originated in the display processor, a collection of about 100 wire-wrap boards and PCBs (printed-circuit boards) plugged into a wire-wrap backpanel. To make matters worse, this system had 1024×1024-pixel resolution, which required four display processors load sharing in the scan-line-interleaved fashion that some commercial graphics chips now use. Thus, the potential offenders included about 400 cards, probably 100 cables, and hundreds of thousands of wire-wrap wires.

Management assigned me to find and fix the problem. We found that the picture from each individual processor

was good on its own, but the combined image possessed the dreaded sparkles. The chief suspect became the fastest clock in the system, which had a blazing 25-nsec period. (Most clocks were 100 nsec.) This clock, which was the basis of all the frame-buffer output and video timing, crossed wire-wrap backpanels eight times, traversed wire-wrap cards 16 times, and traveled through three ribbon cables, with the system rebuffering it several times along the way. By the end, it looked awful.

I figured that a good approach might be to use ECL and differential signaling to control rise and fall times and reduce common-mode noise. Fortunately, we

had a handful of parts available and a built-in -5.2V power supply. Unfortunately, doing the conversion at this late date meant mass rework to get new power to some cards, placing new parts on many cards, and replacing lots of single wires with twisted-pair wires. Nevertheless, from my position in the hot seat, I had almost every company resource at my disposal.

Consequently, our team reworked and tested one processor. The clock looked clean, so we reworked the other three processors in the channel. In the combined system, to our distress, there were severe anomalies in the last processor. We found that, although the clock was clean, its duty cycle was distorted to the point that it was just a thin pulse. Going back to the data book, I learned that the rise time of ECL is typically 0.5 nsec slower than the fall time. Each of the buffers was eating away at that 12.5-nsec high time of the clock. Inverting the clock at strategic points restored balance to the duty cycle and vanquished the sparkles.

I explained the issues and the solution to management. Initially, there was some resistance due to the amount of rework involved. However, I soon had support up and down the chain of command. We completed the sign-off process, which normally took weeks, in hours. Subsequently, a small army of rework technicians attacked the remaining dozen or so processors. As each channel was completed, the technicians would fire it up, verify that it was sparkle-free and still worked, and then wrap up the racks and take them to the loading dock. We started shipping the product on the day it was promised.

Aside from signal-integrity insights, the most valuable lesson I learned was organizational: Though the hot seat is uncomfortable, it bestows unique power to cut through red tape. **EDN**

Reed P Tidwell is a senior staff applications engineer for the Advanced Products Division of Xilinx. Like Reed, share your Tales from the Cube and receive \$200. Contact Maury Wright at mgwright@edn.com.