

A bad-capacitor story ends happily



worked as a design engineer for an optical-telecom company that had deployed 1000 pieces of equipment worldwide. Having so many modules in the field means a trickle of returns, and it was my job to investigate the failures. One investigation taught me a wonderful lesson.

I received a module whose source of failure was easily identifiable: a charred tantalum capacitor. It failed short, making the whole multithousand-dollar module nonoperational. This surface-mount capacitor—with a 7343 footprint and 20V rating—was sitting on a 12V-dc plane. This failure rate of one capacitor in about 10,000 pieces in this time span was well below the statistical prediction. I took a picture of the fallen capacitor and considered the case closed.

In a few weeks, a customer returned a similar module with a charred and shorted capacitor in the same location. Even including this case, the failure rate was still below statistical prediction. I knew there were five more identical capacitors on the board, sitting in parallel on the same 12V-dc plane. In

addition to the module's failure rate, I now had a one-in-six chance with the capacitors. So, I took another picture. I wrote a report to calm upper management, but I had a feeling that I'd better study reliability calculation in general and reliability for tantalum capacitors in particular, and the faster, the better.

In another few weeks, I received another failed module. The same capacitor looked bad. I had by now done my studying and could intimidate other people by saying long and complicated

sentences about reliability, but why was it always the same capacitor? Overvoltage? Spikes? No way. The same plane contained plenty of sensitive stuff that would fry well before the capacitor even felt it. Having nothing better, I clung to the theory of excessive ripple current.

The idea of a temperature rise due to ripple current causing the failure gained traction when all three photos of the fallen capacitors revealed a common condition: almost no solder on each negative terminal. The electrical connection was still good, but there was little solder. The capacitor's positive terminal was fine with a fair amount of curvature-profiled solder. I started to promote the idea that the lack of solder had caused impeded thermal contact, but it was only wishful thinking. I calculated the worst ripple current: 10% of the maximum rating. On an operational board, I got less than 5%.

I had already dismissed other ideas—from excessive humidity to airflow turbulence. Suddenly, the picture of the layout popped up in my mind. The layout sections for the five good capacitors were identical: Vias were close to both terminals going down to an internal layer. The bad capacitor had a via at the positive terminal, but, at the negative end, there was a heavy trace going *inside* the footprint, *beneath* the capacitor, and *only then* outside. That's when I knew how to fit together all the pieces of the puzzle.

On the positive terminal, the solder stayed where it was supposed to, clinching the terminal to the PCB (printed-circuit board). On the negative side, however, during assembly, the melted solder drifted under the capacitor and solidified, lifting the negative end and bending the capacitor just enough to create a microcrack—a capacitor's well-known nemesis. I never felt as much excitement writing a technical report as I did the next day. **EDN**

You can reach Samuel Kerem at skerem@gmail.com. For more Tales from the Cube, go to www.edn.com/tales.