Double, Double, Toil and Troubleshooting: An Engineer's Book of Essential Spells

In this hair-raising book of spells, engineers describe some of their most challenging experiences finding, fixing, and banishing a technical glitch—feats one might even call witchcraft!!!

WITCHES WITH THEIR WARTY GREEN FACES and ominous cackling—like those huddled around a bubbling cauldron in Shakespeare's Macbeth—may not seem to have all that much have in common with engineers.

But engineers do have this kind of uncanny—almost magical ability—to troubleshoot problems, whether it's coaxing dead products back to life, tracking down bugs hidden deep within software code, or sussing out the source of spurious signals that wreak havoc on nearby equipment. The list is endless.

Like modern-day engineers, witches don't just give up when problems crop up or things don't live up to their expectations. No sir, they stir a boiling cauldron while chanting spells or draw giant pentagrams in chalk, applying their magic powers to banish what isn't working and get what they want.

So, in honor of witches and their special holiday—Halloween—*Electronic Design* recognizes this special ability of engineers and witches with real-world stories of how engineers applied their magical powers to troubleshoot (or maybe put a hex on) a vexing technical problem.

TABLE OF CONTENTS

1. A Spell to Cast Out Audio-Video Sync Problems
2. A Spell to Banish a Phantom Bug in an Emulator Tool
3. A Spell to Rehabilitate a Dysfunctional Power Supply
4. A Spell to Reanimate a Disabled Communications System
5. A Spell to Make Right Wildly Erroneous Position Data
6. A Spell to Bring a Dead Transmitter Back to Life
7. A Spell to Defeat a Power-Sucking, Battery-Killing Firmware Flaw
8. A Spell to Vanquish Intermittent Interruptions in a Data Stream
9. A Spell to Eradicate Erratic Sensor Readings
10. A Spell to Locate the Hidden Short

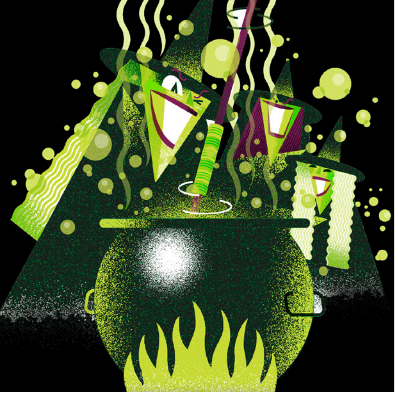

1. A Spell to Cast Out Audio-Video Sync Problems

Sifting through gigabytes of compressed data, an unemployed software developer finds the one bad value.

By JEFF DAVIES, Consumer Electronics Hardware and Software Designer

Back in the early 2000's I was a DVR fanatic. My TiVO and ReplayTV collection and I were mentioned in a *New York Times* article on DVR users.

In 2003 the ReplayTV 5000 series DVR was released. In addition to the automatic commercial skip found on earlier models, this new model could send shows to a ReplayTV in another room or anywhere in the world.

Even with increased sales from the new products, 2003 was a bad year for the company. TiVO was the number one DVR brand and broadcasters were suing over commercial skip and show sharing features. It wasn't long before legal costs overwhelmed their parent company, SONICBlue, which filed for bankruptcy protection.

Although the brand was acquired by Denon & Marantz Holdings under the DNNA brand, ReplayTV was struggling to survive.

During the ReplayTV 5000's first year, users reported losing audio sync in some recordings. Video would momentarily freeze and after resuming the audio was seconds behind. It was only a matter of time before I, too, experienced the same problem.

After talking with tech support, I learned there were no developers left to fix the problem. I was an unemployed software developer taking care of two young kids. I had free time and experience with audio and video compression, so I decided to see if I could find the cause and fix it.

The ReplayTV recording format was known after hackers figured out the disk partition table and file system. Developers created utilities to access the files and those apps were shared on sites like www.avsforum.com.

That meant that by connecting the DVR drive to a computer, a technical user could extract recordings using command line utilities When show sharing was available, a utility, called DVArchive, allowed everyone to move recordings to a PC without pulling the drive.

DVArchive was a Java application that cleverly mimicked the ReplayTV networking protocol. A PC running DVArchive looks like a DVR to a ReplayTV, which allowed the two to share shows.

A recorded show is a large MPEG2 file. I figured the problem might well be caused by a hardware issue like a bad disk or MPEG encoder/decoder problem; a software bug; or a bad data stream.

MPEG video was designed to gracefully play through data loss and recover completely. Since playback wasn't recovering it didn't seem like a decoding error. This could be as simple as one bad key bit, byte, or value in gigabytes of data. Where to begin?

My *brilliant* idea was to ask users with bad recordings to send their recordings to my DVRs. After moving the files to a PC with DVArchive I had 15+GB of files to examine!

MPEG files are interleaved streams of audio, video, and proprietary system data. I trimmed each recording down to a few seconds before and after the glitch, hoping the problem existed nearby. I could only basically hope that was the case, because these trimmed clips played fine on my DVD players and with Cyberlink's PowerDVD.

I wrote a program to parse the MPEG stream, so I could examine all data packets. After weeks of work I made little progress and the decoded data looked fine.

Because these recordings are generated in a closed system, I ignored much of the data frame headers except data type and size fields. I ignored the decode timestamp (DTS) and presentation timestamp (PTS) that are embedded in every data frame, which tell the decoder when to decode and present the data. I found these values may be ignored by MPEG decoders (including those mentioned above), which correctly play these clips.

Running out of ideas, I decided to check the time stamps, notwithstanding that I had absolutely no idea what constituted a "bad" value. Since encoding and playback rates are consistent, I decided to check offsets between PTS and DTS values in successive data chunks.

I got extremely lucky that among the thousands of values displayed a single negative PTS offset stood out. Could this be the cause?

I confirmed a bad PTS was set for seconds before the prior PTS. This jump backward was found in all shared recordings. Since a PTS is derived from the previous PTS, any error will propagate through until the end.

The fix was simple; I needed to recalculate the bad audio PTS and then update the remaining times. I shared my streamfix tool with users on avsforum.com. Users began fixing their recordings with my tool.

This troubleshooting escapade took me a lot of time, but it turned out to be well worth the effort as I enjoyed helping the ReplayTV community!

I shared my findings with ReplayTV, but unfortunately they never fixed their problem.

JEFF DAVIES is a consumer electronic hardware and software designer and developer currently working on intelligent lighting products. He loves tearing things apart to rebuild them to perform better than they did when new.

Back to table of contents

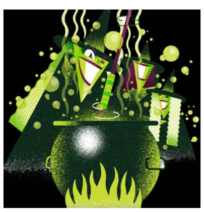

2. A Spell to Banish a Phantom Bug in an Emulator Tool

The customer was convinced he found a bug in a vendor's in-circuit emulator tool, but suspicions turned to his own board design.

By ERIC OVERTON, CEO, Focus Embedded

Twenty years ago, I worked for a company that designed in-circuit microprocessor emulators (ICE's), since JTAG debugging was relatively new. Ours ran well, but every so often we'd get a call from a customer who was convinced he'd found a bug in our tool.

We got one such call from a customer who was debugging new code running on a board that otherwise worked. Just before he got to that new section, the code would run off into empty memory space. He blamed the emulator, because when he put an actual microprocessor chip into the socket where the emulator had been, "printf" statements to an onboard UART indicated the bad jump happened well before the system hit his new code.

Further, the last assembly language instruction that executed correctly before the emulator ran out into the weeds was the same one every time. (It was a CMP operation that left a zero flag either set or unset, and the next instruction was a conditional branch based on the zero flag. The compare operation always executed correctly, but the conditional branch always went off into Neverland.)

He also noted that this new design (it was the controller for an elevator) was based on a series of similar products, and none of those had ever behaved strangely. There were thousands of those in the field in high-rise buildings all over the world.

To see if we could replicate the problem. we asked him to send us a copy of his code, so we could run it on a few bits of similar target hardware. The customer obliged, sending us electronic copies of source code, object modules, and final compiler result.

Notwithstanding the old adage, "The customer is always right," we looked first to see if some compiler switch had been set wrong or something were amiss with his "make" process, but he had done everything correctly. We also confirmed that the customer's code ran just fine under emulation on our target.

At that point, we asked him to send his actual target board.

When we got his hardware and repeated the experiment, we found we could reliably reproduce his problem. We also found that (as he indicated) it never occurred with an actual microprocessor installed. Further, the next bad opcode fetch was always from the same general area of memory.

Maybe he had found a real emulator bug – although since there were no breakpoints being set and the bad code region we went into was not where the ISR for the "trap" instruction was (where we put our monitor code), we couldn't see any connection there. And the same code ran reasonably on other target boards.

Oddly enough, it was visual inspection that got us to the answer. My co-worker and I both came to the same realization, "That's an awfully big board to have such small inductors in its switchmode power supply section."

So, we pulled a few components and hooked the board's power rails to a large linear benchtop supply – at which point both emulator and actual microprocessor chip behaved correctly.

The customer had been a victim of his own success. Previous models of this elevator controller had used the same onboard switchmode power supply to good effect. And with each new generation of controller, somebody copied it into the next design. But somewhere along the line, folks stopped bothering to calculate power budgets.

Clearly what happened is that with their new prototype design, the extra 8pF of pin loading of an emulator head (above that of a "real" processor chip) was just enough so that a saggy power supply couldn't flip all the address bus bits from 1 to 0 when the code hit a conditional branch that led from high memory to some area much lower.

The customer then recalculated his power budget based on all the extra features he'd put in the controller over the years. After that, with a beefier supply, the emulator ran perfectly and a whole lot of field failures of marginal boards with marginally undersized power supplies were averted.

We then got to help him with his custom modifications for adding a real time clock and astronomical tables to his code for a special export model going to an apartment complex in Jerusalem, where the residents (all ultra-orthodox Rabbis) wanted the elevator to change modes between sundown Friday and sundown Saturday (so the occupant wouldn't have to do the "work" of pressing a button to go to a particular floor, but instead would stop at every floor on the way up and down).

Back to table of contents

ERIC OVERTON has been working full-time in the embedded systems design business since getting his MSEE degree in 1989. When he's not designing electronics, he is restoring old cars or playing guitar in a rock.

3. A Spell to Rehabilitate a Dysfunctional Power Supply

Will the fourth engineer brought in to fix a firmware issue succeed and secure a job offer?

By KEITH ABELL, Senior Firmware Engineer, Intellipower Inc.

Early on in my career, I was hoping to get hired as a full-time employee at a start-up company. After a lengthy interview process I agreed to the most challenging "job interview" ever: The hiring VP of Engineering proposed the following to me, "Solve our firmware issues in 90 days and you will work yourself into a permanent position. If not, you will be developer number four to fail."

I'm always up for a good challenge, and so I agreed to his terms.

The company developed electronic security systems for buildings. The door electronics included badge readers, door locks, and the “door connector,” which sent the badge read to the controller via RS485 in order to get a yea or nay to open the door. The communications were faulty.

The design included a small switching power supply chip to convert the wall wart 30V DC to the 5V Vdd required by the Microchip processor brains that controlled the door electronics.

I happened to stumble onto some minor—oops--major issues related to the 5V Vdd rail. The microcontroller chip had an absolute max voltage of 7.5V. A scope on the 5V rail showed pulses of 8V every few tenths of a msec—things were not particularly stable.

But that was just the tip of the faulty iceberg!

While examining the switching regulator, I discovered it had both thermal and short-circuit protection in place to prevent damage by shutting down the regulator.

One of my early tests was to short circuit the supply. Poof! A puff of smoke and the smell of burning parts filled the air. The supply was required to deliver at least 1 Amp under operation.

At this point, it was time to compare the circuit design to the recommended circuit and graphs for capacitor and inductor choices: Two caps--one on the input, one on the output--and a single inductor in the circuit.

For the load the device was to carry and the input voltage range, the caps were of the incorrect value and voltage. The inductor not only had the incorrect inductance, but also was rated for 0.75 amps, which meant it would fry in a short-circuit situation long before the regulator protection mechanisms would kick in. Not to mention it would have no chance to deliver the required 1 Amp. <Sigh>.

After I added the proper caps and inductor for the job at hand, the spikes on the rail were gone and the rail was at a proper 5V. To test the fix, I introduced a short before I went home. I removed it the next morning--and Voila! The regulator recovered and the 5V rail was properly restored.

As it turned out, there was more to the communications problem: Transport layer issues (missing stop bits); protocol issues (causing multiple bus errors because the door connector kept ignoring repetitive commands; interrupt race conditions, etc.

But I made so much progress in a month-and-half that the company offered me the full-time position. I resolved all issues in 86 days, and at the annual holiday party I was awarded a NovusEdge autographed shirt with RS485-86 blazoned over the left breast pocket.

It was only at the party I found out there was concern that the communication problems would cause the company to fail.

KEITH ABELL has been a firmware engineer at various high-tech companies for all of his career (35+ years) after serving in the US Army in the Corp of Engineers as Combat Engineer Platoon Leader and later out of the Division Engineers office, 1st Cav Division.

Back to table of contents

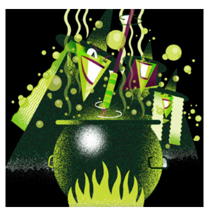

4. A Spell to Reanimate a Disabled Communications System

A “fire” onboard a guided missile frigate provides clues to what felled all communications from HF to UHF.

By JIM MUEHLBERG, Electronics Engineer, National Radio Astronomy Observatory, Front End Local Oscillator Group

Thirty some years ago I was in the Navy. I was an electronics technician or “ET” on the guided missile frigate Samuel B. Roberts, FFG-58. My specialty was maintaining all the HF, UHF and cryptographic gear associated with the communications gear.

We were on deployment somewhere in the Atlantic. When you’re in the Navy, “exercise” has two meanings - one is physical, the other is like practicing for battle. The radiomen were doing an exercise where they had to maintain constant contact with NAVCAMS LANT, or Naval Communications Area Master Station Atlantic (now you understand why we use acronyms) for an extended period of time.

This particular exercise was called “HF Long Hauls. “This was an HF communications channel, somewhere around 2 MHz. It was using FSK, (Frequency Shift Keying) and required a constant carrier. So, the HF transmitter was constantly “keyed down” at about 1KW power.

The transmitters easily tolerated this, so I wasn’t too concerned. But after beginning the exercise, it was noticed that all communications from HF to UHF were disabled, including satellite reception. At first, we did not make the connection, so we looked for common threads, namely the rubidium time standard that fed all comm gear. It was fine.

Baffled, we muddled around until the sun went down. As it grew dark, we heard word of a “fire” near one of the three 35-ft HF whip antennas. I ran up and found a shipmate with a CO2 bottle about to douse what appeared to be an electrical fire, although it was hard to tell what was happening in the dark.

I quickly realized what was going on and asked the radiomen to un-key the HF transmitter being used for the long-haul exercise. Immediately, the “fire” went out. Also, all the communications were restored to normal.

In the Navy’s never-ending battle with corrosion, they employ many seaman to chip, grind, and paint the ship. In order to do a thorough job, one industrious sailor had disconnected a bonding strap from the base of the antenna. Left hanging about 2 inches from the hull, it made a great spark gap.

As soon as power was applied to the antenna, an arc was struck. The resultant broad band RFI (remember spark gap transmitters?) wiped out reception on all of our communication channels.

A wire brush and a bolt cured the problem. I was the hero for a moment, on a ship that turned out to have many heroes, but that’s another story. This event led to a complete bonding and grounding survey (and resultant training for the painters). There are many unintended antennas on the topside of a Navy ship.

It’s good to exercise--that problem would have been much more serious had we been in a real situation!

Jim Muehlberg is a retired Navy Chief and Senior Engineer at the National Radio Astronomy Observatory where he currently oversees maintenance and repair local oscillators in the Atacama Large Millimeter Array Radio Telescope.

Back to table of contents

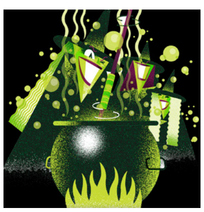

5. A Spell to Make Right Wildly Erroneous Position Data

A radar tracking system for a guided missile goes haywire and all the usual troubleshooting measures turn up nothing.

By NEIL P. ALBAUGH, Analog Design Consultant, Diamond Bell Technology

In the Cold-War days of the early 1960s, I was assigned to a Corporal guided missile unit that guarded the Fulda Gap, a likely Soviet armor invasion route. We spent many days in the field, responding to “Alerts,” which sent our units into the German countryside to conduct preparations to fire a missile.

Fortunately, the order to fire never came.

Central European weather was not easy on equipment, so my repairmen were busy. The Corporal missile guidance system was a primitive all-analog system; the ground guidance computers consisted of 26 chopper-stabilized op amps!

The missile trajectory was tracked by a radar, which fed the data to this computer to be used to guide it to the target. Position information was derived by precision potentiometers on the azmuth (AZ) and elevation (EL) axes of the tracking antenna and these voltages were routed to the computer through a large cable to the radar van and then to the guidance van.

For such an early system, it worked pretty well but it depended on very skilled crews.

One afternoon on an alert in the forest of Wasserlos, the radar operators reported a problem with the position data from the antenna--the position readout voltages were wrong--very wrong! They did the usual troubleshooting, checking the cable connectors to make sure they were tight and even replacing the antenna-to-radar-van cable with a spare without success.

Our Ordnance detachment was called in to track down the problem. After listening to the operators' description of their attempts to troubleshoot their situation and seeing the crazy position voltage readings for ourselves, we decided to start at the source, the position readout potentiometers themselves.

By disconnecting the connector from the antenna pedestal, we could measure the resistance across the pot and to the wiper arm as we manually moved the antenna in AZ and EL. Okay, that eliminated the antenna as the source of the problem so we moved along, replacing the connector and pulling the connector off the other end, a large feedthrough bulkhead mounted in a below-floor-level well on the bottom of the radar van. Checking the resistances with a VOM at this end showed the proper readings which eliminated the cable as the problem.

Plugging that connector back in showed that the problem was still there! Subsequent checks eliminated the radar van to computer van cable path so there wasn't much left to check.

Out of desperation more than anything, I removed the connector from the bulkhead connector inside the van itself and measured the resistance back to the antenna potentiometers. These readings were not correct, in fact they were crazy!

What was going on? The only thing between the antenna cable connector, which I knew was okay, to the pins of the bulkhead connector was the connector itself. What could go wrong with a connector? Disconnecting everything I checked the pins for an open or a short--nothing!

Then I switched the VOM to a higher range and checked the resistance across the pins that carried the AZ pot wiper voltage: 20k ohms! Across the EL pot read 20k. What? In fact, between any pin and any other pin in this big multi-pin bulkhead connector, the VOM read 20k.

This was really strange, I had to find out what was causing this weird resistance reading. With help from the operators, I removed the bulkhead connector and disassembled it. Inside there was a rubber seal that all pins projected through and underneath that was a phenolic block. Under the rubber seal there was a thin layer of black mold! This was the culprit that caused all the problems.

Scrubbing the connector and seal clean and giving it a coat of DC-3 silicone grease solved the problem. We were back up and operational-- Western Europe was safe for another day.

From that experience I learned about "sheet rho" and that the resistivity of black mold was 20k/square.

At the time of this incident, Neil Albaugh was Section Chief of the 157th Ordnance Detachment in Babenhausen, Germany. After a long and rewarding career in industry, he is now an analog consultant.

Back to table of contents

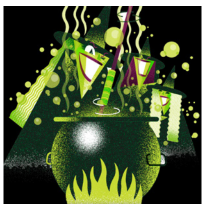

6. A Spell to Bring a Dead Transmitter Back to Life

Moisture inside and outside of a malfunctioning transmitter provides a useful clue.

By JIM SWARTOS, Electronics Engineering Instructor

A friend of mine had been working for nearly three days straight, trying to get a 20 KW transmitter back on the air after a major storm had passed through the area. Since there was no back-up generator, the transmitter building had been "chilled" to about 20 degrees Fahrenheit and was slow to warm up…

That is when he called me for help.

When I arrived, I noticed that there was a layer of condensation on the surface of everything that was metal. Guess what? There was a layer of condensation on the inside of everything too!!!

We spent the next hour using hairdryers and paper towels to soak up the moisture and dry things out.

At that point we tried to power up the transmitter. The initial "Start-up" circuits worked, but as soon as we tried bringing up the "RF Power Out" the "Shut-Down" protection circuits would spring into action.

Over the next couple of hours we troubleshot circuits, correcting problems that may have been caused by trying to bring up a "wet box," by replacing some burned components, fuses, etc.

While we had one large printed circuit board on the bench, I tapped a large power resistor with a wooden dowel. It "exploded" into a bunch of glass shards. My friend looked at me with eyes wide and yelled, "You broke it!!!"

I asked him if the transmitter had been experiencing any sporadic issues. He said that there had been "ghostly problems," but that he was never able to fix them. I explained that a wire wound, ceramic power resistor can develop a small break in the wire (as it expands and contracts, during the heating and cooling phases).

If the resistor had been "good" (as in no faults), it would not have broken. I demonstrated, tapping another resistor with the dowel and was rewarded with a "ringing" sound. I tapped another, and it exploded into pieces.

We found four more bad power resistors by tapping them. When they are cold, testing these resistors will normally give a good reading on a multimeter. But as they heat up and expand, the break in the wire will cause the resistor to "go sporadic" and cause all kinds of mayhem.

We completed repairs and the 20 KW transmitter came back to life and worked dependably for many years. My friend added a wooden dowel to his tool kit.

JIM SWARTOS started working as an Electronics Technician in 1968 and has been in the broadcasting field most of his career along with a stint teaching electronics. He now works full time for a city in its Communications Shop.

Back to table of contents

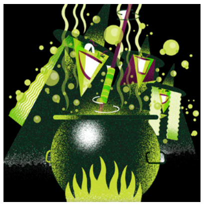

7. A Spell to Defeat a Power-Sucking, Battery-Killing Firmware Flaw

Something funny in the firmware was causing sensors in a home security system to communicate at a frenetic pace.

By AMIR MARANDI, Specialist, Senior Design Engineer; Lighting Division, Eaton

My group was working on a ZigBee-based home security system when some testers began reporting that their door/window sensors had quit working.

The sensors were battery operated. I checked the returned sensors and discovered the batteries were all dead. Not good, since the battery should have lasted for at least a few years.

I replaced the batteries and returned the sensors to the testers, but soon the same people were returning the sensors. It was puzzling, as only this group was affected--not all the sensors were failing in the field.

I put a setup together to measure the battery output current with a small series resistor. I connected it to a Tektronix DSO scope to display the voltage signal across the resistor, which was directly related to the current passing through it and consumed by the sensor. The scope had a built-in function to measure the area under the curve and by dividing that value by the series resistor value, I was able to calculate the energy consumption in mAs.

I also had a ZigBee sniffer to monitor the communications between the sensor and controller.

I quickly found out that the bad sensors were doing more communications than the good sensors. That made a certain kind of sense, as the more communications, the faster the battery drain. Looking at the ZigBee communications log, I discovered that the bad sensors were sending an "I am alive" message every 10 seconds instead of every 5 minutes. I checked the firmware versions of a good and bad sensor; both had the latest update. It did not make any sense. Checking the source code did not reveal anything.

Then someone brought in another dead sensor to be checked. After replacing the battery, I noticed that it behaved exactly like the other bad ones, the only difference was that it had an older version of the firmware.

I surmised that the sensor was probably dead before the firmware upgrade happened. So I upgraded its firmware, but that did not fix the problem.

I did a factory reset to the initial settings, and the bad unit now behaved like a good unit. Most of the sensors had their labels on them and turned out that all the bad unit were built prior to the date that the new firmware was released. Our testers had a mixture of sensors, some built before the release of the new firmware and some after.

The sensors with the old firmware had been upgraded automatically to the new firmware, but the problem was that the upgrade was not correcting the 10 seconds interval to 5 minutes.

Only a factory reset would do that. After instructing the tester to do a factory reset and repeat the pairing, the problem was fixed.

Since I was able to calculate energy consumption by the sensor, I went through all the possible scenarios and captured energy used by different events. With some assumptions of how many times a day the sensor was triggered, I was able to calculate energy usage per day and calculate the battery lifetime.

AMIR MARANDI is an experienced lead hardware design engineer for complex electronic designs.

Back to table of contents

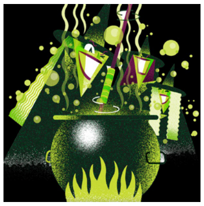

8. A Spell to Vanquish Intermittent Interruptions in a Data Stream

An engineer turns to prayer when all else fails while troubleshooting a faulty auction clock.

By JOHANNES VAN KEMPEN

About fifty years ago, in a time well before desktop computers, digital interfacing was done with relays and transistors.

Oscilloscopes had two or four carrying handles and the multi/AVO meter was the main tool for troubleshooting. A computer hard-drive was the size of a washing-machine. Data input and storage was usually on punch-cards.

At the time, I did service and maintenance on auction installations in western Europe. The auction-clocks were interfaced with relays and transistors to a punch-card machine. This was usually a very reliable process but not on one particular day there were intermittent interruptions in the data stream.

The fault was located in the interface box and all questionable relays were replaced by re-conditioned, cleaned adjusted and tested ones. All transistors and diodes were looking okay with the ohm-meter. The intermittent fault was still there.

The transistors, diodes etc. were mounted on a Vero-board with a blueprint on top. I re-soldered the entire board and checked wiring and connectors. The intermittent fault was still there. Really spooky! After many hours, alone and at night, I had run out of ideas.

So I said a prayer and took the box out again.

Carefully checking everything on the board again, I bumped a transistor, which promptly bent over on two legs. Bingo!! It turned out that the developing liquid (ammonia) in the blueprint had eaten up the steel transistor leg, rusting it through the gold layer where the leg was bent in the board.

Needless to say, we did things differently after that

Now retired, Johannes van Kempen's engineering career spanned the gamut from working as an engineer at an instrument factory in the Netherlands to serving as a transmitter engineer with Trans World Radio in Swaziland.

Back to table of contents

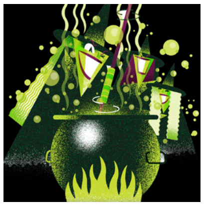

9. A Spell to Eradicate Erratic Sensor Readings

An engineer begins to question his sanity when he can't locate the source of intermittent errors that are getting worse.

By ROBERT CAMPBELL, Programmer and Digital Design Engineer

About 5 years ago I took over an IoT embedded sensor project that involved custom hardware sensors with a custom wireless bridging device.

We had a couple of potential customers that we were testing with, and everything was going great. All of a sudden, we started getting bad readings every once in a while. They were intermittent but kept getting worse.

We would get the devices back and test them, every time would find that everything was okay. It made no sense.

At some point while questioning my sanity I realized that most of the time the problem was happening during the night at the customer sites. I watched to pattern long enough to also realize it was happening when it was cold.

Everyone told me I was nuts. I eventually stuck a device in the office kitchen freezer and behold, the error was reproducible at last. (The rest of the team had to try it several times themselves before they believed it.)

The previous engineers had hand rolled their own bi-direction single wire bit bang interface. There was a timing bug where a line check took place almost exactly at the moment of the line being pulled down instead of waiting for a reasonable amount of time.

It worked when it was warm, but when the temperature dropped to around freezing or below the communication failed and all bits came back the same.

I almost left engineering due to that bug. I threw away the bit bang interface and wrote a software UART driver and never looked back.

BOB CAMPBELL is a self-taught engineer with 25 years programming and about 5 years of digital electronics design experience. He works primarily on IoT gadgets.

Back to table of contents

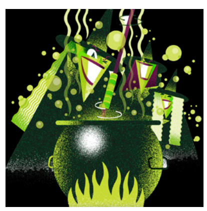

10. A Spell to Locate the Hidden Short

The live concert was about to start; would the young recording engineer find the short in the recording machine in time?

By MICHAEL SIRKIS, Engineer at Radio Systems Corp.

In 1975 I was working for Broadcast Electronics (BE) of Silver Spring, MD, manufacturer of endless loop tape cartridge machines for radio broadcast.

At the time. BE was owned by Filmways Corporation. Filmways was the company that produced TV shows such as the Beverly Hillbillies. In addition to BE, Filmways also owned Wally Heider Recording, which had studios in Los Hollywood and San Francisco.

Although well respected for its studio work, Heider was known as the go-to studio for remote live recording. Whenever Heider was doing a live recording near the Washington, D.C. metro area, tickets and backstage passes were offered to a few BE employees.

I was fortunate enough to be on the short list--as was my friend Tom, who also worked in the engineering department at BE. Tom and I attended several concerts as the guests of Heider Recording.

The remote recording studio was

either constructed in a backstage room at the venue or in a large truck just outside. Packing and shipping a recording console, four multitrack tape machines, microphones, cables, monitor speakers, amplifiers and a mountain of tape for construction into a recording studio 3,000 miles away is quite a task.

One evenng Tom and I arrived for a concert but first went to the remote studio--built in a truck for this particular event--to say “Hello” to the guys from Heider. Instead of the usual laid-back atmosphere, it was a scene fraught with incredible tension as one of the multitrack tape machines was not powering up.

Normally, the four tape machines are operated as pairs in ping-pong mode. Recording is started on the first pair, and when the reel of tape on the first pair is halfway through the second pair is started. The inputs on pair one and pair two are wired in parallel.

In addition, one track on each machine is dedicated to control, allowing the machines to be synced together at a later time during the mix. During the live recording there would be no mixing, instead the goal is to use as many pickup and microphones as the tape machines would allow.

Losing one tape machine would render the same effect as losing two machines. The number of inputs would have to be reduced to fit onto only one machine.

As the plan for reducing and rewiring the inputs was being discussed, I asked if a service manual was available – it wasn’t. In addition, with the exception of fuses there were no spare parts, boards or modules.

The only test equipment available was a Simpson 260 VOM. The head technician (from Heider) had traced the issue to the power supply by comparing the dead machine to a working one. The power supply was a multl-voltage output supply, and one of the outputs was at zero volts.

With nothing to lose and with some linear supply design experience under my belt, I said, “I don’t think you have a dead supply.” This brought a cold stare from the tech who was several years my senior (I was 22 at the time).

I quickly explained foldback current protection, and that I believed the problem was a short. The supply was simply protecting itself.

The supply output in question was feeding several wires that snaked their way through the machine. I asked the tech to remove one wire at a time and to cycle the power to the machine.

When the fourth wire was removed, the power supply started outputting the correct voltage. The first three wires were reattached, and the power supply continued to operate normally.

Pulling gently on the fourth wire and following it through the wiring harness brought us to a control panel that was full of illuminated switches. These switches were used to select which tracks were in record mode and which tracks were locked out.

The control panel switches were wired point-to-point to a harness that controlled the machine. I didn’t see any broken wires or shorts. I then remembered that sometimes an incandescent bulb will fail and short. After removing a number of bulbs, we found the culprit.

The wire was reattached to the power supply, the control panel was reassembled, and the machine powered up normally. All four machines were used to record the concert.

At the end of the concert, Tom and I were asked to go backstage for a few minutes. That is when we met the members of Earth Wind and Fire, who were elated that the concert had been successfully recorded.

Michael Sirkis has been developing products for the radio broadcast market for over 40 years. He has worked for Broadcast Electronics, Ampro/Scully, Fidelipac, and Radio Systems.

Back to table of contents