CIRCUIT CELLAR®
THE MAGAZINE FOR COMPUTER APPLICATIONS

FEATURE ARTICLE

Norman Rogers

Killing the EMI Demon

Emissions regulations can make it expensive to get your microprocessor boards to the market. In an effort to help us reduce EMI testing costs, Norman reveals the EMI reduction tricks Rabbit Semiconductor developed for its second-generation processor.

Some government bureaucracies are fairly nice, even if you do have to wait for two hours to meet with them. Other bureaucracies are horrible. The Immigration and Naturalization Service comes to mind. To meet with them, you have to line up the night before outdoors in the cold. The FCC is probably closer to the INS than to the SSA. Why should you care about the FCC, assuming that you're an electronics engineer who doesn't own a radio station? The answer is found in the *Code of Federal Regulations* (Chapter 1, Part 15, Subpart B—Unintentional Radiators). [1] You can read the regulations online at www.access.gpo.gov/ecfr/. These are the regulations that say you have to prevent excessive radio emissions from digital devices, which usually means a microprocessor-based system.

If you read these regulations closely, you'll realize that they are really stupid. There's no nicer way to put it. For example, you're prohibited from operating a device unless you first test it for compliance with the regulations for unintentional emissions; however, it isn't feasible to develop a digital device without operating it, and it's illegal to operate it without first testing it. So, it's actually impossible to legally develop a microprocessor board.

In addition, the regulations blithely outlaw vast swaths of the cottage electronics industry. If you build a one-of-a-kind microprocessor board and sell it for, say, $500, you're in violation of the regulations unless you first spend $5000 or so to have it tested. Thus, small engineering consultantcies are granted the privilege of conducting their business in the same dignified manner as pirate taxi operators and undocumented farm workers.

This is the "we're not taking any garbage" school of regulation. Nobody gets off on a technicality because everyone is guilty if they do anything. However, if you take a closer look at the regulations, you'll discover that regulations apply to the little guys who don't have political pull. When important industries come calling, it's "Yes sir," "No sir," and "How can we be of assistance?" Home appliance makers and automobile manufacturers, to name two important industries, are exempt from the regulation in Part 15. Hmm, is that why my microwave oven interferes with my cordless telephone?

The laws of government regulatory evolution are at work here, favoring a situation where the powerful are unregulated, the weak are regulated, and the regulations generate a supportive constituency of people partnered up with the regulators. A huge testing industry has grown up around the FCC's Part 15 regulations. The industry is comprised of the test labs that are charging you $500 per hour. These people are the FCC's loyal courtiers, and they are a dependable lobby in favor of more regulation.

The FCC's method for measuring the level of potential interference is based on the use of a quasi-peak filter, which was invented in the 1930s to rate interference for AM radio. There's actually no reason to suppose that the quasi-peak filter is relevant to modern communication technology; it probably would be relevant for AM radio, but the FCC doesn't regulate emissions in the AM band (unless they're over 30 MHz).

The FCC recognizes class A devices that are intended for use in offices and class B devices intended for use in homes. The regulation is strictest for class B, specifying a signal level that's low compared to television signal levels. But as cable and satellite television spread to more than 80% of U.S. homes, the FCC's testing procedure becomes more and more irrelevant.

Does the FCC really care about interference? Does it care if its regulations actually work to prevent interference? I believe the FCC cares that people think the regulations work. If the FCC really cared, they could have done a lot more to make the regulations less of a burden and more cost effective. Of course, the FCC's friends in the test labs would not appreciate such cost effectiveness. Our cost is their revenue.

A simple regulatory change would be to officially recognize that one device or 1000 devices don't deserve the same heavy-handed regulation as devices produced by the millions. There is no reason at all to require the testing of low-production devices. The cost of regulation is burdensome, and low-production devices will never create much of an EMI problem because they are used in such low numbers. There are catch-all regulations and laws that outlaw devices that create interference problems that would still apply for the rare extreme case.

Another simple change would be to borrow an idea used in building codes. Every pipe or beam doesn't have to be tested before it can be used. The building code permits the use of components that meet certain specifications. For example, pipe made of a certain material with a certain wall thickness may be approved for the distribution of water. The same technique could be used for the regulation of unintentional radiation from computer systems. For instance, devices with a ground plane and a clock speed of less than 30 MHz, as well as meeting certain other specifications, might be exempted.

This approach could be a lot cheaper for regulation when low production volumes are involved. Would this approach be perfect? Of course not, but neither is the current approach, which involves many inaccuracies, dubious assumptions, and measurement compromises.

I'll put forth another suggestion for rationalizing the regulations. How about allocating spectrum for EMI? Most EMI comes from microprocessor clocks, so it's parked in a very narrow band around the harmonics of the clock. The harmonics of 6 MHz represent an interesting series because the harmonics fall smack in-between all the VHF TV channels, except for channels 5 and 6 (72 to 88 MHz). The harmonics at 90, 96, 102, and 108 MHz fall exactly between FM stations. If a clock frequency of 18 MHz is selected, the harmonics miss channels 5 and 6 as well as all the other VHF TV and FM stations (if you divide 6 MHz by 13 you get a clock frequency that's good for generating data rates close to standard data rates).

Under a plan like this, devices using a clock that is a harmonic of 6 MHz would be regulated only on certain harmonics because most of the harmonics won't interfere with anything. Over a period of time, services conflicting with the harmonics of 6 MHz could be moved to other frequencies and the regulations could be harmonized with radio spectrum allocation in other countries. Life could get simpler. Digital devices would be cheaper and there would be less EMI. Are any bureaucrats listening?

Can small companies, being more numerous than big companies, influence the regulations? The answer is, not much. There aren't trade organizations representing small digital electronics companies, and without professional representation in Washington, it's unlikely that their voices will be heard. It may be that the people writing these regulations are not actually aware that small companies exist. Lobbyists for small companies never come to visit.

THE STRUGGLE WITH EMI

When Rabbit Semiconductor introduced its first microprocessor, the Rabbit 2000, the company thought that the customers wouldn't have too much trouble with EMI. The maximum clock speed was fairly low at 30 MHz, and Rabbit gave suggestions in the design manual on how to avoid EMI problems. But, it turned out that the customers were mostly smaller companies with limited expertise in this fairly exotic area. They were focused on the innovative products that they were designing, not gobbledygook government regulations.

Many of the customers took their product along with $5000 to a test lab, but they subsequently found themselves going back to the drawing board. Executives at Rabbit decided to do everything they could to make it impossible for the next-generation processor, the Rabbit 3000, to flunk the radiated emission tests. In the process, they even found some solutions that could be compatibly retrofitted to the first-generation micro-

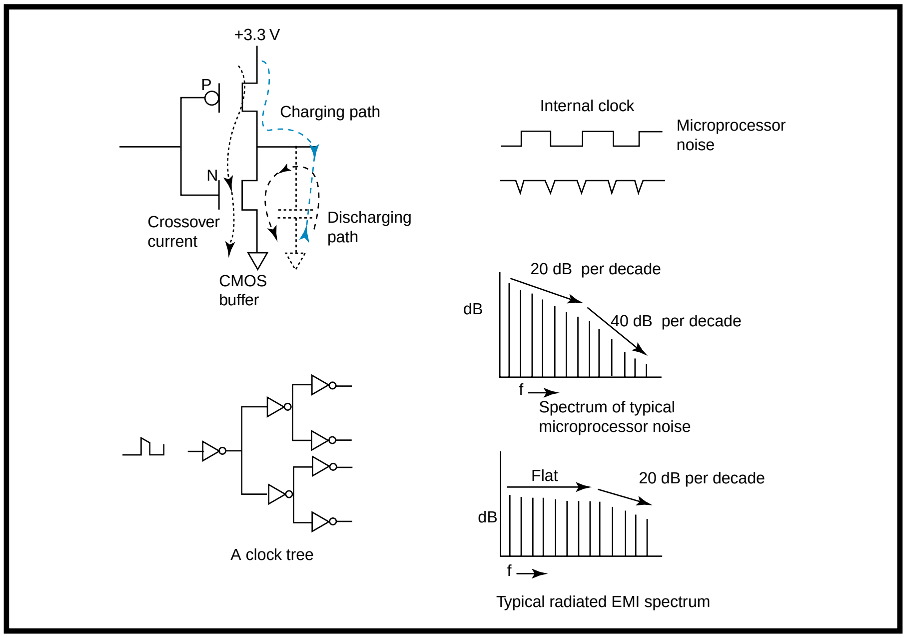

Figure 1—*Here you can see the sources of EMI in a typical microprocessor and the resulting spectrum.*

processor to make that device much better from the EMI perspective. The cumulative effect of the improvements turned out to be so effective that Rabbit met the goal to make it nearly impossible to flunk the government EMI tests with a Rabbit 3000-based product, even at 50-MHz clock speeds.

SOURCES OF EMI

EMI on a typical microprocessor board is related to the clock. If the clock is a square wave, it contains frequencies at the clock frequency and harmonics. A perfect square wave clock would have harmonic frequencies at f, 3 × f, 5 × f, 7 × f, and so on. For a perfect square wave, or any string of pulses with a fast rise time, the strength of the harmonics declines inversely with frequency. So, the eleventh harmonic would be one-eleventh as strong as the fundamental frequency. This corresponds to a decline in harmonic amplitude of 20 dB per decade.

Real time clocks are not perfect square waves, and pulses do not have infinitely fast rise times. As a result, the higher harmonics of any real waveform start dropping faster than 1/n at higher frequencies, generally dropping as $1/(n^2)$, or 40 dB per decade, after the frequency is high enough. You can see this in Figure 1. The antenna efficiency of PC board structures or cables increases 20 dB per decade as frequency increases and wavelength gets shorter and closer to the size of structures found on typical PC boards. As a result, the beginning part of the radiated spectrum tends to be uniform, the 20 dB per decade decline in harmonic strength being balanced by the 20 dB per decade increase in antenna efficiency, until a high enough frequency is reached where the curve takes a bend and harmonics start declining at 40 dB per decade zone (see Figure 1). Above this frequency, the radiated spectrum starts declining by 20 dB per octave. But, the amplitudes of the real harmonics of a real device are often quite irregular because of resonances that weaken some and reinforce others.

What is not usually understood is that the biggest source of EMI is not the clock directly, but a train of pulses generated on both edges of the clock when current surges into the microprocessor for a nanosecond or two when the clock transitions up or down. This pulse train has a frequency that's double the clock frequency. It seeps out of the processor chip into the power supplies and generally infects the board with high-frequency EMI. It also gets into the output lines emanating from the processor package; therefore, it's further spread around the board and to cables and devices connected to the board.

The current surges on both clock edges are related to the clock tree. The clock tree is a system consisting of a branching network of buffers that distribute the internal clock around the silicon die. Because these buffers drive considerable capacitance and have both polarities of the clock present, there is a surge of current on both edges of the clock. This occurs as current flows into the chip to charge up the capacitance in the part of the clock tree that is transitioning from 0 V to the power supply voltage. On-chip devices, such as flip-flops, also contain internal gates and buffers where both polarities of the clock are present and contribute to the current surge.

An additional current surge is related to the crossover current when both the N and P transistors in a CMOS buffer are momentarily conducting during a logic transition. The silicon chip tries to suck in the required current to service these fast transients through its power supply pins. However, these connections have inductance created by the bond wires and lead frame, so the voltage drops briefly on the die, creating an on-chip power supply voltage drop with an amplitude on the order of a few tenths of a volt and the duration of a nanosecond or so.

If this same on-chip power supply drives the output buffers that carry signal lines out of the chip, these lines will also be infected with the fast pulses present in the power and ground supplies. This is because the power supply noise is directly transmitted through the buffer power inputs to the output lines. The on-chip current surges create fast noise that passes out through the power supply pins to the power and ground planes on the PC board, further spreading the infection.

The amplitude of the harmonics of the periodic noise pulses, at least at lower frequencies, declines inversely with frequency (1/f). Unfortunately, the effectiveness of a short antenna, such as a PC board trace, increases directly with frequency (~f). The result is that the radiated EMI tends to be flat across the spectrum. Fortunately, the amplitude of the harmonics starts declining more rapidly than 1/f; it's more like $1/(f^2)$ at some higher frequency determined by the finite rise time of the pulses in the pulse train. The balance of these countervailing effects is such that the most trouble is often found in the area of 100 to

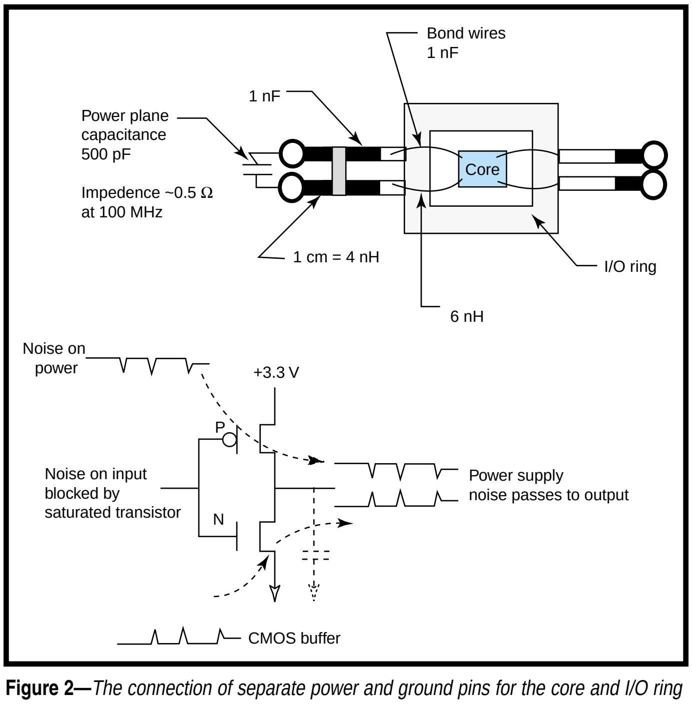

Figure 2—*The connection of separate power and ground pins for the core and I/O ring of a processor is shown here. A PC board filter blocks core noise from power planes. You can also see how I/O buffers spread power supply noise.*

300 MHz for lower-speed 8- and 16-bit microprocessor boards.

Decoupling capacitors and the intrinsic capacitance of the power and ground planes can be used to short circuit or filter noise on the power supply. However, this technique loses effectiveness above 100 MHz, because the decoupling capacitors have inductance of about 1 nH, giving an effective resistance of about 0.5 Ω at 200 MHz. The large currents involved will develop millivolt-level voltages across such capacitors.

REDUCTION TRICK #1

A key feature used in the Rabbit 3000 to mitigate the problem of noise on the I/O lines is the provision of two sets of power supply pins. One set is used for the processor core; the other is for the output drivers that are located in the I/O ring on the periphery of the die (see Figure 2).

If the I/O buffers are supplied with the same power that is made dirty by the fast transients in the processor core, every output pin of the processor will spread EMI. The EMI that tries to come out of the power pins for the core can be blocked by a combination of decoupling capacitors and PC board trace inductance. This keeps the PC board power planes a relatively clean source of power for the processor I/O ring. The design team figured this feature decreases EMI amplitudes by 10 dB, which is a factor of three in EMI electrical field strength measured by the prescribed calibrated antenna. This is a lot because it's common to flunk the tests by 5 dB.

REDUCTION TRICK #2

Most microprocessors have I/O and memory devices connected to the same bus with distinct control signals for the devices. Generally, there is a lot more activity at a higher frequency for the memory devices. The Rabbit 3000 has an option to use separate pins for memory and I/O devices, both address and data. The advantage is that the physical scope of the high-speed memory bus is limited to the memory devices. A separate address and data bus handles I/O cycles and has a much lower average operating frequency. In particular, the address lines toggle only during I/O bus cycles, greatly limiting the emissions from the I/O bus. This avoids the situation where the fast-toggling address and data lines of the memory bus have to be run all over the printed circuit board of a large system. This scheme also limits the capacitive loading on the memory bus, which does not have to extend to numerous I/O devices.

REDUCTION TRICK #3

A line spectrum is the spectrum generated by a square wave clock or by a train of short pulses. All of the energy is concentrated in a narrow spectral line at the harmonic frequencies.

When the FCC EMI measurement tests are performed, the spectrum analyzer measures the amplitude of the signal from a 120-kHz wide filter that is swept across the frequencies of interest. With a line spectrum, all of the energy in a single line passes through the filter, resulting in a strong signal. If the energy in the line could be spread out over a wider frequency, say 5 MHz, only one-fortieth the energy would pass through the 120-kHz wide filter, considerably reducing the reading (by 16 dB in amplitude for one-fortieth of the energy). This is what a clock spectrum spreader does. It modulates the clock frequency by a little so as to smear out the spectral line in frequency. The idea to do this for the purpose of reducing EMI was patented by Bell Labs in two patents during the 1960s.

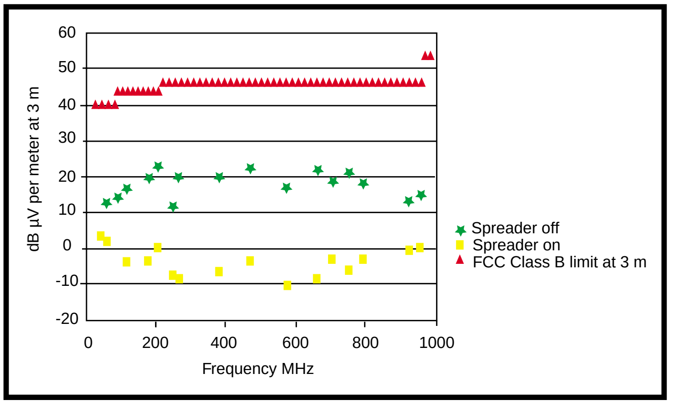

Figure 3—*Take a look at the actual EMI measurements for a Rabbit 3000 RCM3010 core module board with clock spectrum on and off. Peak values are shown as dots. Units are in amplitude dB relative to 1 µV per meter at a 3-m distance.*

There are numerous ways to modulate the clock frequency. One method is to use a voltage-controlled oscillator and phase-lock loop so that the frequency sweeps back and forth at a low modulation rate (e.g., 50 kHz). Another method is to insert random delays or dithers into the clock. These methods are all covered in the original Bell Labs patents. The Bell Labs people were probably interested in EMI because telephone switches involve a large amount of equipment in a small space. In addition, it's conceivable that the early computerized switches suffered from EMI problems.

We installed a clock spectrum spreader in the Rabbit 3000 based on a combination of digital and analog techniques. The spectrum spreader reduces FCC-style EMI readings by around 20 dB, which is a lot.

A control system makes sure that the modulated clock edge is never in error by more than 20 ns compared to where the clock edge would be if it were not modulated. This prevents disruption in serial communications or other timing functions. For example, a UART operating at 460,000 bps can tolerate about 500 ns of clock edge error before it will be near to generating errors. This is far less than our 20-ns worst error in clock edge position.

THE CONTROVERSY

Perhaps it's because the spectrum spreader solves EMI problems (it gives an advantage of 10 to 20 dB) that it's not well accepted in some quarters. Those in the anti-spectrum spreader crowd argue that the use of spectrum spreaders is a form of cheating because the radio frequency energy is still present but just spread differently in the spectrum. It's assumed that changing the spectrum from narrow band to broadband will simply spread the interference around or even make the total interference greater.

People holding competing points of view have performed experiments in an effort to demonstrate that spectrum spreaders either increase or decrease interference. But what both sides have failed to realize is that there are different types of spectrum spreaders that have different effects on different forms of communications signals. It's important to realize that traditional radio communication is being gradually displaced by new techniques rooted in digital technology.

The newer digital techniques are generally broadband and transmit redundant information. For example, frequency hopping is a technique in which the transmitter and receiver hop from one frequency band to the next in a predetermined pattern. Interference on one or a few bands will not have a noticeable effect because the missing data will be retransmitted or recovered by the use of error correction techniques.

Another technique is Orthogonal Frequency Division Multiplexing (OFDM). This is used for digital radio and television, particularly in Europe. OFDM is also used for DSL Internet service over copper telephone wires. With this technique a wide band is divided into hundreds or thousands of narrow band channels, each of which carries a slow data signal. The bands are typically 1- to 10-kHz wide.

All of the data signals are combined and error corrected to create a fast datastream. Because as much as half of the transmitted information is redundant, this technique can suffer considerable interference and retain perfect transmission of the picture or sound. A big advantage is that it's resistant to fading and multipath interference. OFDM is made possible by cheap digital signal processing that is now available via ASIC's. The digital technology makes possible frequency analysis via techniques such as the Fast Fourier transform. In addition, the digital compression of sound and speech is achievable via techniques such as MPEG.

Most spectrum spreaders modulate the clock by sweeping the frequency or phase back and forth in a regular pattern. The repletion rate is typically in the range of 20 to 100 kHz. However, it's also possible to modulate the frequency or phase by a random signal or a pseudo-random signal that does not repeat, or if it does, it has a long repetition period. If the modulation is periodic so that the clock pattern periodically repeats, then the spectrum will be split into separate spectral lines separated by the repetition frequency. If the modulation is random or has sufficient random noise in the circuit, then the spectrum will be smeared in a continuous fashion. This can make a difference.

In an OFDM TV system, for example, if a TV signal is transmitted in 8000 bands that are separated by 1 kHz and the spread spectrum clock splits the original clock into separate spectral lines that are separated by 100 kHz, then only one-one hundredth or perhaps one-fiftieth of the channels in the OFDM will be interfered with. This is a degree of interference that can be easily handled by the error correction facility. However, if the spectrum is continuously smeared, then it's conceivable that every channel would experience interference and the TV picture would be lost.

The repetition frequency of the modulator can also have an effect on voice transmission. If the repetition frequency is 5 kHz and a harmonic of the clock falls in the traditional FM band, then a 5-kHz whistle will be heard in an FM receiver (this is something the ear is sensitive to). If the repetition frequency is 50 kHz, then the whistle may still be there, but it will be beyond the range of human hearing.

MORE REDUCTION TRICKS

There were other tricks that helped to reduce EMI on the Rabbit 3000. For instance, some of the internal clocks are gated, so they're only enabled when needed. The use of gated clocks reduces the amplitude of the current surge into the chip.

The external processor bus is designed so that it isn't required to run the clock around the PC board. The clock is available at a dedicated pin, but in most systems it is turned off because it's not needed. When it's needed (e.g., to provide a clock for an FPGA), it can be supplied at full or half the internal clock frequency.

Additionally, the external processor bus cycles are not all the same length. This breaks up the periodicity associated with the external bus.

An internal clock doubler allows the external crystal oscillator to operate at half frequency. This reduces EMI by lowering the frequency of the external clock, which is physically larger with large currents and thus in great danger of generating EMI.

THE RESULTS

Figure 3 shows the EMI measurements on a Rabbit 3000 RCM3010 core module, which is a small microprocessor board operating at 29.5 MHz. The EMI is virtually undetectable; it's far below the noise floor of the spectrum analyzer as used in the normal FCC-mandated measurements. We were able to use a special measurement technique to ascertain where it was and plot it on the graph.

Norman Rogers is the President of Z-World, Inc. and Rabbit Semiconductor, which is a division of Z-World. He holds a BA in physics from the University of California at Berkeley and an M.S. in physics from the University of Hawaii. His technical interests include C compilers and microprocessor architecture. Norman may be reached at nrogers@zworld.com.

RESOURCE

Rabbit Semiconductor, "PC Board Layout Suggestions for the 3000 Rabbit Microprocessor, TN221.

REFERENCE

[1] *Code of Federal Regulations*, Government Printing Office, Washington, D.C.

SOURCE

Rabbit 3000 RCM3010 core module
Rabbit Semiconductor
(530) 757-8400
www.rabbitsemiconductor.com