2020
ARRL/TAPR
Digital Communications Conference

Frequency Estimation: for Research, for Competition, for Fun!

by David Kazdan AD8Y
Case Western Reserve University
Case Amateur Radio Club W8EDU
HamSCI: The Amateur Radio Science-Citizen Investigation

Thanks to TAPR for interesting work since the 1980s, and to ARRL for getting me started with all of this in the 1960s

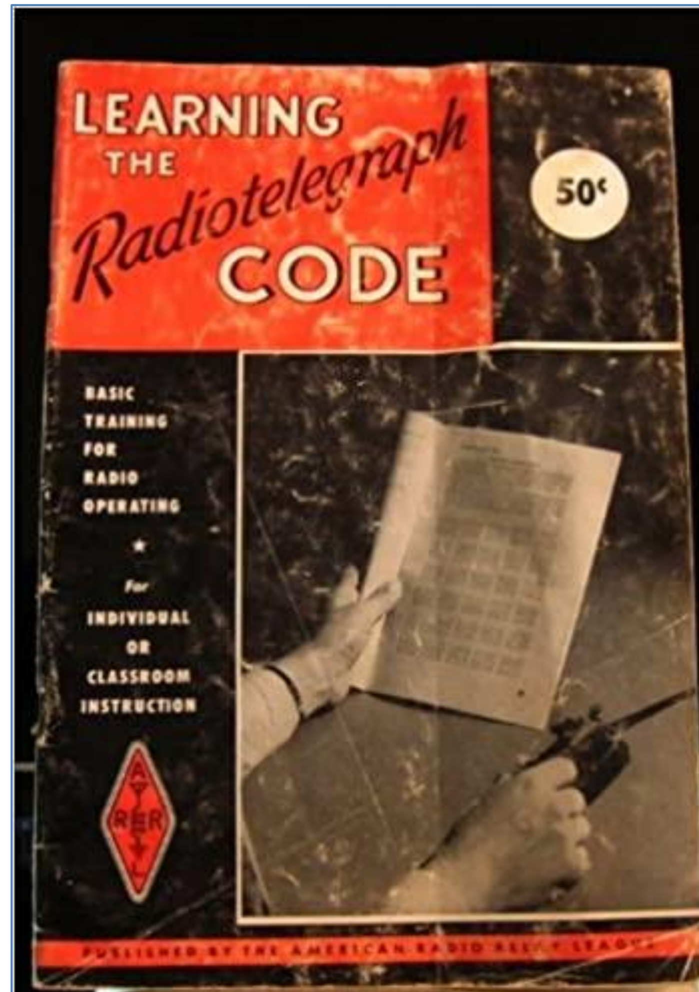

Fifty years since Novice license!
Wow!

And thanks to CARC and HamSCI for making it all interesting now

w8edu.wordpress.com

http://hamsci.org

And to NIST and WWV for making it possible. Standards matter.

National Institute of Standards and Technology
U.S. Department of Commerce

WWØWWV

Fort Collins, Colorado
WWV Centennial 1919-2019

Governmental & Military

CHU

means
National Research Council Canada Time Signal Shortwave Radio Station

by acronymsandslang.com

And to the other national physics laboratories around the world!

We all estimate frequency every day

Hearing a voice or a piece of music and recognizing it is an extraordinary piece of time-frequency analysis and frequency estimation in the presence of noise and interfering signals.

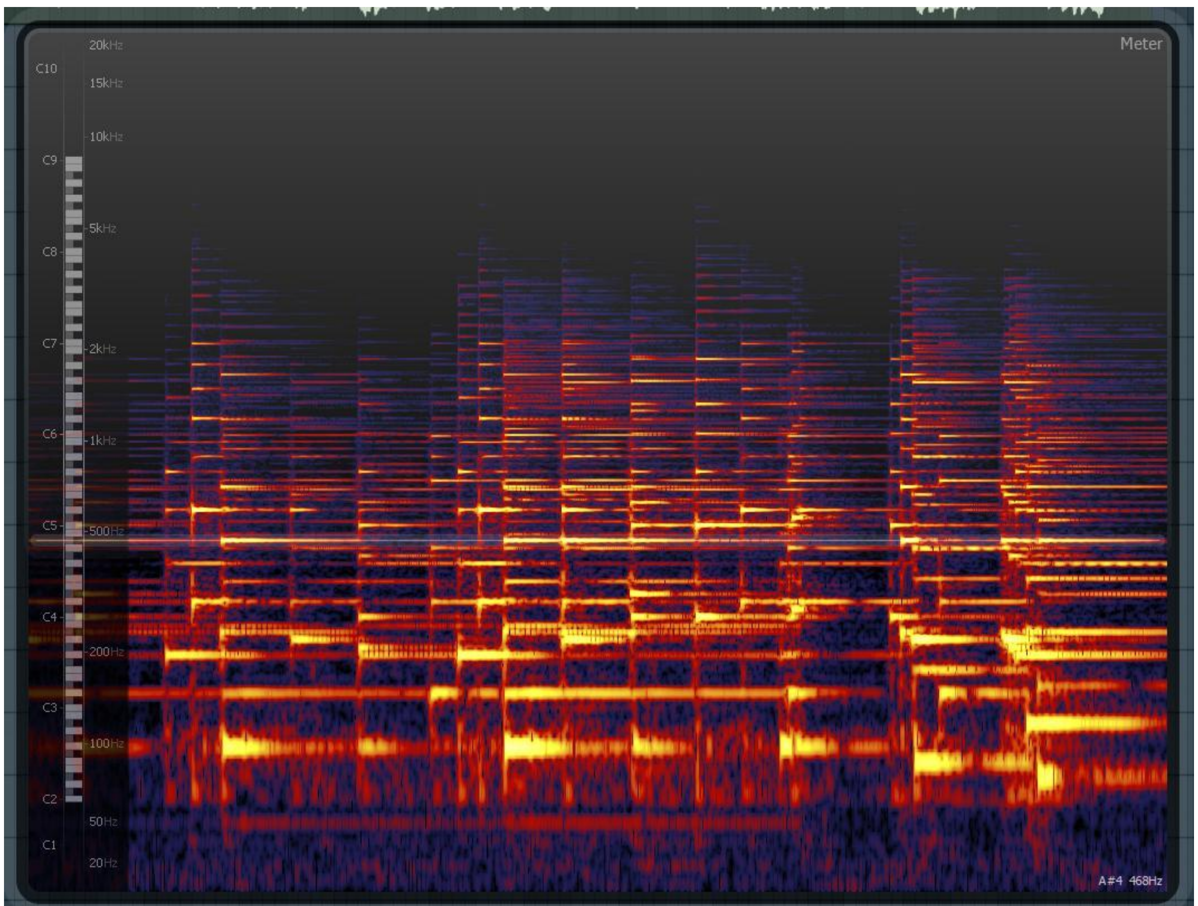

Psychoacoustics and vibratory sense are their own universes.
Not today.

Horizontal "waterfall" of piano music.

Picking one CW signal out the many on the 1970s Novice bands with a 1970s receiver is more akin to today's problem.

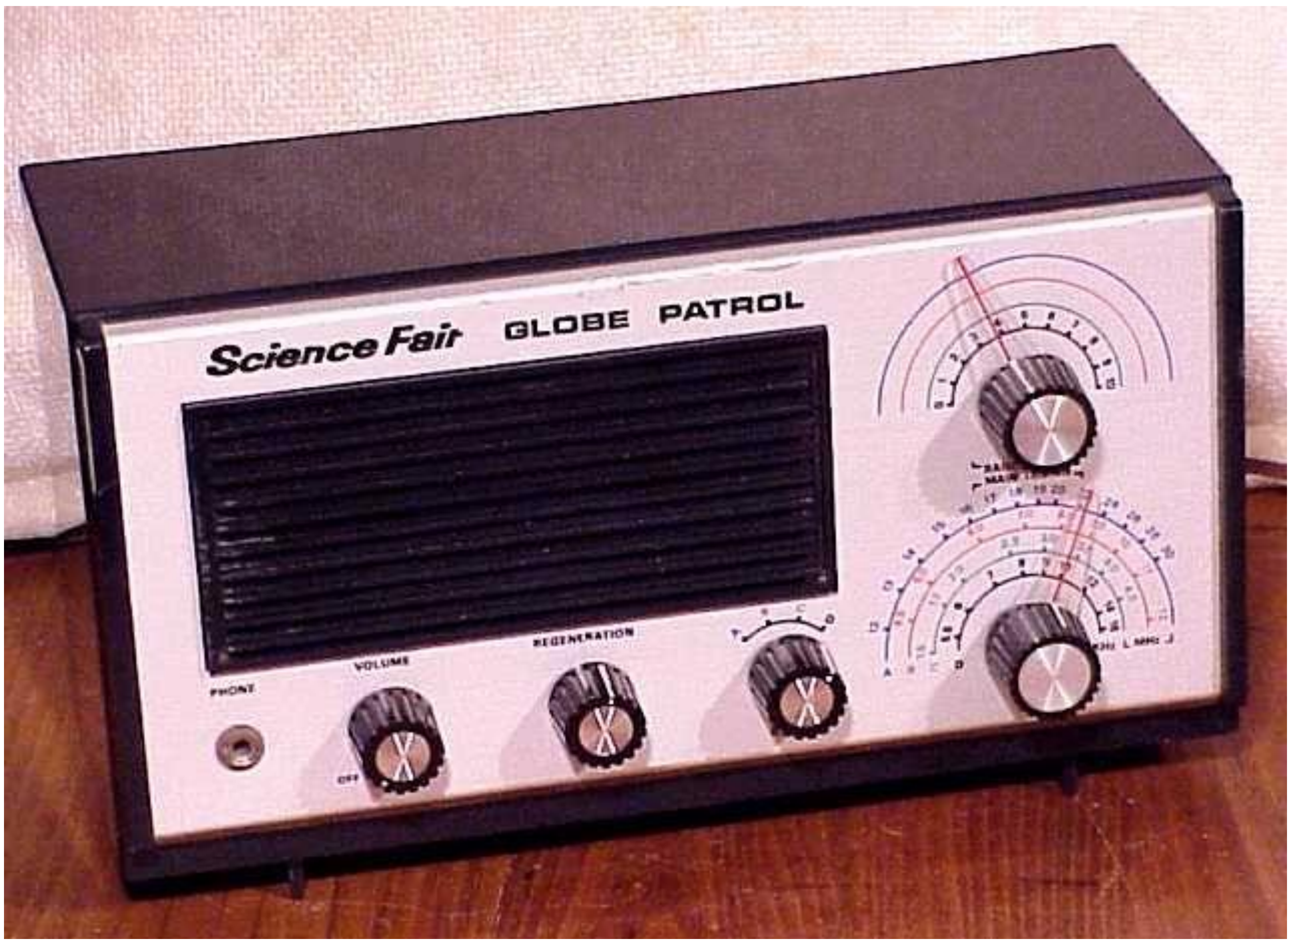

We've all been there.

Understanding a signal in noise is a statistical problem.

An Assertion:

If every college statistics class included a
Frequency Measuring Test,
students would pay attention.

The FMT. Around 120 enter each session.
The winner is crowned ***Nerdo di tutti i nerdi***.

We love it at W8EDU.

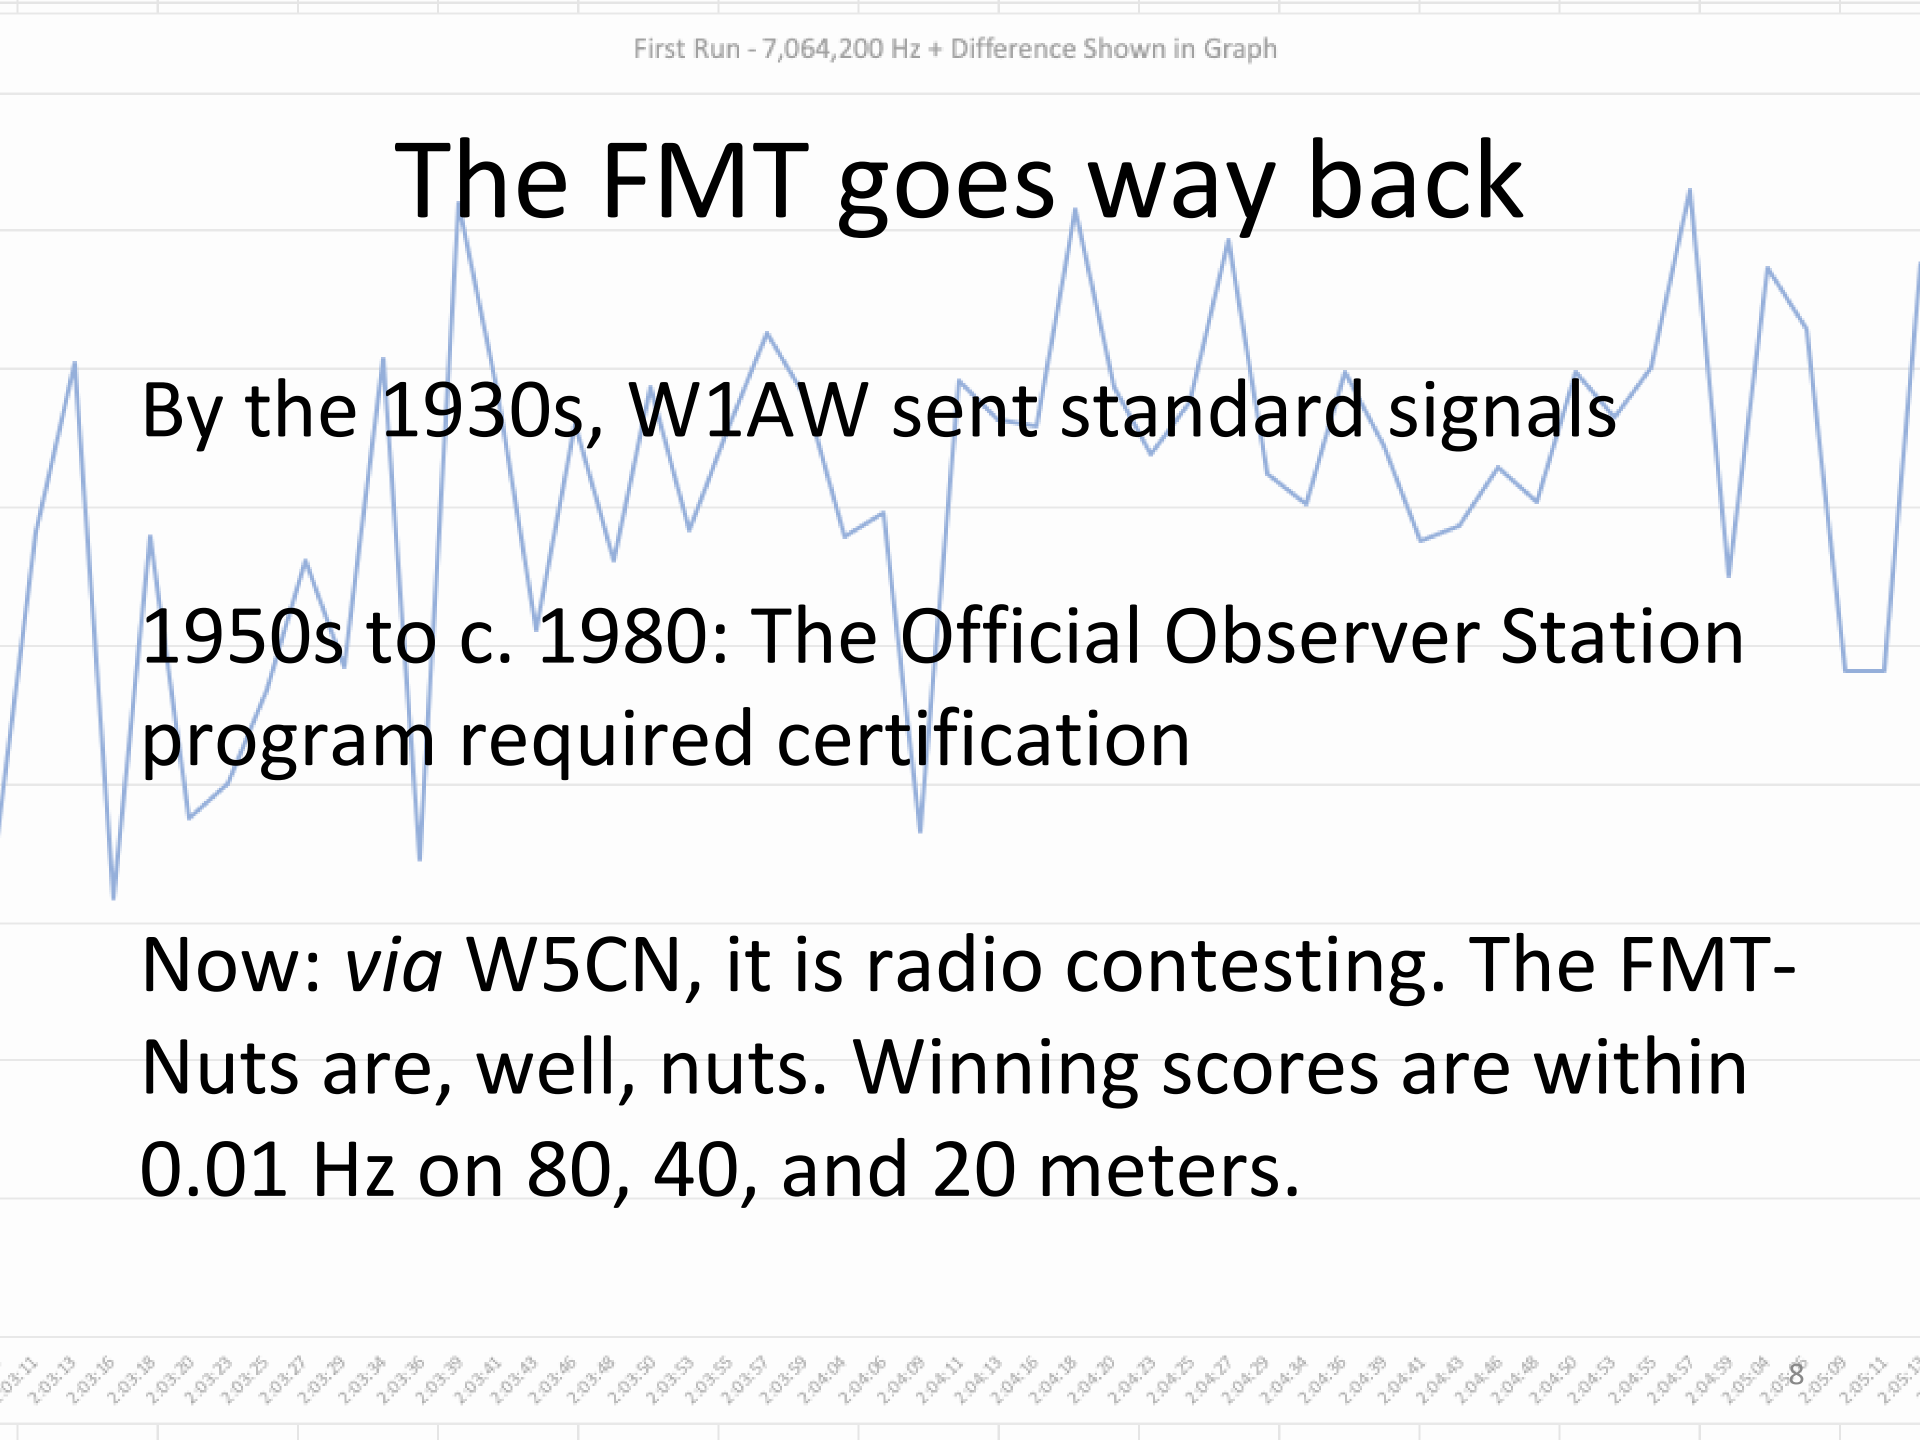

The FMT goes way back

By the 1930s, W1AW sent standard signals

1950s to c. 1980: The Official Observer Station program required certification

Now: *via* W5CN, it is radio contesting. The FMT-Nuts are, well, nuts. Winning scores are within 0.01 Hz on 80, 40, and 20 meters.

The only apparent FMT rule: Keep the referee anonymous

This helps prevent unfortunate incidents.

FMT is fun, but for HamSCI, frequency estimation has real significance

- Beacon signals are launched with very high accuracy and stability. WWV, CHU, MIKES, the rest of them are all very accurate.
- Bounced off the ionosphere, they undergo linear and nonlinear distortions, Doppler shifts, multipaths, etc. and get noisy.
- HamSCI wants to read the ionosphere's EKG: Given received frequency changes, what's going on 200 km above the earth?
- We need fast, automated measurements with high resolution. Power! Statistical power!

Estimating frequency doesn't sound that hard:

Given a signal of single frequency and constant amplitude,

Measure its frequency! Set a timer and count squiggles! How hard can it be?

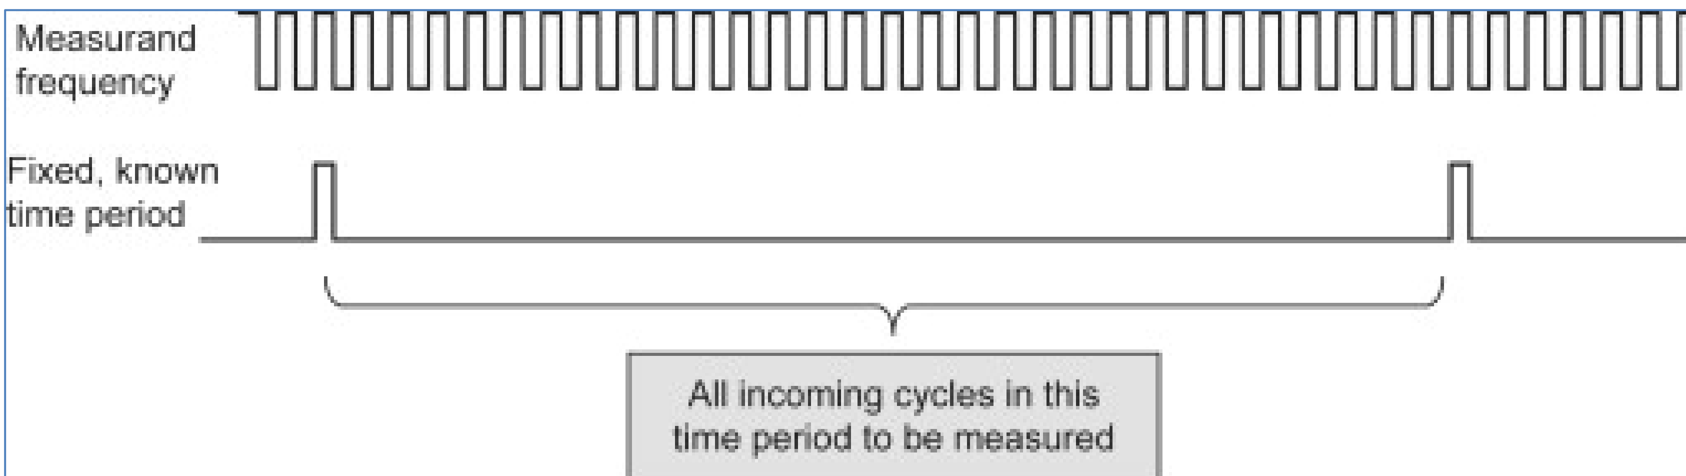

Estimating frequency doesn’t sound that hard. It can be.

Given a signal of single frequency and constant amplitude

and

Given that the signal is contaminated with noise,

What is the best statistical estimator of the signal’s frequency?

Earliest methods measured wavelength as surrogate for frequency

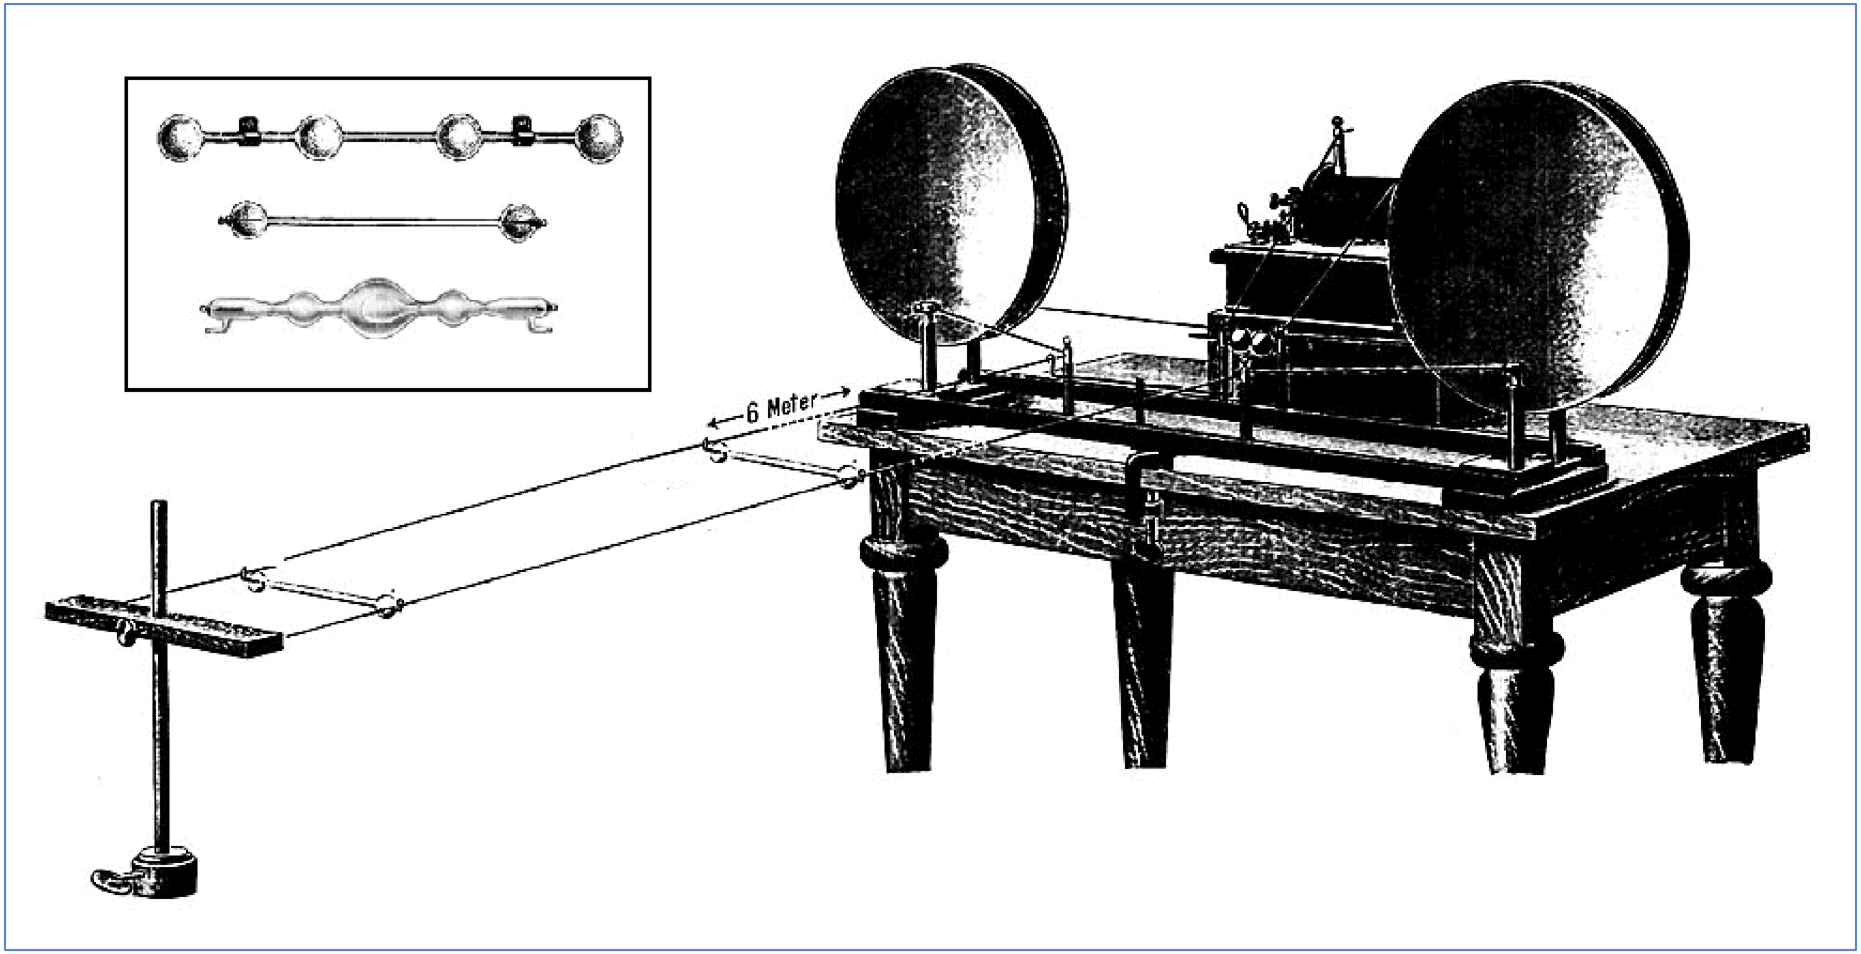

Lecher line: High SWR by design! Light bulb is used to probe for voltage nodes.
It's the reason we refer to bands by wavelength.

Tuned LC circuits as "absorption wavemeters," were next

Heterodyning receivers followed. The Army Signal Corps BC-221 Frequency Meter

Frequency Meter Set SCR-211-(*) in a Wooden Cabinet

Used to be used a lot for FMTs. Could get within about 10 Hz on a good day.

With modern signal analysis, how do we begin?

We have oscilloscopes, spectrum analyzers, sampling systems, and computers that the Signal Corps didn't have.

We have atomic clocks and ways to disseminate atomic clock level of accuracy (GPSDOs)—Signal Corps only had crystals

(Crystals helped, though: book here)

Noisy CW. Blue is the received signal. Red is filtered. Can you measure the frequency? Accurately? Quickly?

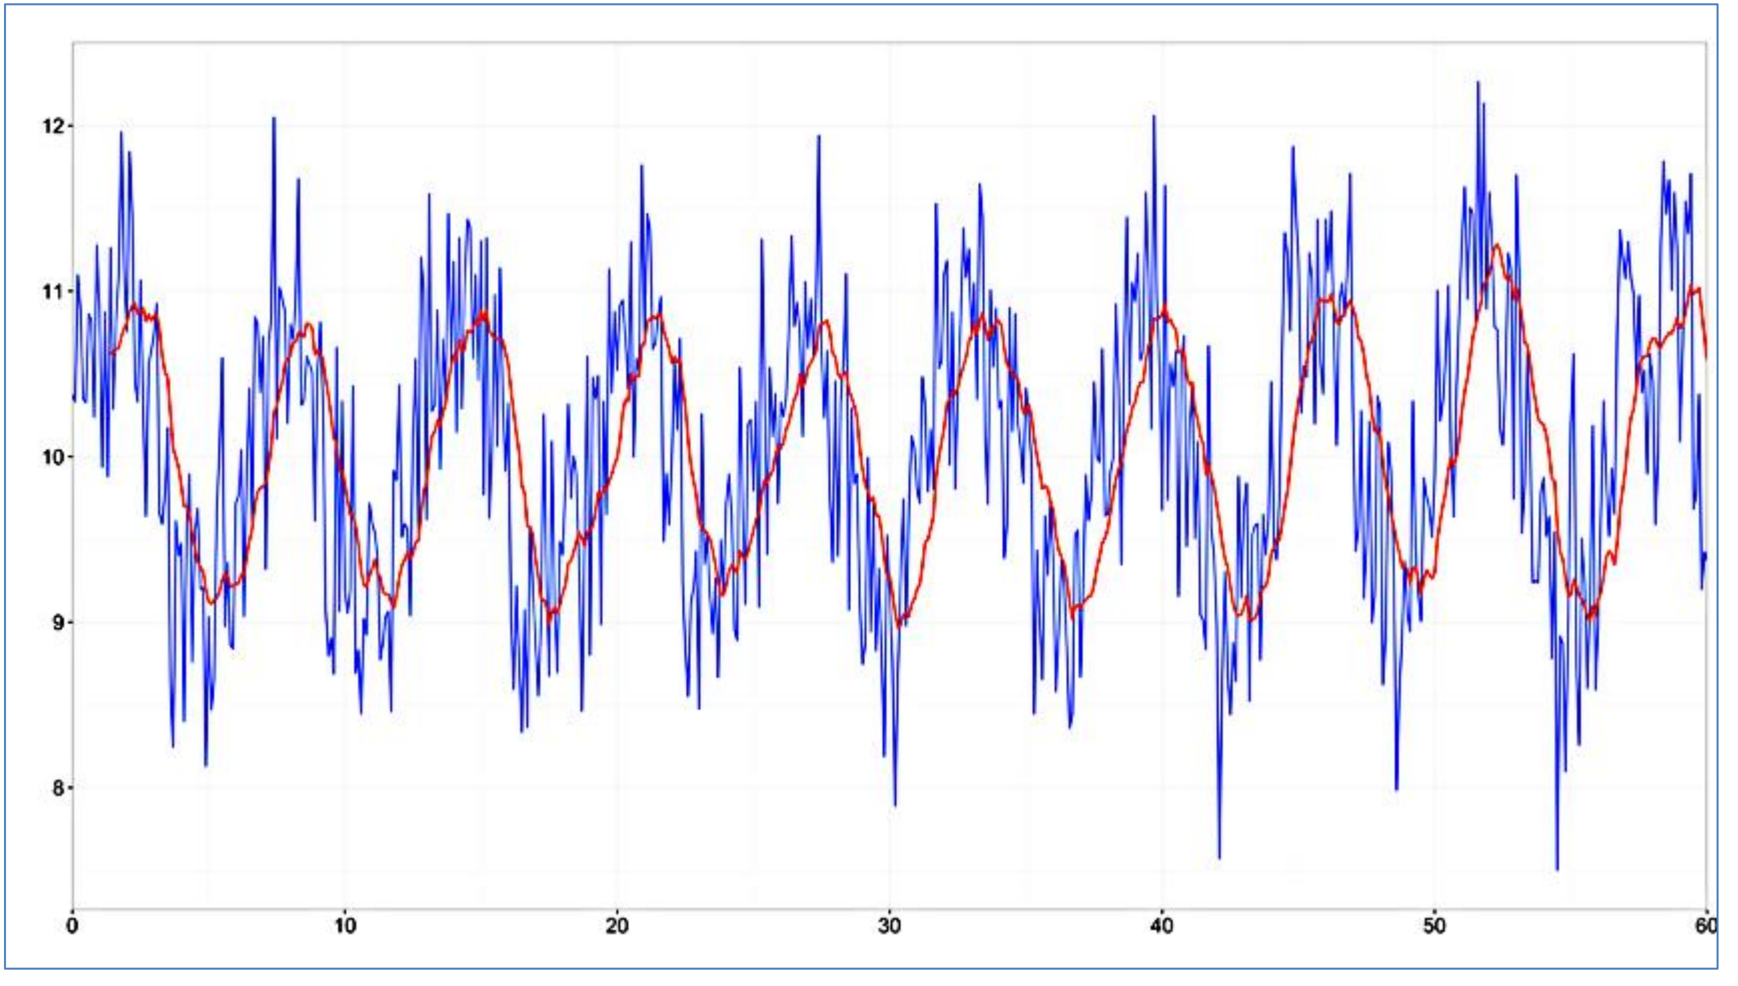

How to begin estimating the underlying frequency of the signal, $s(t) = \mathrm{A} \cdot \sin(2\pi f t) + AWGN$

A first approach is to examine the signal's zero crossings, rather like an FM receiver.

That's what a frequency counter does.
It's a reasonable looking digital device:

Signal goes in the coax. Frequency gets displayed.

And it don't mean a thing if it ain't got those
Nixie tubes. Have we won the FMT yet?

Frequency counters are conceptually straightforward. Nice tutorial here

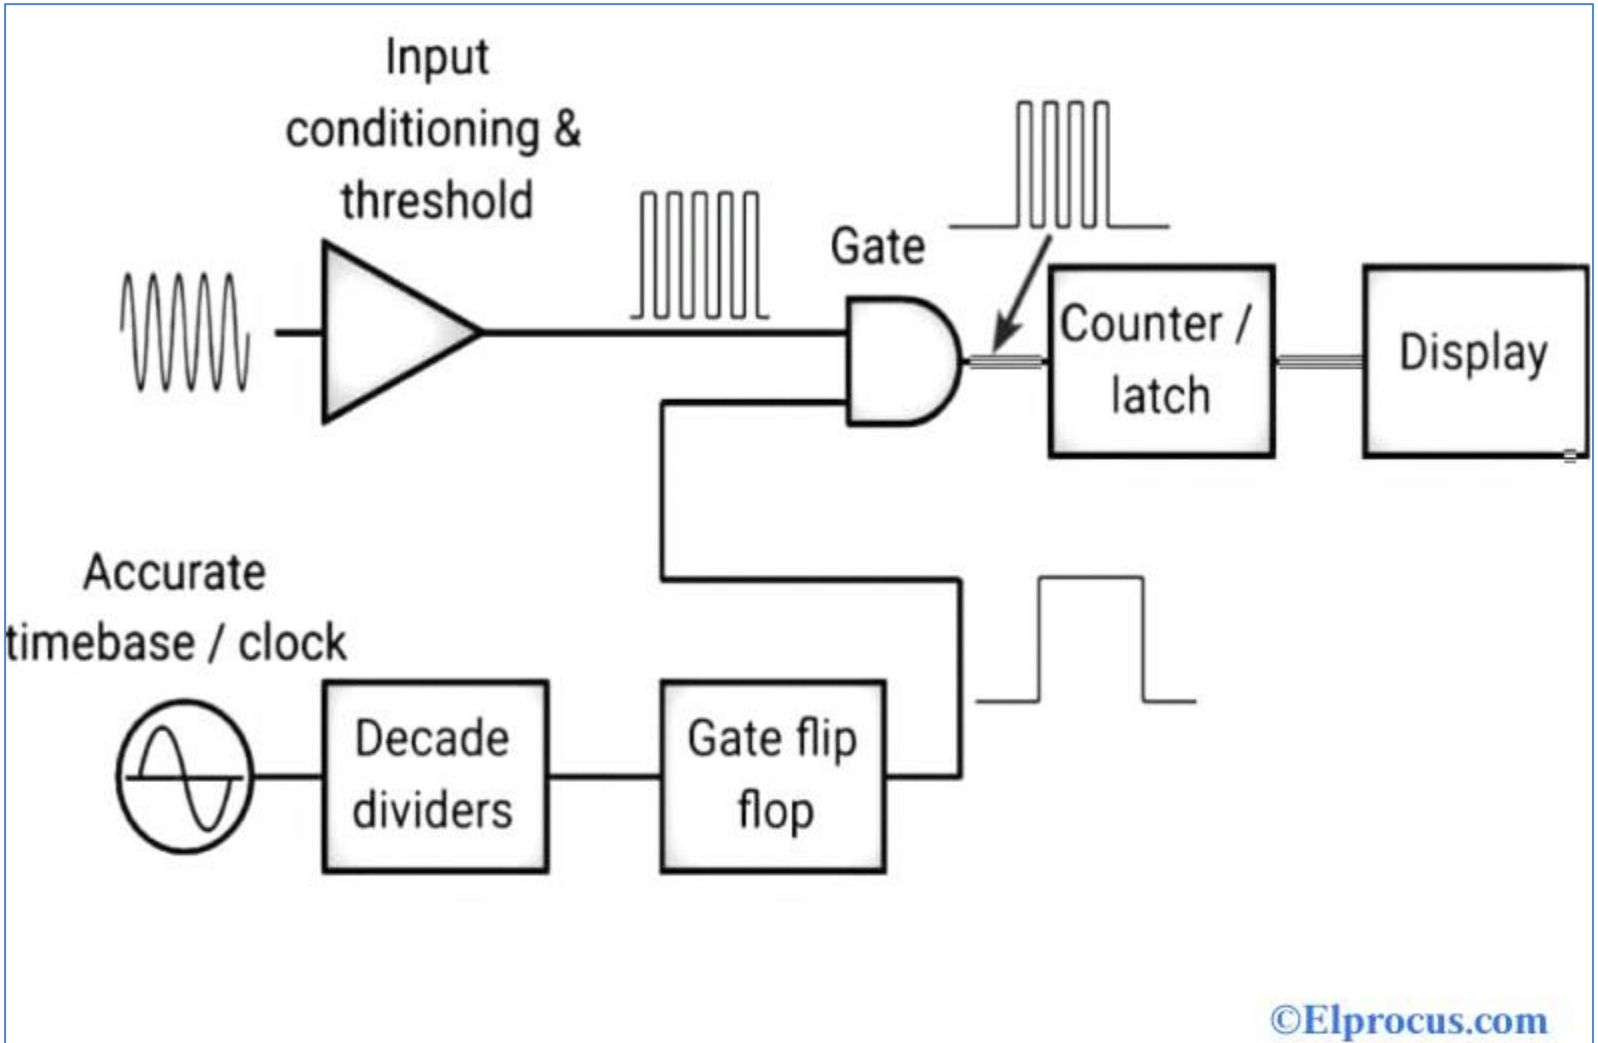

Specifically, a "gate" frequency counter.

What’s the problem?

1) They measure zero-crossings per gating, so noise will distort the measurement *high.*
2) They have an unforgiving tradeoff of gate time vs. precision *in linear frequency, not in fractional error.*

For example, a one second gate time gives one Hertz resolution *and no more.** There are no fractional zero crossings!

HamSCI says that’s not good enough. So does Connie.

----

* OK, frequency multiplying front ends are possible. Read up on those on your own time. They’re fussy.

The problem gets worse.

Say Connie is transmitting on 40 meters, 7 MHz.

1 Hz precision will give about 0.14 ppm error.

Does your filter have adequate *Q* for a frequency counter? Noise will ruin the measurement.
Q of 100,000 or 1,000,000 is hard to attain and to set: You can’t find the signal!

It gets yet worse. You don't just want to do well in FMT, you want to win.

The 1 second gate time gives 1 Hz precision.

10 seconds will give 0.1 Hz.

100 seconds will give 0.01 Hz

Connie goes key-down for 2 minutes in the FMT. You won't be able to find the signal, set your filter, and begin your measurement quickly enough to average 100 seconds. This won't work well.

You can finesse the *Q* problem by adding a frequency down-conversion. Measure at a lower frequency.

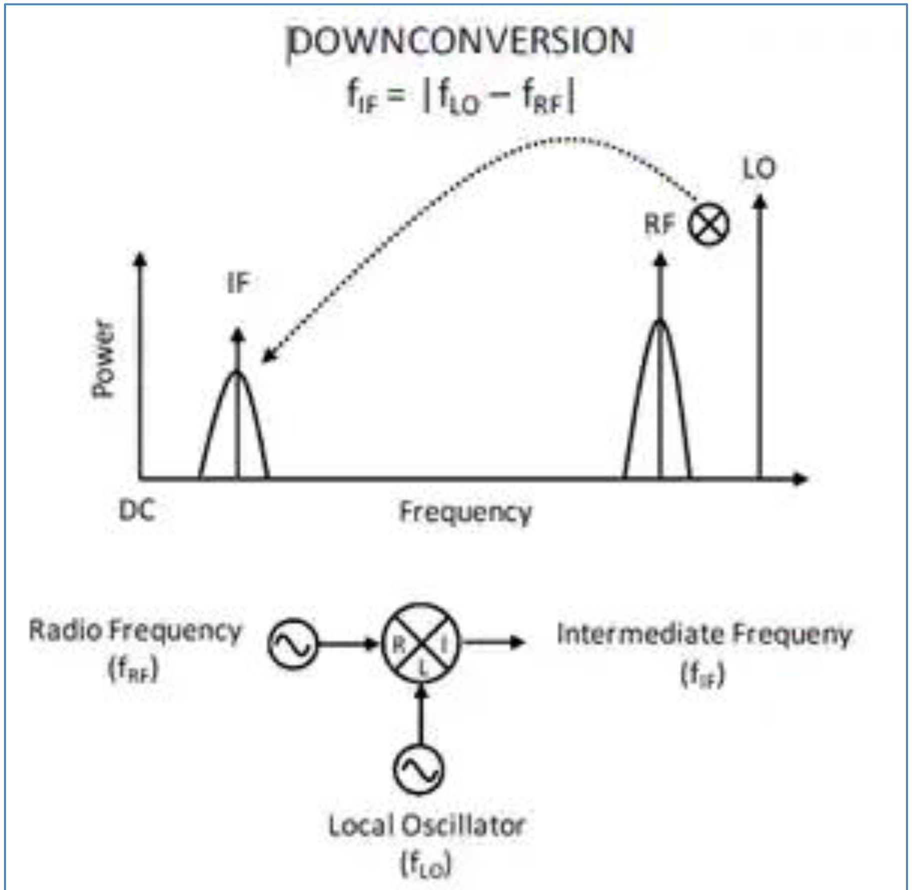

Mixer stage:

Generate a high-accuracy local oscillator. Subtract its frequency off the unknown.

That’s heterodyning—like the BC-221! Can improve noise response, *but not precision* for this linear measurement.

The noise filter is much easier to build, but we’re not yet winning the FMT.

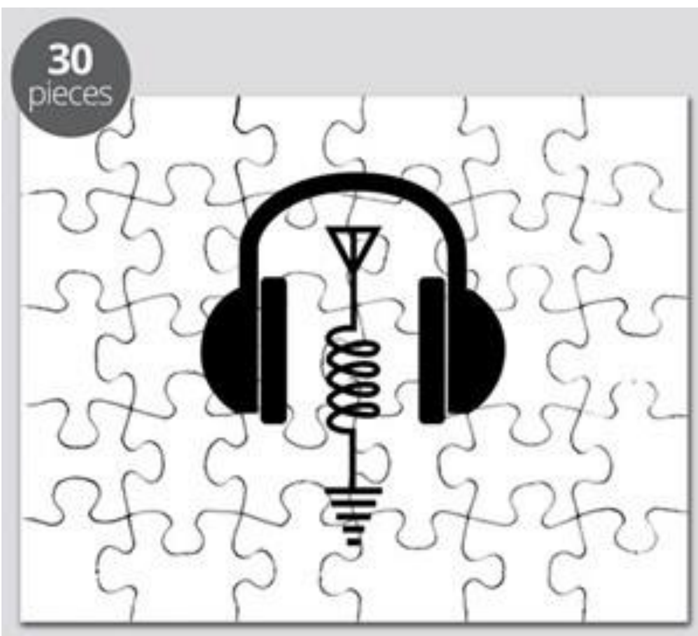

Ham pithing?
Looks uncomfortable.

fldigi *analysis* mode uses a sampled, digital approach to this:

The incoming low-IF signal is mixed by a computer-generated “guess” signal (the cursor).

The phase difference between the two signals is accumulated over one second.

That is the frequency difference between the cursor and the signal. fldigi does the math and records the result.

It’s not a bad approach but turns out to be more noise-sensitive than other methods.

fldigi *analysis mode*: Dial set at 9999 MHz. WWV's carrier at the "low IF" of 1000 Hz.

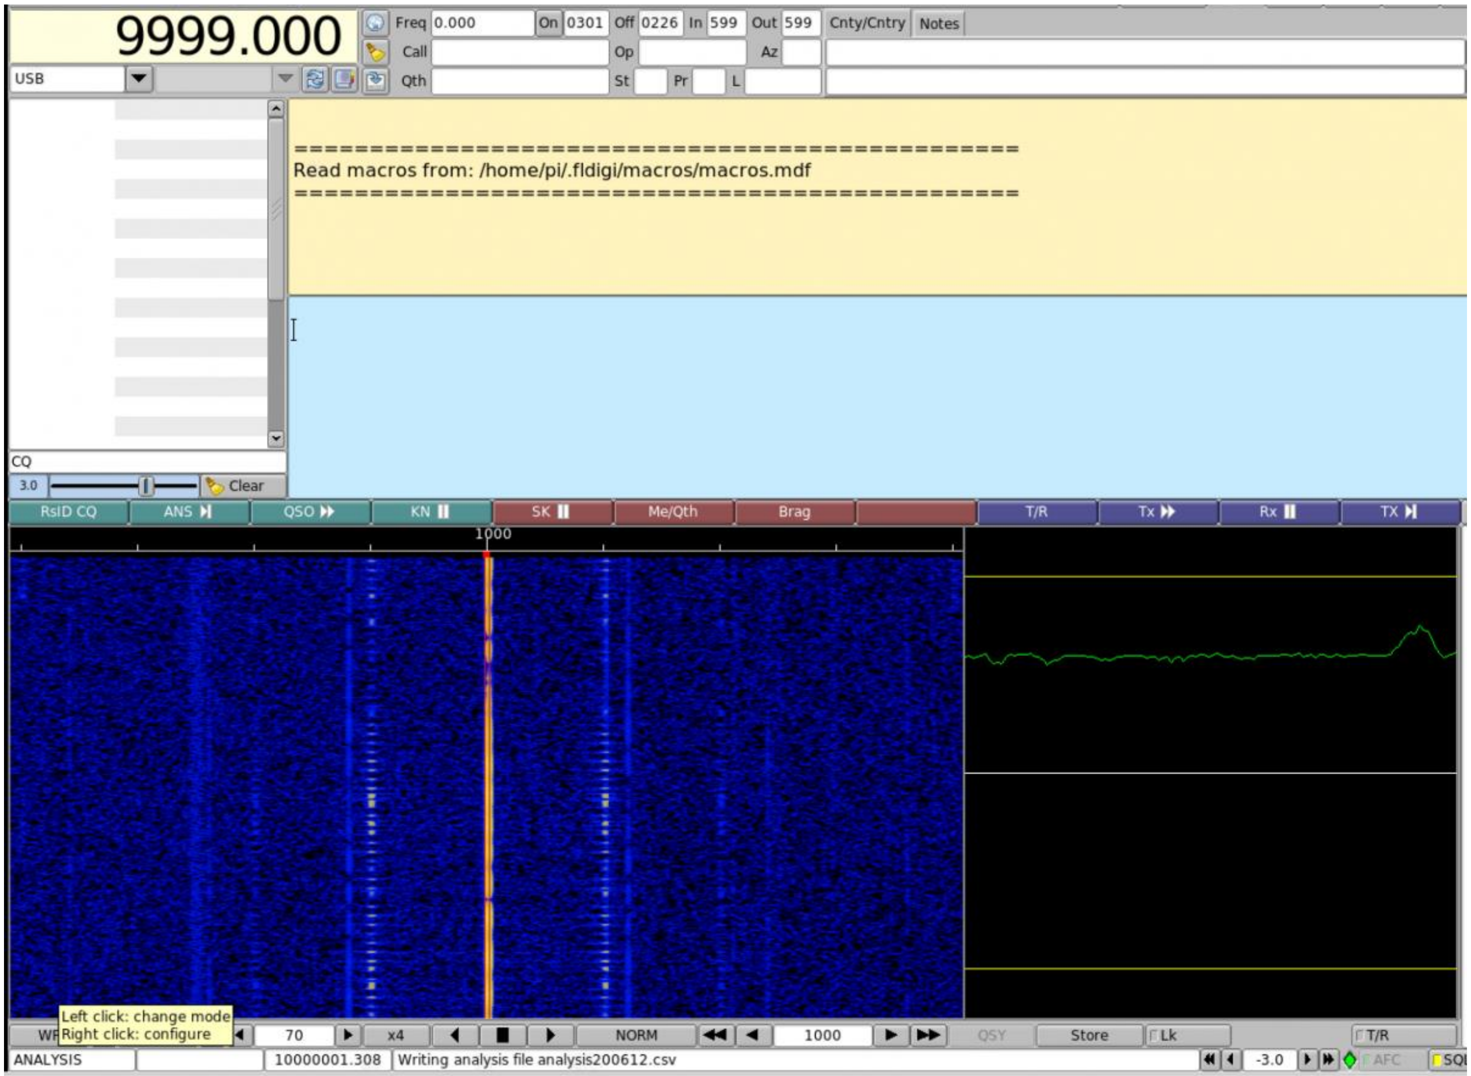

The waterfall is from a windowed DFT.
Not the analysis!

W8EDU is pleased to gloat that we have (nearly) won the FMT this way.

So has AD8Y. Nearly.

Next: Frequency domain methods

General idea:
Collect sampled data for a period of time,
do a discrete Fourier transform (by FFT),
perform computations in frequency domain,
find the peak…and win FMT?

not just yet…

The DFT is a bunch of simultaneous filters

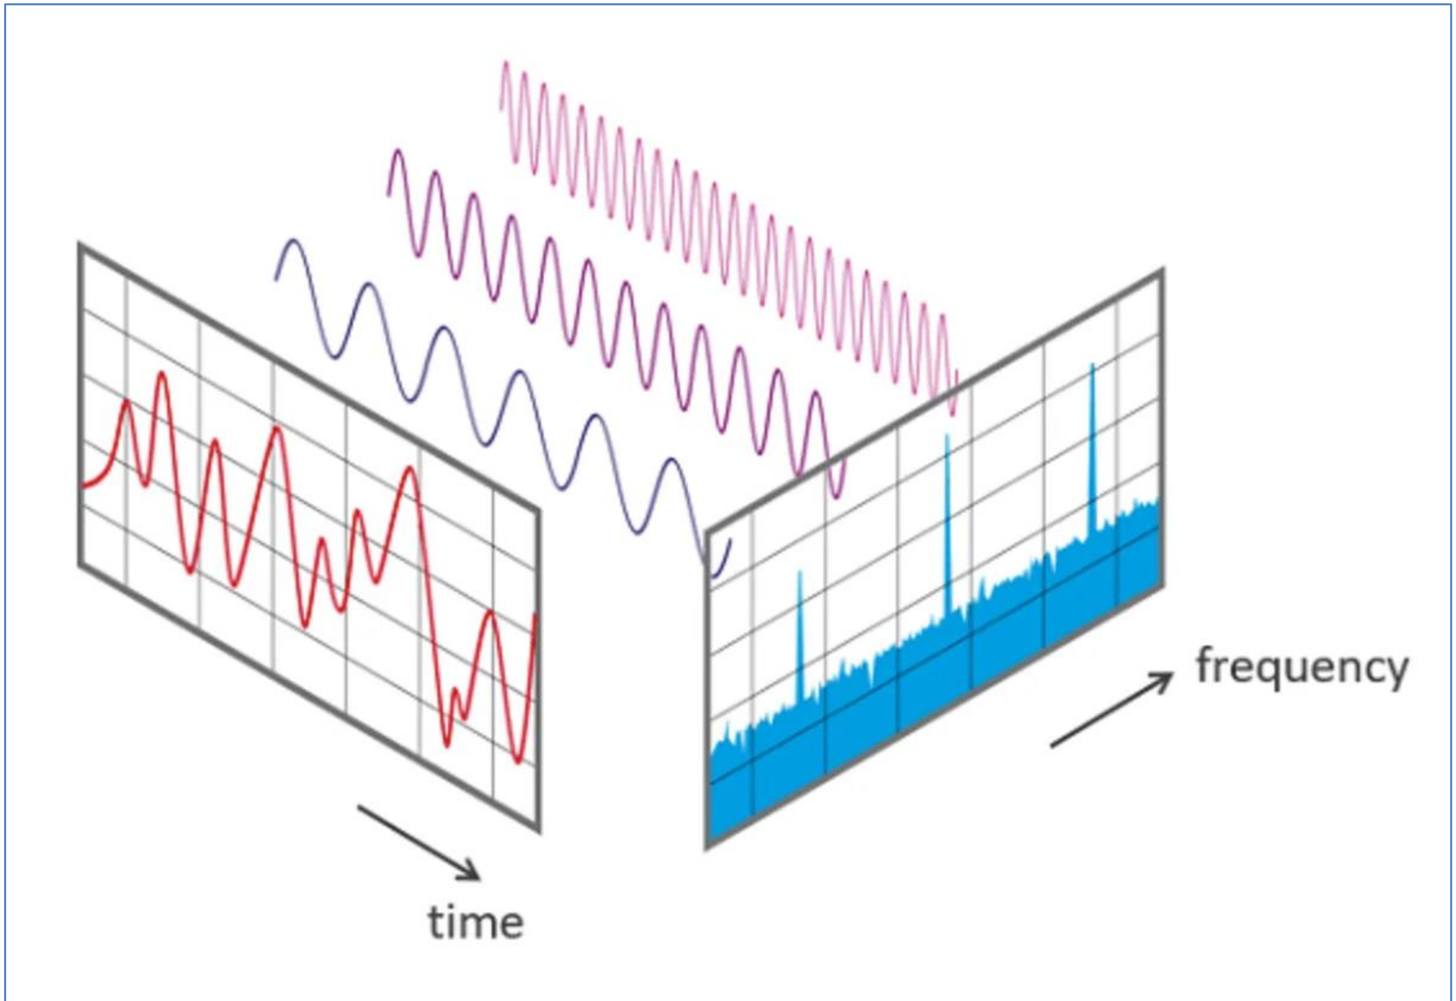

Here, three sinusoids plus noise are transformed.

Just read off the spikes?

The “bins” are too wide!

The DFT filters: a digital inner ear

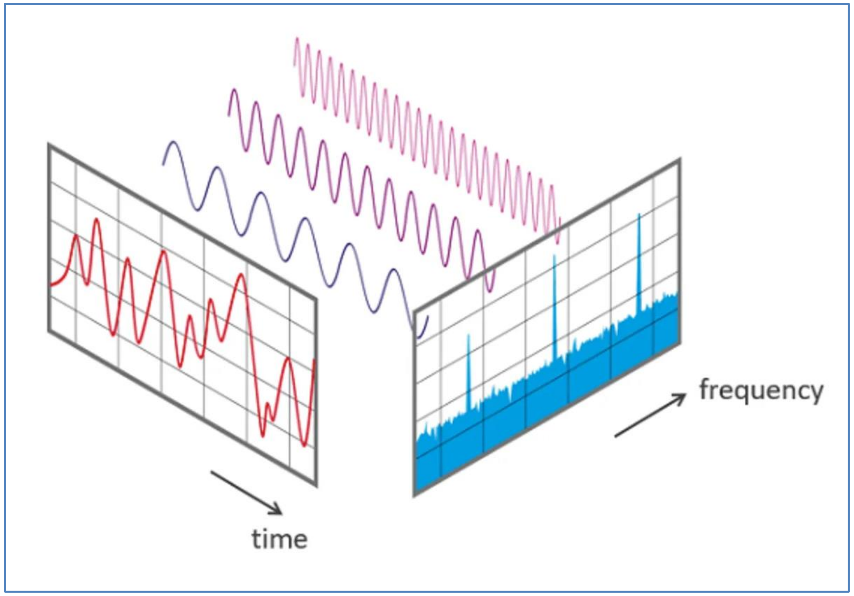

Substitute "frequency resolution" for "*Q*" and you have pretty much the same problem.

1 second of data collection gives 1 Hz precision of measurement.

10 seconds gives 0.1 Hz, and so on

Or does it?

The *linear* operation of the DFT retains enormously more information about the signal than the *limiting amplifier* ahead of a frequency counter. An FM detector also uses a limiting amplifier—to sharpen zero crossings and discard amplitude information.

Maybe there are ways of extracting some of the DFT's retained information?

You can take longer DFT sampling epochs and have narrower “filter” bins

That’s what you saw earlier from Steve Cerwin WA5FRF: long DFTs with a lot of overlap, a moving average.

It’s a variant on taking longer gate times with a frequency counter but with multiple, simultaneous linear filters. The resolution does go up, but so does time-averaging.

Short-time effects get smudged, detail is lost.

That FMT win is remaining elusive.

From WA5FRF's *QEX* article

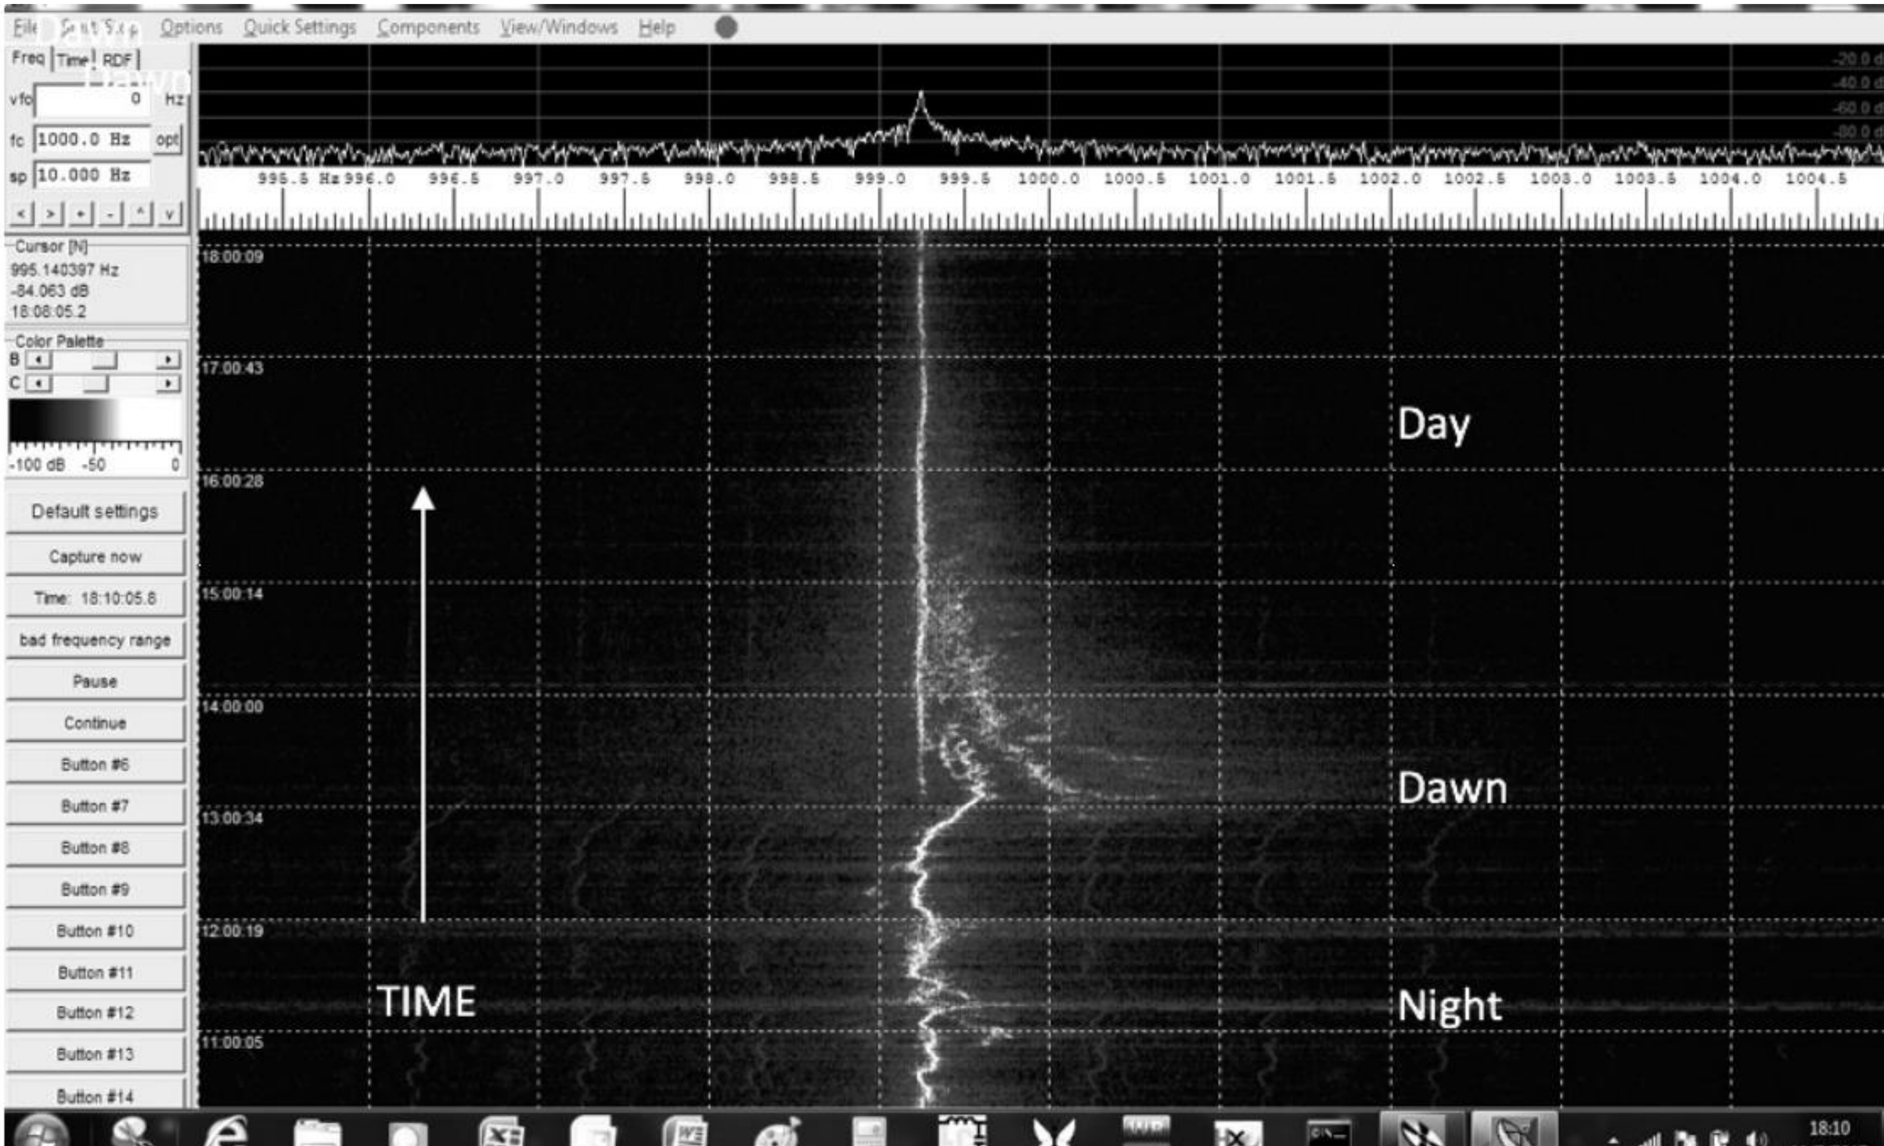

An alternative DFT approach: nonlinear interpretation of the linear transform. It's a data compression, fast but at a statistical price: multiple frequencies may be conflated into a single one by the processing.

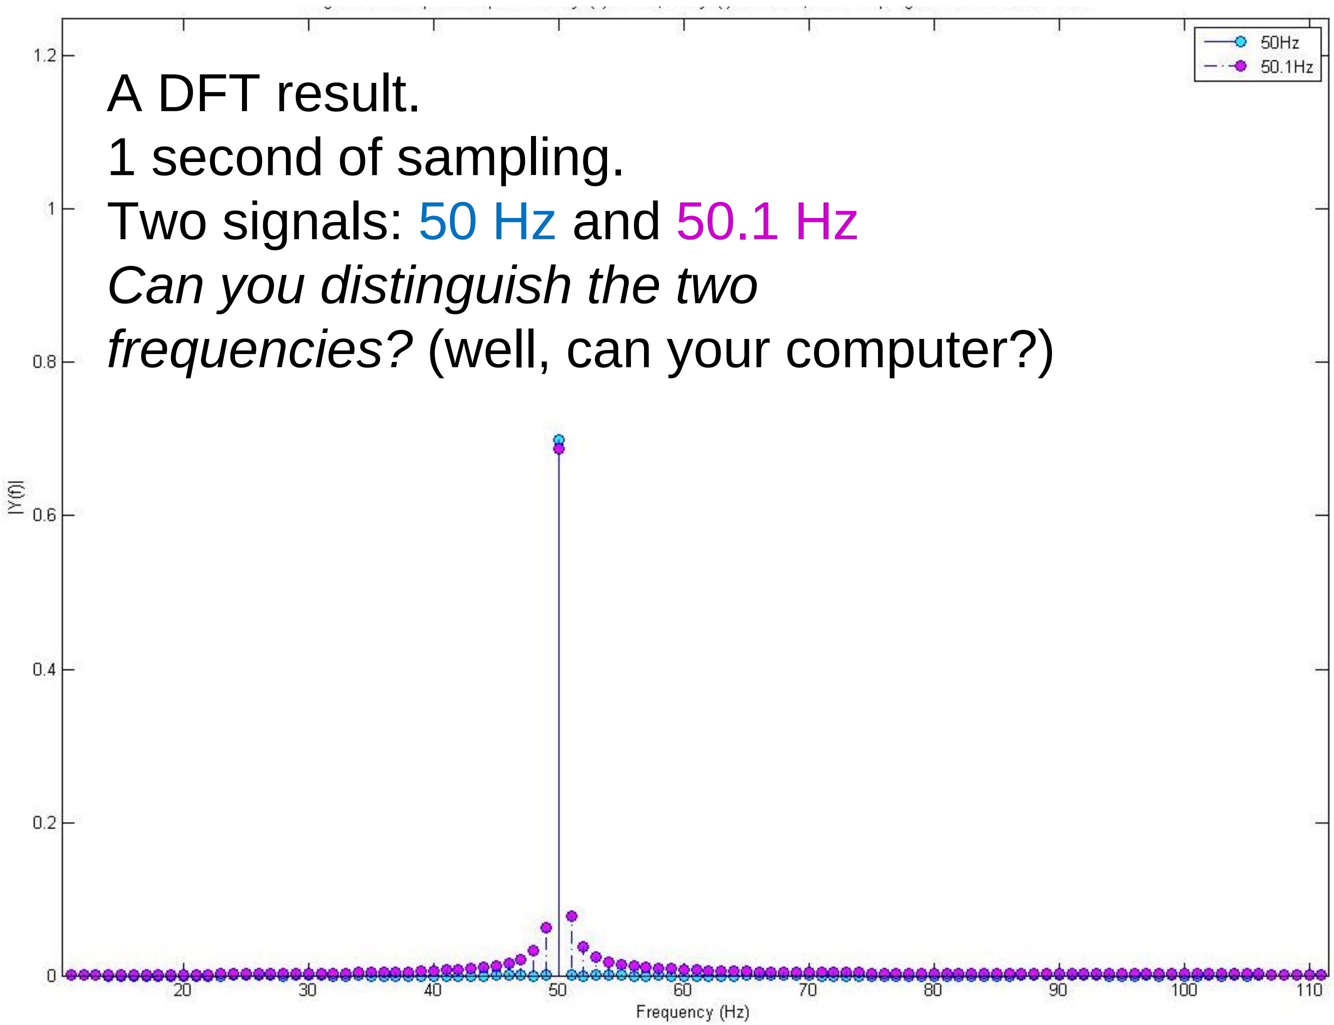

One second of sampling, so the bins are 1 Hz.

Two sine waves at 50 and 50.1 Hz have dominant "bins" but 50.1 has others with energy, too.

Note the asymmetry: the actual 50.1 Hz frequency isn't centered on a bin's frequency.

Here’s an example of nonlinear DFT usage: Maybe we can get around the resolution limit.

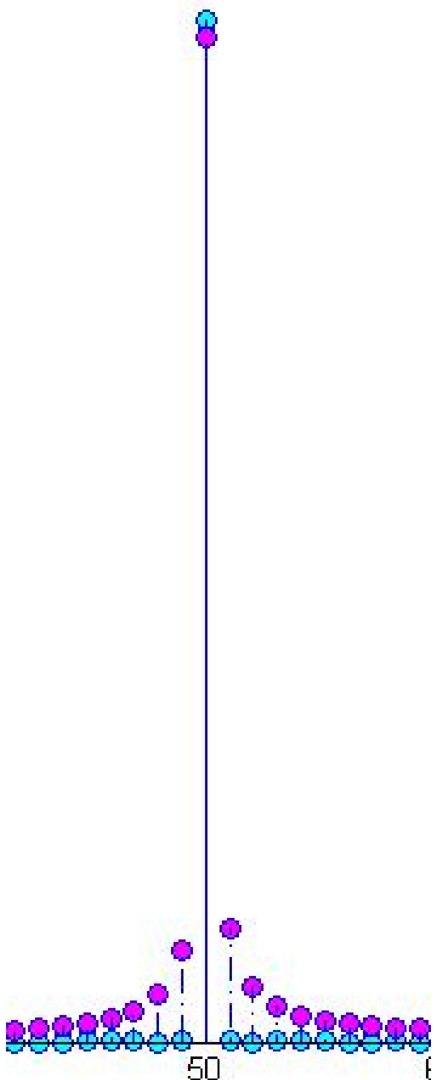

Zoomed in, notice that the 50.1 “spills” toward the right. It’s not centered on one of the frequency bins, which range from zero to 1023 Hz in 1 Hz increments, since one second of data were sampled at 1024 samples per second.

You could draw a spline curve that includes the points and estimate its actual peak. A parabolic spline with the peak and one point on each side is worth a try. This method isn’t much used but it does work.

The new fldigi *FMT mode*
(lots of code work by David Freese W1HKJ)

Uses the *ne plus ultra* of the nonlinear, single-frequency DFT techniques, based on this:

Paper

A highly accurate DFT-based parameter estimator for complex exponentials

Jeffrey Tsui and Sam Reisenfeld

76 1/2006 JOURNAL OF TELECOMMUNICATIONS AND INFORMATION TECHNOLOGY

This method takes a set of samples (typically for us, one second) and performs a nonlinear, recursive estimation that gives frequency, phase, and amplitude very close to the statistical-theoretical best possible.

fldigi *FMT mode*

permits comparing an unknown signal to a local standard, gives very clean graph. Has brought the FMT winning bar down to nearly 1 milliHertz in the past few months.

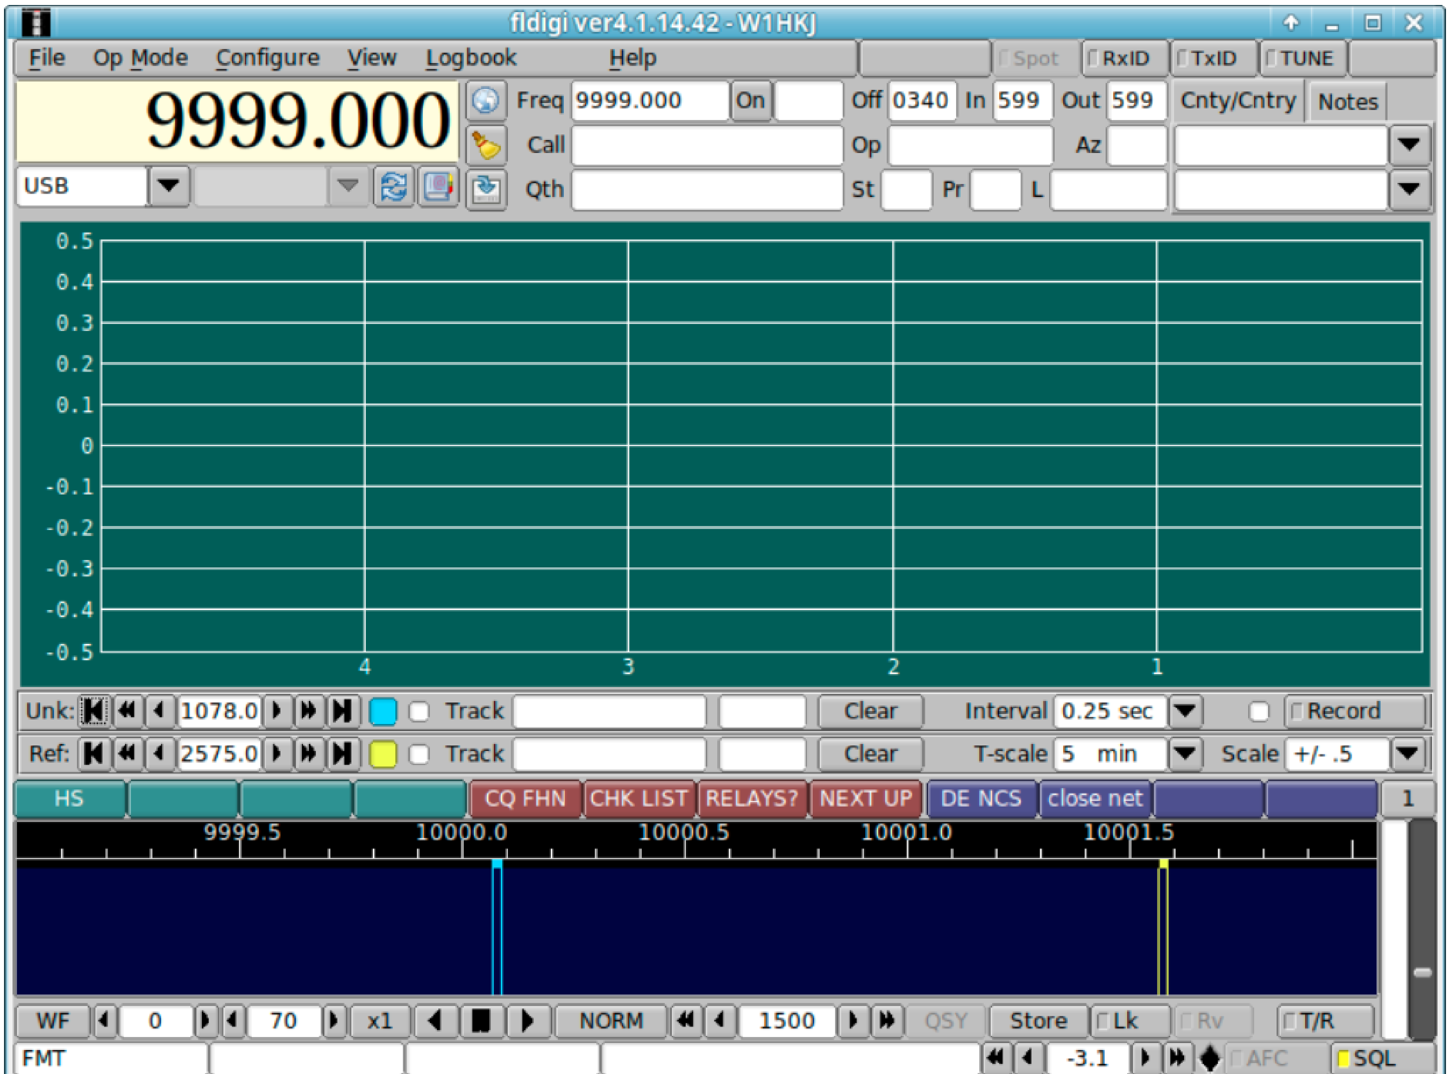

We have won the FMT! It's still hard work.

How to estimate frequency?

Make sure that you ask the right question:

Is the signal noise-free, stable, and of high amplitude? Use a frequency counter.

Is it noisy but stable? Might have multiple frequencies? Use a DFT technique with long time-averaging.

Is it noisy, only short-term stable, but expected single-frequency? Use a nonlinear DFT technique. *May use several simultaneously.*

You can’t have it all, all the time

If your signal is of low amplitude, noisy, unstable, and maybe not of single-frequency:

You’ll have to accept the tradeoffs of frequency precision *vs.* time resolution. Short transients will be lost to the analysis. Rapid changes will be averaged out and invisible.

Get to work and do some research, write a paper on the subject, be famous.

That’s the measurement life!
Jerome Lettvin: “More data, more noise.”

The Wisdom of RAL, K8DTS:

"Everything must be done first."

A motto of W8EDU.

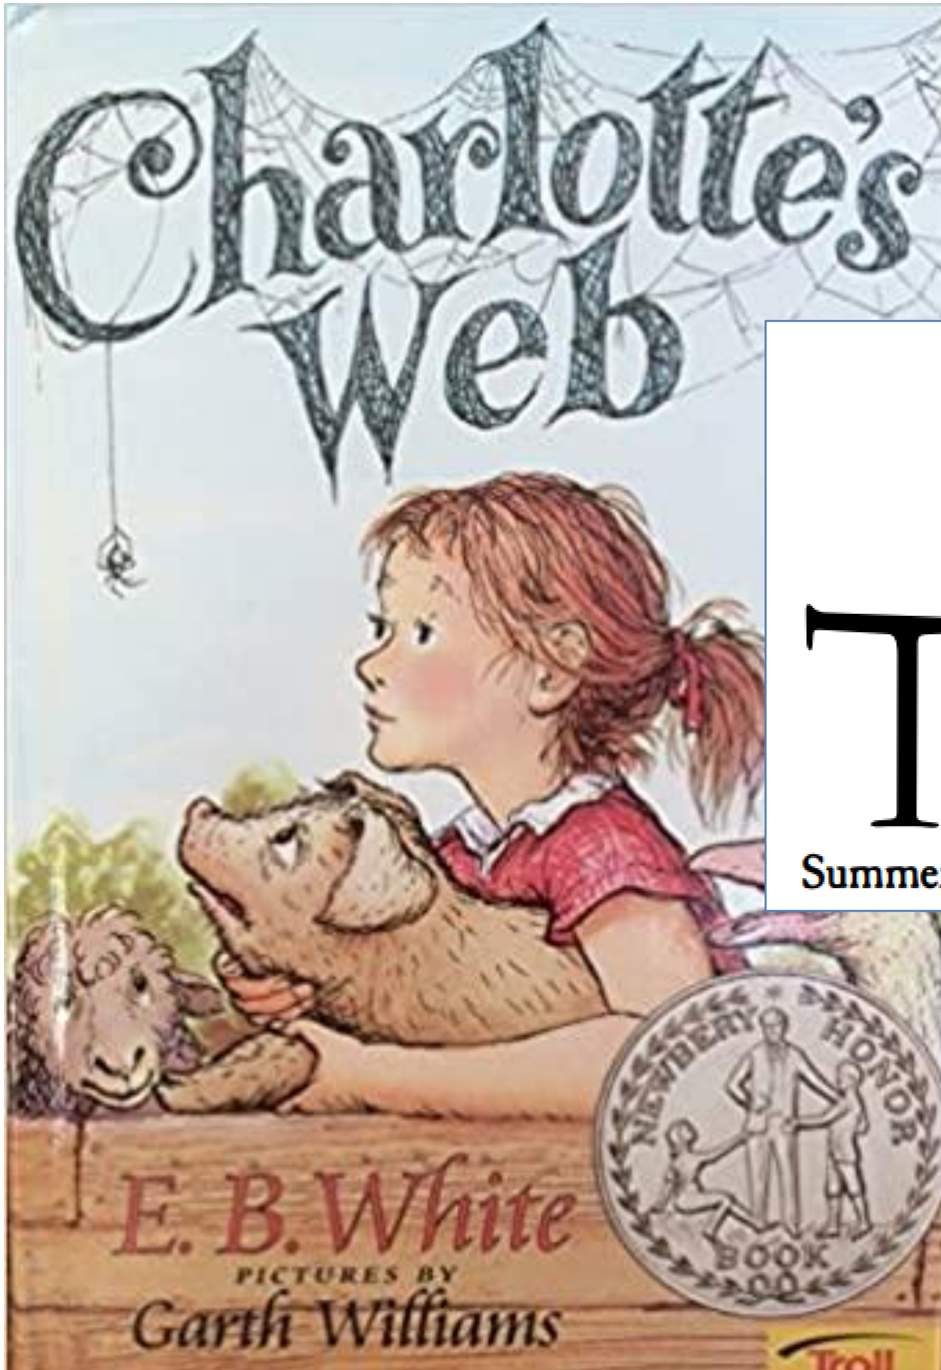

The Crickets

THE CRICKETS sang in the grasses. They sang the song of summer's ending, a sad, monotonous song. "Summer is over and gone," they sang. "Over and gone, over and gone. Summer is dying, dying."

Our WWV and CHU Doppler curves are changing noticeably now. Everyone stay safe and healthy. Enjoy the conference!