JFET Push-Push Frequency Doubler for 5MHz Input

C. Steinmetz 1/26/15 © 2015

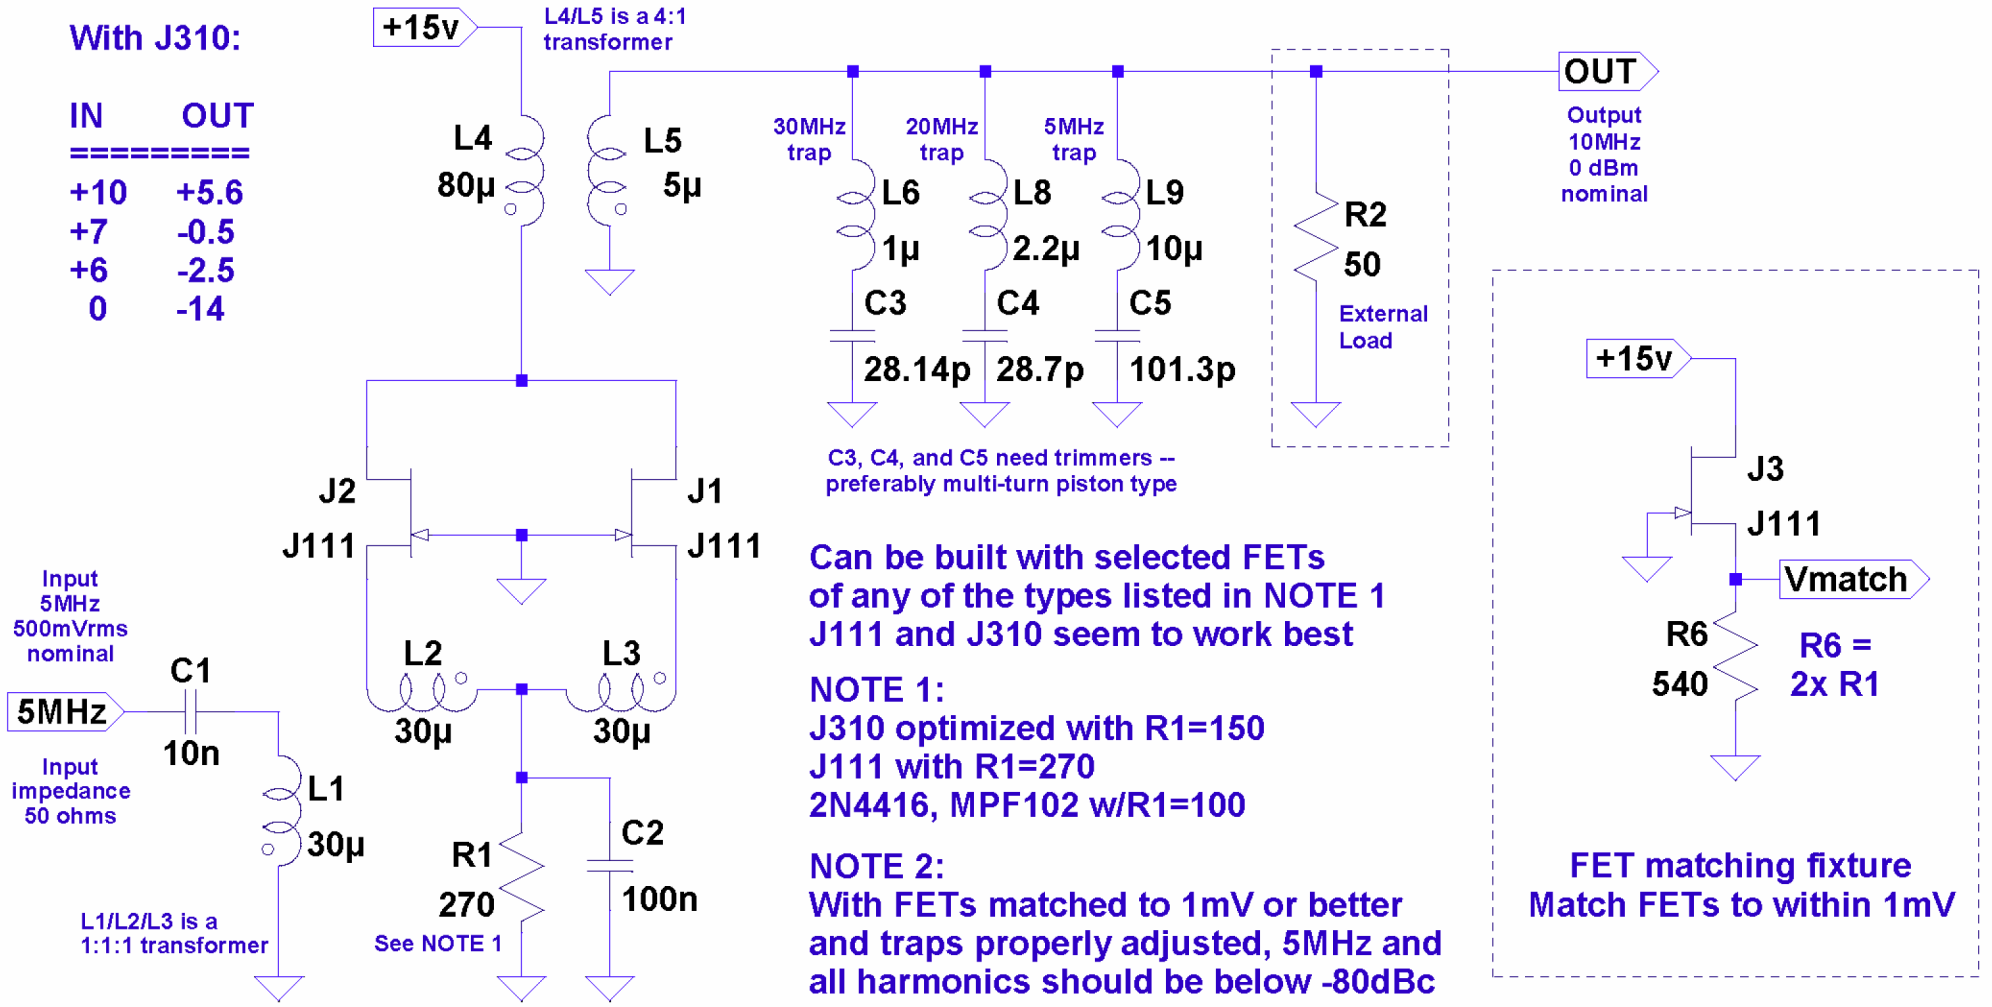

Discussion of the "barely Class A" push-push JFET frequency doubler -- my responses to issues raised by members of the "Time Nuts" email reflector, January 27 through February 4, 2015

Andrea wrote:

> *Now I have some 5MHz DOCXO. I have started to experiment with them and I would like to build a frequency doubler.*
>
> * * *
>
> *By the way, I see that really many of the 10MHz reference out there, are in effect doubled 5MHz ones so build a doubler seems reasonable for me.*

One thing to watch for is the 5MHz leakage component. If you are going to use the 10MHz standard for time-nuts experiments, the 5MHz component needs to be WAY down (< -80dBc) or you will get funny periodic ripples in stability plots. Despite having two 5MHz traps, one recently published design suppresses the 5MHz component only about 52dB below the 10MHz output, and the 20MHz and 30MHz components are also only -50 to -55dB.

For this reason (and some others, see discussions over the last several months in the archives) I prefer a doubler built with a quadrature hybrid coupler and a balanced mixer. There is a write-up here:

<http://www.ko4bb.com/manuals/download.php?file=02_GPS_Timing/4_App_Notes_and_Articles/Frequency_doubler_quadrature_DBM.pdf>

I recently revived an old, stalled project to develop a JFET push-push doubler for use at 5MHz (see schematic below). [Schematic is on preceding page]

FETs with very high transconductance and very small pinchoff voltage (what a tube designer would call a "sharp cutoff" characteristic) (e.g., 2SK369, BF862, etc.) are attractive on first look because they can operate with lower conversion loss or even some conversion gain. However, they are not well suited for doubler duty for two reasons: (i) their characteristics have a very short range of 2nd-order curvature, so in order to keep noise down they must be driven into regions of higher-order distortion and therefore generate lots of spurious energy; and (ii) they are devilishly hard to match well enough to suppress the input frequency feedthrough. Note that you also need to put enough voltage on the FET drains to get them well into the saturation region -- a Vcc of 5v is not enough. Again, the penalty is lots of spurious energy. So, the lower conversion loss of sharp-cutoff FETs is not the benefit it might at first appear to be -- it is much easier to add gain after the doubler than to remove unwanted spurious mixing products.

The design below uses medium-cutoff FETs and a Vcc of 15v (I found that J111 and J310 work best and can be matched sufficiently with a one-point match; 2N4416 and others also work, but are fussier and would benefit from a 2- or 3-point match). At an input of 500mVrms, their long 2nd-order characteristic is used efficiently to generate 10MHz with relatively little spurious energy.

I had no problem finding one or more FET pairs matched to within 1mV, given 20 devices from the same lot (YMMV). With properly adjusted traps at 5, 20, and 30MHz, all harmonics and subharmonics were below -80dBc. The inductors can be commercial RF parts with Q of 100 or so (I used some high-quality through-hole RF inductors I had on hand -- I doubt SMD inductors will work). The trap capacitors should be C0G/NP0 ceramics for the bulk of the capacitance, plus very small trimmers (I used 27pF, 27pF, and 100pF plus 0.2--6pF glass piston trimmers). I wound the two transformers on Mix-61 toroid cores (each winding is 20 turns on a FT37-61 core -- the inductance is a little lower than called out). Mini-Circuits parts (or equivalents) may also work, but I have not tried them.

==========

Bruce wrote:

> *Using the square law characteristic will inevitably increase the phase noise floor particularly in the flicker region with respect to just using the switching characteristic of a JFET, diode or BJT (non saturated).*

Even FETs with very large transconductance and low pinchoff voltage do not switch cleanly in a push-push mixer -- there is always plenty of square-law behavior to add flicker noise. So the difference is not really all that great. Furthermore, in practice whatever difference there is will be swamped in many (if not most) cases faced by amateur time nuts by the much larger phase noise of their sources.

> *The only viable solution is to use better filtering of the output of a switching multiplier.*

For the reasons given above, I disagree that a switching multiplier is the only viable solution for amateur time nuts. For NIST, possibly. But not for amateur time nuts. Given a dirty multiplier such as a push-push doubler with 2SK369s or BF862s, the only way to get the fundamental and distortion products at 30MHz and below down to suitable levels for critical work is to use a high-Q 10MHz filter plus a series of traps. The high-Q 10MHz filter brings its own low frequency stability problems.

Notwithstanding theoretical objections, the described circuit works very nicely and performs substantially better than most time nuts need. I presume that when one is willing to afford sources that require better performance, one will (1) know that, and (2) also be willing to afford an expensive doubler that meets that need.

==========

Gerhard wrote [regarding using series crystal filters to clean up the output of a frequency doubler]:

> *It is a different game when you want to notch away sub/harmonics.*

One problem with using crystals as traps (notch filters) is that the series resistance of a crystal is several orders of magnitude higher than that of a good series-resonant LC -- generally in the 50-100 ohm range. So, although the notch is very narrow, it will not be very deep unless it is in a high-impedance circuit. For example, in a 50 ohm circuit (50 ohms looking each way, so 25 ohms at the node) you will be lucky to get 3dB of attenuation. To get 40dB of suppression, the nodal impedance would need to be at least 5k ohms, perhaps even >10k ohms -- and the high impedance adds noise, which means there is a phase noise penalty. Another problem is that the narrow notches are prone to sliding off frequency with small temperature changes.

Also, while a 5MHz trap crystal will almost certainly be a fundamental-mode resonator, that will probably not be true at, say, 30MHz -- so a 30MHz trap would most likely have a notch at or near the desired output frequency.

Made with good, high-Q RF inductors (forget SM parts), an LC trap is generally preferable to a crystal trap. There is still some temperature sensitivity, but the greater width is much more forgiving. At the same time, the Q is high enough that you don't have to worry about effects 5MHz away when you are trapping frequencies of 30MHz and below.

==========

Andrea wrote:

> *Let me sum up everything and please correct me:*
>
> *the square-law characteristic of devices should be avoided, so the configuration of the doubler must be some sort of "ideal" full wave rectifier*

I disagree strongly with this, at least where push-push JFET doublers are concerned.

If you look at the schematic Bruce posted on his site, which uses a pair of J310 FETs driven into the pinch-off region, it runs the FETs from 0 to about 21mA. My circuit, when using J310s, runs the FETs from about 1mA to about 16mA. In both cases, when the FETs are conducting they are operating as common-gate linear amplifiers, NOT as switches.

In either case, when one FET is drawing low (or zero) current, the other one is drawing high current. The theoretical noise improvement due to running the low-current FET past the pinch-off point is, in practice, totally swamped by the noise from the other FET.

In order to realize a useful reduction of noise, the FETs would have to switch hard, from "off" (beyond pinchoff) to "full on" (with Vgs=0) -- but JFETs don't work like that, unless you drive the gates hard with square waves (that is how commutating mixers such as the ones designed by Ed Oxner and the later "H-mode" mixers work). See below for a schematic of an Oxner mixer using a quad JFET (but note that commutating mixers generally use MOSFETs).

When my circuit is normalized for 50 ohm output (by using a 4:1 transformer at the output -- which is the preferred method of driving 50 ohms with it) and the bias and drive are adjusted for the same currents as Bruce's circuit, the models predict almost identical noise from the two circuits. As a real-world check, I adjusted the bias conditions and drive on my breadboard doubler to give FET currents from 1 to about 22mA, and the measured noise decreased by a fraction of a dB. (The spurious distortion products rose slightly, but not nearly as much as when one drives the FETs beyond pinch-off.)

So no, running the FETs in Class AB or B does NOT confer a material noise advantage compared to running them "barely Class A," as my design does. It does, however, create an exponential explosion of odd-order distortion products that must be removed if the circuit is to be useful for time nuts purposes.

So in my view, the "barely Class A" push-push JFET doubler is clearly superior to its Class AB or B cousin.

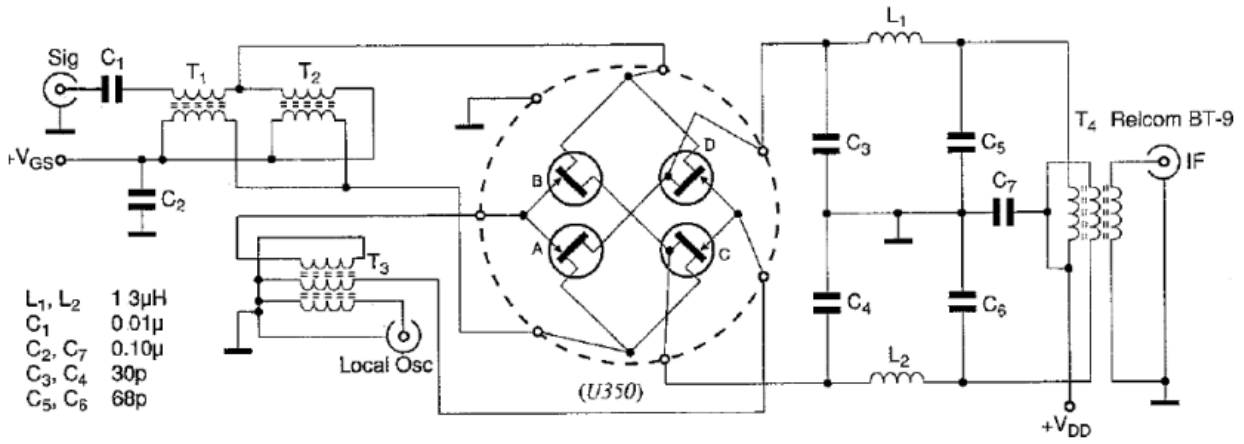

Ed Oxner JFET commutating mixer

==========

Andrea wrote:

> *I see. This configuration is in effect a common gate B-class (or AB, or "barely A") amplifier and the rectification is a side effect.*
> *But, what is the advantage between it and a couple of diode-connected transistors with a full A-class (more linear, so less spurs) amplifier in front of it?*

It is the rectification that causes the gross nonlinearities, not the amplification. So no matter how linear an amplifier you make, the diodes (or Class B or AB amplifier) will cause gross nonlinearities that we do not want. Furthermore, transistors have both even- and odd-order distortion products, while JFETs have predominantly second-order products. So JFETs naturally tend to produce the second harmonic, while transistors also produce the odd-order products we are trying to avoid (as well as higher even-order products).

> *I know that the circuit originates at NIST and thus there surely IS an advantage. Are it trading more spurs (that you can cancel out with filtering) for less phase noise (that you cannot recover anymore)?*

I do not know precisely how the NIST circuit is biased, and as far as I know it is not general knowledge among time nuts -- so any substantive response would be conjecture. I don't even know if NIST still uses it. NIST historically settles on something that works well enough, then sticks with it for a long time (until the phenomena they are trying to measure get distinctly better than their instruments). NIST has lots of considerations besides pure performance, such as power consumption and fitting into old form factors, so they do not necessarily have the best possible solutions, even when they have just designed the next generation. What we know for sure is that the JFET push-push doubler worked well enough for NIST's purposes when it was designed. That does not mean improvements weren't possible.

> *Adding negative feedback linearize further the "barely Class A" amplifier; so, it's good to sacrifice part of the gain of the push-push stage to reduce flicker noise (and thus add less phase noise) and at the same time spurs.*

But it is the natural second-order distortion of the JFETs that makes it a particularly good way to build a push-push doubler. We don't *want* to linearize it!

> *If it's so, why use a nonlinear (or barely linear) gain stage to rectify? Using just one stage means in general less phase noise output (but with probably more spurs that can be filtered out), versus a more stage linear amplifier (perhaps with strong negative feedblack) followed by a rectifier?*

The "barely Class A" push-push doubler does not rectify the signal -- it creates the second harmonic largely because of the device characteristic. The design goal is to map the bias and input to the portion of the FETs' characteristic curve that has the best fit to a second-order transfer function, while at the same time holding noise down below the noise budget. That is why medium-cutoff FETs like the J111 and J310 are the best choices, not sharp-cutoff FETs like 2SK369 and BF862.

==========

Andrea wrote:

> *But, what is the advantage between it and a couple of diode-connected transistors with a full A-class (more linear, so less spurs) amplifier in front of it?*
>
> *If it's so, why use a nonlinear (or barely linear) gain stage to rectify? Using just one stage means in general less phase noise output (but with probably more spurs that can be filtered out), versus a more stage linear amplifier (perhaps with strong negative feedblack) followed by a rectifier?*

I replied:

> *The "barely Class A" push-push doubler does not rectify the signal -- it creates the second harmonic because of the primarily second-order transfer characteristics of the JFETs. The design goal is to map the DC bias and the input signal to the portion of the FETs' characteristic curve that has the best fit to a second-order transfer function, while at the same time holding noise below the design requirement.*

Perhaps some pictures would be helpful (see below). Figure 1 (top) shows an ideal full-wave rectified sine wave, similar to what is produced by a full-wave diode rectifier, a bipolar transistor push-push doubler, or a FET doubler driven into pinchoff (Class B). Obviously, it is extremely rich in harmonics. The second harmonic of the output (doubled) frequency is only 14dB below the desired signal, and a series of even harmonics stretches as far as the eye can see, diminishing only very slowly with increasing harmonic number. (In practice, there will be a HF rolloff that makes things slightly better. However, there will also be odd-order components, which an ideal full-wave rectifier would not produce.)

Figure 2 (bottom) shows waveforms from the simulation of my "barely Class A" push-push doubler, using a matched pair of J111 FETs (J310s perform almost identically, with the appropriate change in the bias resistor). I purposely introduced a 10mV gate voltage imbalance in the simulation to model imperfect matching. The red and magenta traces are the currents in the two FETs, showing a primarily second-order transfer characteristic. When these currents are added by the push-push connection and

put through a 4:1 (turns ratio) transformer into a 50 ohm load, the green trace results. This trace shows the simulated raw output, without any traps. Obviously, this is very much closer to a clean 10MHz signal than the rectified signal in Figure 1.

The 5MHz component is ~40dB below the desired 10MHz signal. This depends strongly on how well the FETs are matched and on the layout and shielding. J111s or J310s from the same lot, matched to within 1mV, should do better than this (the 5MHz component from my breadboard circuit is below -45dBc, without any traps). The other visible distortion products, and their levels, are:

15MHz -75dBc
20MHz -45dBc
25MHz -100dBc
30MHz -75dBc
35MHz -100dBc (all figures are approximate).

The breadboard circuit performs similarly (the 15MHz component is about 10dB lower from the breadboard, so I needed traps only at 5, 20, and 30MHz to get all spurious responses below -80dBc).

As I noted before, the "barely Class A" circuit is not materially noisier than a FET push-push doubler that is run into Class AB or B, but it has MUCH lower spurious outputs and, therefore, does not need the sort of aggressive filtering the Class AB/B circuits need, avoiding the increase in phase noise and other problems associated with aggressive filters.

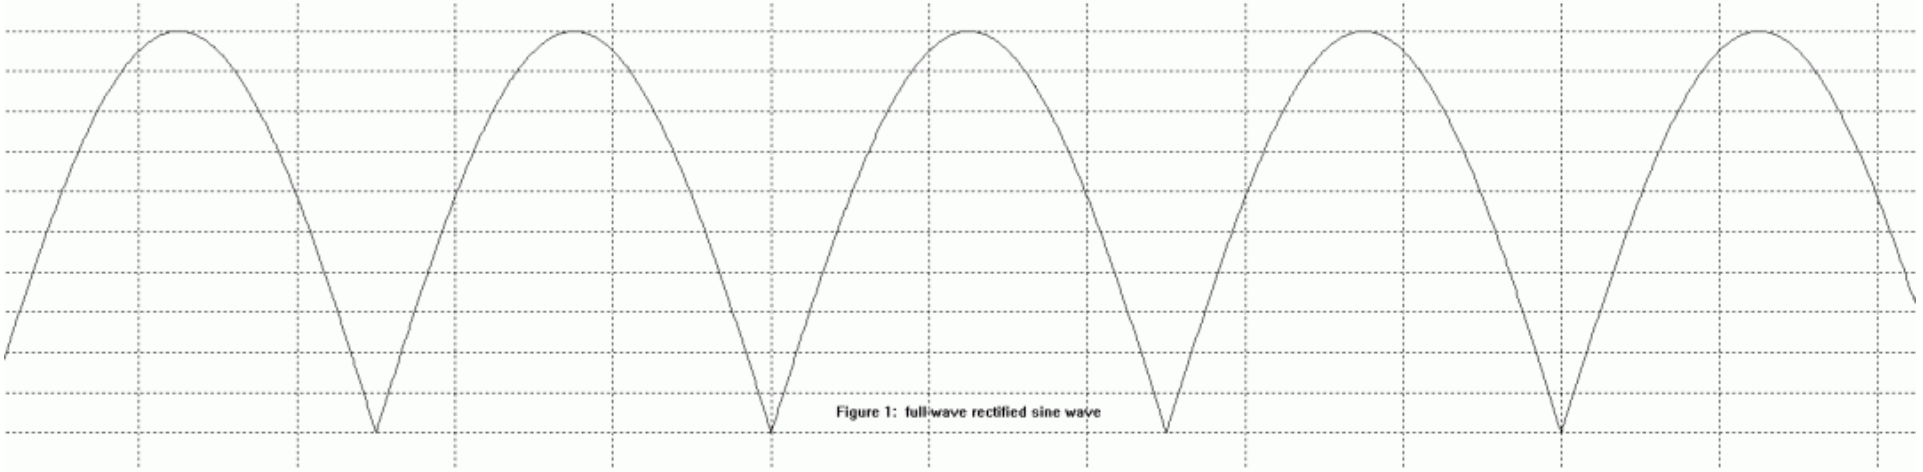

Figure 1: full-wave rectified sine wave

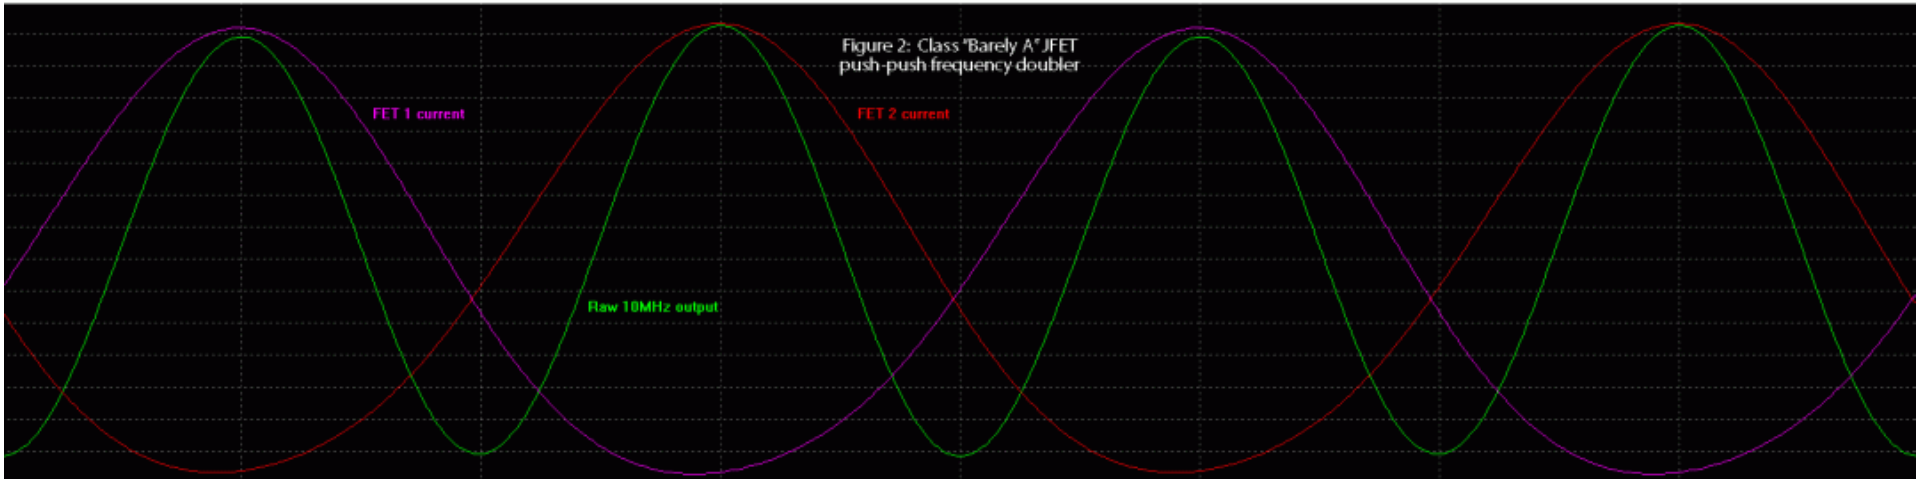

Figure 2: Class "Barely A" JFET push-push frequency doubler

==========

Bruce wrote:

> *Whilst the output signal of the barely class A JFET amplifier has a lower unwanted harmonic content and thus requires less filtering to achieve a given suppression of unwanted harmonics and/or subharmonics, the question of the flicker phase noise penalty incurred by the barely class A amplifier approach remains unresolved.*

I posted the resolution a few days ago.

As I said then, I adjusted the bias and input parameters of my breadboard doubler to match the conditions under which the FETs are operated in the doubler posted on your site, and measured the change in noise (including in the flicker region). The noise decreased by a fraction of a dB. Accordingly, I conclude that the barely class A doubler's noise, including flicker noise, is within a fraction of a dB of a Class AB doubler using the same FETs that you consider optimized.

I also explained then why this result should come as no surprise (one FET in a Class AB or B doubler will not be contributing noise when it is cut off -- but that coincides with the other FET being at or near full current, so the total noise is dominated by the noise of the full-current FET and the benefit due to the cut-off FET is insignificant).

There may be quieter FETs with lower flicker noise corners available that have similar medium-cutoff characteristics and are, therefore, suitable for this use -- but for the reasons I have given, I believe that similar relative noise relationships between barely Class A and Class AB doublers using such FETs would hold for them, as well. NOTE: For anyone simulating JFET circuits, be aware that many available JFET models do not model flicker noise at all, and many of those that do are wildly inaccurate at simulating noise in the flicker region. As always, there is no substitute for building and measuring the circuit.

==========

Andrea wrote:

> *Can you put in another graph the calculated difference to a pure sine wave?*

I'm not sure what you mean by "the calculated difference to a pure sine wave." I already reported the amplitudes of all of the visible spurs (that is, the ones above the noise floor), which define the departure from a pure 10MHz sine wave. I am attaching below the simulated spectrum analysis from which I took those reported amplitudes, if that helps. (There is no new data here, it is just graphical rather than tabular.)

Again, this is from a simulation, and I purposely introduced 10mV of gate imbalance to model imperfectly-matched FETs. It shows the raw output from the doubler, with no traps installed. The breadboard circuit performs similarly, although the 5MHz and 15MHz components are about 10dB lower from the breadboard than they are shown in this simulation (this depends on how well matched the FETs are -- I was able to get a better balance in real life than the imbalance I purposely introduced for this simulation).

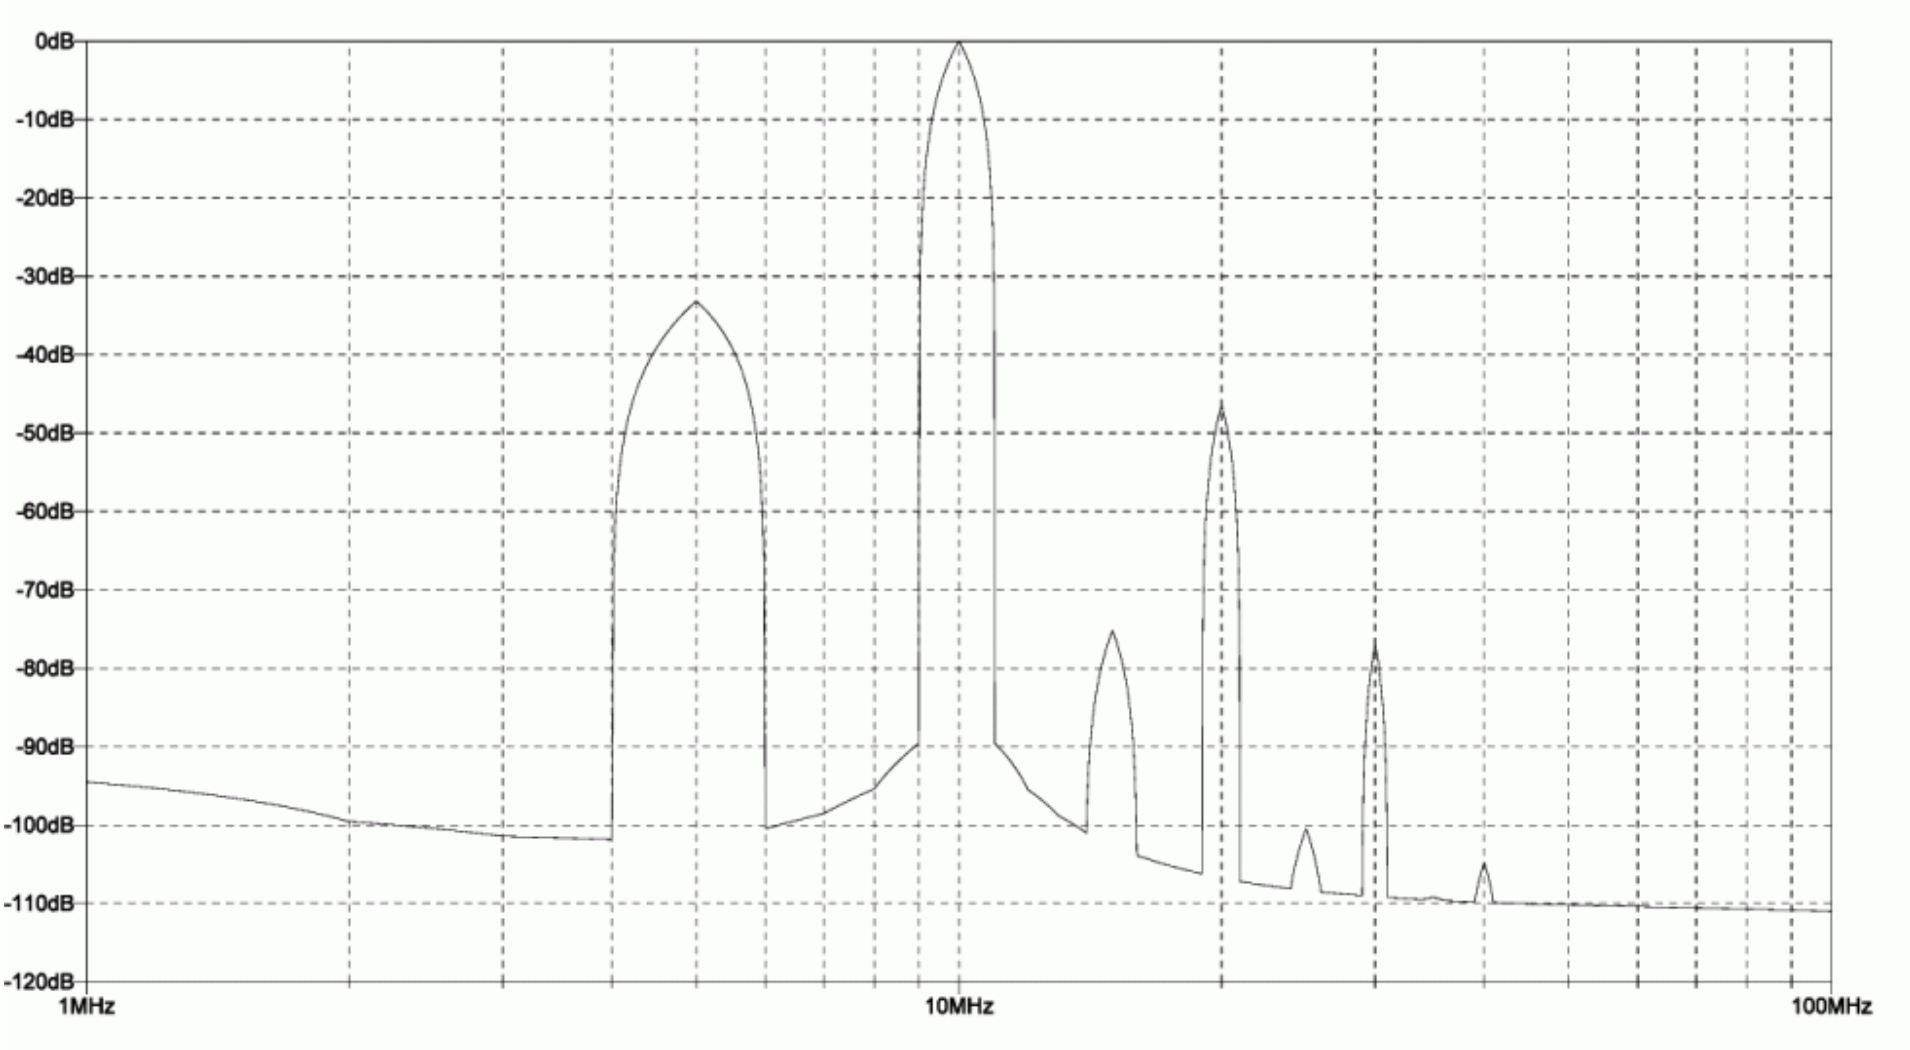
0dB
-10dB
-20dB
-30dB
-40dB
-50dB
-60dB
-70dB
-80dB
-90dB
-100dB
-110dB
-120dB
1MHz
10MHz
100MHz