ELECTRONOTES
1016 Hanshaw Road
Ithaca, NY 14850

APPLICATION NOTE NO. 422

Mar 6, 2015

NOTCH FILTER AS A "WASHED-OUT" COMB

INTRODUCTION:

We recently reviewed notch filters [1] and thought of them as a class of devices that were capable of blocking one particular frequency, usually completely blocking it, with a necessary partial rejection of frequencies on either side of the notch. This notch was usually the direct result of a "zero" (pair or zeros) placed on the appropriate frequency axis (jω–axis in s-plane or unit circle in the z-plane). The sharpness of the notch could be controlled by placing a stable pair of poles near the notching zero. In the continuous-time case (s-plane) we always had those poles around somewhere anyway. Old stuff.

The digital case could be similar. Indeed, with such methods as "Bilinear-z" we could mimic the analog system as a prototype. We also could make notch filters by just placing zeros without any supporting poles. (To be technical, the poles were relegated to z=0). It only took a length three filter (Finite Impulse Response of FIR) to produce the pair or zeros. Without sharpening feedback-produced poles, the resulting notch was no very sharp. If we still insisted on FIR, but wanted a sharper notch, it was quite possible to use a longer FIR filter and the notching zeros were supported by a "ring" or zeros. In this case, the role of the pole was played by one pair of zeros moving off the ring – indeed becoming the unit-circle notching pair (see below).

A relative of the longer FIR notch is the so-called "comb filter" – a filter with multiple equally-spaced unit-circle zeros [2] which has some obvious applications such as the removal of harmonic frequencies (integer multiples) [3]. Such filters typically have a very simple impulse response: for example: h(n) = [1 0 0 0 ……..0 0 0 1] and thus are seen summing an input with a long-delayed input. This is related to constructive and destructive "interference" ideas in physics. In the case where the time between the input and the delay corresponds to half a cycle, the frequency is cancelled, as is the case where the delay is 3/2 of a cycle, 5/2 of a cycle, etc. Thus we have a periodic notch or comb-filter.

There is the usual limit on input frequency of being half the sampling frequency. In the case of the comb just described we can think (most directly) that there is an array of equally spaced zeros that is allowed to extend to half the sampling frequency. Sometimes we think of it as being allowed to continue around the unit circle multiple times. In as much as some original realizations of comb filters involved so-called "bucket-brigade" devices (often called, in error, analog delay lines) these notions seemed natural. The sampling was there, but hidden in the large number (like 32 to 512 or 1024) of samples (buckets of charge passed along) between input and output. We had no access to these in general. They worked for a variety of "phasor" devices, artificial reverberators, etc., or as simple comb arrays.

Necessarily as multiple notches were crowded into the range from zero to half the sampling frequency, the individual notches got sharper. We just had to worry that the notches were multiple! This was a fundamental difference from the usual one-frequency notch. Here we want to look at the manner in which upper notches may be "washed out" leaving mainly an original notch. This is a curiosity, but is interesting and fun. Further it may help us understand certain naturally-occurring processes where a delayed path is not as clean as the h(n) above, but may involve multiple paths of differing lengths. Examples such as the so-called "jetsounds" effect (like listening to a jet engine as a plane moves relative to a sound-reflecting runway) and auditorium reverberation may be likely examples.

A PROTOTYPICAL EXAMPLE

Here is the general notion. In Fig. 1 we suggest a sound source (Jet) that is presumably broadbanded to some degree. There is a distance away a microphone in the vicinity of a runway, presumed reasonably flat. Sound from the jet reaches the microphone directly (blue), but let's also assume there is a component that reflects off the runway to the microphone (green) which is necessarily a longer path. The summation constitutes a delay/add situation, although additional details could be very complex. In the simplified picture, the blue and green paths constitute a comb filter. Again, the details are complex, including different attenuations, frequency responses inside the paths, additional paths (the mountain and the UFO in the figure) and the fact that the plane is moving.

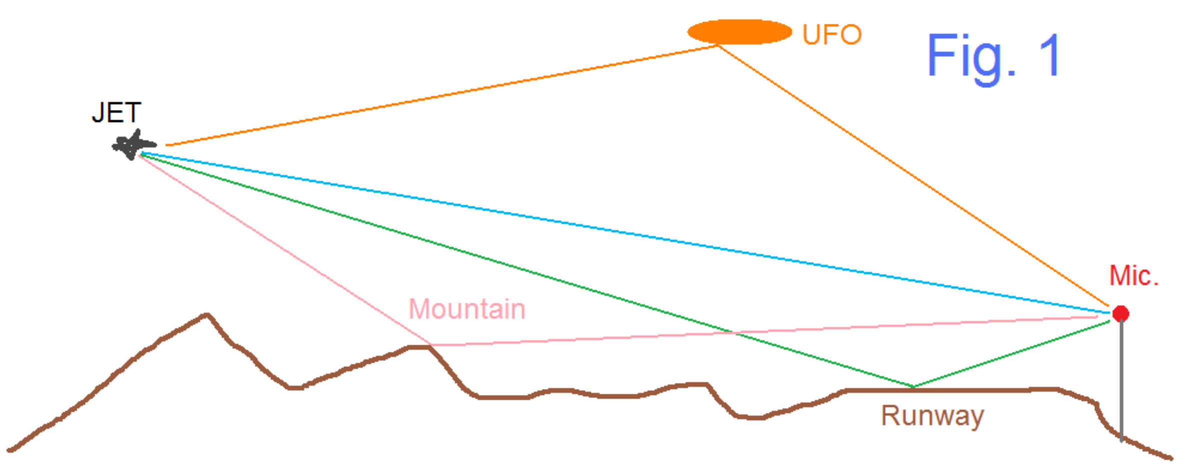

Fig. 1

We intend this setup to be considered general, but at the same time point out that this corresponds to a familiar special effect of electronic music – the phasing, flanging, or "Jetsounds" effect formed by moving notches, the origin of the term now suggested.

If we were assigned the task of determining the frequency response of the setup, it is straightforward in principle. We should consider a particular sinewave frequency we believe to be involved, calculate the various paths (and any complications) and then it is a matter of trig identities. Finding the response amplitude at that frequency is a matter of examining a full cycle. Note that this is not necessarily an LTI (Linear Time-Invariant) system because things may be moving, but it should be valid as a quasi-stationary finding.

A TOY EXAMPLE

At this point we can both simplify and give a specific example, just to see what is happening. So instead of the somewhat arbitrary multi-paths of Fig. 1, we will consider the simple FIR system of Fig. 2. This familiar network is NOT a model of the postulated physical situation of Fig. 1. What we show is a long delay line consisting of M delays (pink) followed by N delays (green). Potentially this can form a length M+N+1 FIR filter. Here we suggest that the very first "tap" (red, with a multiplier 1) is directly from the input itself. The input is then delayed by M unit delays (pink) and then this is followed by N+1 potential taps (blue - all shown as having a value "a", but could be general). It is important that we understand what this does NOT represent.

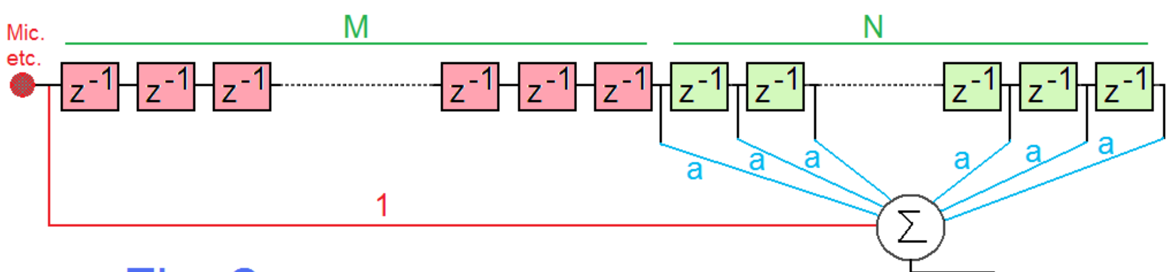

Fig. 2

The pink delays are NOT the delay of a direct path (blue) of Fig. 1. Indeed, the input to the pink delays <u>should be understood to correspond to the exact arrival time of the direct component</u>. The red path of "1" duly records this arrival. So what are the pink delays? They represent <u>the delay between the direct arrival and the first indirect component</u>. The green delays are a very rough notion of a rapid arrival of multiple multi-path signals. (Note that if a=1/(N+1), the tapped green delay line is a familiar moving-average filter.) So, recognizing this is only by analogy, the red path might be the direct sound of a jet. The first blue path, following the last pink delay might be the reflection off a runway. The second blue path from the first green delay (probably with a different "a") might be a reflection off a barn. That sort of arrangement. Our representation of multi-path components as equally spaced and equal weights is our idea of a reasonable starting point.

At this point we consider the two possible impulse responses, the first being:

h_0=[1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1] which is a basic length-17 comb filter impulse response, and:

a=1/11

h_1=[1 0 0 0 0 0 0 0 0 0 0 a a a a a a a a a a a]

which replaces the tap on the end with 11 taps all of the same value (1/11) and centered about the end tap of h_0. The new FIR is longer: 17 + (11-1)/2 = 22. This is perhaps easiest to see from the plot of Fig. 3 where h_1 is shown by the red stems and h_0 by the blue stars. Fig. 4 shows the corresponding frequency response magnitudes, and Fig. 5 the corresponding zero plots for the two filters.

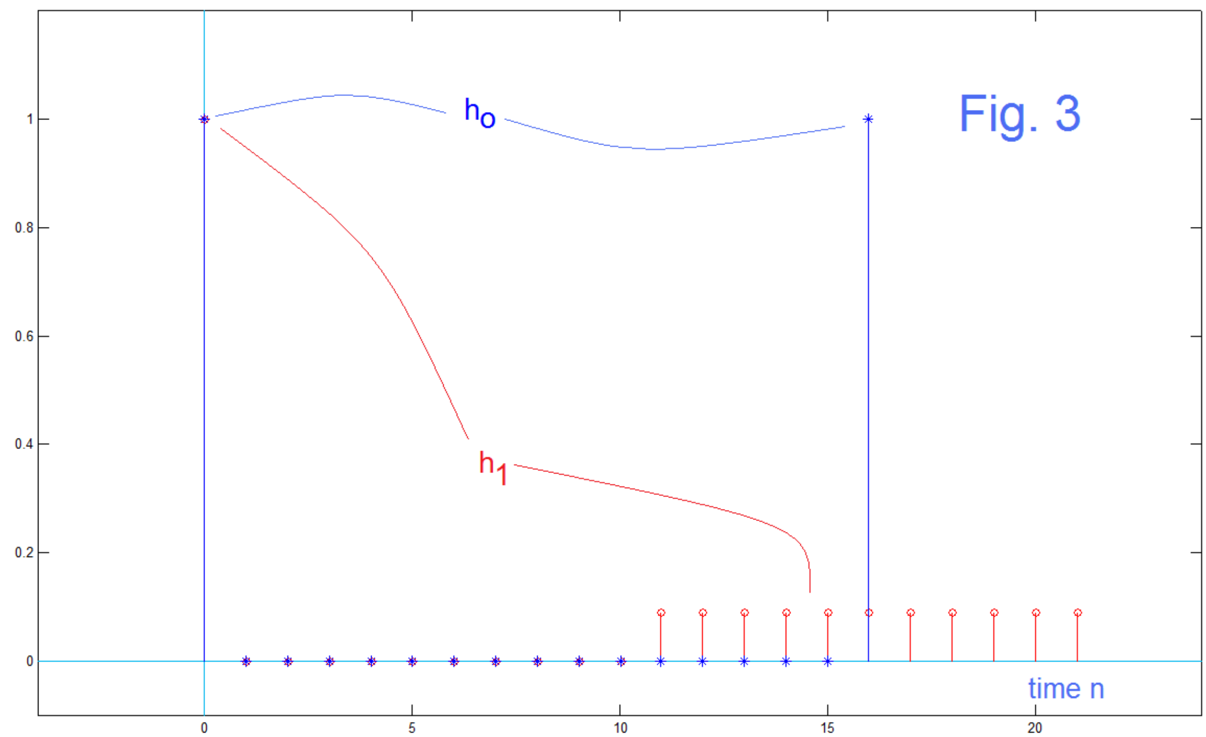

The motivation here was intuitive, that a comb filter is not a notch, and that if we had instead a whole array of notches <u>with different spacings</u> we could perhaps more or less preserve the lowest notch of the comb while higher frequency notches might tend to cancel (Fig. 6 which follows). That is, the notch at the lowest frequency might remain while the notches for the higher frequencies will cancel in the summation. Here the impulse response, being a summation of unit sample responses, IS the sum of each and every impulse response value. The corresponding frequency responses are likewise a sum. In doing summations, it is imperative to do the phase right (Matlab's ***freqz*** will do this for is). It is easiest at this point to look at the results.

Fig. 4 shows the frequency response magnitudes of the two impulse response, the dashed blue is the response for h_0 and is the expected comb with its multiple notches. The red curve for h_1 shows an imperfect null for the lowest notch, with a flatter response for the washed-out upper notches – exactly what we suspected would happen.

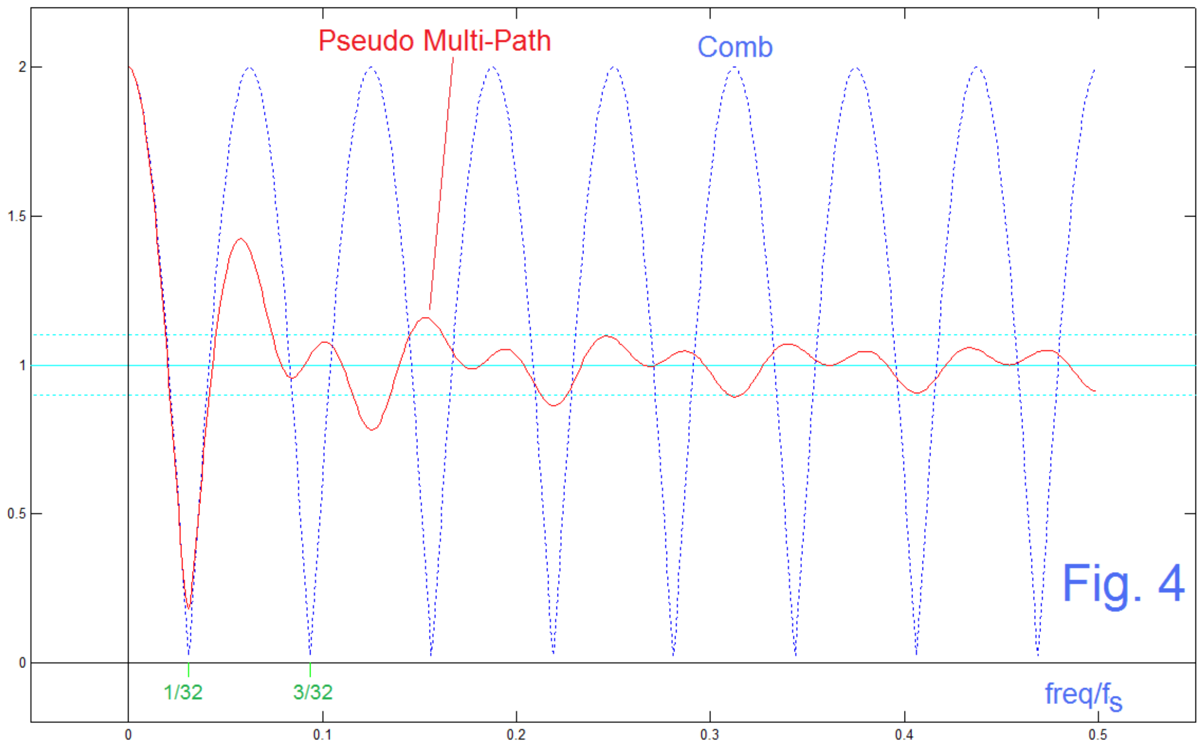
Pseudo Multi-Path
Comb
Fig. 4
1/32
3/32
freq/fs
0
0.1
0.2
0.3
0.4
0.5
2
1.5
1
0.5
0

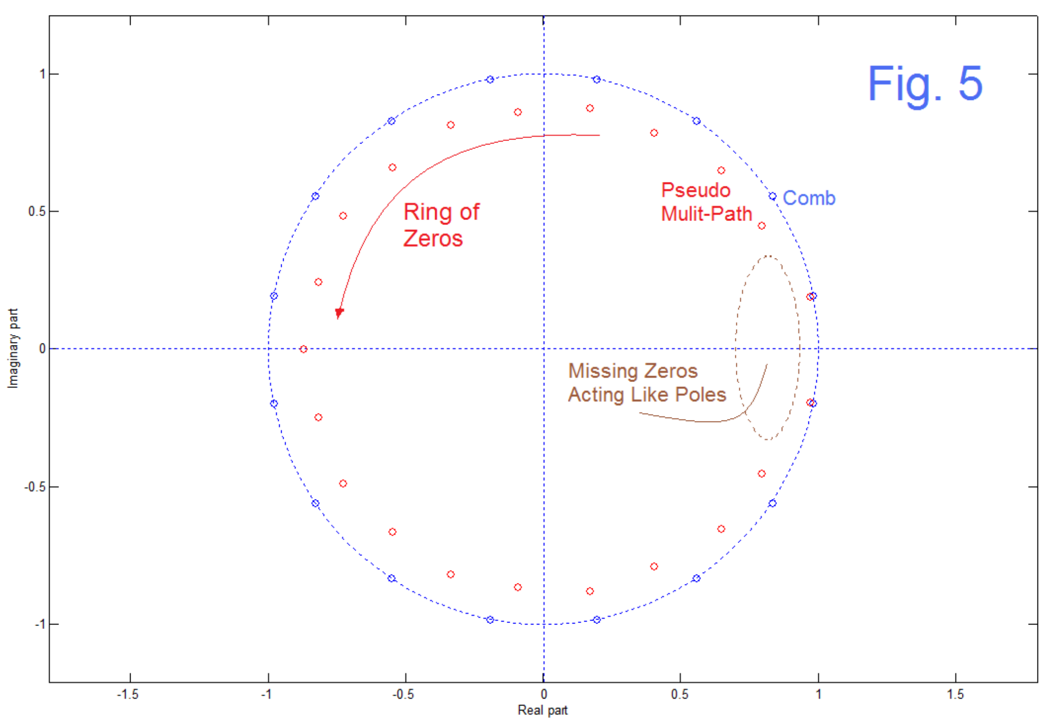
Fig. 5
Ring of Zeros
Pseudo Mulit-Path
Comb
Missing Zeros Acting Like Poles
Imaginary part
Real part
-1.5
-1
-0.5
0
0.5
1
1.5

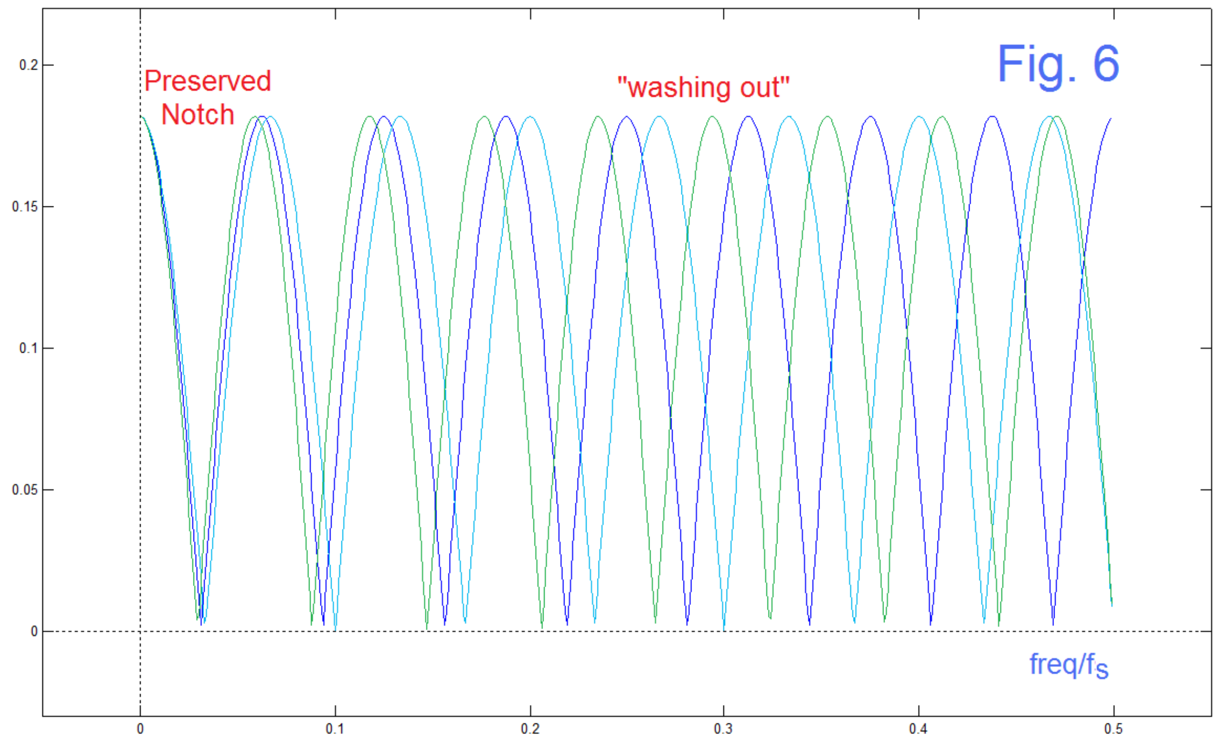

We are not bragging about the quality of the notch in Fig. 4. Instead we are indicating how a multi-path summation could result in a single notch-like response instead of a full comb. Below we will review the simplest way of obtaining a proper FIR design for a similar notch. Fig. 5 and Fig. 6 provide insight into how the response of Fig. 4 came about.

Fig. 5 shows the positions of the zeros for the two filters. The length-17 impulse response h_0 has 16 equally spaced zeros, all on the unit circle – the exact comb with eight notches. The pseudo multi-path case h_1 had no unit circle zeros (red) although the notch for the lowest frequency is close – having moved inside the circle only slightly (the deepest point in that notch is down below about 0.2). Being length 22, h_1 has 21 zeros total. Note that the zeros all tend to be inside the unit circle, somewhat equally spaced on a reduced circle, except for the two notch-producing zeros. The more familiar case of a notch that would be similar to Fig. 4 (red) would be a conjugate pair of zeros slightly inside the unit circle (which we do have) and a pair of poles, slightly further in at about the same angle, thus sharpening this notch. That is, it would be a 2nd-order IIR (Infinite Impulse Response) realization. Here instead we have a ring of relatively close zeros with one pair of zeros missing; the missing zeros having an effect similar to a pair of poles [4].

Fig. 6 shows the way in which notches with adjacent delays tend to cancel in the upper frequency regions, as compared to the lowest frequency notch. Here we show notches at delays 15, 16, and 17 (light blue, blue, and green) as individual responses. In Fig. 4 (red) we have 11 delays from 11 to 21 all summed. (The individual notches in Fig. 6 of course sum the delay with a zero-delay weighted at a=1/11, not at 1.)

A DESIGN USING “FREQUENCY SAMPLING”

Above we were not really treating the sum of multi-paths as though it were supposed to be designing a notch filter. Rather we wanted to emphasize that the response turned out to seem to want to have a notch-like character. Here we will want to try to intentionally use a standard FIR design procedure and see if anything like the features we observed as “accidental” appear. Likely the most basic FIR design procedure is the use of the inverse FFT (or similar) known as Frequency Sampling (FS). This procedure is straightforward, but there are pitfalls to be avoided [5, 6].

The simplest FS procedure here is to begin with a length 32 vector of samples for a frequency response magnitude as (presented in Matlab code):

```
A=[2 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 ];
```

where we have chosen the DC (k=0) value as 2 to match the original response, and set the notch to 0 at k=1 and k=31, and made all other samples 1. We then establish the correct phase delay:

```
n=[0:31];
A = A.*exp(-j*2*pi*n/32)
```

and take the inverse FFT (and take off a very tiny accidental imaginary part):

```
hfs = ifft(A)
hfs = real(h)
```

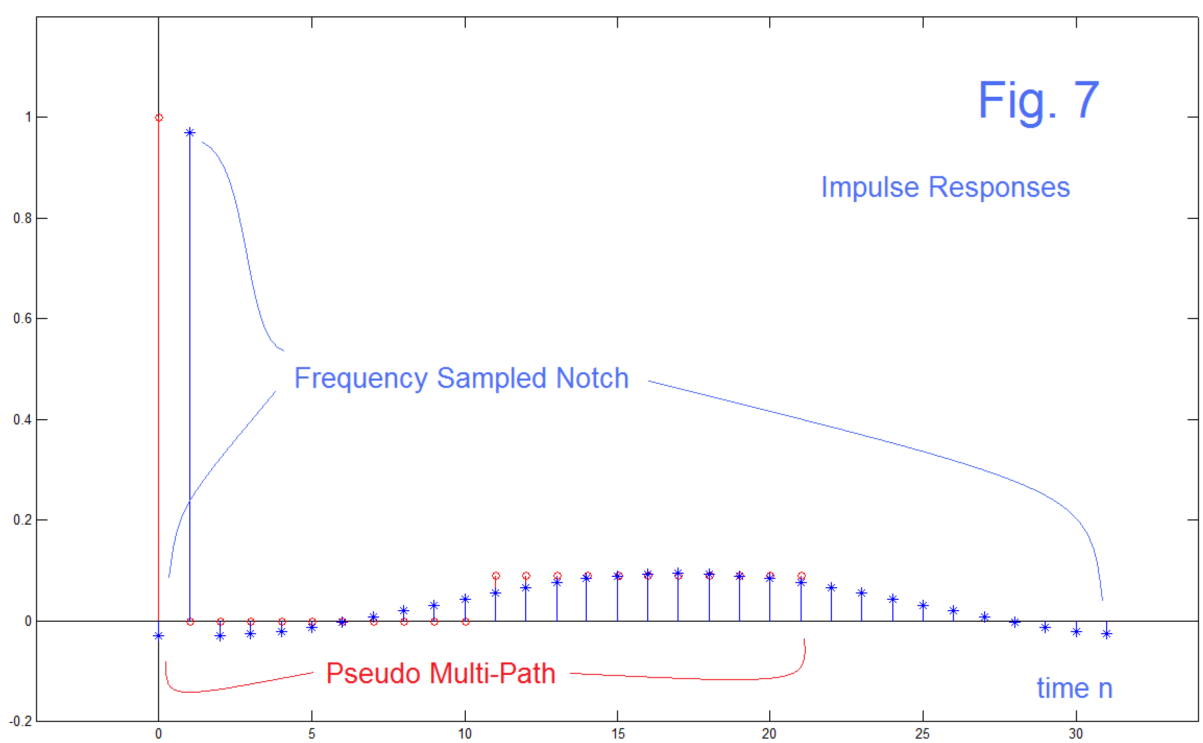

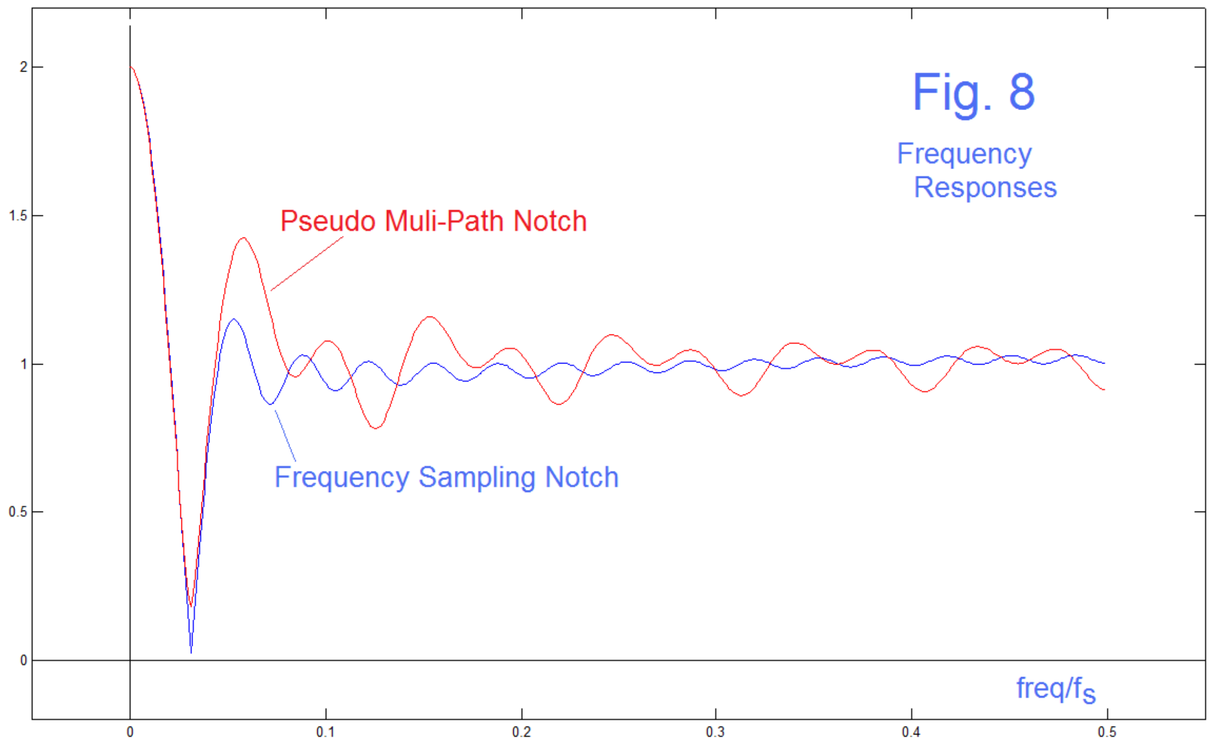

Fig. 7 shows the resulting impulse response, Fig. 8 the resulting frequency response, and Fig. 9 shows the zeros of the FS design (blue), all as compared to the pseudo multi-phase (red) of Fig. 4.

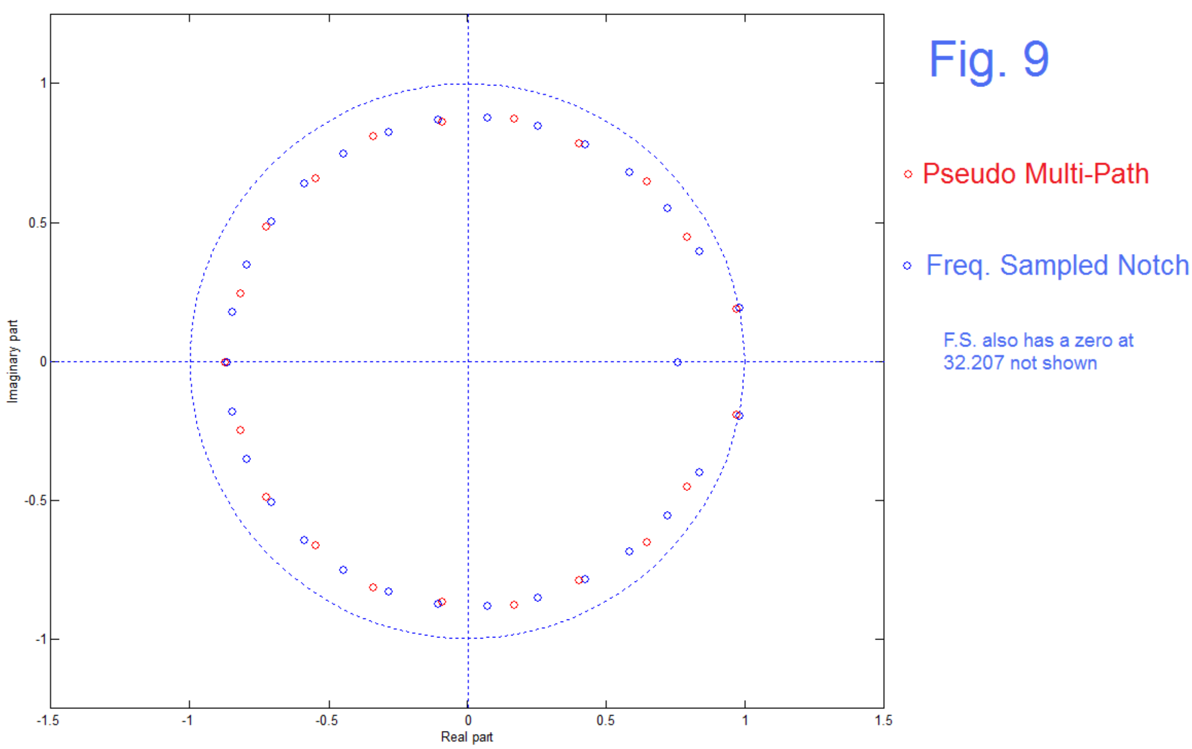

We see that this comparison of the pseudo multi-path result is quite similar to the FS design. The FS has more zeros due to the somewhat trial-and-error selection of the exact parameters, but it otherwise well matched in frequency response (Fig. 8) and in having the same general "ring of zeros" (Fig. 9). This we anticipated. The shape of the impulse response (Fig. 7) was something we could not guess as well. But we see that the general profiles of the taps are not dissimilar.

All and all, the notion that a system with multiple paths could lead to a single non-comb-like notch seems to be promising.

REFERENCES

[1] B. Hutchins, "Notch Filtering – Some Notes," Electronotes Application Note No. 413 July 30, 2014
http://electronotes.netfirms.com/AN413.pdf

[2] B. Hutchins, "Analog Delay for Musical Engineering," ***Electronotes***, Vol. 7, No. 56, Aug. 1975

[3] B. Hutchins & W. Ku, "An Adapting Delay Comb Filter for the Restoration of Audio Signals Badly Corrupted with a Periodic Signal of Slowly Changing Frequency", ***J. Audio Eng. Soc***., Vol. 39, No 1/2, Jan/Feb 1982
http://electronotes.netfirms.com/AES5.PDF

[4] B. Hutchins, "FIR Filters With Sinusoidal Tap Weights, Electronotes Application Note No. 238, Nov. 18, 1981

[5] B. Hutchins, "Basic Elements of Digital Signal Processing: Filters Part 2" ***Electronotes***, Vol. 20, No. 198, June 2001, (Generalized Frequency Sampling, Part 3b, pp 12-20)
http://electronotes.netfirms.com/EN198.pdf

[6] B. Hutchins, "A General Review of Frequency Sampling Design", Electronotes Application Note 337, March 1996
http://electronotes.netfirms.com/AN337.pdf