Moving Average

- All weights of the moving average (MA) are set to one:

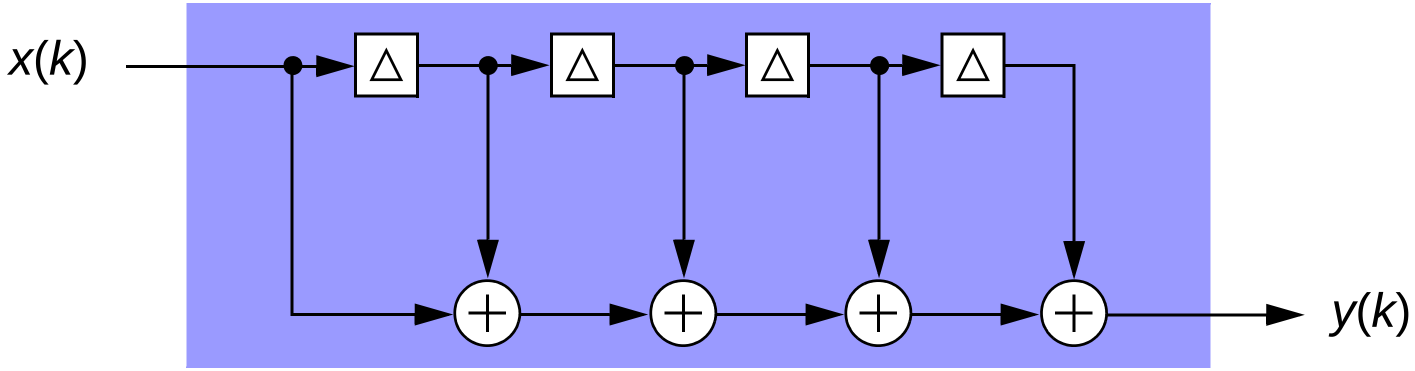

- Simple "low pass" characteristic
- Low cost - no multiplies required.

This filter might preferably be implemented use a power of two number of weights - why?

Top

Notes:

This filter is a very simple low pass characteristic.

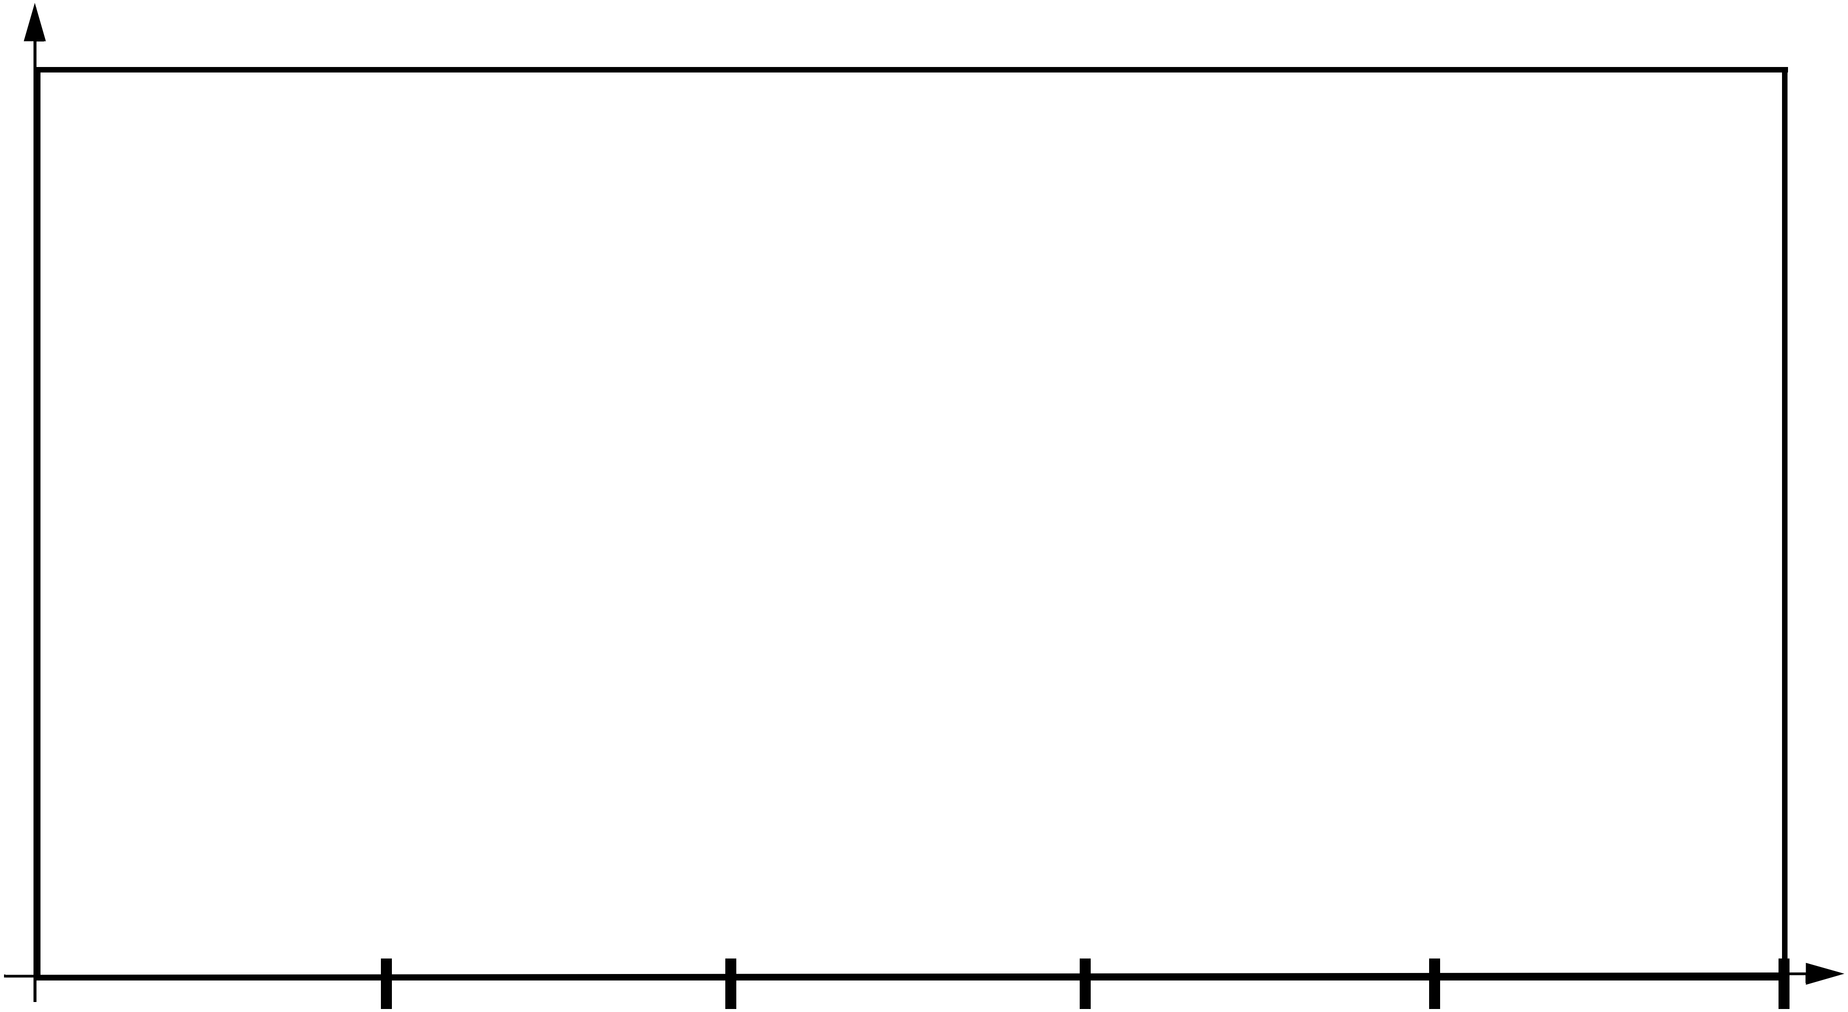

Differentiator

- Two weight filter, with values of 1 and -1:

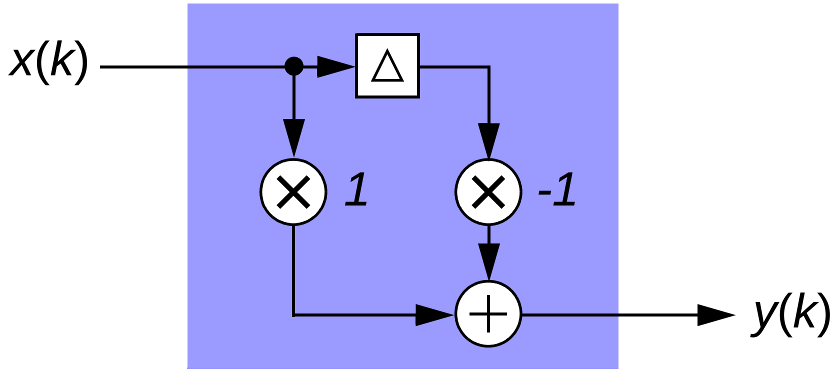

- Simple “high pass” magnitude response with no multiplies required.

- Output is: $y(k) = x(k) - x(k-1)$ and in the z-domain:

$$Y(z) = X(z) - X(z)z^{-1} \qquad \Rightarrow Y(z) = X(z)[1 - z^{-1}]$$

and hence the differentiator transfer function is:

$$H(z) = \frac{Y(z)}{X(z)} = 1 - z^{-1}$$

Top

Notes:

Inputing a constant value, ie. DC or 0 Hz will result in no output appearing after an initial transient. Hence there is a spectral zero at 0Hz, i.e. a spectral zero is where the gain is precisely 0 in a linear scale, and if we attempt to represent in a log scale: $20\log 0 = -\infty$.

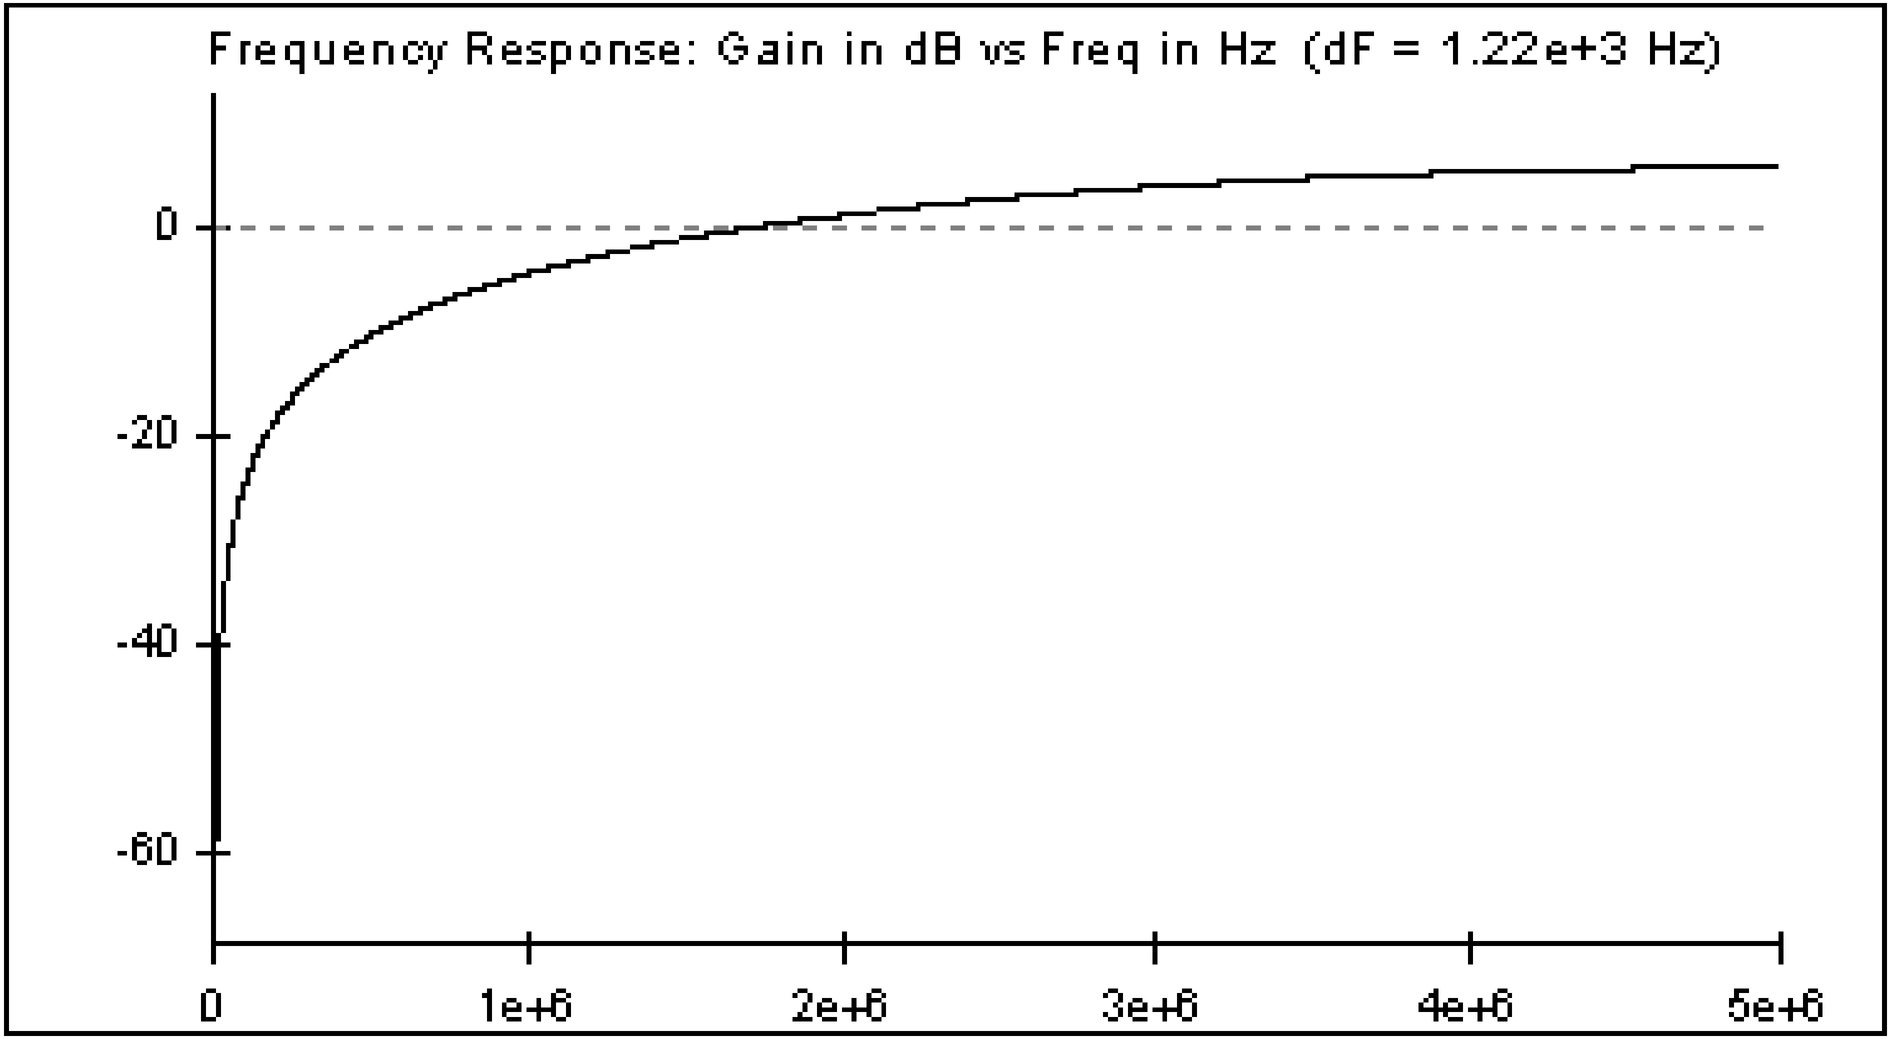

Developed by: www.steepestascent.com

Integrator

- Integrator is a single weight IIR filter:

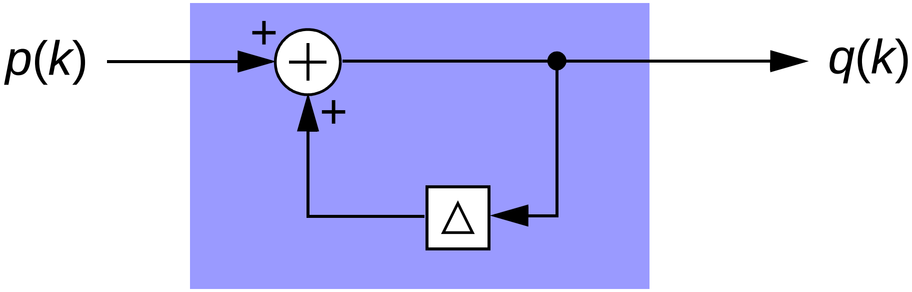

- "Low pass" (infinite gain at DC) with ***no*** multiplies required.

- Output in the time domain is: $q(k) = p(k) + p(k-1)$ and in the z domain:

$$Q(z) = P(z) + Q(z-1) \qquad \Rightarrow Q(z)[1 - z^{-1}] = P(z)$$

and hence the integrator transfer function is:

$$G(z) = \frac{Q(z)}{P(z)} = \frac{1}{1 - z^{-1}}$$

Top

Notes:

If a feedback weight of b is introduced, where $|b| < 1$ this is often refered to as a leaky integrator. Generally speaking for DSP for FPGAs/ASICs we will not be concerned with leaky integrators. If $|b| > 1$ then the filter would have a pole outside of the unit circle and would be diverging or unstable. .

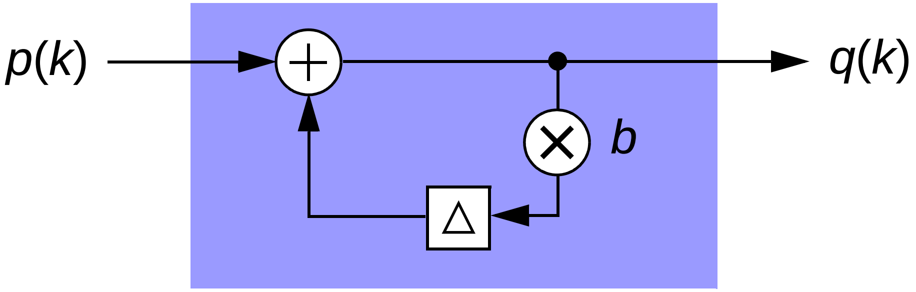

An integrator and a differentiator are clearly perfect inverses of each other. From a spectral point of view it is interesting to note that the differentiator has infinite attenuation at 0 Hz and the integrator has infinite gain at 0 Hz,.... and any engineer knows infinity multiplied by zero, might just be 1 in many cases!

$$G(z)H(z) = \left(\frac{1}{1-z^{-1}}\right)(1-z^{-1}) = 1$$

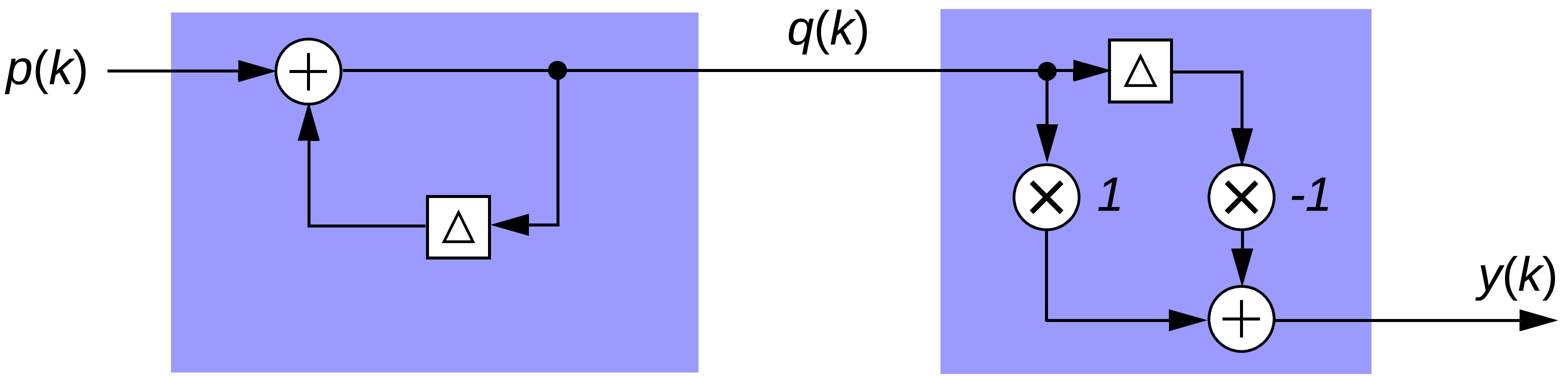

$$y(k) = q(k) - q(k-1) = [p(k) + q(k-1)] - q(k-1) = p(k)$$

Comb Filter

- Weights set to 1 and -1 at either end of the filter.

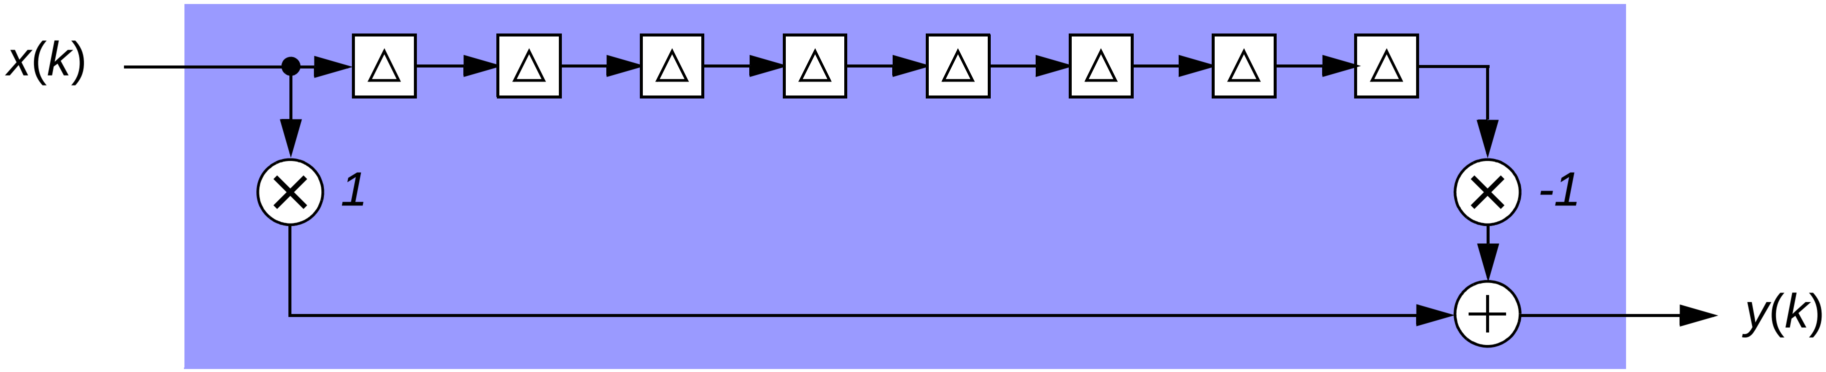

- Simple multichannel frequency response - ***no*** multiplies required.
- Using the z-notation to represent the 8 delays we can show as:

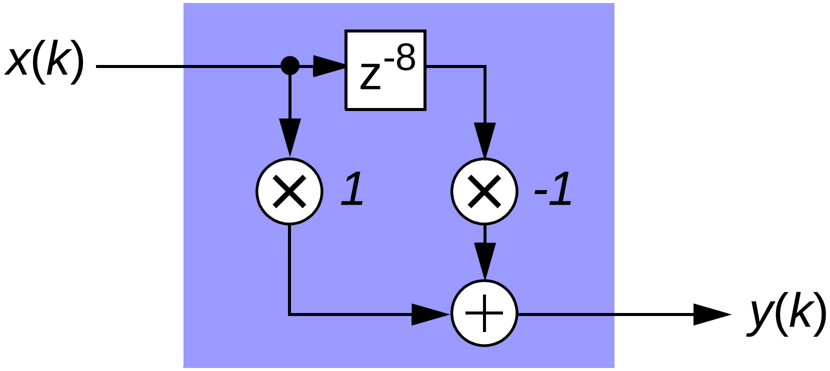

Top

Notes:

A comb filter with *N* sample delays (or *N*+1 weights) will have N evenly spaced spectral zeroes from 0 to $f_s/2$. Therefore the 8 delay comb filter above will have 4 spectral zeroes from 0 to 5 MHz, at spacings of 1.25MHz, when the sample rate is set to f_s = 10MHz.

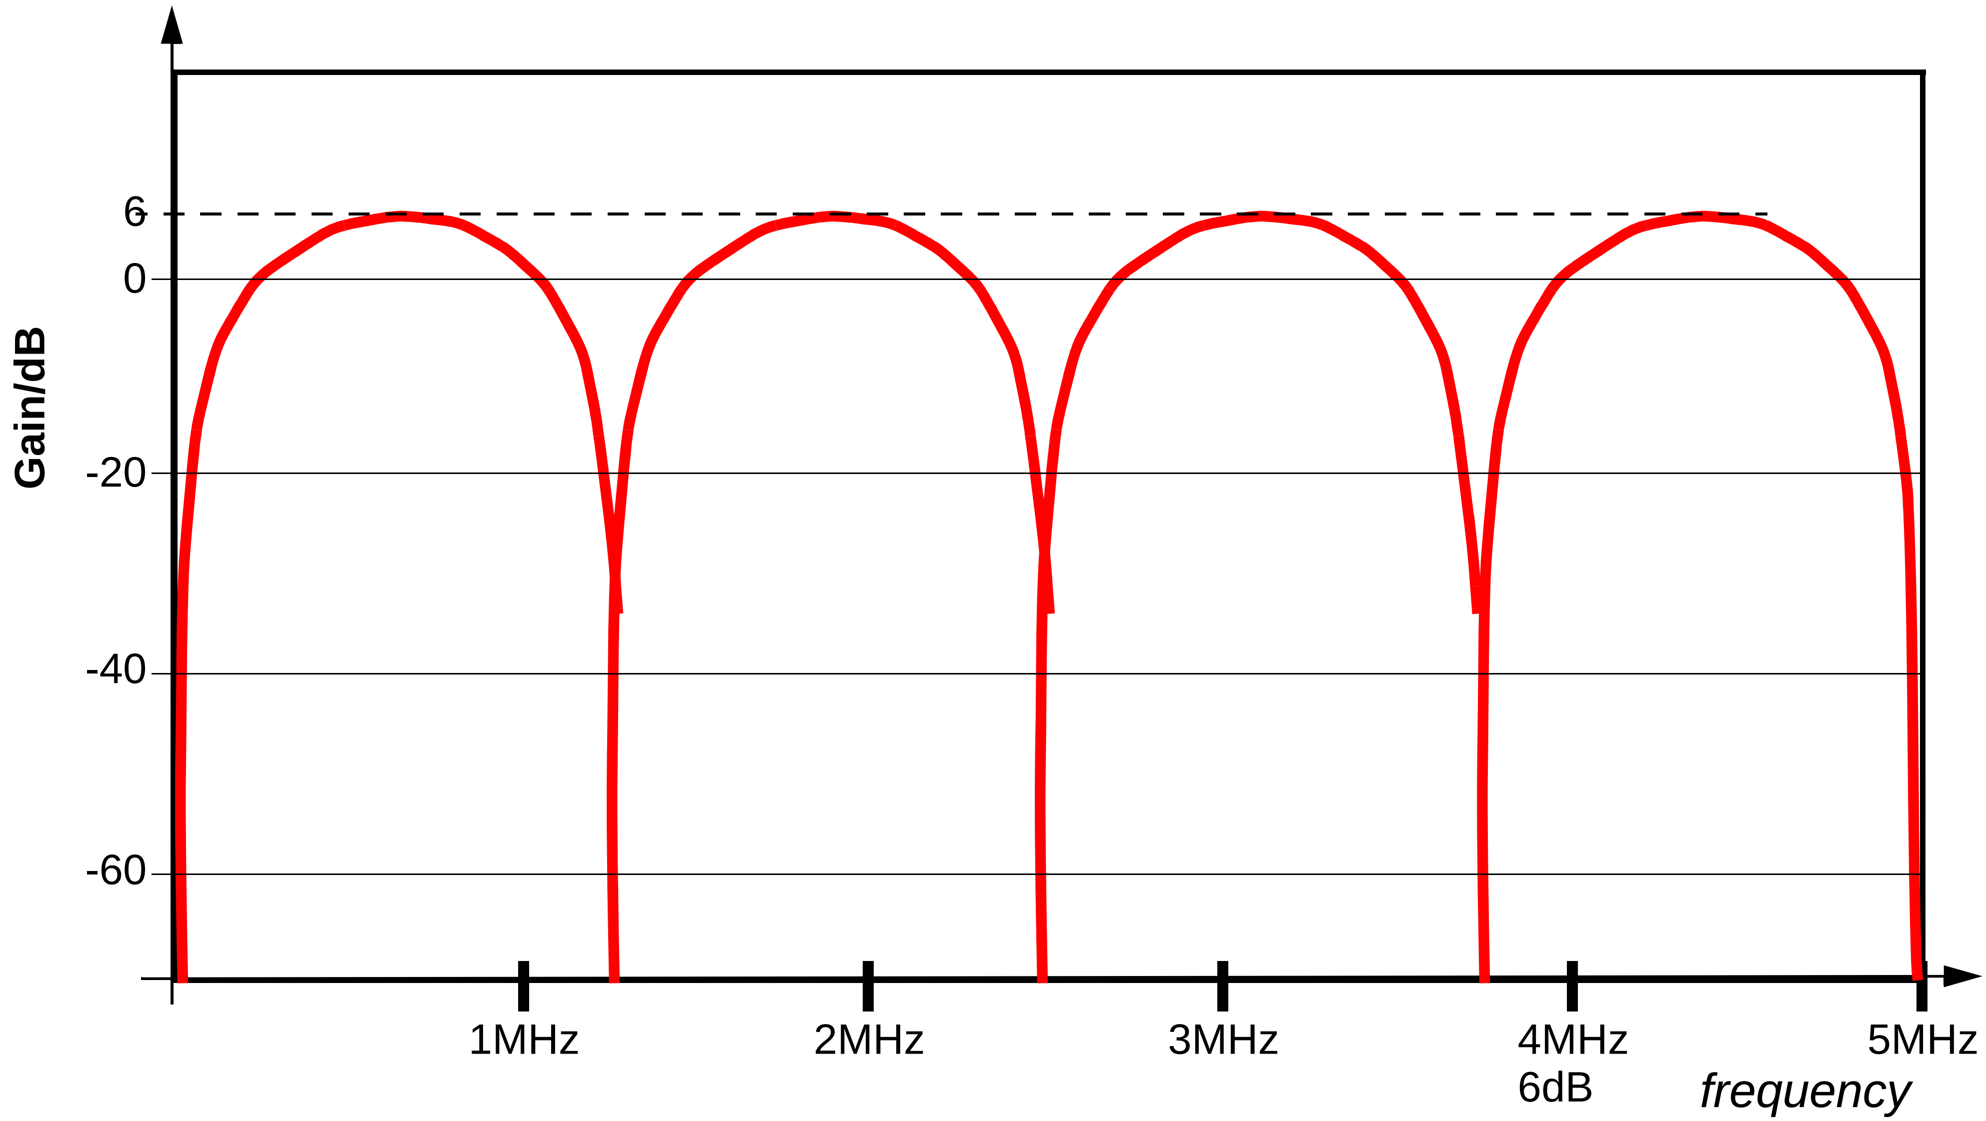

Eight Weight Moving Average

- Consider again the moving average (MA); all weights of “1”

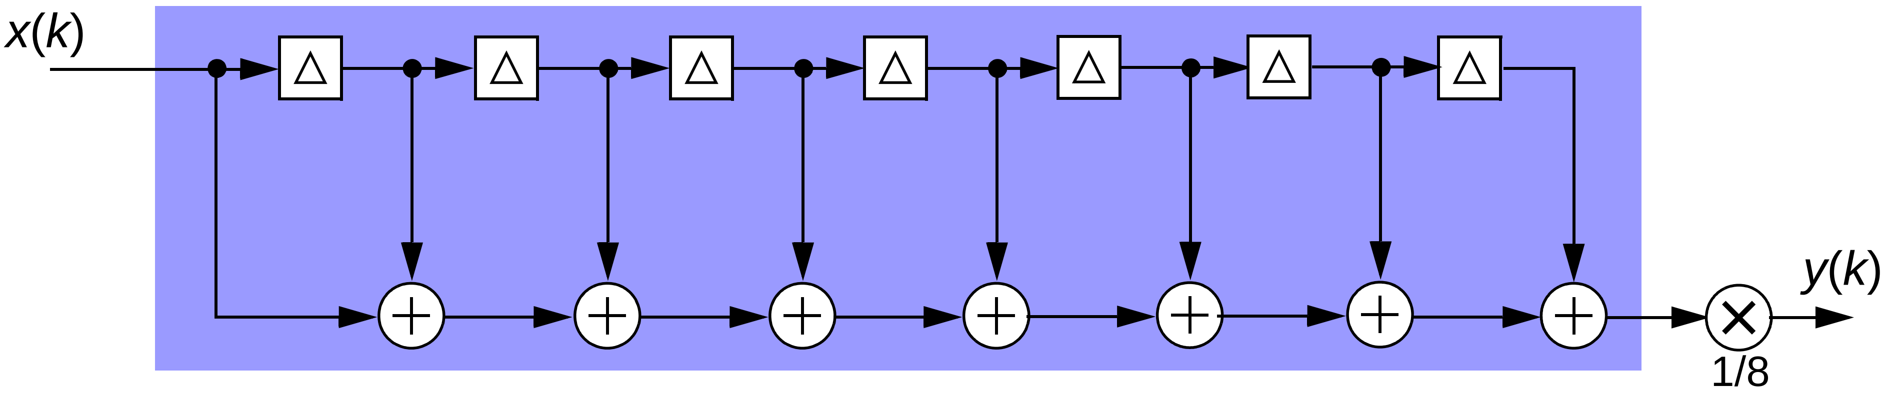

$$H(z) = 1 + z^{-1} + z^{-2} + z^{-3} + z^{-4} + z^{-5} + z^{-6} + z^{-7}\frac{1}{8}$$

- True moving average if we scale the output by $\frac{1}{8}$ (left shift 3 places)- equivalent to all weights being 1/8.

- In the spectrum the moving average filter has *N-1* spectral zeroes from 0 to f_s. In our case $N = 8$, we can see 4 spectral zeroes from 0 to $f_s/2$.

Top

Notes:

To allow a numerical representation, we choose $f_s = 10,000,000$

We can see four spectral zeroes between 0 and $f_s/2$, i.e. 8-1=7 spectral zeroes between 0 and f_s.

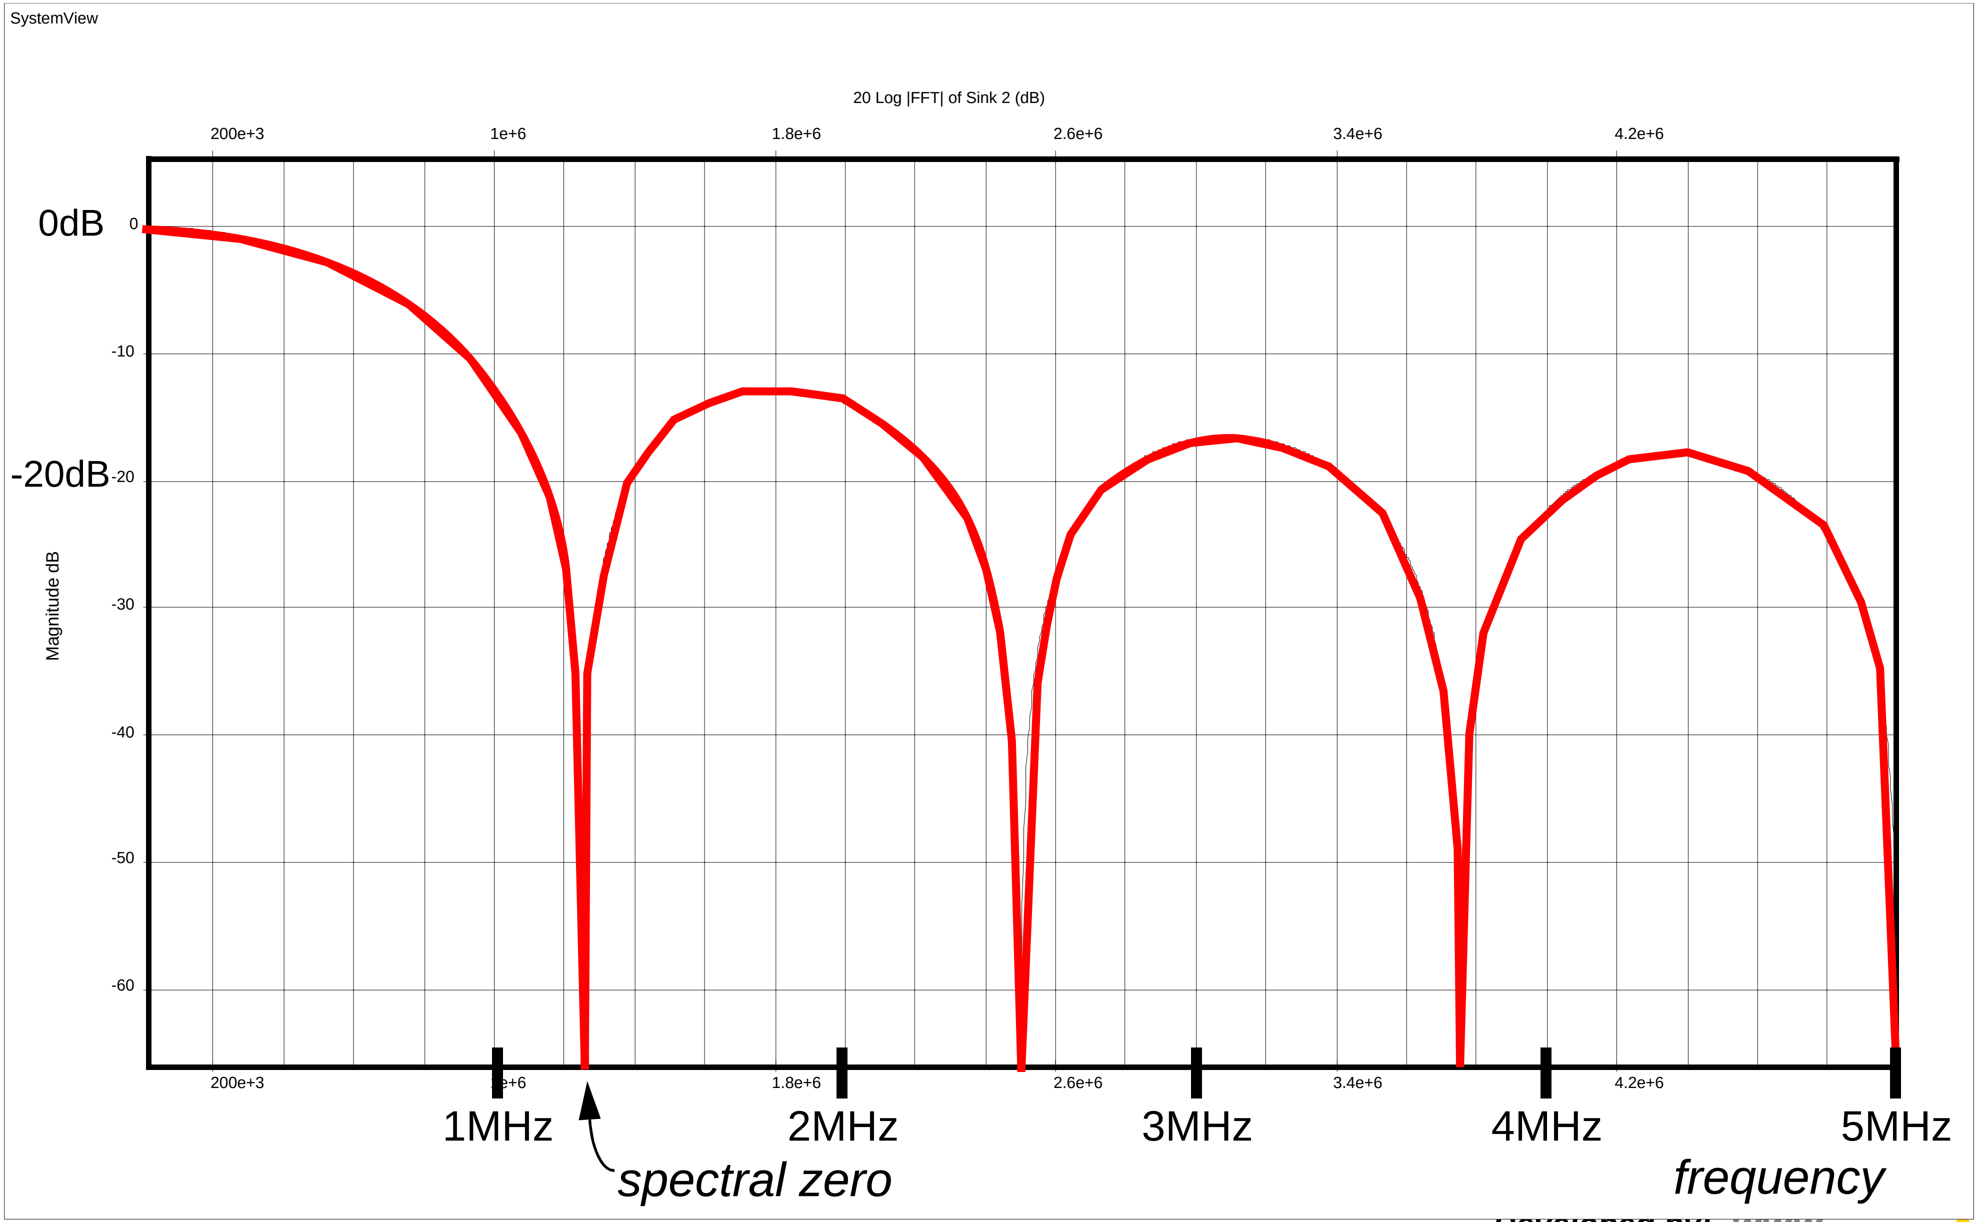

Nine Weight Moving Average (MA)

- All weights are "1"

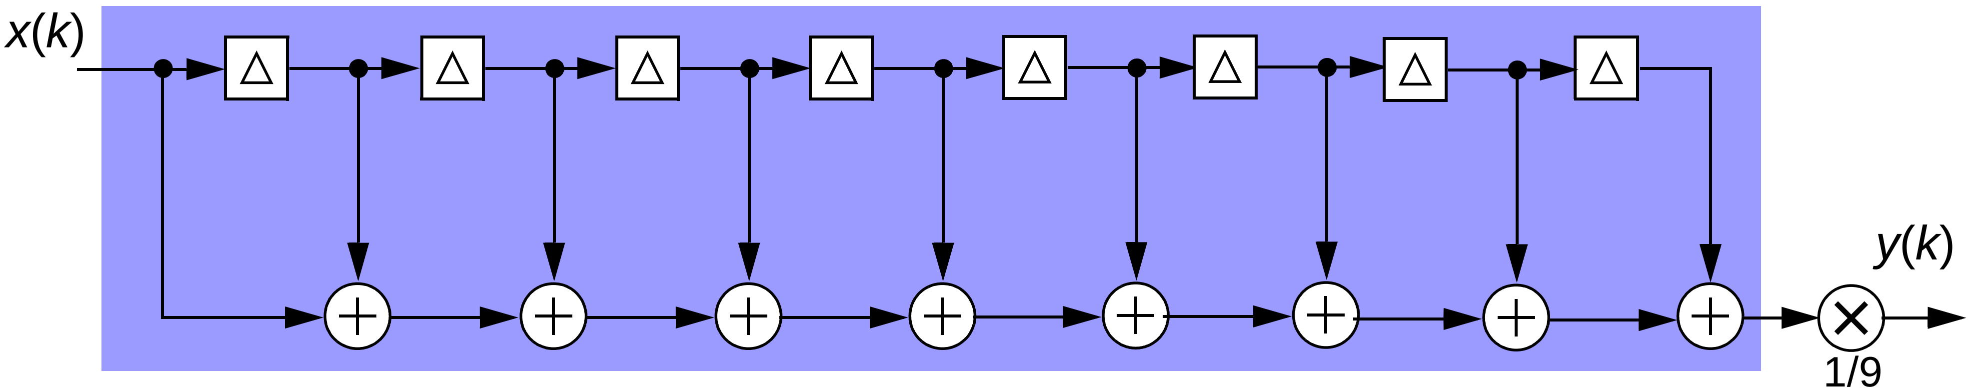

$$H(z) = (1 + z^{-1} + z^{-2} + z^{-3} + z^{-4} + z^{-5} + z^{-6} + z^{-7} + z^{-8})\frac{1}{9}$$

- Multiplying by 1/9 is not so convenient.....

Notes:

For ease of numerical representation, we choose $f_s = 10,000,000$

We can see 4 spectral zeroes between 0 and $f_s/2$, i.e. 9-1=8 spectral zeroes between 0 and f_s.

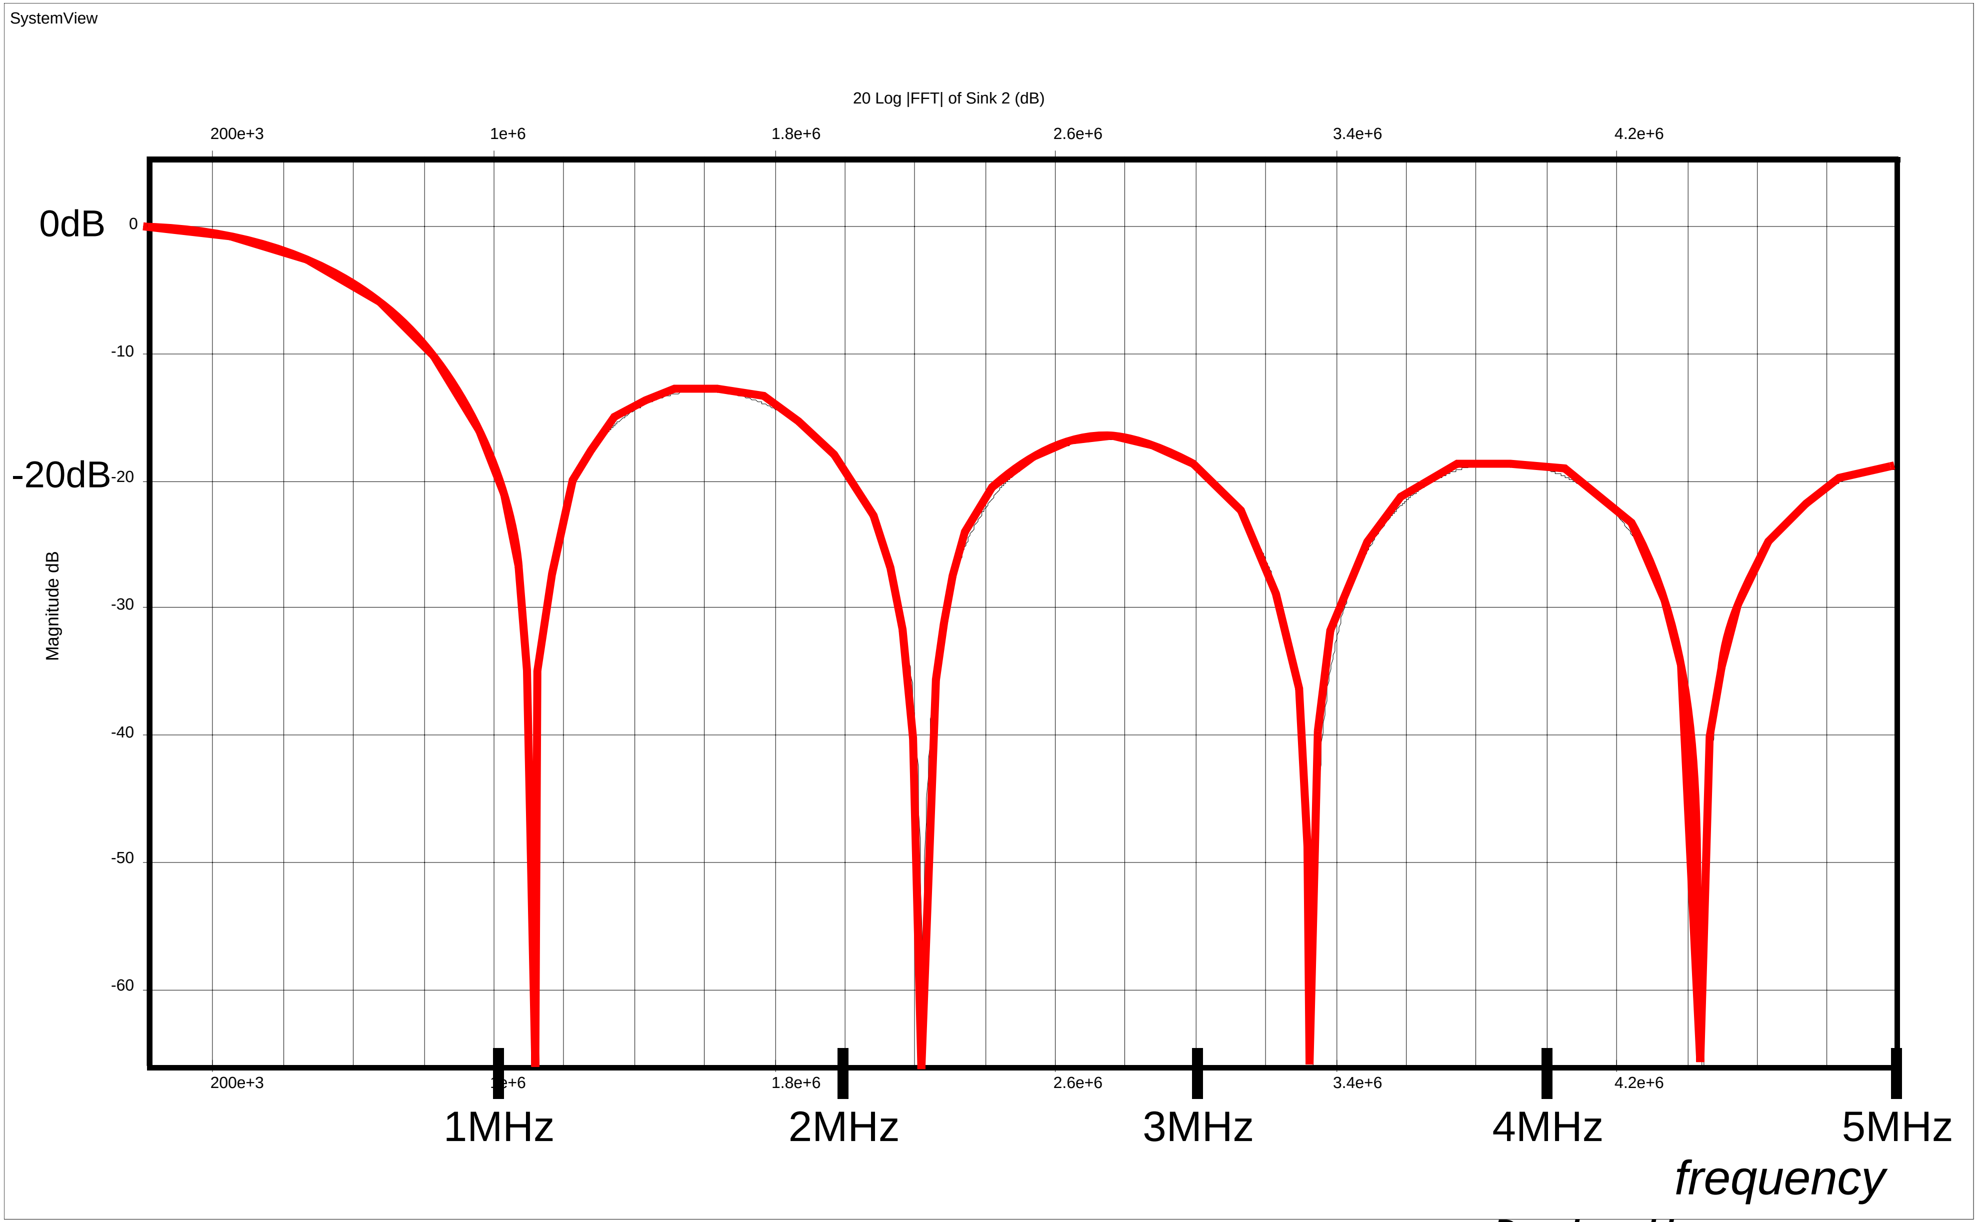

Cascade Integrator Comb (CIC)

- Generate a MA impulse response with CIC structure (see Slide 7.6)

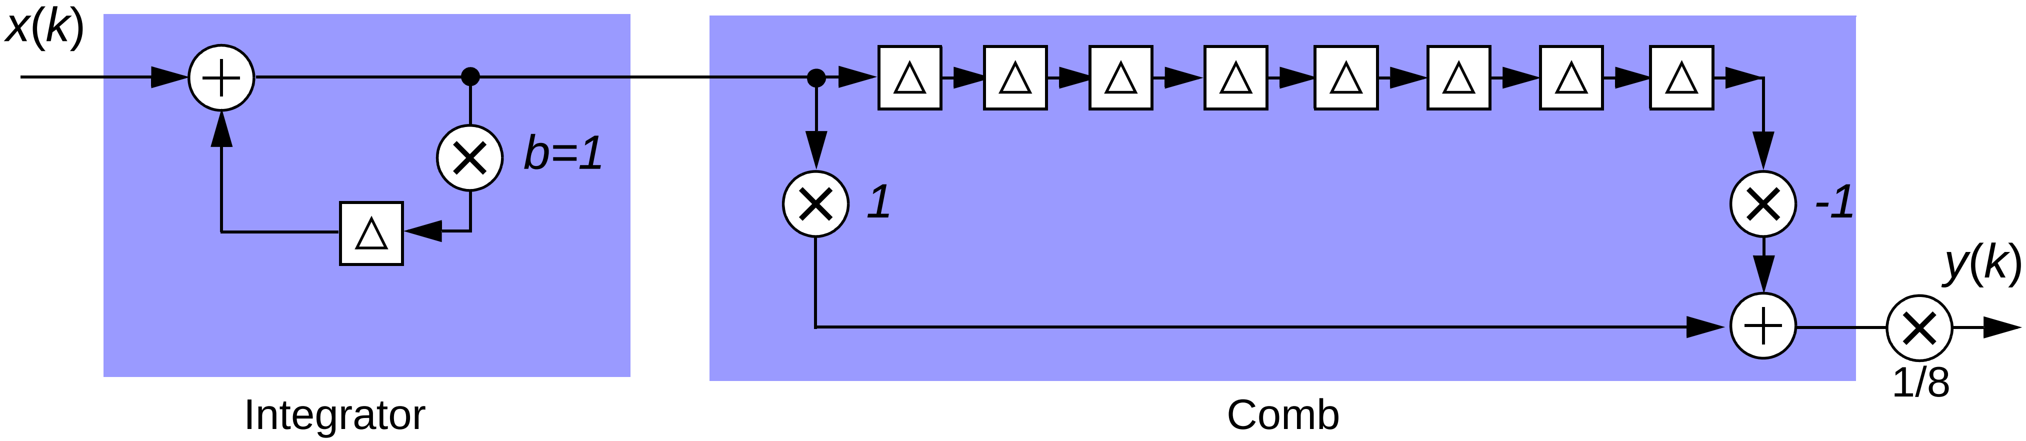

$$H(z) = \left(\frac{1}{1-z^{-1}}\right)(1-z^{-8}) = \frac{1-z^{-8}}{1-z^{-1}}$$

- Note that: $\frac{1-z^{-8}}{1-z^{-1}} = 1 + z^{-1} + z^{-2} + z^{-3} + z^{-4} + z^{-5} + z^{-6} + z^{-7}$

- i.e ***an integrator and*** ***M comb weight CIC*** = ***M-1 weight MA***

Notes:

$$\frac{1 - z^{-8}}{1 - z^{-1}} = 1 + z^{-1} + z^{-2} + z^{-3} + z^{-4} + z^{-5} + z^{-6} + z^{-7}$$

$$1 - z^{-8} = (1 + z^{-1} + z^{-2} + z^{-3} + z^{-4} + z^{-5} + z^{-6} + z^{-7})(1 - z^{-1})$$

$$\begin{aligned} 1 - z^{-8} = {} & 1 + \cancel{z^{-1}} + \cancel{z^{-2}} + \cancel{z^{-3}} + \cancel{z^{-4}} + \cancel{z^{-5}} + \cancel{z^{-6}} + \cancel{z^{-7}} \\ & - \cancel{z^{-1}} - \cancel{z^{-2}} - \cancel{z^{-3}} - \cancel{z^{-4}} - \cancel{z^{-5}} - \cancel{z^{-6}} - \cancel{z^{-7}} - z^{-8} \end{aligned}$$

It is interesting to note that the integrator has infinite gain at DC and the comb filter has zero gain an DC!

CIC Advantages: CIC has Only two additions compare to 8 additions in MA.

CIC Disadvantages: CIC requires 9 storage registers, and MA requires only 7 storage register.

Integrator Overflow

- The integrator of the CIC has infinite gain at DC (0 Hz).
- Therefore consider the input of a ***step signal*** to the CIC:

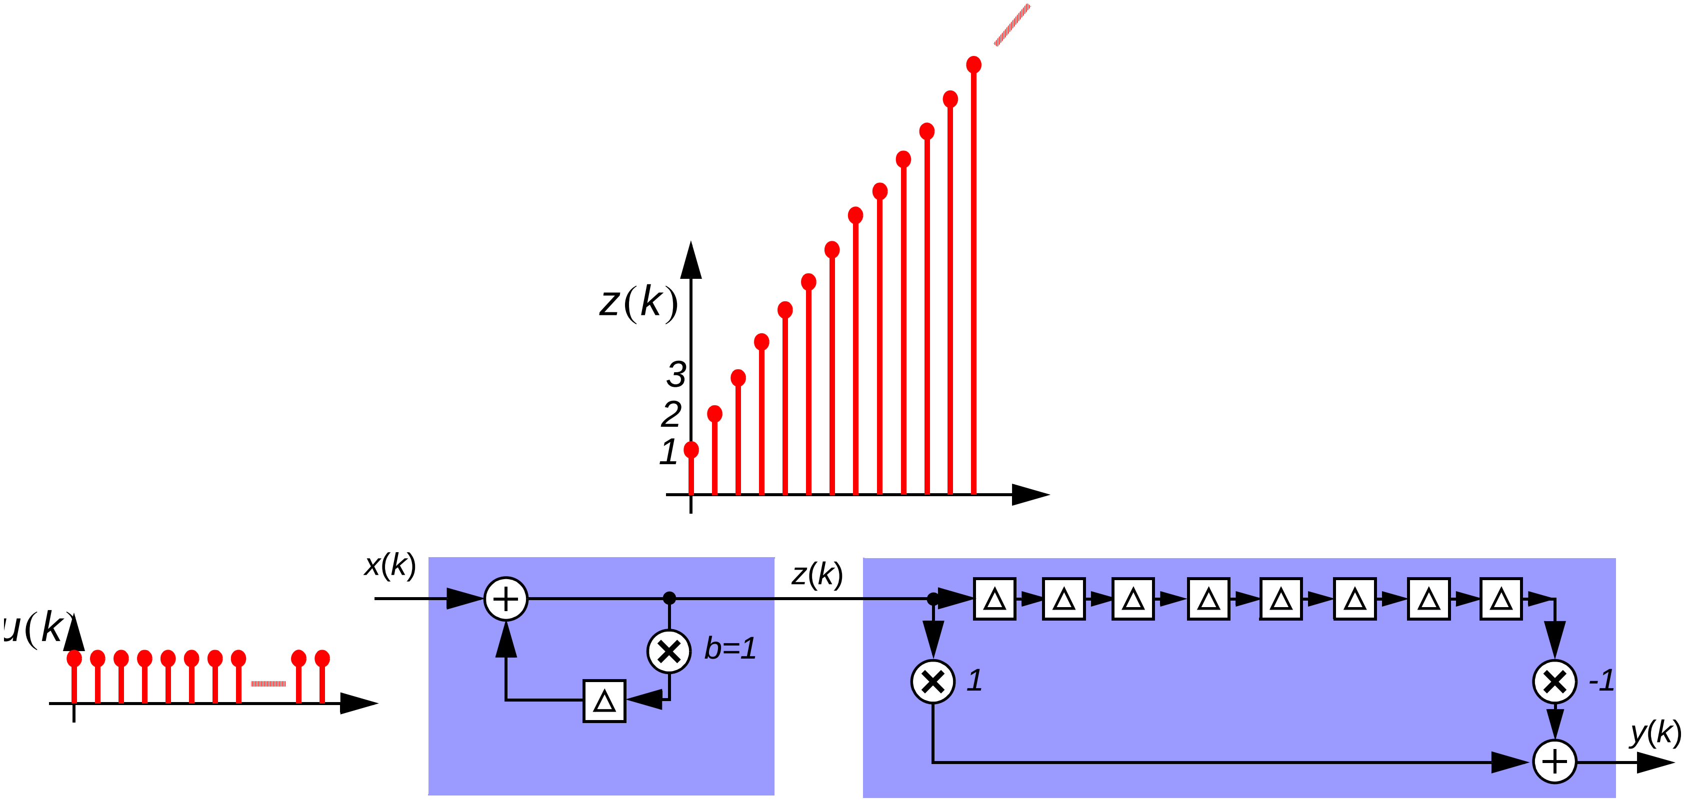

- The integrator output "grows" unbounded for the step input.

Notes:

Eventually the integrator output will overflow.....

To address this we can use modulo arithmetic.

Cascade of CICs

- We can cascade CIC filters to produce "better" low pass characteristics:

- Cascade of 5 CICs of 8th order MA filters:

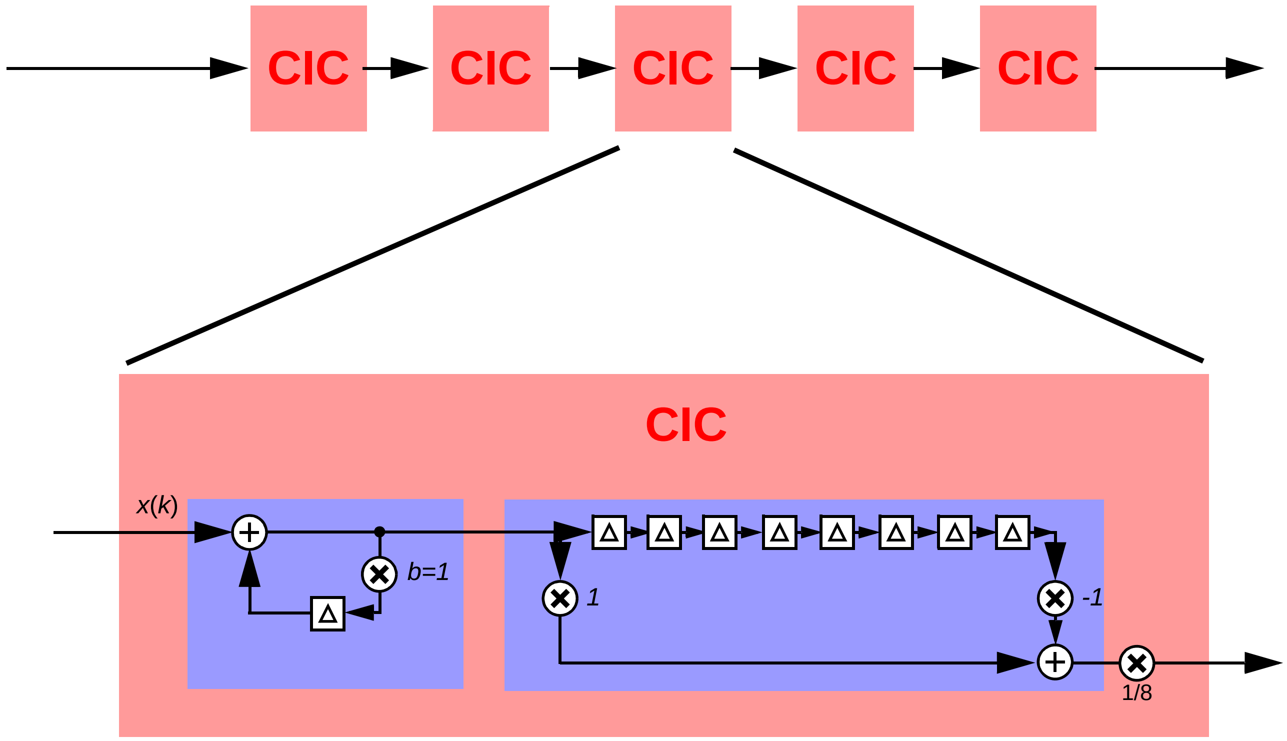

- Note however the baseband droop is "worse".

Top

Notes:

Plots of CIC and cascade of 5 CICs for 8th order moving average.

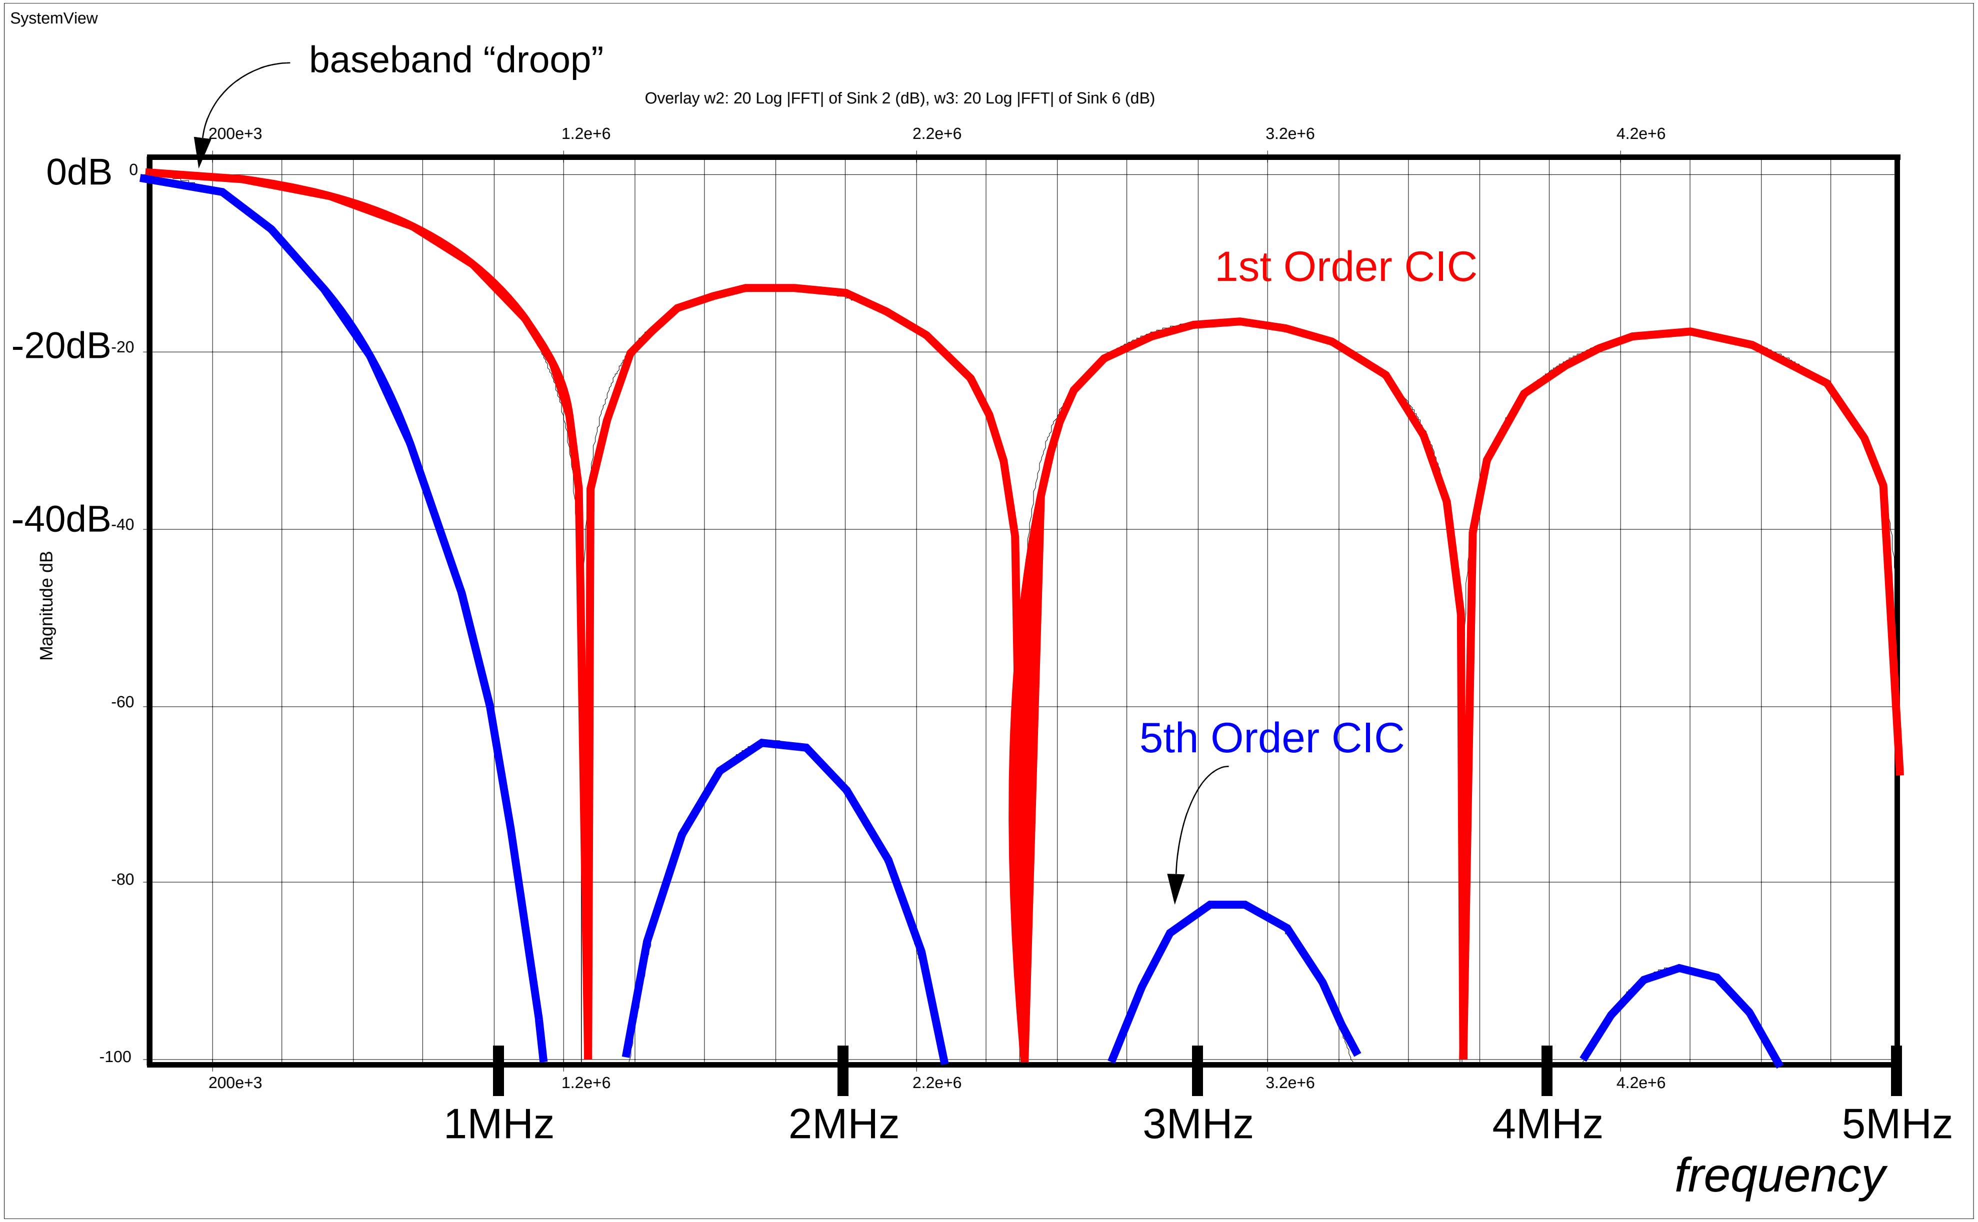

Developed by: www.steepestascent.com

Recovery of an IF modulated Signal

- Consider the following scenario:
  - Signal of interest centered at $f_c = 2.5\text{MHz}$
  - Signal bandwidth = 100kHz
  - Sampling rate, $f_s = 10\text{MHz}$

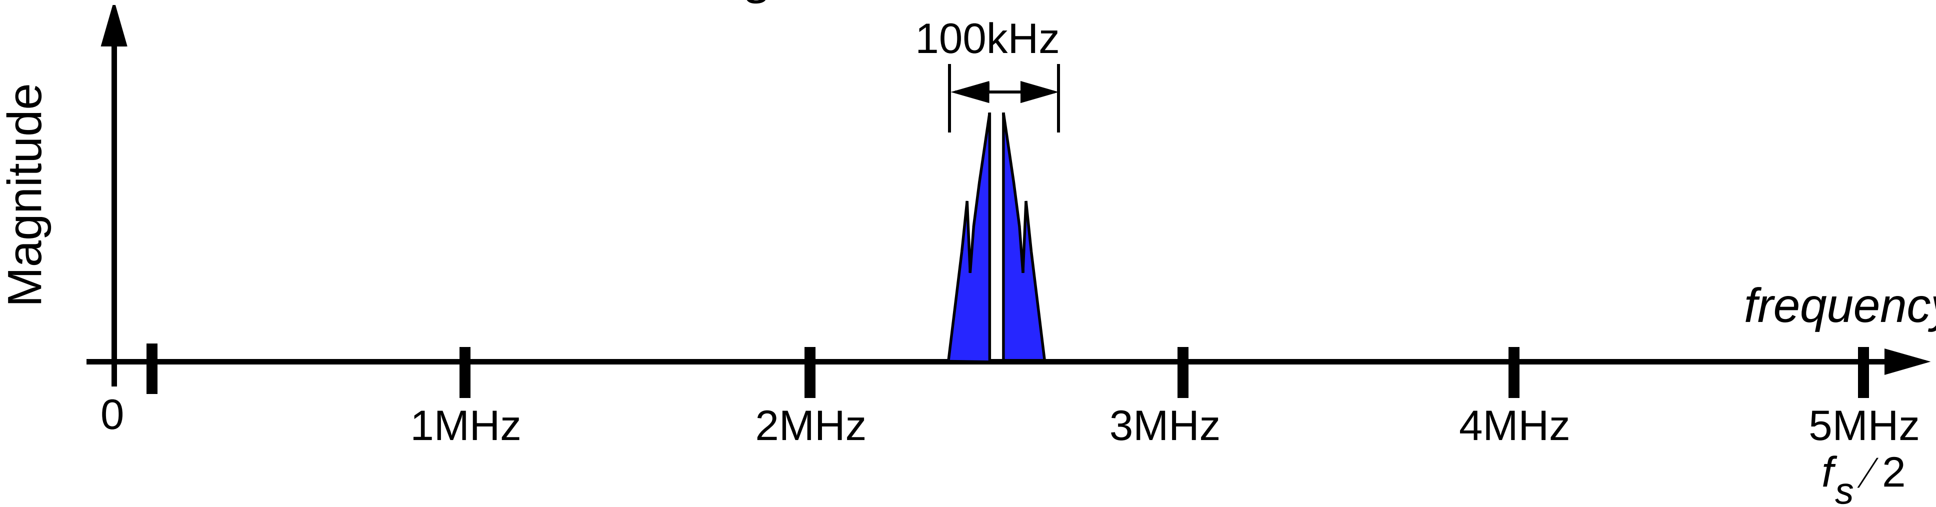

- Requirement is to recover the IF signal at baseband frequencies using as low computation as possible.

Top

Notes:

This bandpass signal has been created by simple amplitude modulation:

Amplitude of a “high” frequency carrier sinusoid is varied in proportion to the amplitude of signal with lower frequency components.

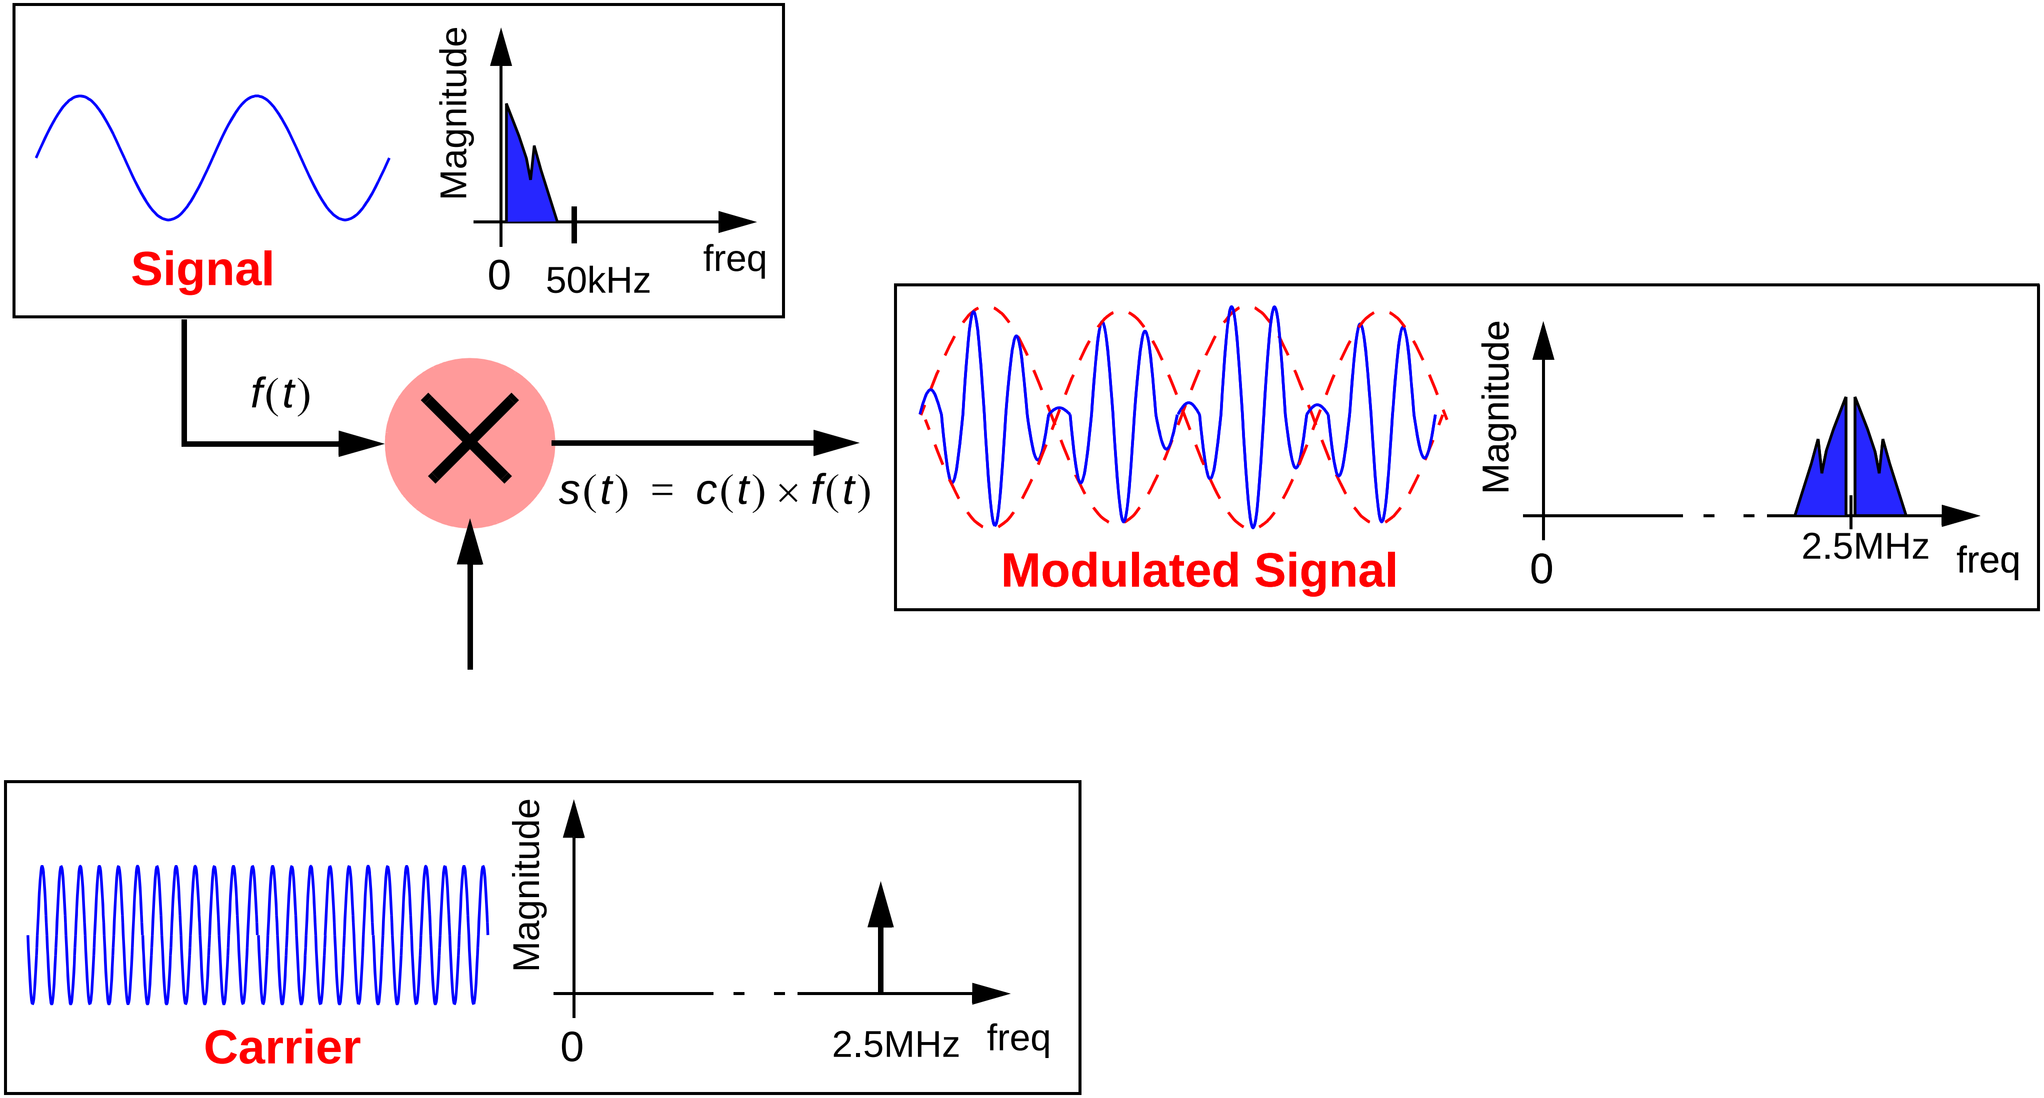

Recovery of an IF modulated Signal

- When the signal is received, the spectrum outside of the 50kHz band of interest is likely to be occupied with other signals and noise:

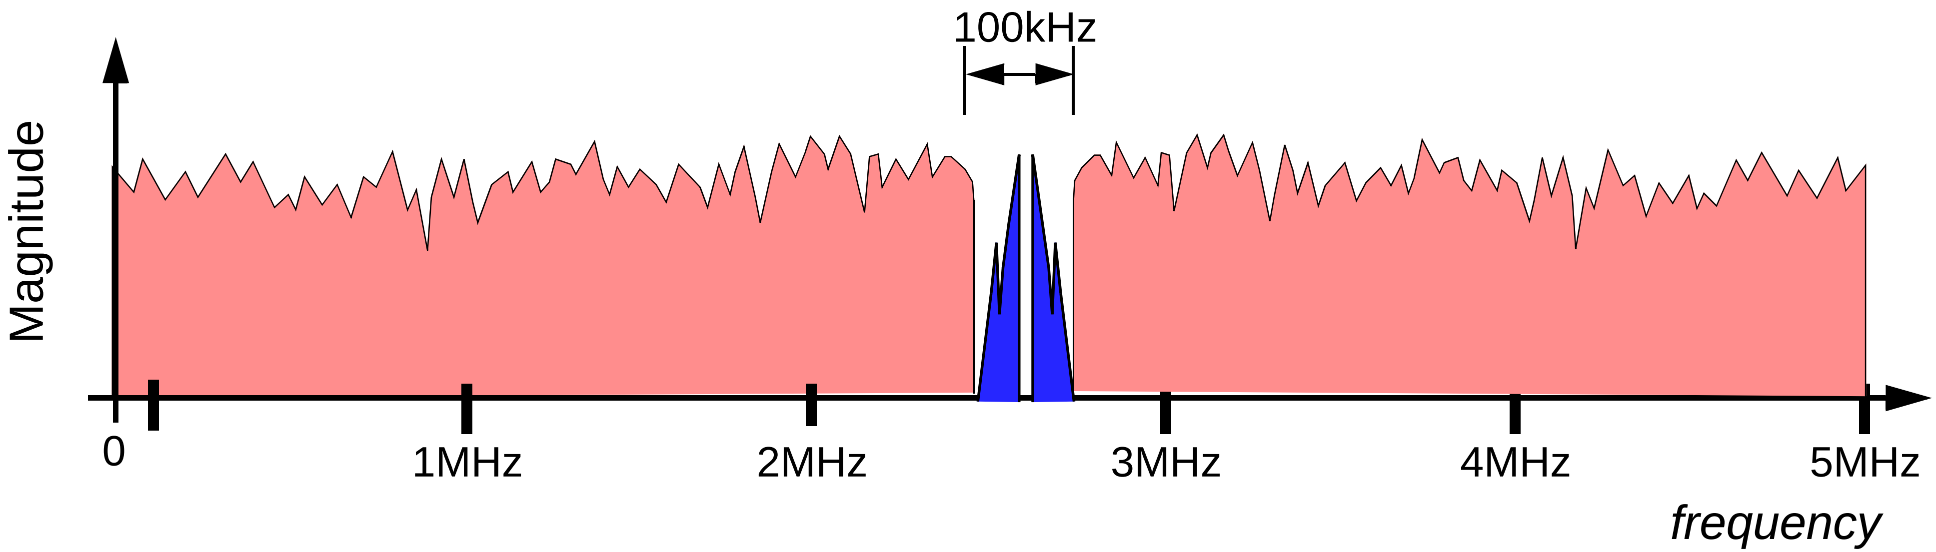

- To recover we require to demodulate to baseband and then low pass filter to recover the signal.

Notes:

Demodulation of Signal

- Sampling with a high frequency ADC we can first digitally demodulate the signal:

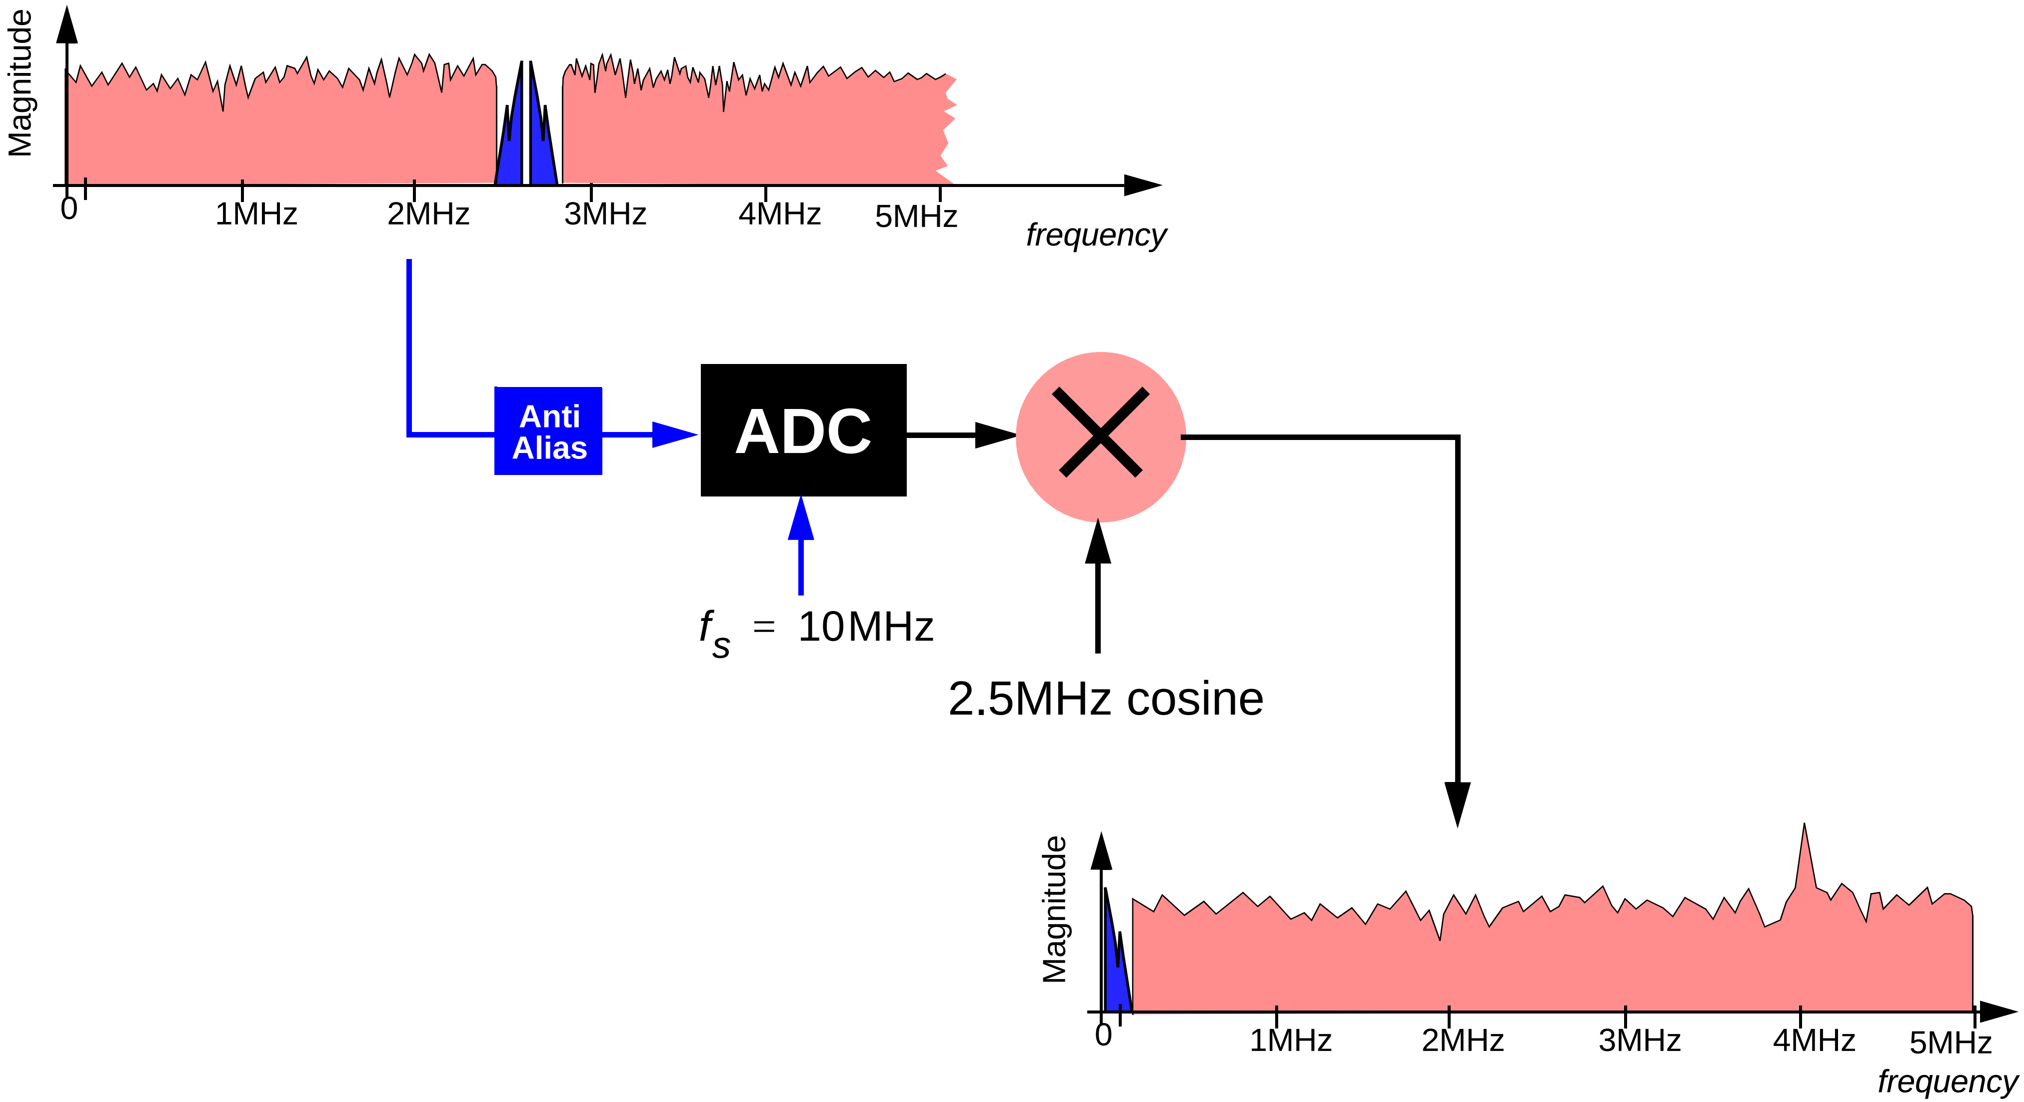

Notes:

Top

Demodulation of Signal

- ....then low pass filter:

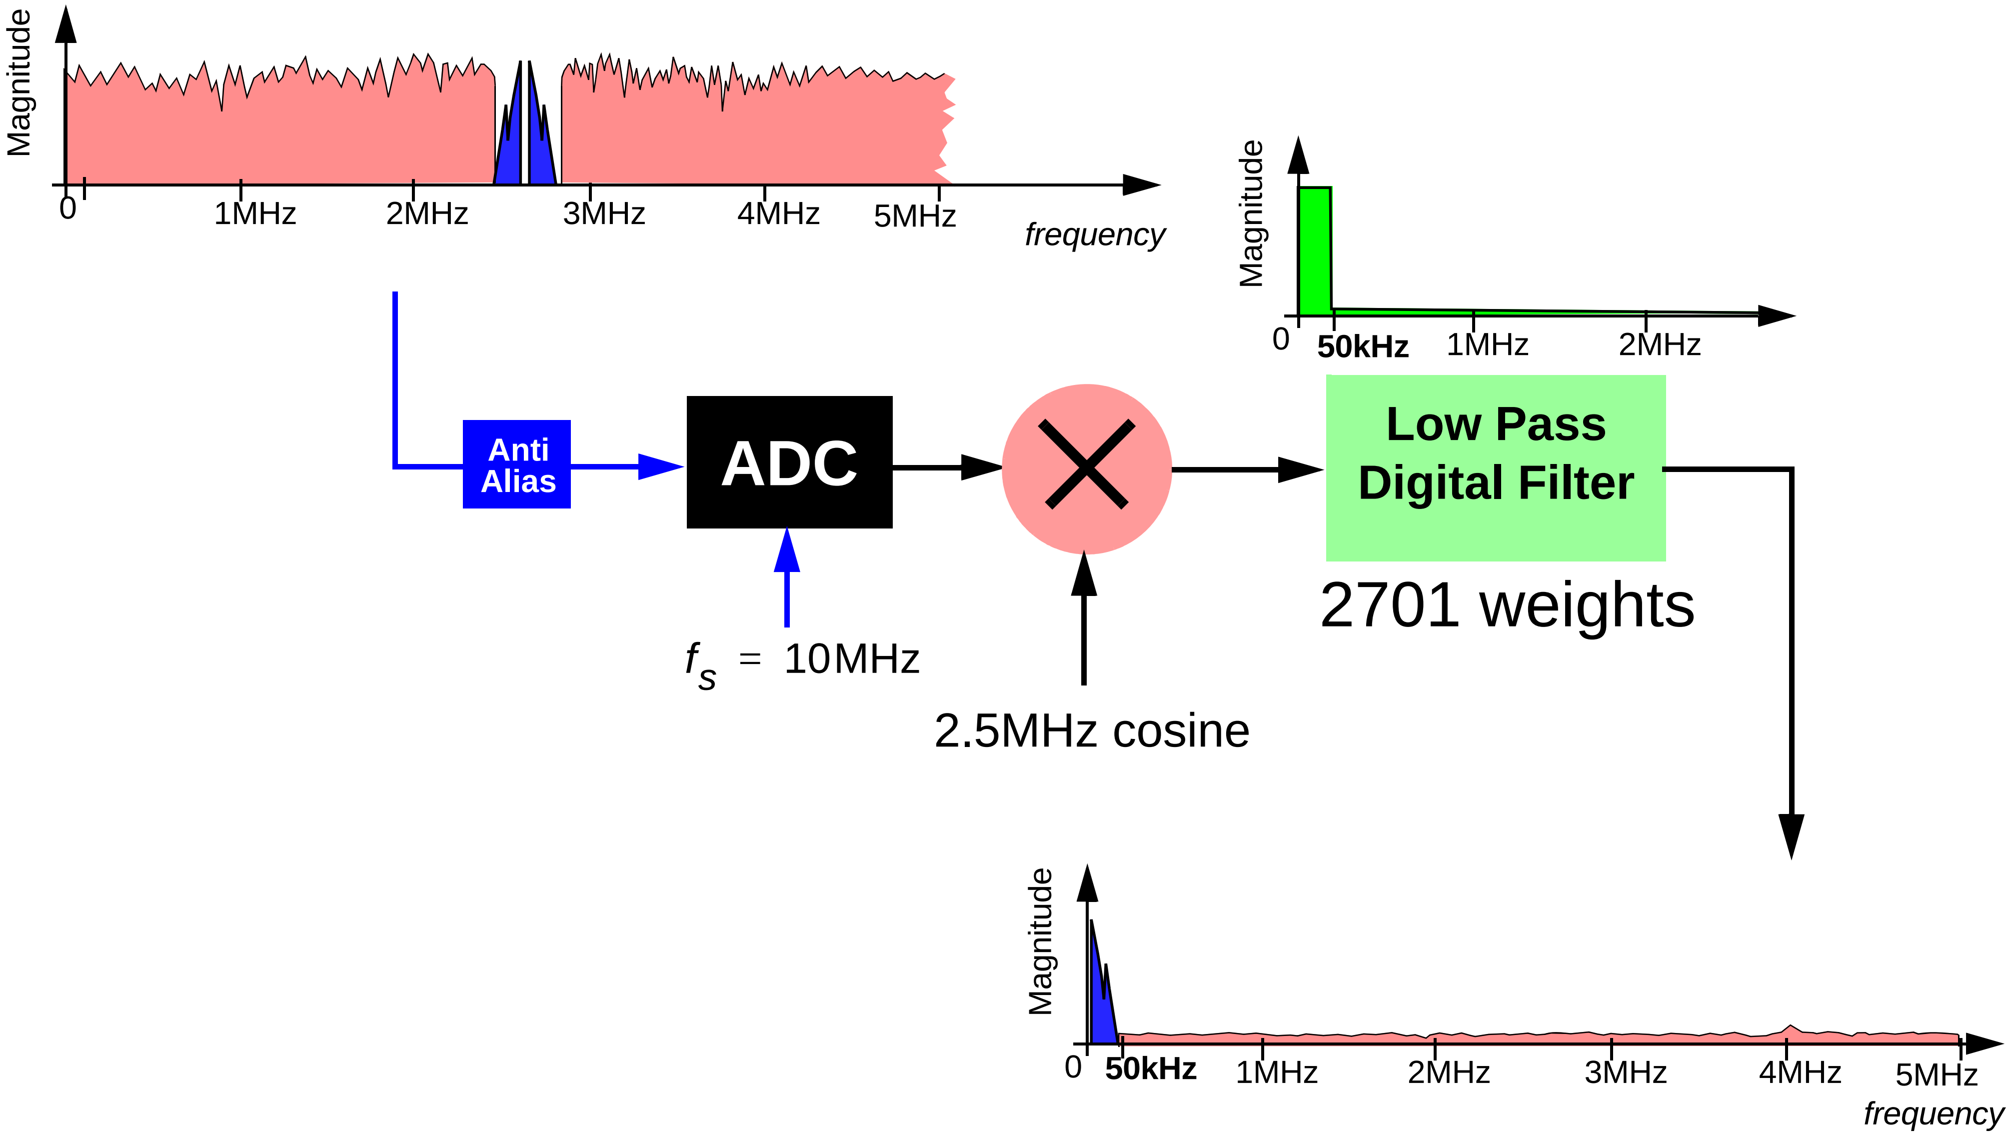

Top

Notes:

Cost of Digital Filter

MACs/sec = 10,000,000 x 2701 = 27,010,000,000 = ***27 billion MACs/sec***!

But remember the Downsampling...!

- ....then low pass filter:

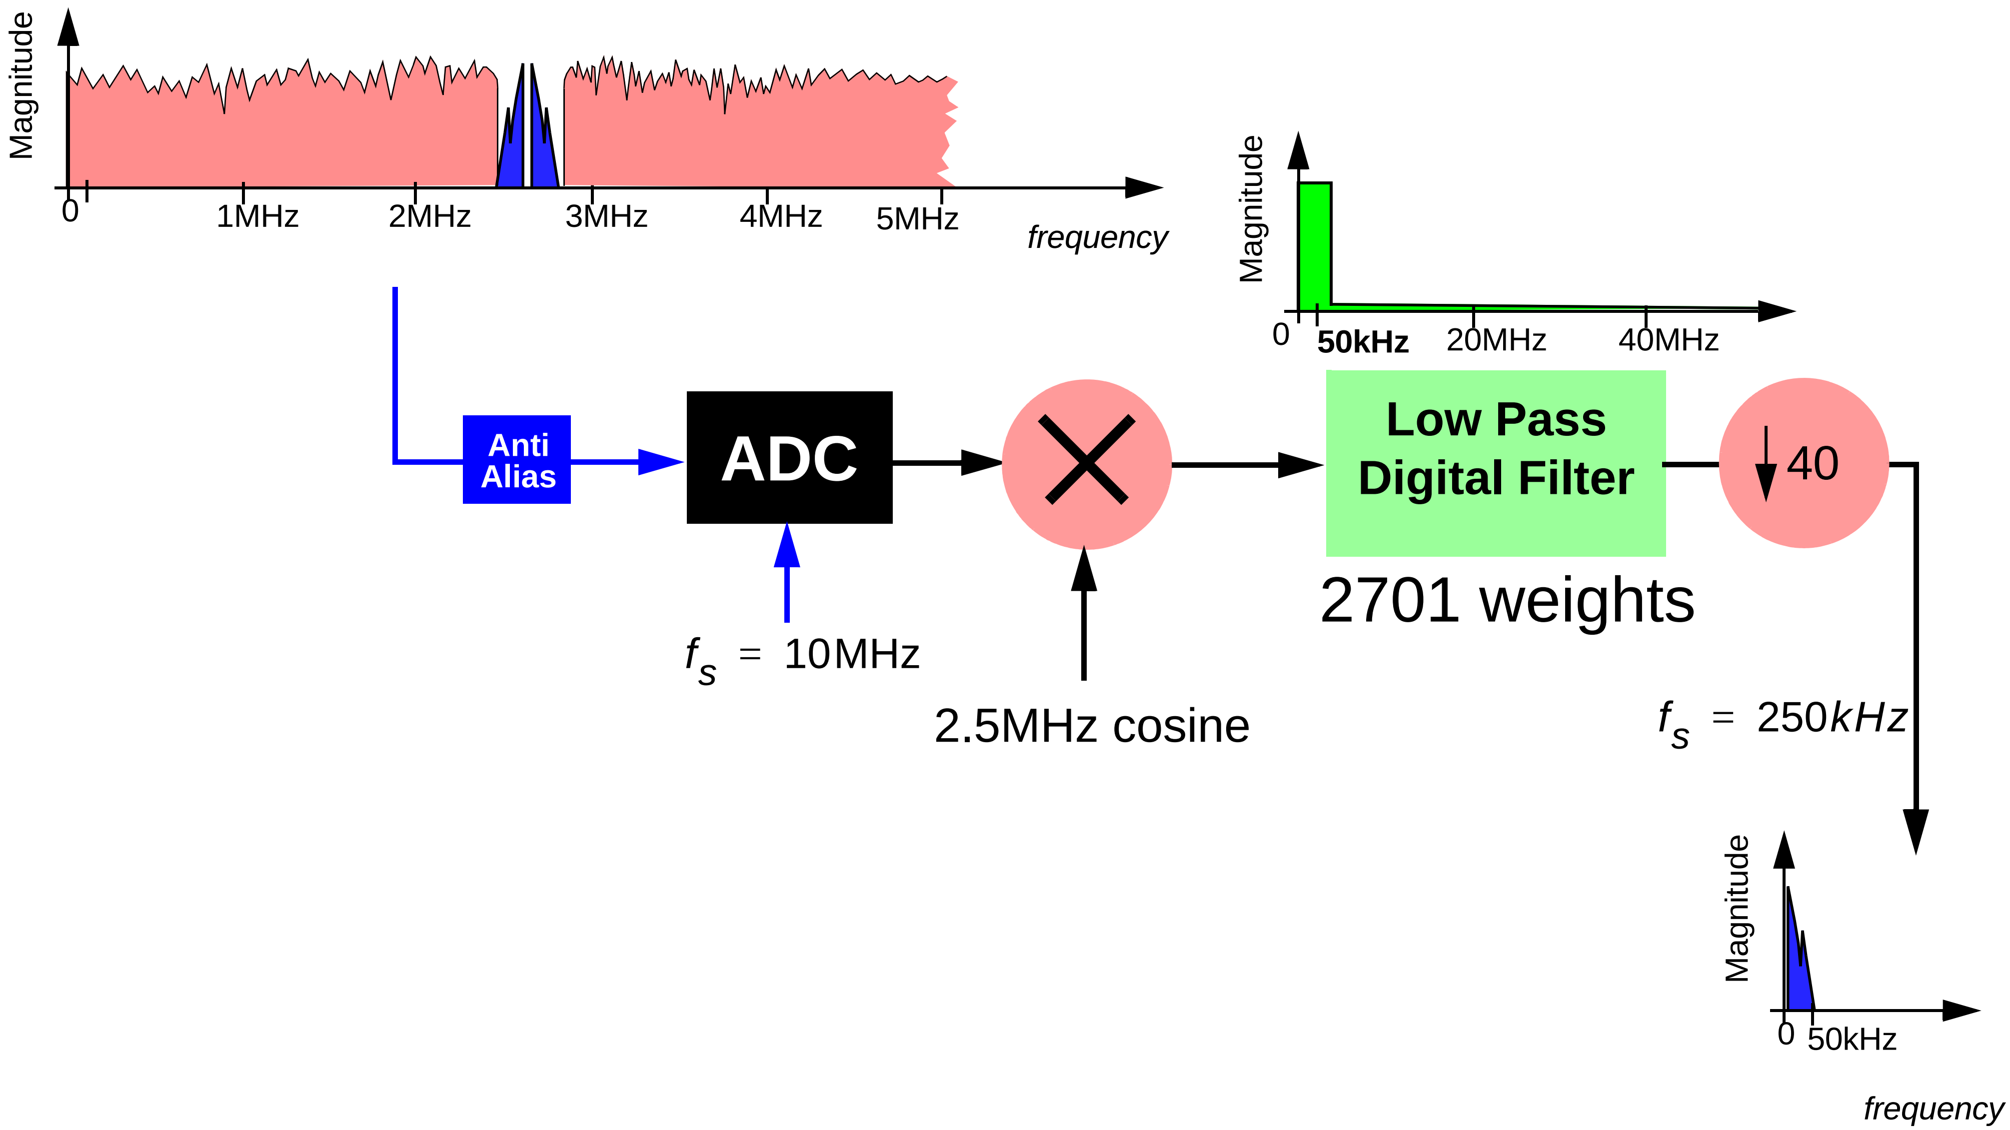

Top

Notes:

In this example the final required sample rate is 250kHz and hence as we have bandlimited we can now downsample by a factor of 40.

Cost of Digital Filter

MACs/sec = 10,000,000/40 x 2701 = 270,100,000 = ***675 million MACs/sec***!

CIC stage for Decimation

- Consider now designing the low pass filter to extract 0 to 50kHz using a cascade of low cost simpler filters. Is there a cost saving?

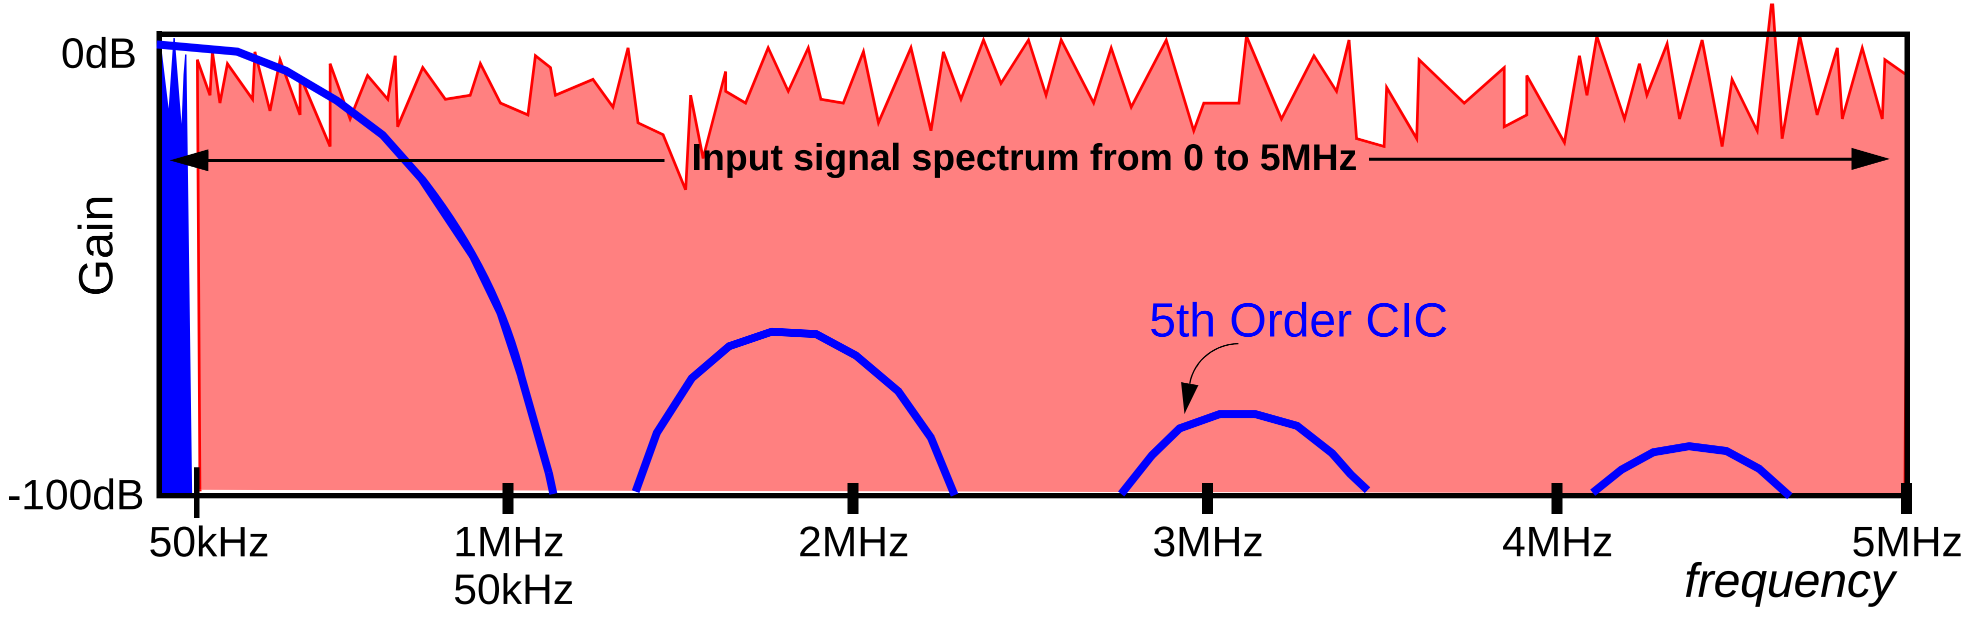

- If we low pass filter the signal of interest with the 5th order CIC then downsample by 2 to 5MHz, then the aliasing of higher frequency signals comes from frequency regions where the energy is very low.

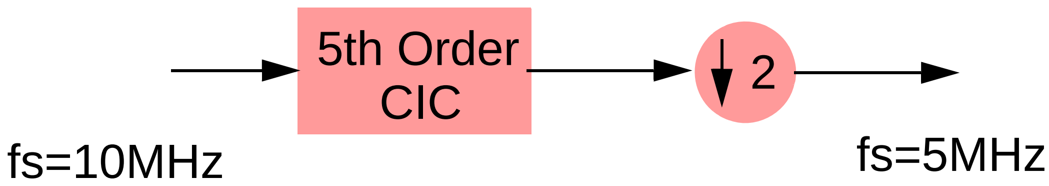

Notes:

The output spectrum almost leaves the 0 to 50kHz signal untouched in and attentuates the signal energy above 50kHz as below.

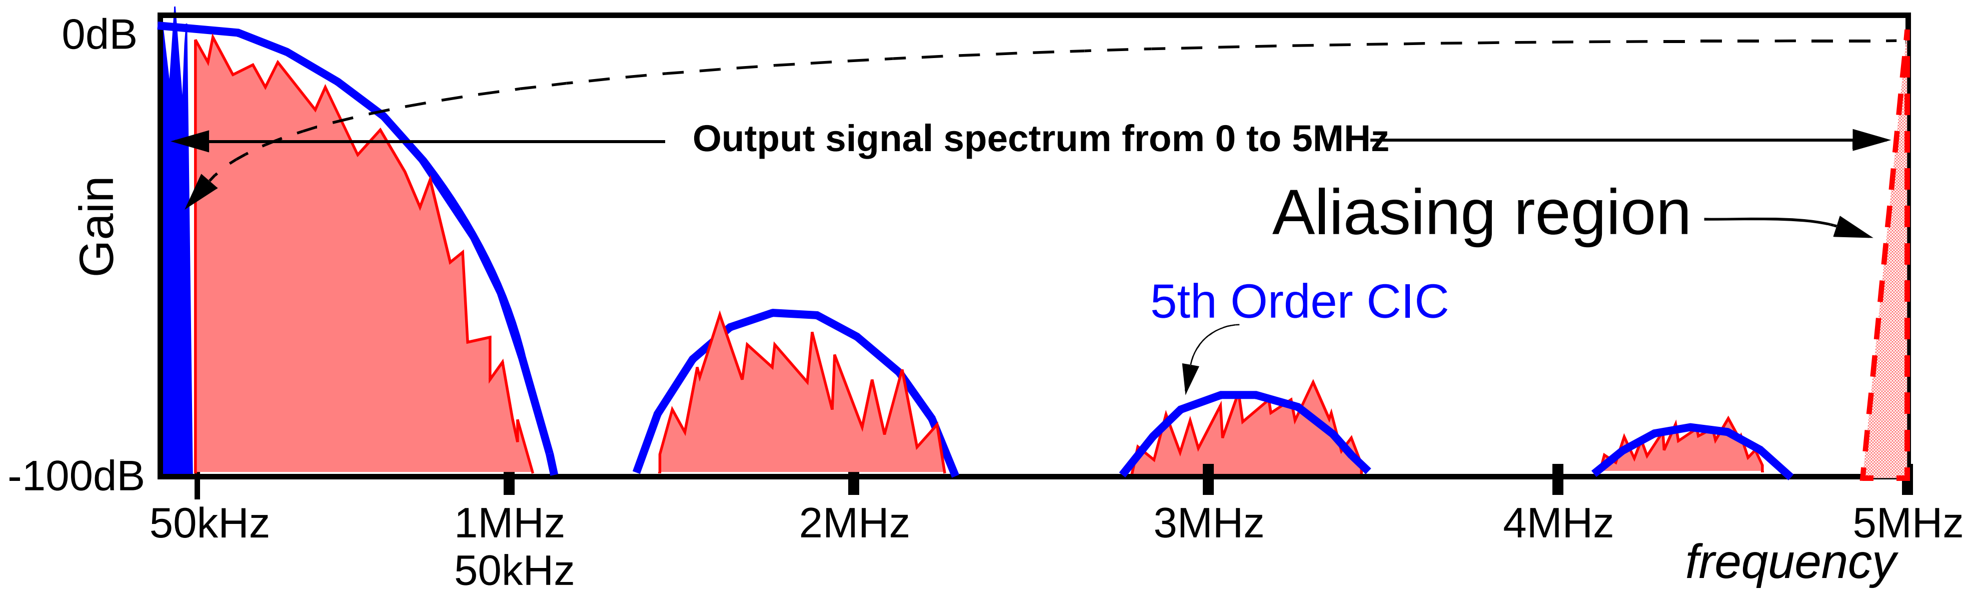

Final Stage Decimation

- Noting the other *empty* spectral regions we could downsample by 4:

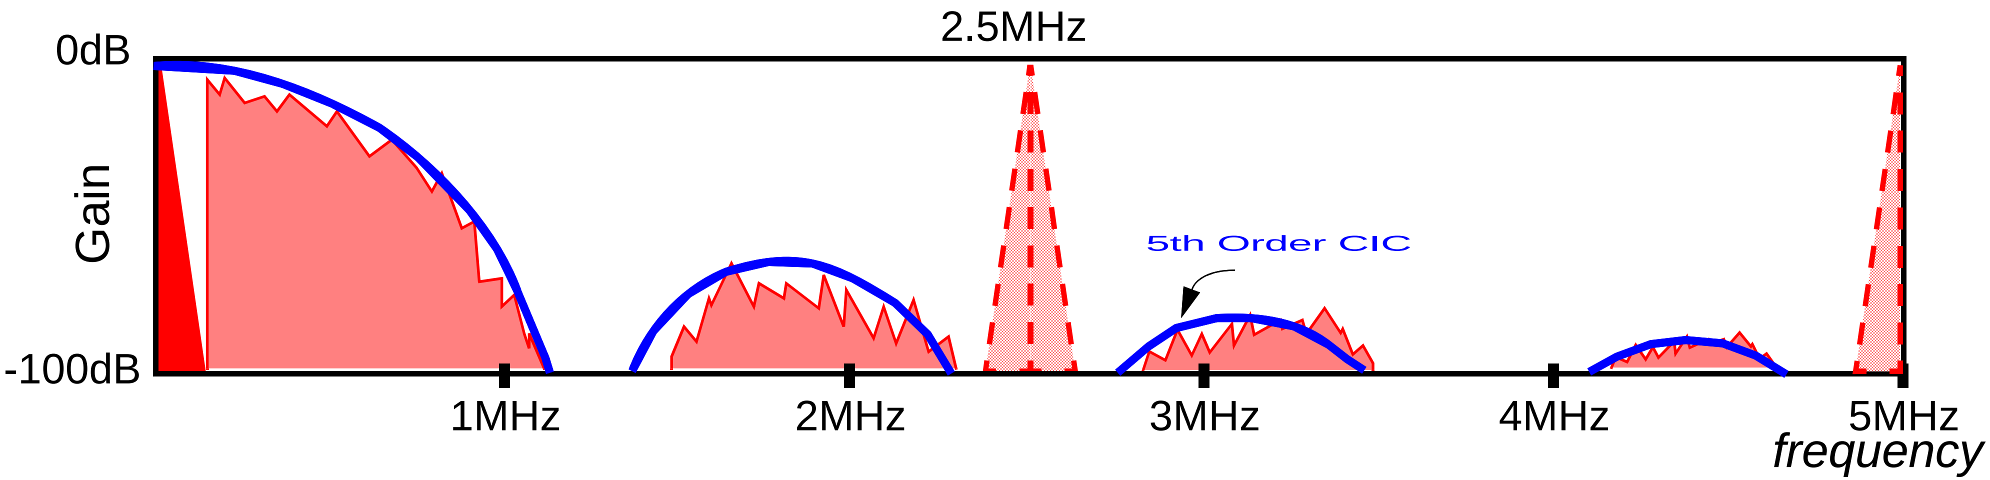

....in fact we could probably downsample by 8:

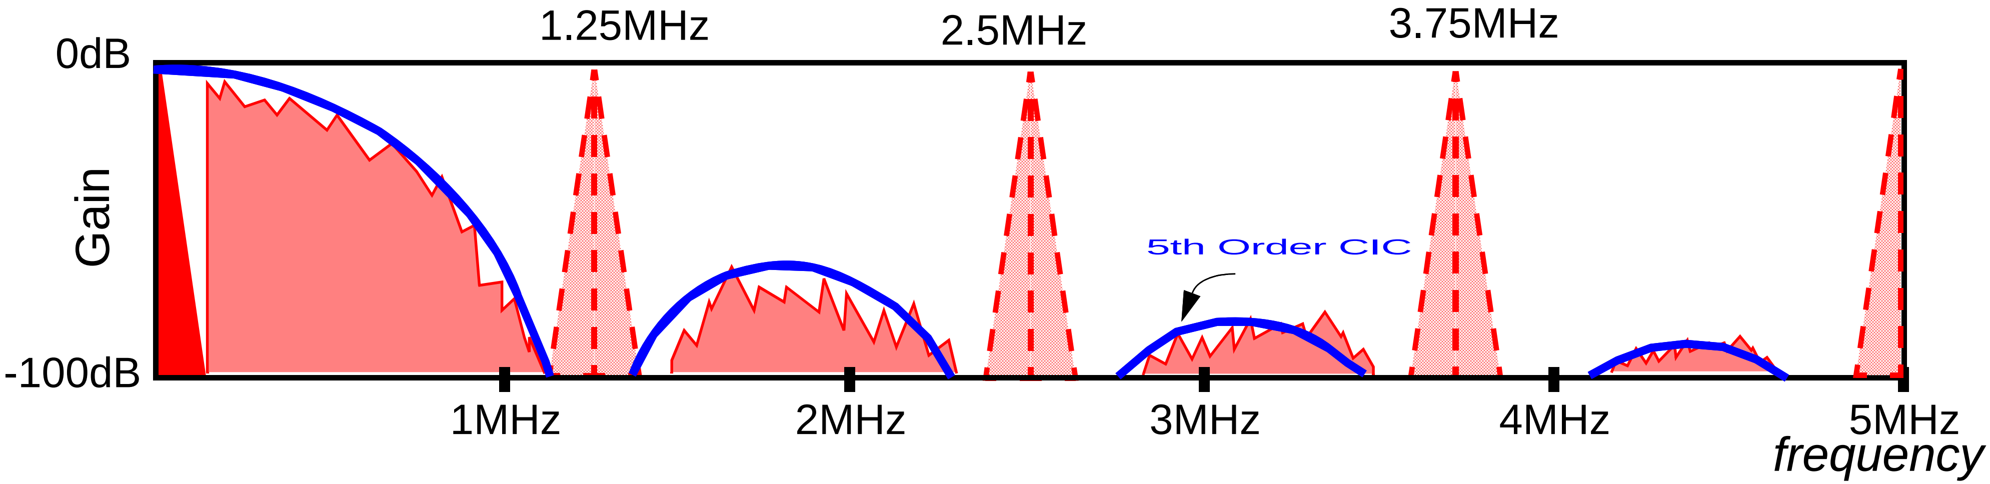

Notes:

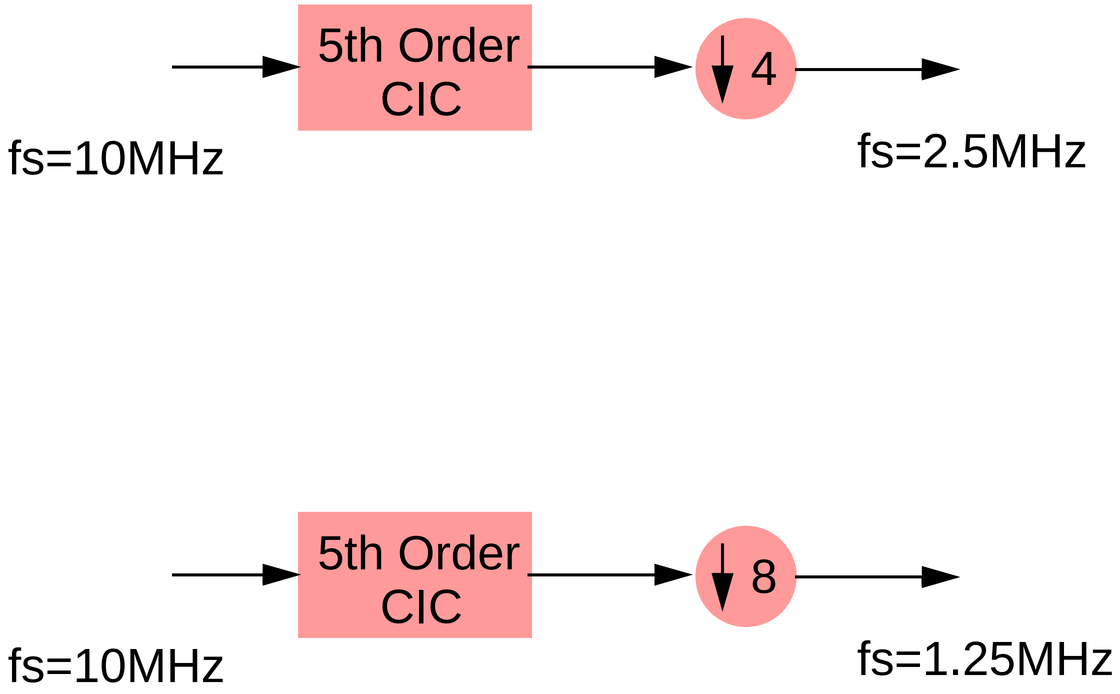
5th Order
CIC
4
fs=10MHz
fs=2.5MHz
5th Order
CIC
8
fs=10MHz
fs=1.25MHz

Downsampling Values

- We can then perform a final stage of decimation using a standard low pass filter:

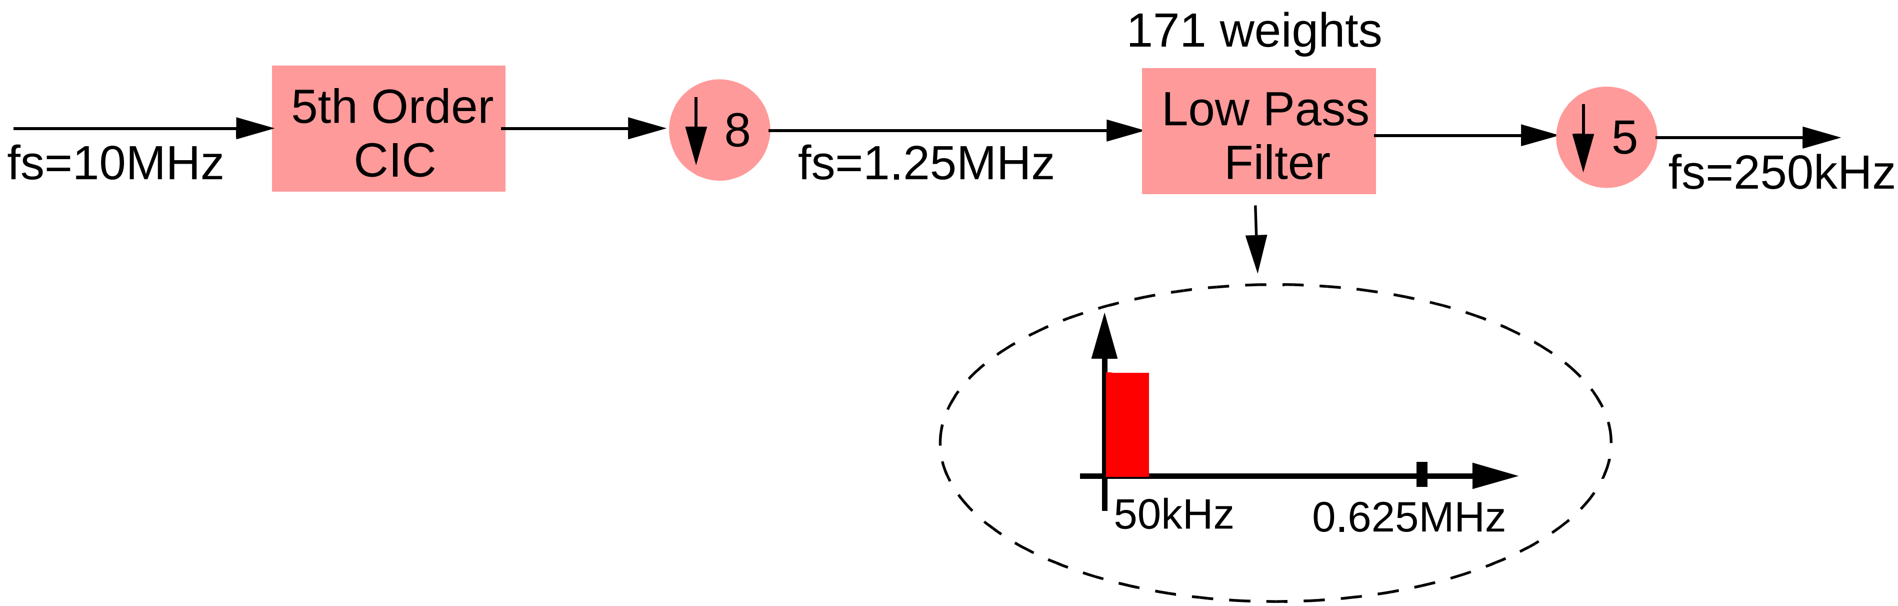

- Therefore we are now *anticipating* that the above staged decimation is similar to the one step decimation presented earlier (and shown below):

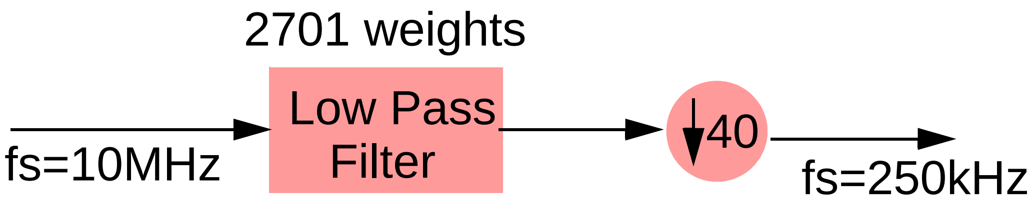

Notes:

Cost Comparison

- One stage ***Low Pass Filter*** decimation:
  - 2701 weights, 10MHz sampling, Downsample 40

$$\frac{2701 \times 10,000,000}{40} = 675.25 \text{ million MACs/sec}$$

- ***5th Order CIC and low pass*** (at f_s = 1.25MHz)
  - 171 weights, 1.25MHz sampling, Downsample 5

$$\frac{171 \times 1,250,000}{5} = 42.75 \text{ million MACs/sec}$$

  - 5 CICs with 2 adds each at 10MHz = 100 million adds/sec

- ***Computation reduced by a factor of almost 16!***

Notes:

CIC Droop

- One difference we have ignored so far is the “droop” at low frequencies of the CIC low pass filter:

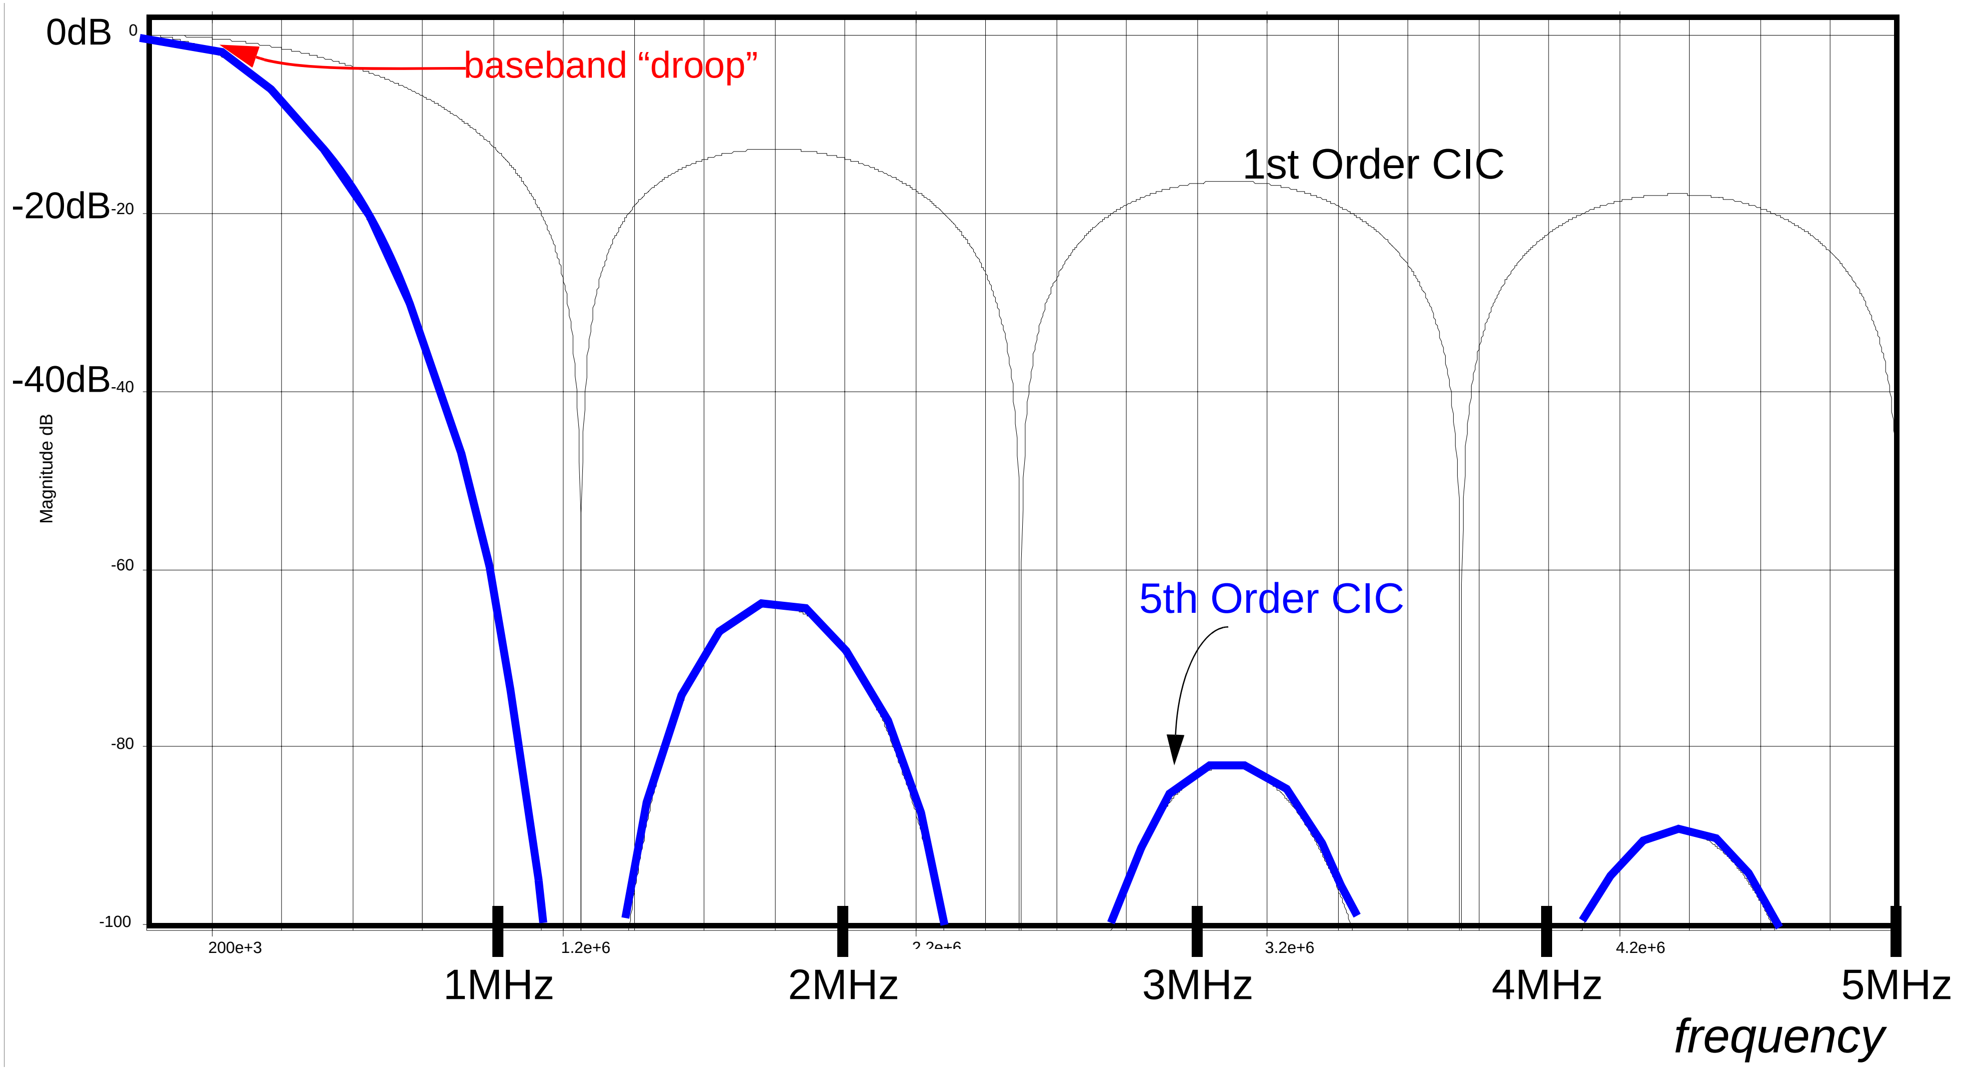

Notes:

Careful viewing of the spectrum shows that the droop is around 0.5dB:

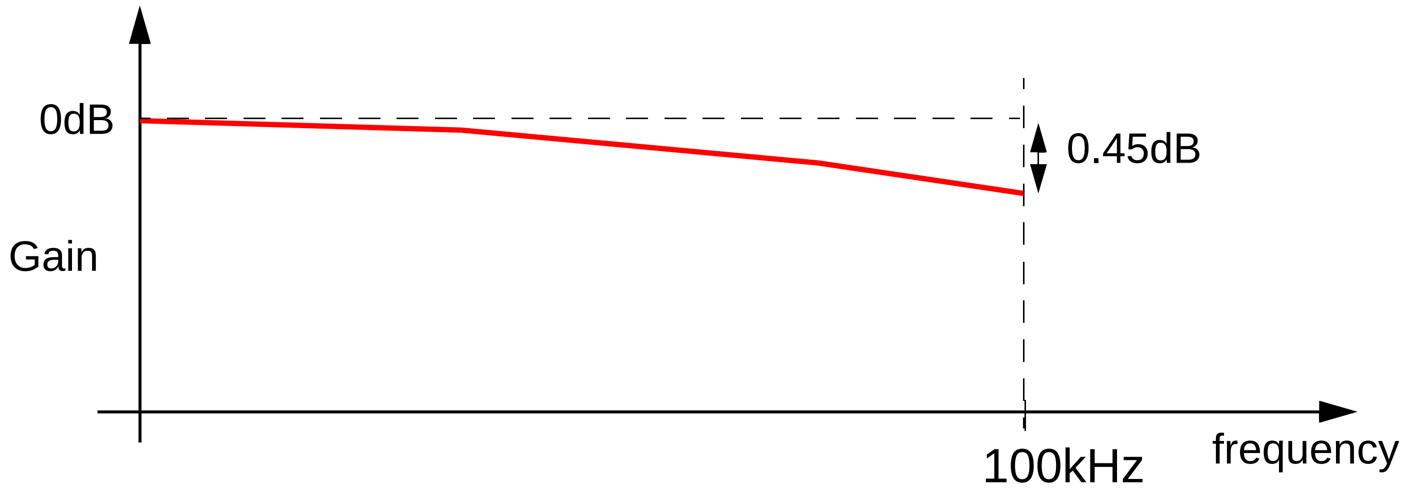

Correcting the Droop

- So how do we correct the droop?

- Incorporate a “lift” in passband of the final stage low pass filter:

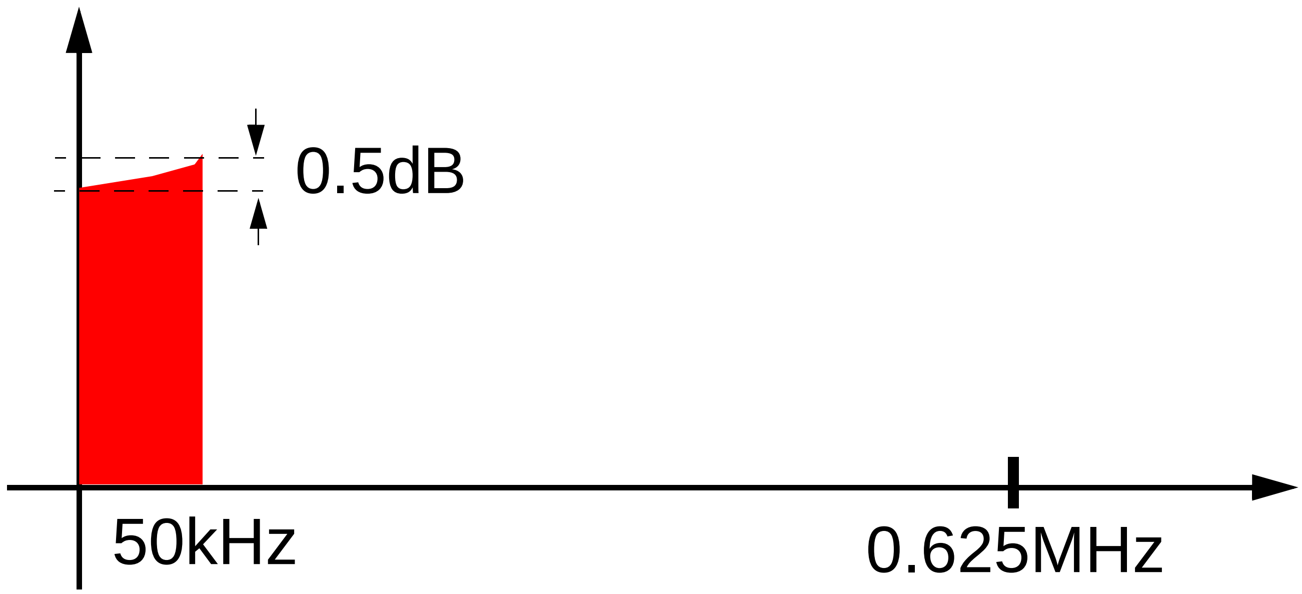

- ....therefore the decision of this final stage filter must be done very carefully to correct for the droop.

Notes:

CIC Implementation

- Consider the 3rd order CIC cascade shown below with a final stage downsampler:

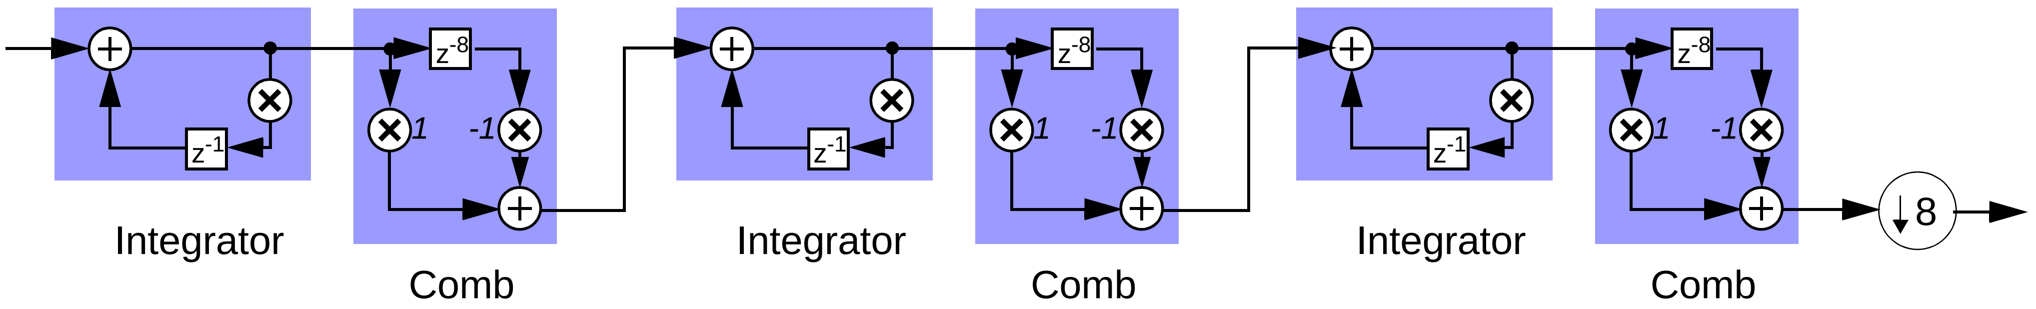

- We can rearrange the order of the filtering:

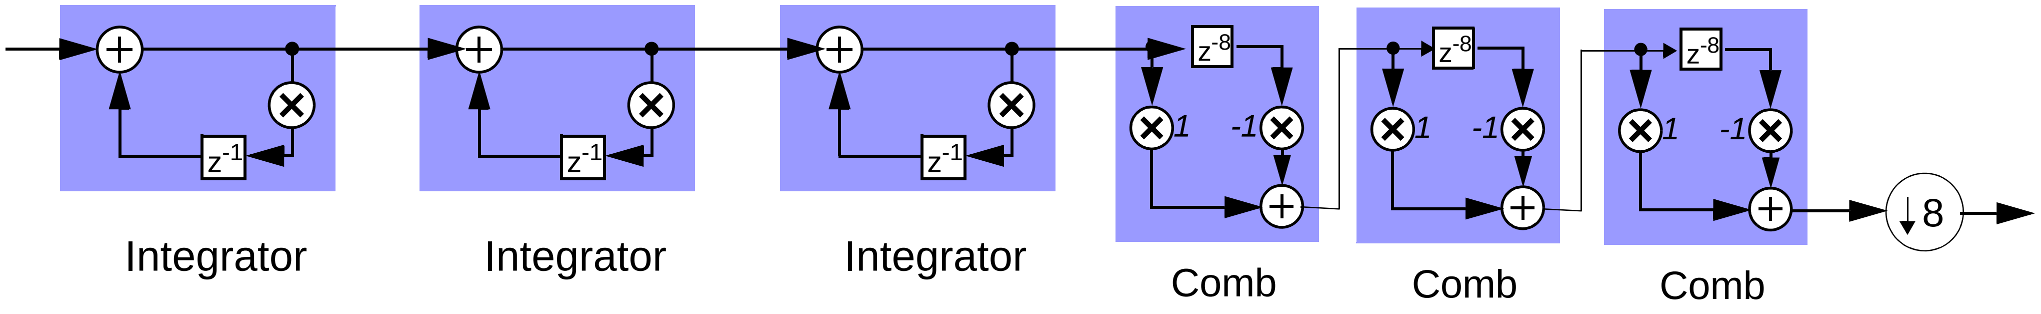

Notes:

We can represent the comb filter more compactly using the z^{-1} notation for a delay:

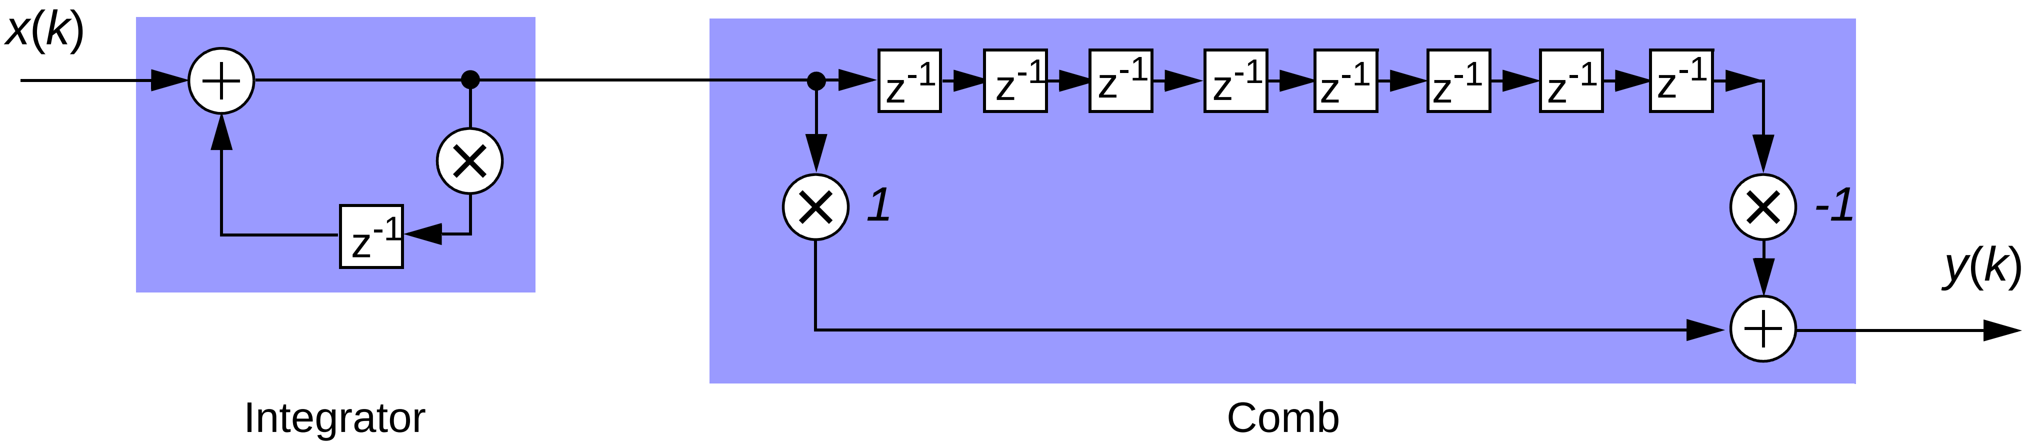

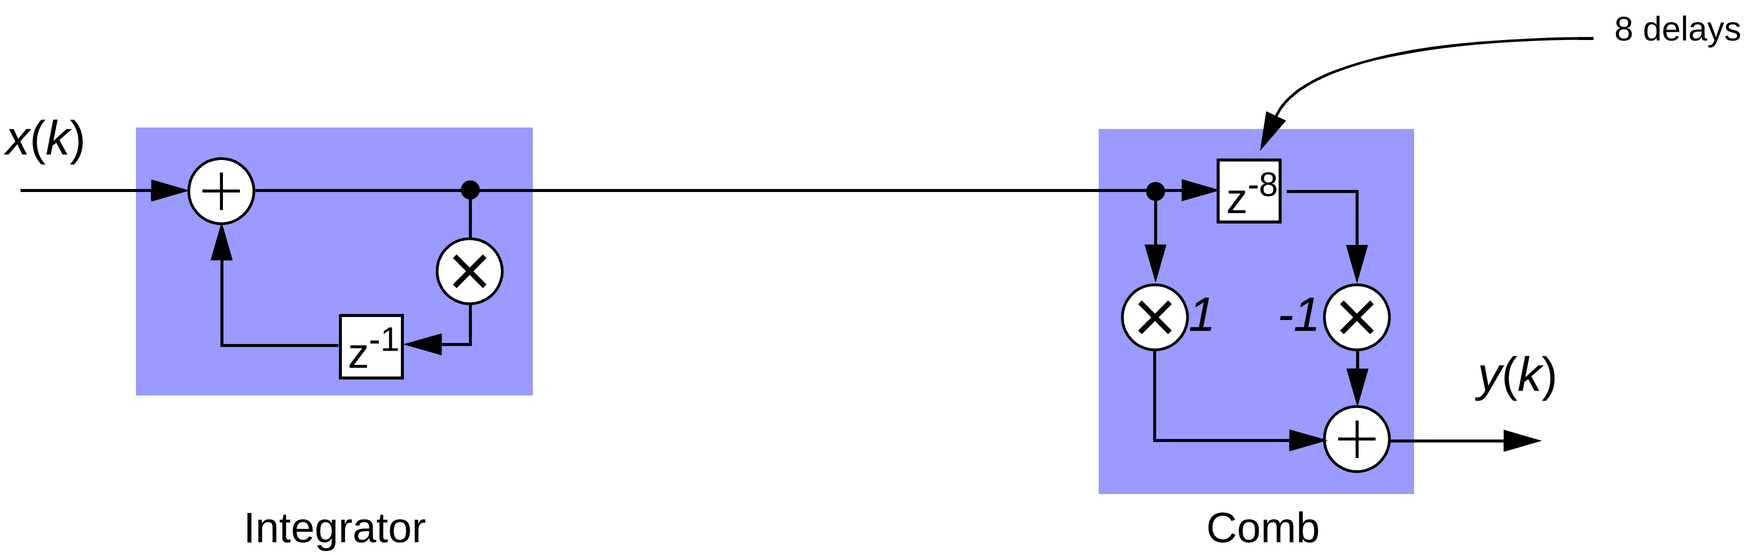

CIC Implementation

- Based on the noble identity we can move the downsampler to before the comb filters:

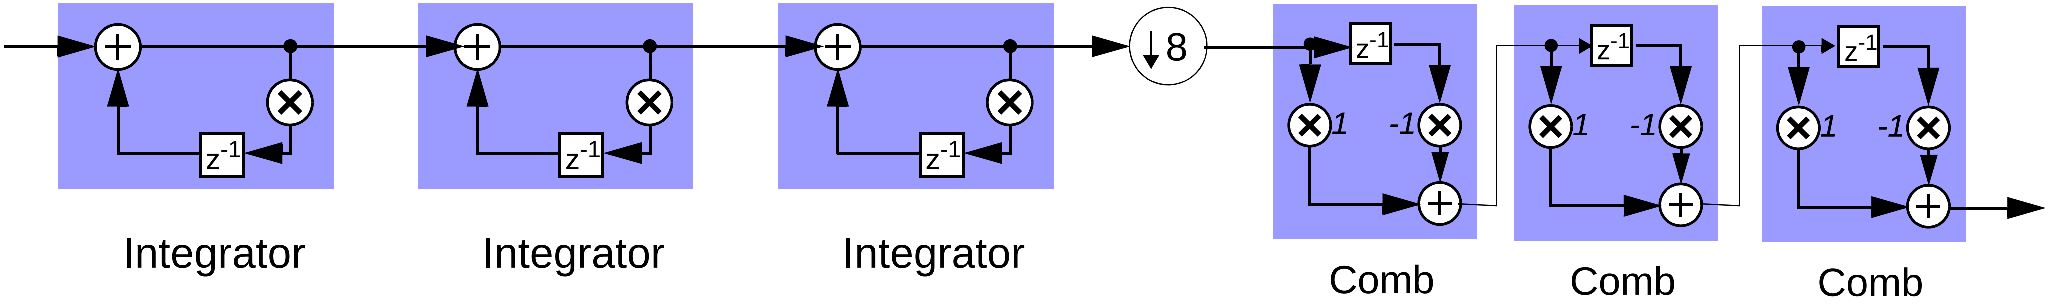

- Hence the comb filters now run at the downsampled rate, and require fewer registers for implementation.

Top

Notes:

The noble identity allows the two systems below to be demonstrated to be equivalent.:

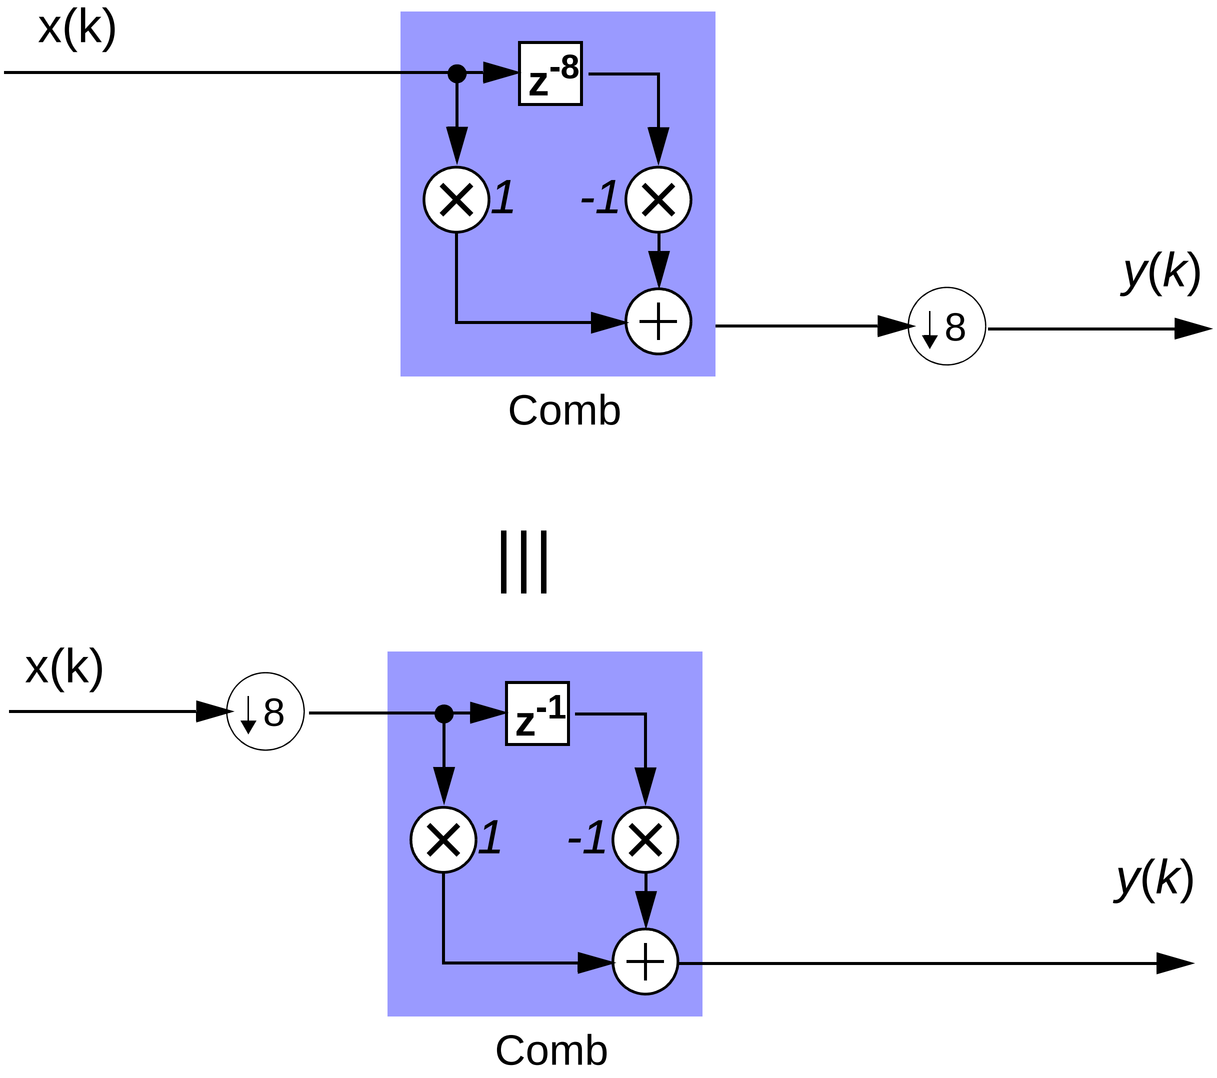

CIC Filter study

- In the next few slides the Cascaded Integrator-Comb filter is examined in more detail. This will cover the following areas:

  - Introduction to the CIC filter and some examples of where it may be used

  - An examination of word length growth in CIC filters and how 'bit-pruning' may be used to reduce resource consumption

  - The Sharpened CIC (SCIC) filter structure: how it differs from the CIC filter and where its use is appropriate

  - The Interpolated Second Order Polynomial (ISOP) filter: an alternative to the SCIC filter for compensating for CIC filter passband droop

  - A discussion of the costs and benefits of CIC and SCIC filters compared with non-recursive, 'moving average'-based filter structures