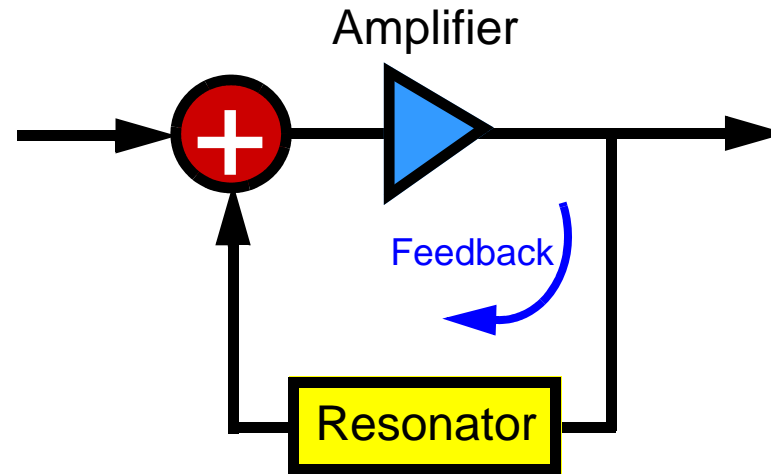


Oscillators

- Analogue oscillators are realised from feedback loops with **gain = 1**. The feedback path contains a resonator (such as a quartz crystal or ceramic) to control the frequency of oscillations.



- Digitally, an equivalent can be formed from a “marginally stable” **IIR filter**, as discussed in the *Digital Filtering* chapter.
- NCOs can also use the **CORDIC** algorithm to rotate a vector.
- A further option is to store samples of a sine wave in a **look up table**, and read them out at a rate commensurate with the desired frequency.

Notes:

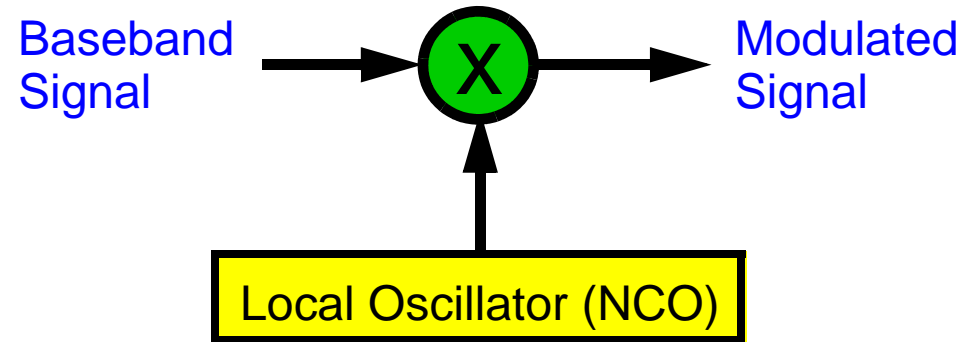
Applications of NCOs

- NCOs are used extensively in DSP.
- Perhaps the most common application area is digital communications. Modulators and demodulators both multiply signals with sinusoids to move between baseband, Intermediate Frequency (IF) and sometimes even Radio Frequency (RF).
- You will find NCOs in the following systems (and many others!):
 - Cellular telephones and basestations
 - Radar systems
 - Digital televisions
 - GPS satellites and handsets
 - Wireless LAN

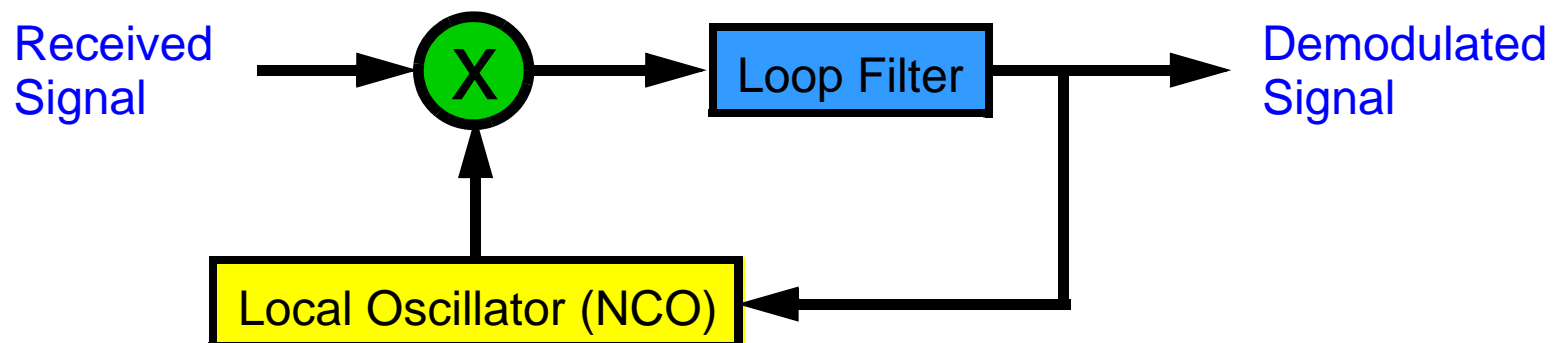


Notes:

Modulation involves multiplying the baseband signal with a sinusoid at the carrier frequency (this is usually the IF frequency). The frequencies which the transmitter NCO must generate are limited by the frequency channels of the communications system.




Demodulation is more difficult as the phase and frequency of the carrier must be recovered. The carrier frequency of the received signal may be different from that generated by the local oscillator in the receiver, due to device variations and the Doppler effect. For this reason, the NCO in the receiver must be capable of realising a range of frequencies with fine resolution, in order to match the frequency of the incoming signal.



Lookup Table NCO

- The construction comprises a LUT of sine wave samples, and an accumulator which generates the address.



- The LUT has $N = 2^n$ entries, where n is the number of address bits generated by the accumulator.
- The precision of the LUT output is L bits. This parameter is completely independent of n .
- The step size, μ , determines how quickly the address is accumulated, and hence the frequency of the generated sine wave.

Notes:

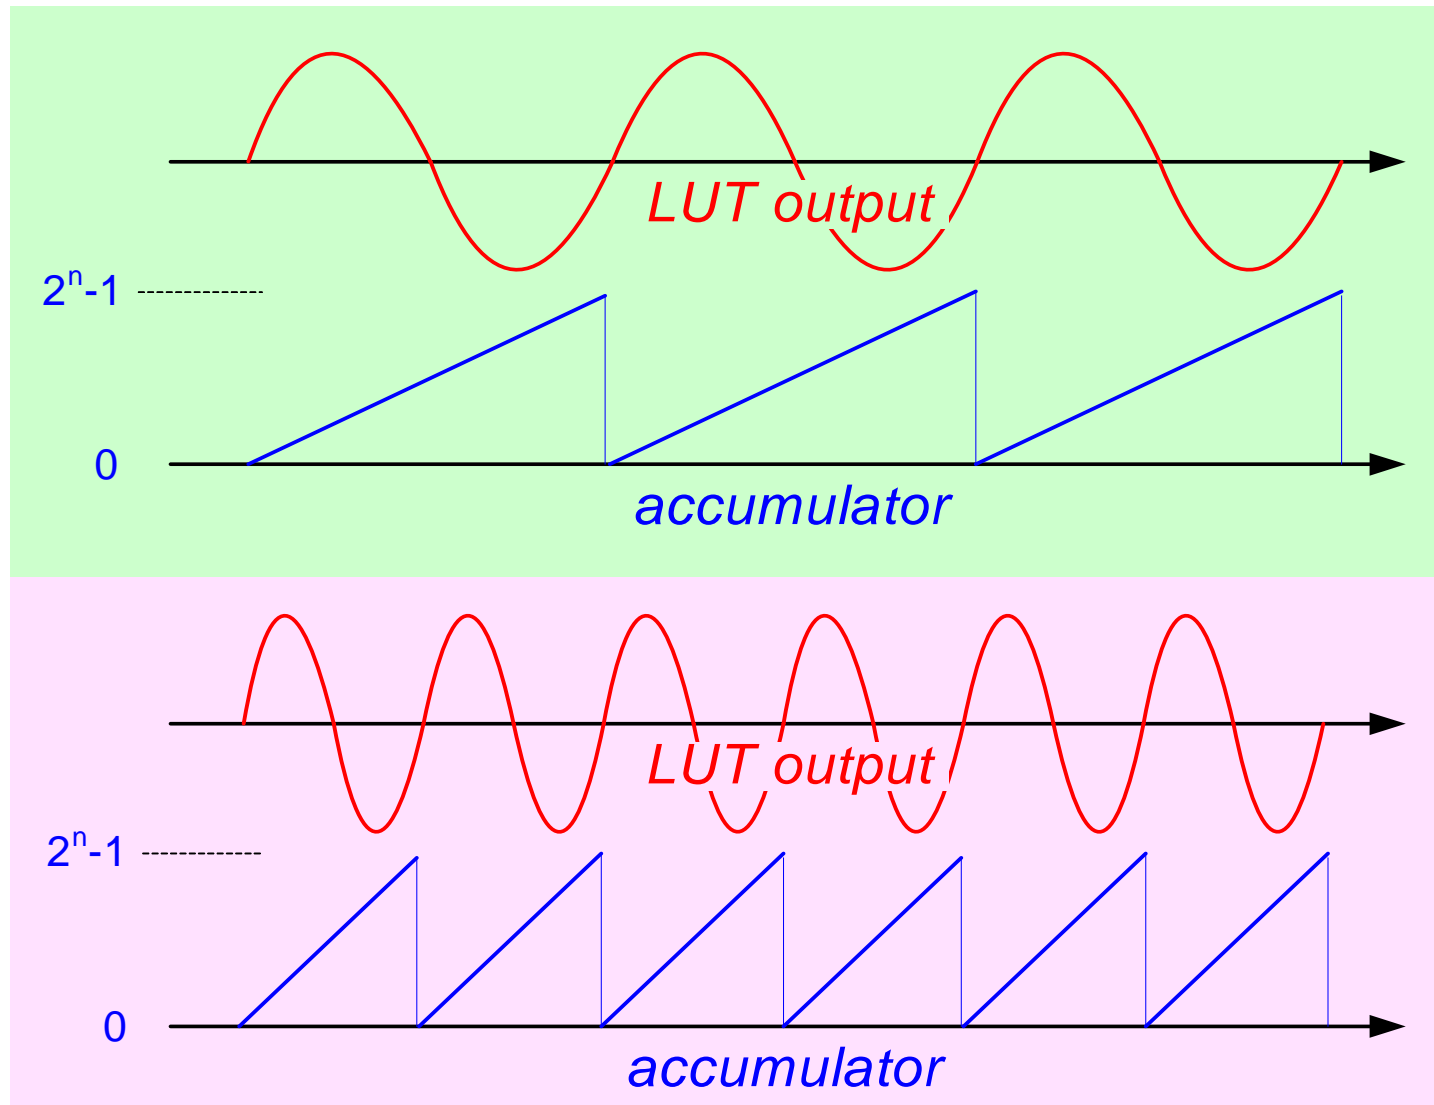
The accumulator wraps round when the maximum value of 2^n is reached.

It is also possible to save hardware in the LUT by storing only one quarter of a sine wave. Some simple translation of the address may be used prior to the LUT, in order to access the required sample and invert if appropriate. This approach results in the number of ROM addresses being reduced from 2^n to 2^{n-2} .

In the discussion which follows, we will assume that the full sine wave is stored, and therefore that the number of entries in the LUT is 2^n .

Frequency Control

- The accumulator cycles through the LUT addresses more quickly with a larger step size (bottom), producing a higher frequency sine wave.



Notes:

The step size, μ , is derived from the following system parameters:

N : the number of entries in the lookup table (taken as 2^n , even if a quarter wave lookup table is used)

f_s : the system sampling frequency

f_d : the frequency of the desired sine wave

where:

$$\mu = N \frac{f_d}{f_s}$$

If, for example, there are 8 address bits (giving a lookup table size of $N = 2^8 = 256$), the sampling frequency is 10MHz and the desired frequency is 2.5MHz, the resulting step size is...

$$\mu = 256 \frac{2.5MHz}{10MHz} = 64$$

But what if we wanted to create a 2.4MHz sine wave? Here, $N = 256$, $f_s = 10MHz$ and $f_d = 2.4MHz$, giving....

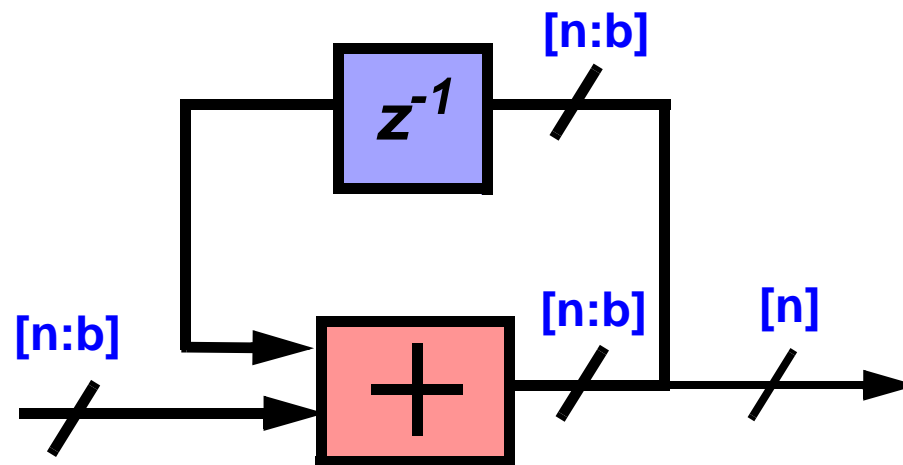
$$\mu = 256 \frac{2.4MHz}{10MHz} = 61.44 \dots \text{ How do we include this fractional part?}$$

NOTE: Choosing $\mu = 61$ and applying to the rearranged version of the equation above for the actual frequency, f_a , we achieve

$$f_a = \mu \frac{f_s}{N} = 2.383MHz \dots \text{ which is not what we wanted at all!}$$

The Accumulator in Detail

- The accumulator includes a fractional part, which is not used when referencing the LUT, but which allows greater frequency control to be achieved.
- The step size is a fixed point number, comprising n whole bits and b fractional bits. The entire word may be written in the form $n:b$.
- For fixed n , increasing b provides greater precision in the step size, and hence the synthesised frequency.



- The cost of using fractional bits is increased accumulator complexity.

Notes:

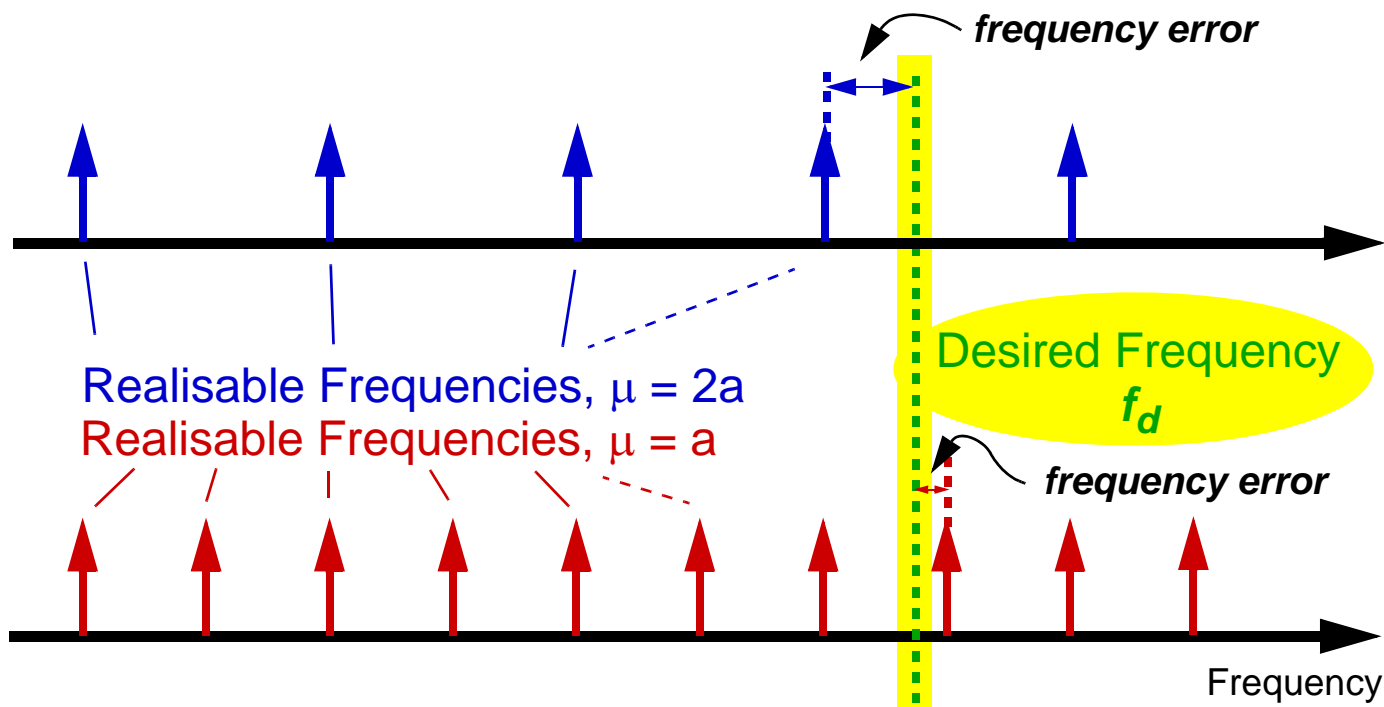
This might lead us to think that more accumulator bits is better, but from an implementation point of view the wordlength should be limited to the minimum required for the desired resolution. Using more bits in the accumulator not only costs more, but may also limit the clocking rate at which the design can operate.

Returning to the previous numerical example, suppose that we now have 4 fractional bits available. This means that we have 8 whole bits and 4 fractional bits, so we can represent the desired step size of 61.44 more accurately. In fact, we can represent 61.4375 with this level of quantisation. Therefore we can realise an actual frequency of

$$f_1 = \mu \frac{f_s}{N} = 2,399,902 \text{ Hz} \quad \dots \text{ much closer to our target!}$$

Frequency Resolution

- If quantisation of the step size is too coarse, the frequency which is actually realised may suffer from a large frequency error.
- The number of fractional bits in the accumulator should therefore be chosen according to the maximum acceptable frequency error.



- The maximum error between the desired and synthesised frequencies is half of a frequency interval.

Notes:

Returning to the equations introduced previously, it is useful to define the minimum possible step size increment. We will denote this by $\Delta\mu$.

$$\Delta\mu = \frac{1}{2^b}$$

The step size resolution translates directly into the frequency resolution of the synthesised sine wave, and may be represented by Δf_a .

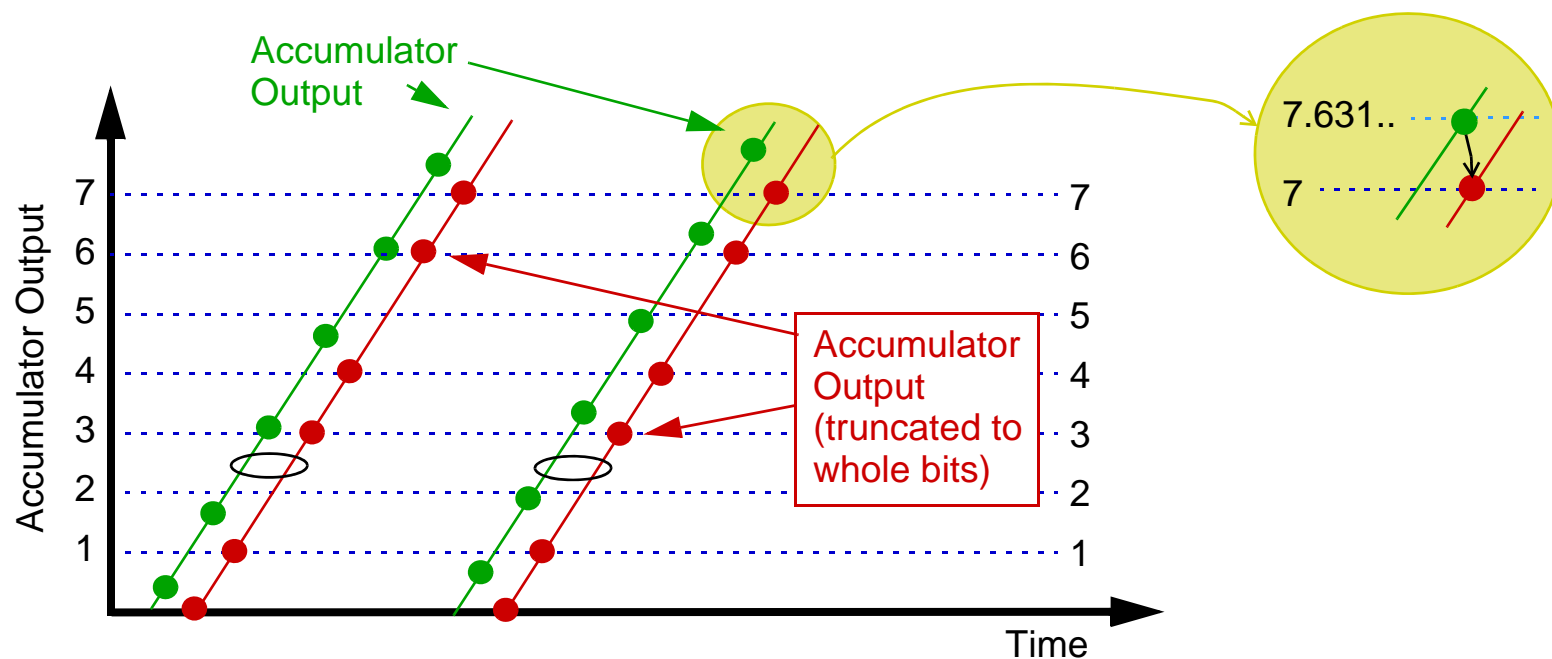
$$\Delta f_a = \Delta\mu \frac{f_s}{N} = \frac{f_s}{2^b N}$$

As an example, consider the previous scenario where a 256-entry lookup table was used in a system with sampling rate of 10MHz. By implication, $n = 8$. Therefore we can calculate the frequency resolution attainable when b bits are chosen for the fractional part.

Fractional bits (b)	Step Size Resolution ($\Delta\mu$)	Frequency Resolution (Δf_a) (Hz)
0	1.0	39062.5
1	0.5	19531.25
2	0.25	9765.625
5	0.03125	1220.703125
10	0.0009765625	38.14697265625
20	0.00000095367431640625	0.037252902984619140625

Addressing the Lookup Table

- When addressing the look up table, only the whole part of the accumulator output is used. The fractional bits are disregarded.
- This results in a **truncation** of the phase, as in this simple example with 3 whole bits (resulting in 8 addresses), and arbitrary fractional bits.



- Phase truncation is detrimental to the spectral purity of the synthesised sine wave.


Notes:

Notice from the diagram that the green dots (which represent the output from the accumulator) are evenly spaced. In contrast, the red dots (which represent the truncated output from the accumulator) are not evenly spaced. An unwanted irregularity has been introduced by the act of truncating the addresses, and therefore the synthesised sine wave is not accurate. As a result, the sine wave is not spectrally pure.

In addition, some degree of amplitude quantisation is applied to the samples stored in the LUT, depending on the number of storage bits available. This causes a further degradation to the ideal sine wave output, which can be observed in both the time domain and the frequency domain.

LUT resolution and depth are therefore the key factors influencing unwanted spectral components.

Spectrum: Phase Truncation Spurs




Notes:

In this example, the spurs visible in the FFT plot are due to phase truncation (amplitude quantisation has been reduced to a negligible level). The relevant parameters used in this example are:

Sampling frequency :	100kHz
Desired frequency:	24.3kHz
LUT depth :	64 addresses ($=2^6$)
LUT resolution :	32 bits
Accumulator wordlength :	6 whole bits, 16 fractional bits
Step size:	15.552001953125

NOTE: 32 bits has been chosen to ensure amplitude quantisation effects do not contribute significantly to the displayed spectrum, and therefore that the spectrum conveys the effect only of phase truncation.

Spectrum: Amplitude Quantisation Spurs




Notes:

In this example, the spurs visible in the FFT plot are primarily due to amplitude quantisation. The level of phase truncation has been reduced to a minimal level. The parameters used in this example are:

Sampling frequency :	100kHz
Desired frequency:	24.3kHz
LUT depth :	4096 addresses (= 2^{12})
LUT resolution :	8 bits
Accumulator wordlength :	12 whole bits, 16 fractional bits
Step size:	995.3280029296875

NOTE: the number of LUT entries has been chosen to ensure that the phase truncation effects introduced into the spectrum are minimal, and therefore that the spectrum shows, as closely as possible, only amplitude quantisation effects.

Spectrum: Amplitude & Phase Effects



Notes:

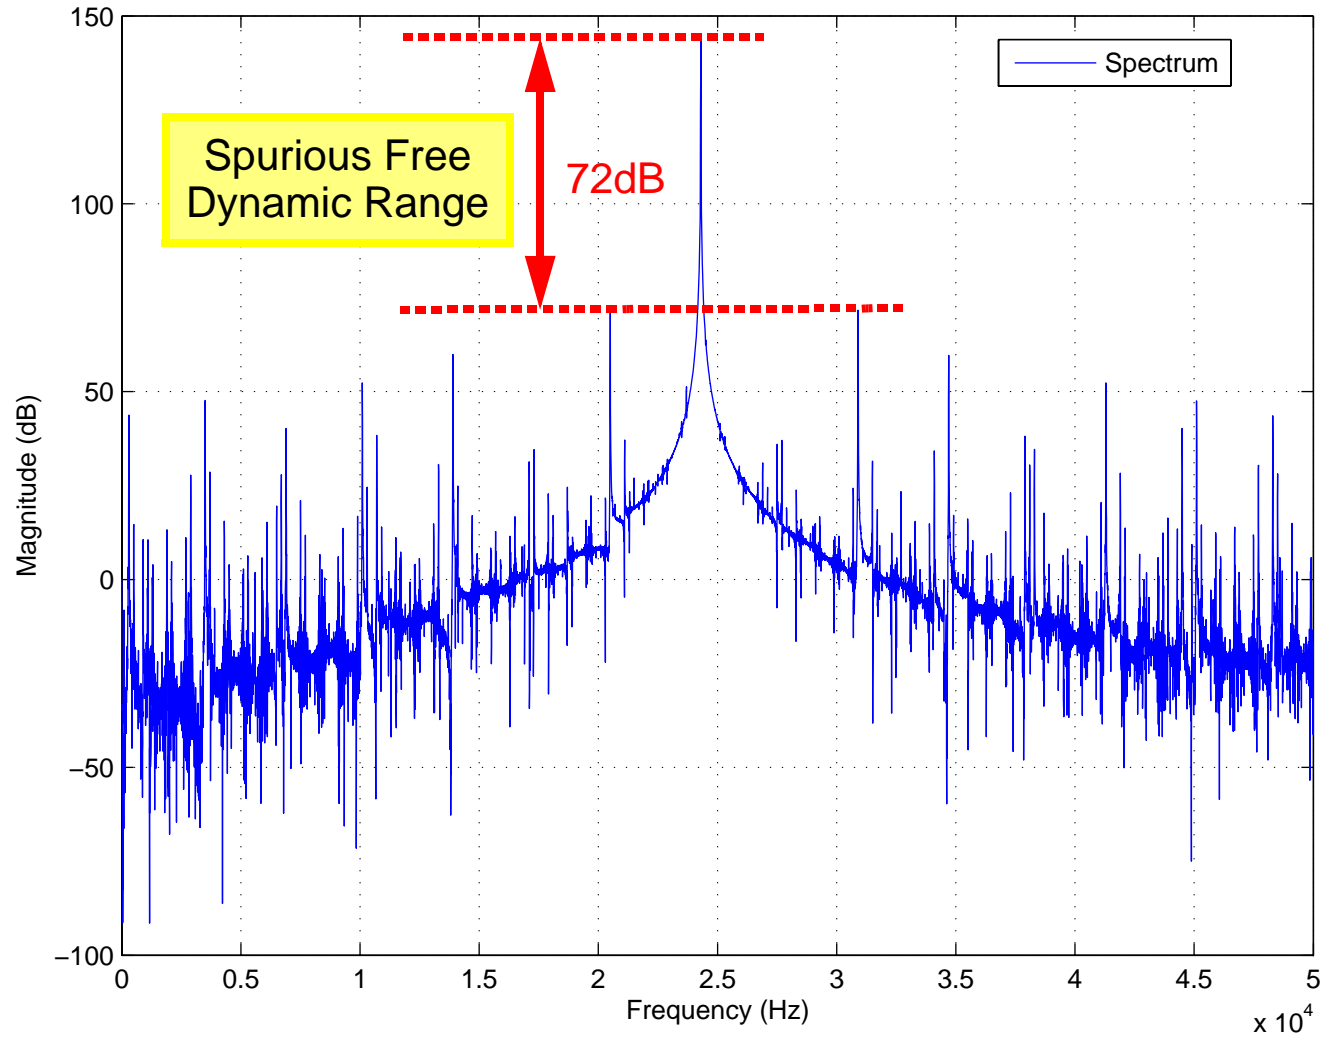
Finally, the effects of amplitude quantiation and phase truncation are combined by creating a small lookup table, with coarse amplitude quantisation applied to the sine wave samples. The parameters used in this example are:

Sampling frequency :	100kHz
Desired frequency:	24.3kHz
LUT depth :	64 addresses ($=2^6$)
LUT resolution :	8 bits
Accumulator wordlength :	6 whole bits, 16 fractional bits
Step size:	15.552001953125

Notice that the spurs arising from phase truncation are more dominant than those arising from amplitude quantisation. These can be derived mathematically as described in the following paper:

H. T. Nicholas and H. Samuelli, "An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase Accumulator Truncation," *Proceedings of the 41st Annual Frequency Control Symposium, 1987*.

Spurious Free Dynamic Range (SFDR)



Notes:

The Spurious Free Dynamic Range is the difference between the spectral peak and the highest spur. In this example, which is reproduced from the previous slide, the SFDR is 72dB.

SFDR is an important consideration when designing radio systems. If the oscillator signal contains significant spurious frequencies, these unwanted components can contaminate the signal being mixed. They may be difficult to subsequently remove, especially when close to the centre frequency.

As an example, the GSM specification requires a SFDR of at least 110dB - much greater than the NCO in this example provides!

Design Parameters

- It has been established that **whole bits**, **fractional bits**, and **LUT output resolution** all influence the performance of the NCO.
- Increasing the number of **whole bits** results in:
 - A larger lookup table
 - Less severe phase truncation and hence frequency spurs
- Increasing the number of **fractional bits** results in:
 - Greater step size precision and frequency resolution
 - A more expensive accumulator
- Increasing the **LUT output resolution** results in:
 - Less severe amplitude quantisation of sine wave samples
 - Greater LUT memory requirements



Notes:


The ultimate would be a large wordlength comprised entirely of whole bits (i.e. no phase truncation!) but this would result in a very large lookup table, and would therefore be prohibitively expensive to implement.

The effect of adding **1 bit** to each of the parameters can be summarised as follows:

Whole bits: $n \rightarrow n+1$


Adding 1 bit to the whole part of the number doubles the number of entries in the lookup table.

If the number of fractional bits is kept the same, the wordlength in the accumulator also grows by 1 bit (see bottom diagram).



LUT Output bits: $L \rightarrow L+1$

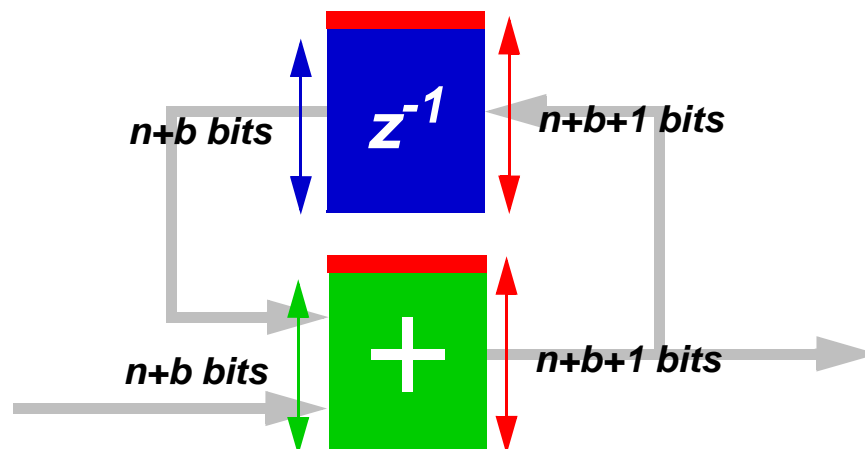
Adding 1 bit to the output resolution of the LUT increases the total memory required by the number of entries (2^n).



Fractional bits: $b \rightarrow b+1$

Adding 1 bit to the fractional part of the number increases the width of the adder and delay element of the accumulator, each by 1 bit.

The increase has no impact on the LUT.



Improving the SFDR

- The SFDR can be improved by increasing the number of whole bits, and also by increasing the resolution of the LUT output.
- However this is not an economic approach!
- To deal efficiently with the problem of spurs, other solutions should be considered:
 - **Amplitude dithering** - Adding a low level of noise to the output of the lookup table, to break up the structure of amplitude quantisation noise.
 - **Phase dithering** - Likewise, adding a low level of noise to the accumulator output, to break up the structure of phase truncation noise.
 - **Bandpass filtering** - Using a filter at the oscillator output to remove spurious frequencies. However, spurs lying close to the centre frequency cannot be removed in this way.



Notes:

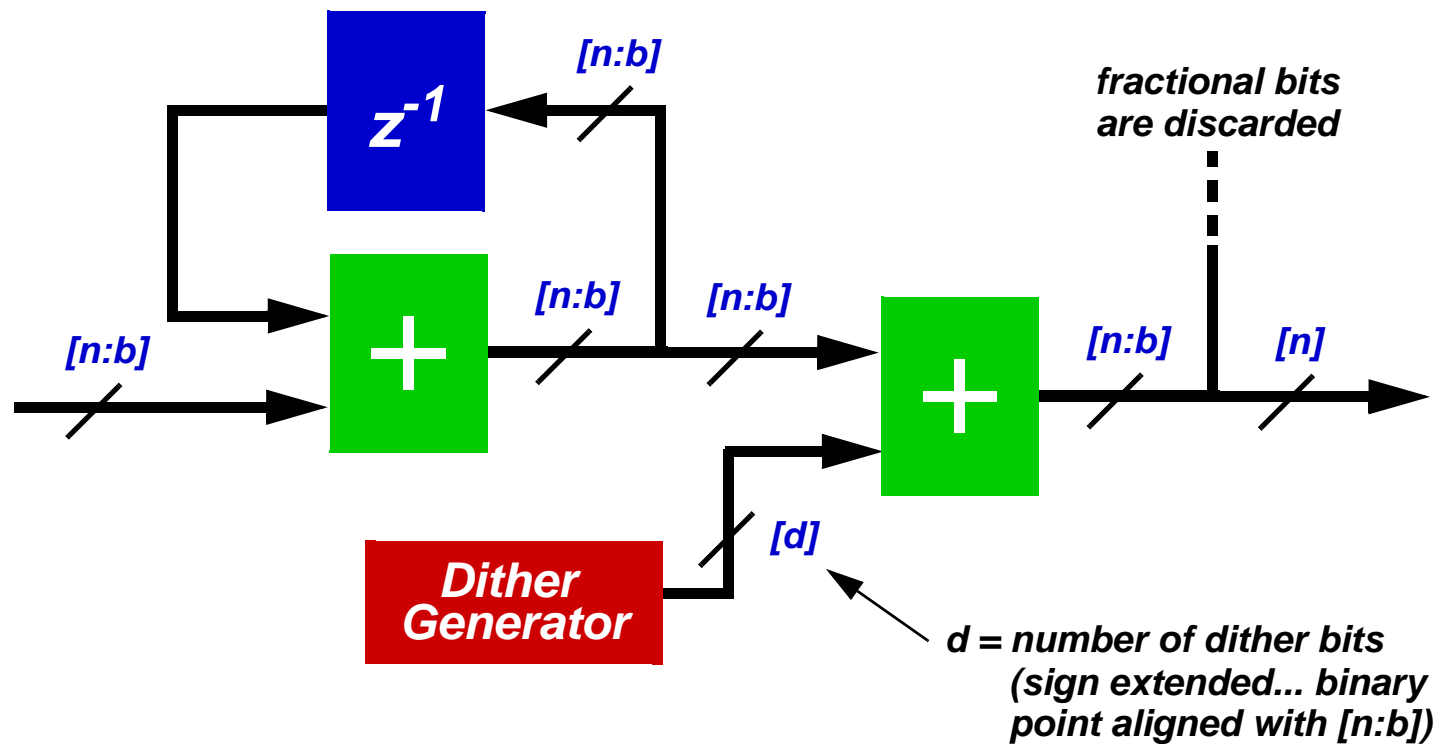
It has been established that phase truncation has the most significant effect on Spurious Free Dynamic Range.

Addressing the issue by expanding the whole part of the word is expensive in hardware terms. For every whole bit added to the accumulator wordlength, the LUT size doubles.

Phase dithering is therefore a favoured method of increasing the SFDR.

Phase Dithering

- As shown, a low-level of pseudo random noise is added to the accumulator output. This happens prior to truncation of the word to the LUT address.




- The number of dither bits can be varied.
- We will now analyse the impact of dither bits on SFDR performance.

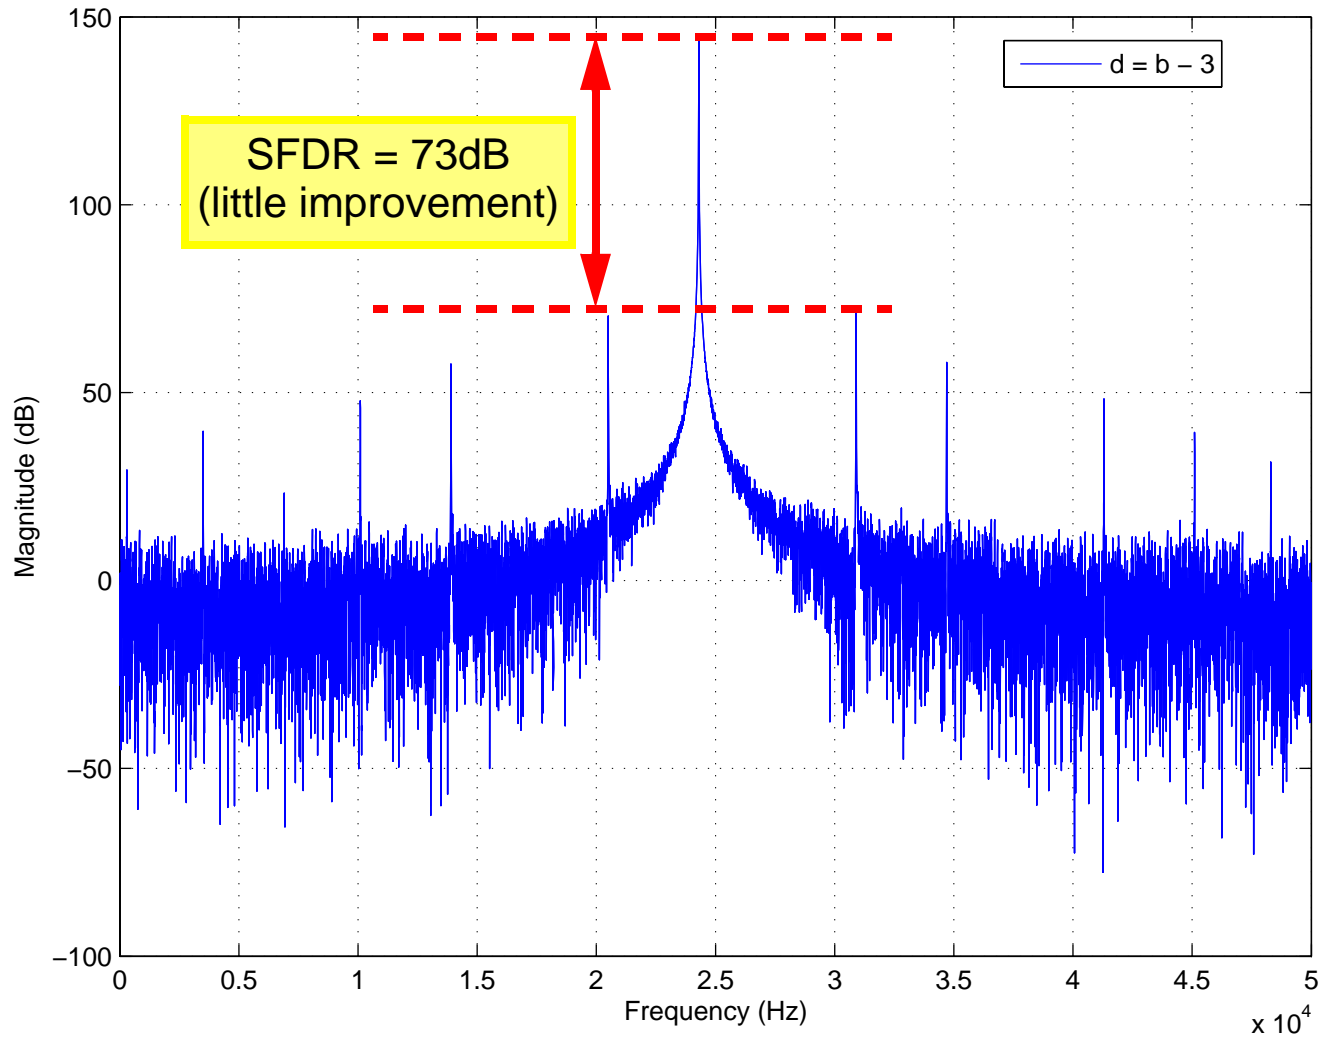
Notes:

Dither may be generated using a Linear Feedback Shift Register (LFSR). This involves a shift register of M taps, with selected taps fed back (usually via XOR gates) to create a sequence which repeats after $2^M - 1$ clock cycles. Therefore, long pseudo random noise (PN) sequences can be created using very little hardware. For example, a 12-element LFSR can create a PN sequence which repeats after 4095 clock cycles. The cost of implementing such an LFSR is simply M D-type flip-flops, and a few combinatorial elements.

To create d dither bits, the outputs from d taps are combined into a vector, where $d \leq M$. This is illustrated for a simple example below. In practice, a longer shift register would be used. The longer the shift register, the more random the output appears.



Too Few Dither Bits



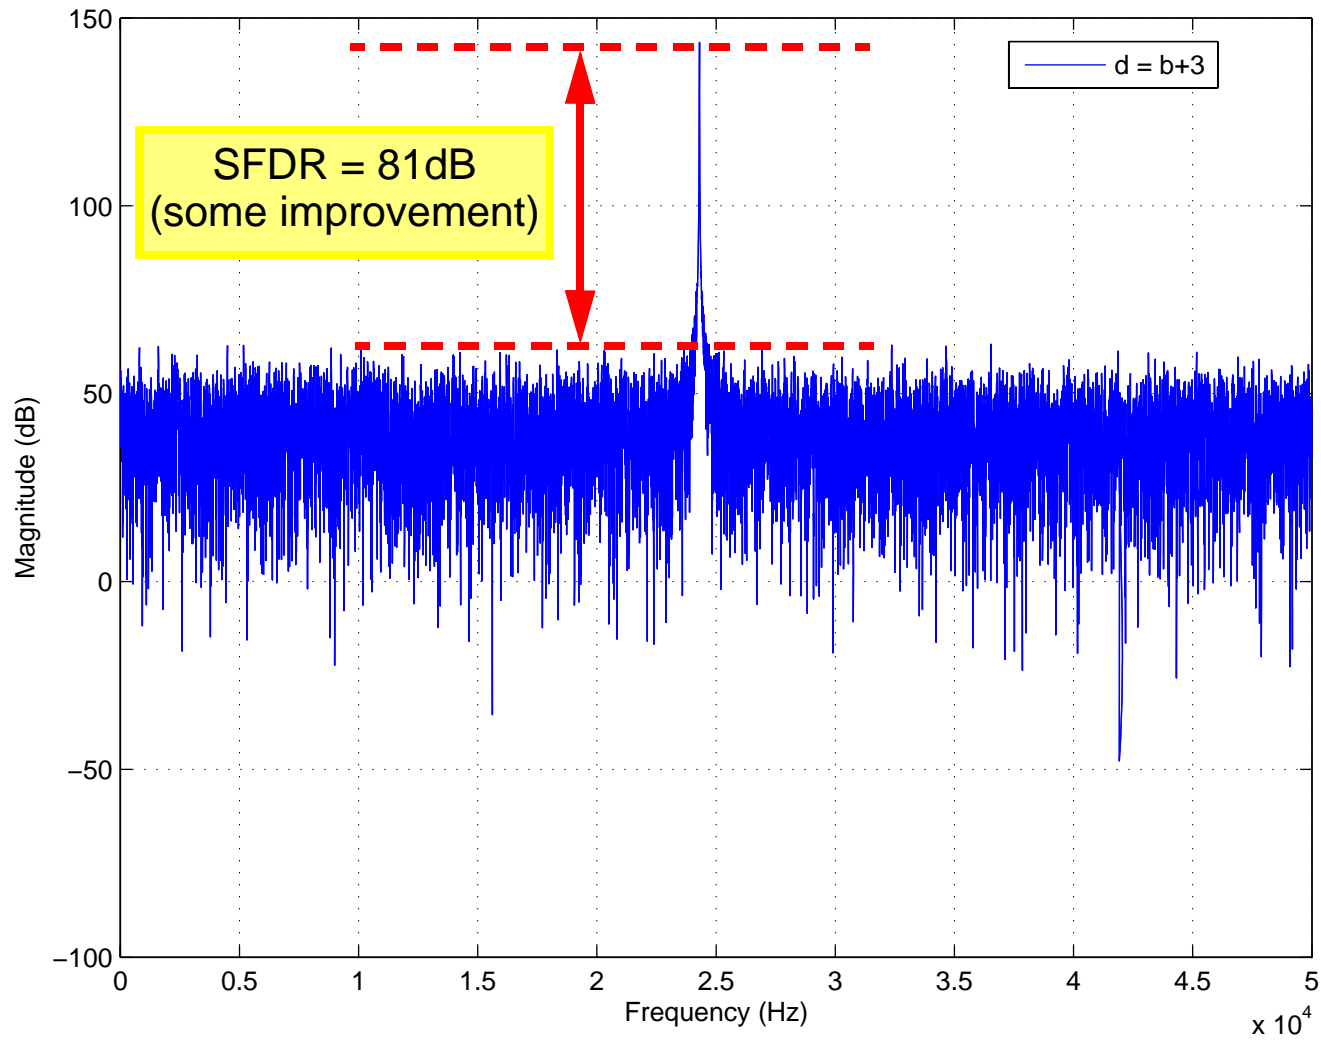
Notes:

This spectrum corresponds to the NCO example given in Slide 11.12.

Adding too few dither bits has little impact on the structure of the phase truncation, so the spurs remain prominent.

In this example, $b - 3$ dither bits are added. The number format is unsigned, so on each sample, the dither added is between 0 and $\frac{1}{8}$ th of one whole bit. SFDR has been improved by about 1dB compared to the original case.

Too Many Dither Bits




Notes:

If too many dither bits are added, the spurs disappear but noise floor of the spectrum is raised. Depending on how many dither bits are added, this might be even worse than the original SFDR.

In this example, $b + 3$ dither bits are added, so the maximum extent of the dither is 8 whole bits. The SFDR has been improved by 9dB compared to the original.

The “Right” Amount of Dither Bits



Notes:

Finally, choosing $d = b$ appear to give the best result. The spurs have disappeared and the noise floor is much lower than in the previous example.

Here, the dither samples are in the range 0 to 1 whole bits, which provides a good level of noise to break up the pattern of the truncated phase.

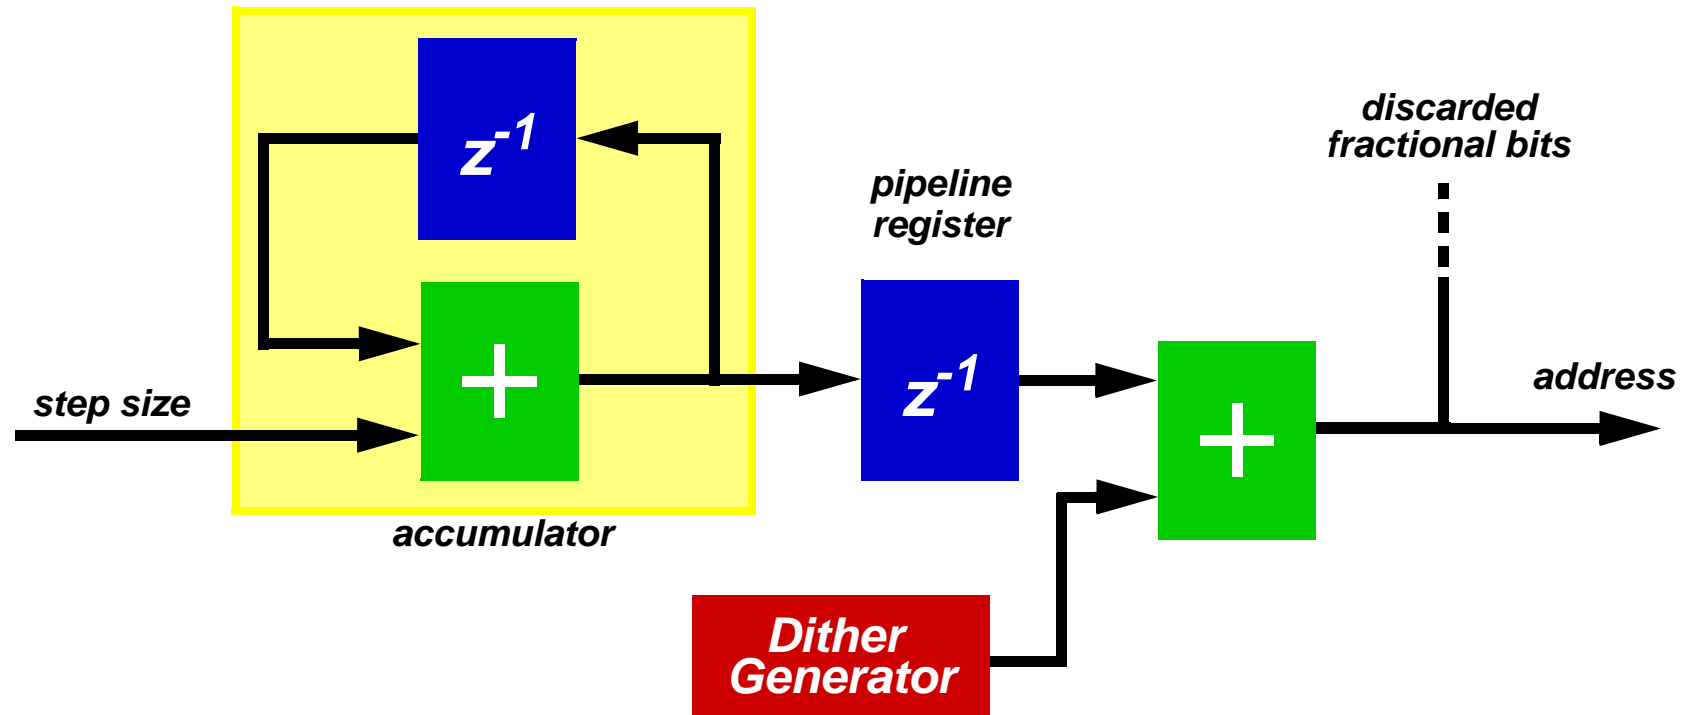
NCO Hardware Considerations

- The adder in the accumulator represents a long combinatorial path, particularly for large wordlengths. This adder is ideally mapped to a DSP48 (in Virtex-4) or DSP48E (in Virtex-5).
- The maximum speed at which the design can be clocked may be limited by this adder.
- If adding dither, the critical path contains both adders, thus doubling the problem!
- It is vital to check that the design will clock at the required sampling rate (f_s) - if not, the output frequency will be wrong!
- The resolution of the sine wave samples stored in the LUT may be influenced by the requirements of subsequent DSP blocks.



Notes:

The critical path can be shortened by inserting a pipeline register between the two adders, as shown. This simply adds a delay of one clock cycle to the NCO output.



Conclusions

- Numerically controlled oscillators play an important role in DSP - particularly digital communications - and may be found in a variety of devices.
- NCOs are constructed cheaply in digital hardware. The most common form of NCO comprises an accumulator and a sine wave look up table, with the output frequency defined by a step size parameter.
- A critical performance criterion is Spurious Free Dynamic Range, which indicates the purity of the synthesised sine wave.
- The frequency resolution is also important. This describes how large a frequency offset may be encountered between the desired and synthesised frequencies.
- Several trade-offs exist between NCO performance and hardware cost.
- Dithering the phase of the accumulator output is a simple, low cost method of improving SFDR performance.

