A Mixed Signal Frequency Synthesiser for Configurable Communication Systems

P. Gray, D.C. Hendry, R. Cambio, W.W.S. Ma
Department of Engineering
University of Aberdeen
Aberdeen AB24 3UE
UK

Email: pete@eng.abdn.ac.uk

Abstract

This paper describes the development and characterisation of a wideband frequency synthesiser. The synthesiser is a development of a design originally described in a US patent filed in 1975. The problems with the original design are discussed and alternative approaches are presented that overcome the limitations of the technology originally available in the 1970s so as to resolve such issues such as clock contention, which were not addressed at that time. The objectives of the research described herein are the development of a wide band frequency synthesiser, the simulation of the synthesiser at frequencies compatible with those of interest to research collaborators, and the estimation of the power consumption of the overall design. The work undertaken is part of Project MuMoR, an IST Framework V funded project, an objective of the project being the development of flexible mixed signal architectures relating to configurable communication systems. The paper describes the approach to frequency synthesis adopted in the research. It concludes that the frequency synthesis technique offers the possibility of synthesising an exceptionally wide band of target frequencies whilst automatically switching the VCO's tuning bands .

1. Introduction

Frequency synthesisers find many diverse applications within the electronics field. Direct analogue synthesis is one of the oldest techniques, originally used in communications and in apparatus such as laboratory signal generators. The availability of crystal-controlled oscillators which allowed the output frequency of such synthesisers to be referenced to a signal source whose frequency could be defined and controlled (for example within a constant temperature environment) with some considerable accuracy was an important factor in the development of such equipment.

1.1 Types of Frequency Synthesisers

Frequency synthesisers can be categorised into three groups: analogue, digital or mixed signal (hybrid). The frequency synthesiser described in this paper falls into the latter category. The majority of frequency synthesisers employed within the field of mobile communications employ a phase locked loop with a divide-by-N in the feedback loop. This allows the channel to be switched conveniently by changing the value of the integer, N.

1.2 Frequency Synthesisers for mobile communications

The integer-N PLL (phase-locked loop - see section 2.1) is a commonly used building block of frequency synthesizers as employed in mobile communications. It has the advantages of low-power operation and low component count, compared to other synthesizer implementations. The PLL-based synthesiser has by its nature a narrow bandwidth and the VCO phase noise in this context is a critical design parameter. In order to keep phase noise low it is essential to limit the conversion gain of the VCO, and to do this band switching is normally employed.

The basic functional blocks of an integer-N PLL synthesiser[1] are outlined in section 2.1.

1.2 Re-configurability of building blocks

In the MuMoR project the synthesisers for the RF frontend must support different local oscillator signals for the WCDMA FDD, WCDMA TDD and GSM 900/1800 modes. It can be demonstrated that a synthesizer count of two is achievable if a combined synthesizer for WCDMA and GSM is used. Within the project, Nokia are currently working towards this goal by employing a reconfigurable multi-mode VCO under development at EPFL and fabricated at ST Microelectronics. In general it is found with this approach that functional blocks within the synthesiser under development are directly applicable for all supported modes without the need for auxiliary circuitry for reconfiguration purposes, made possible by the fact that both WCDMA and GSM require the same channel spacing of 200 kHz.

2. PLL-based Synthesisers

The use of PLLs in frequency synthesis is well-established. PLLs are flexible, but they do have certain drawbacks; in particular they are limited in their ability to continuously synthesise signals over a wide band of frequencies. The oscillator (see below) within the loop needs to be band-switched in order to cover the frequencies associated, for example, with mobile communications. Otherwise, a high conversion gain, K_v, is required, with a subsequent increase of in-loop noise. However, the PLL itself cannot perform this task automatically, and some extra circuitry is needed for this function.

2.1 Integer-N loops

The most commonly implemented PLL synthesiser is the integer-N PLL, as represented by Figure 1. The frequency of a reference signal from a crystal-controlled source is pre-scaled by frequency divider (omitted from Figure 1 for clarity) before being presented to one input to a phase detector. The output of the phase detector is low pass filtered to remove unwanted signal components. The filter's output is a DC voltage, which is used to control the frequency of a voltage-controlled oscillator whose output is the desired synthesized signal. A second frequency divider within the loop can be used to switch the output signal frequency from channel to channel by changing the value of N, the frequency divider's input-to-output frequency ratio.

2.2 Fractional-N PLL

Fractional-N synthesisers offer further control of output signal frequency by dithering the in-loop divider's input-to-output frequency ratio value between N and N+1. Unavoidable in this technique is the generation of spurious sidebands; these require elimination by extra hardware, which pushes up the power consumption figure.

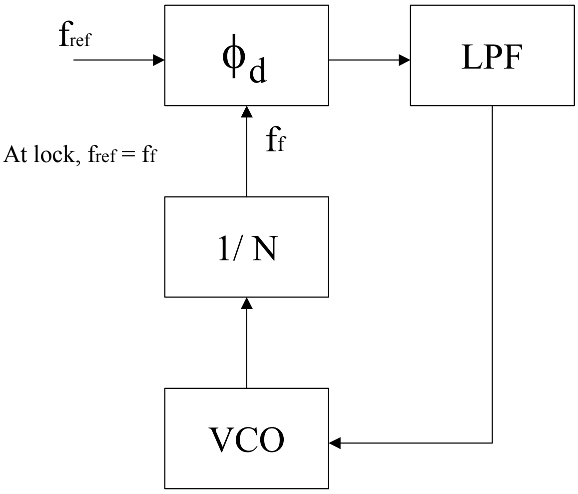

Figure 1: The Integer-N PLL

2.3 Direct Digital Synthesis

Direct digital synthesis (DDS) offers an alternative approach to synthesis; it involves the repetitive generation of sine wave cycles, typically via a look-up table for the function sinθ held in a ROM (read-only memory). PLLs have been designed with a DDS synthesizer incorporated within the loop[2] that offer great flexibility in the determination of output frequency. Again there are power consumption issues, and the upper bound of the output frequency is limited by the clock frequency of the DDS block and the size of the sampling interval of the [sampled] sine wave stored in ROM.

3. A Wide Band Synthesiser

In this paper we present a wide-band hybrid frequency synthesiser. The ability to generate signals over a wide range of frequencies as compared to a PLL-based synthesizer is an important feature. Another feature of the design is the ability to automatically band switch the VCO without the need for extra circuitry. The essential elements of the synthesiser, based on a patent filed in 1975 by R.J. Bosselaers[3], are shown in Figure 2. It is similar in structure to a PLL, but in this synthesiser design the phase detector is replaced by a digital frequency discriminator. In general the reference frequency f_{ref} is not equal to the feedback frequency f_f, so the device cannot be categorised as a PLL. A DAC (omitted for clarity) is needed to convert the digital N-bit signal to analogue form before admitting it to the LPF.

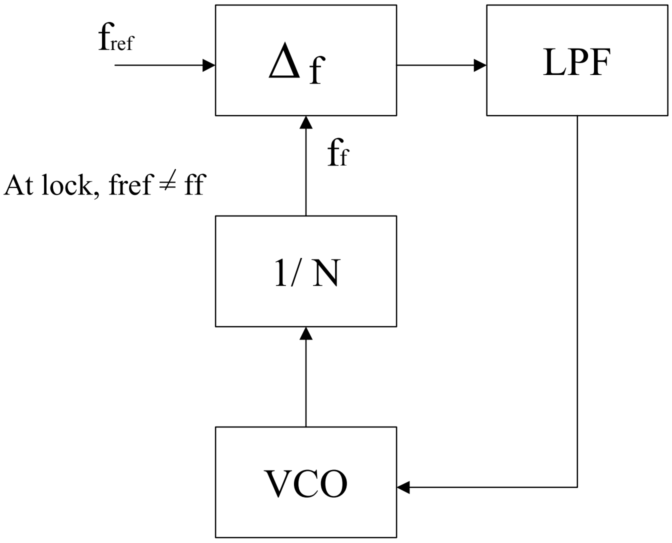

Figure 2**:** The Essential Blocks of the Wide-band Frequency Synthesiser

3.1 Operation of the Digital Circuit

The digital phase detector design developed at Aberdeen is loosely based on the patent of Bosselaers but formulated in a manner which permits synthesis to a standard cell library using typical synthesis tools. A block diagram of the synthesiser is shown in Figure 3. The circuit provides a VCO control output expressed generally as an N-bit integer (N=20 is shown) derived from two clock signals. The first clock signal, f_ref, represents the reference signal (as might be supplied by a crystal controller local oscillator), and the second signal, f_feedback, represents the output of the VCO after dividing by a known integer. The VCO control output is then stable when the ratio of these two frequencies is given by N1/N2. The values of the constants N1 and N2 are supplied as binary integers to the circuit from, for example, a controlling microprocessor.

In the original Bosselaers design a value N1 was added to the result register on a rising edge of the reference signal, and a value N2 was subtracted from the result register on a rising edge of the feedback signal. Clearly, clock contention problems were unavoidable; also two ALUs were needed to implement the design. In the synthesiser design presented here the constant N1 is held in the register K1 (see Figure 3) and added to the result register on the rising edge of the reference frequency, as before. The register K2, however, holds the *difference* between the two constants (N2 – N1).

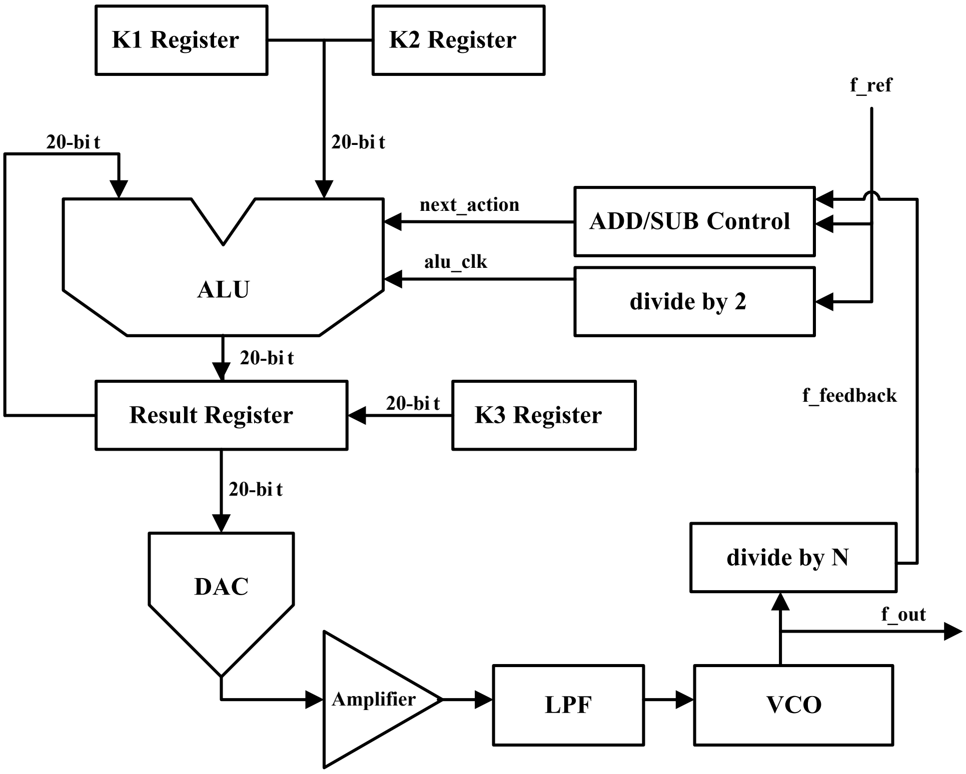

Figure 3: Block diagram of the complete synthesiser

The principle of the circuit is to synchronise the lower of the two frequencies to the higher frequency via the logic block 'ADD/SUB control'. This eliminates a problem associated with the original design. By sampling the lower frequency signal on both the rising and falling[1] edges of the higher frequency signal, the occurrence of cycles of both signals may be detected. The occurrence of these cycles then controls the data added or subtracted from a register containing the accumulated phase difference. In any cycle of the reference frequency clock in which only the reference frequency completes a cycle, N1 is added to that accumulator; whereas in a cycle of the reference frequency clock in which both signals complete a cycle, N1 – N2 is subtracted from the accumulator.

The result register itself is pre-loadable with a constant via the register K3, which can facilitate significantly faster acquisition of the target frequency f_out, as discussed later in this paper.

3.2 Behavioural Model

Prior to modelling with Cadence tools an ADS behavioural model was built. The frequency discriminator block was modelled as shown in Figure 4 using ADS library components. All other circuit elements VCO, frequency divider, single pole filter, etc were also readily available from ADS libraries. The VCO was modelled with a conversion gain, Kv, of 143 MHz/v corresponding to a frequency output range of 4.4 to 4.6 GHz.

---

[1] In which case the 'divide by two' block shown in figure 3 is not needed; it is shown (for generality) as present for an alternative implementation where only rising edges of fref are used as sampling triggers.

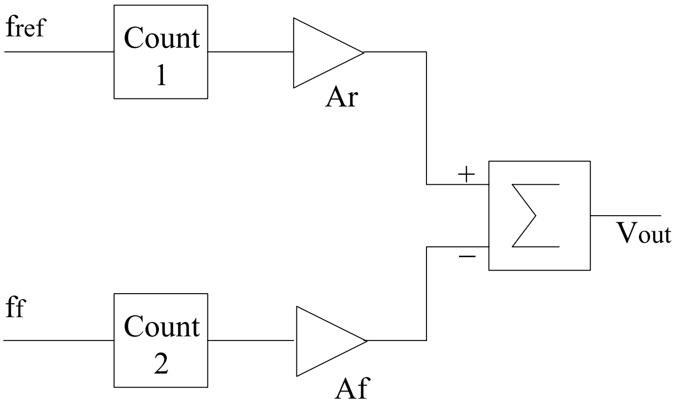

Figure 4: Behavioural Model of Frequency Discriminator

Amplifiers are employed within this model to simulate the constants N1 and N2 used in the synthesiser proper. They are fed from counters; one counter (‘count1’) is triggered by the reference signal while the other counter (‘count2’) is triggered by the feedback signal. Each amplifier supplies a staircase voltage to the difference amplifier; when the two staircase voltages have the same slope the output voltage of the frequency discriminator is constant.

In general, experimentation with this ADS behavioural model showed that the synthesiser ramped exponentially upwards or downwards towards the target voltage required to drive the VCO to its stable output frequency. This behavioural model provided valuable insight into the dynamics of the operation of the synthesiser, and allowed rapid and convenient evaluation of the synthesiser at various different frequencies. Figure 5 shows a typical response of the circuit for Ar = 1 and Af = 0.8.

3.3 Simulation of the Bosselaers-derived Frequency Discriminator

The design has been evaluated using a typical standard cell library and experiments show that an operating frequency of at least 100MHz (logic clock) is feasible. It is likely however that an operating frequency this high will not be required, and a lower operating frequency selected to provide a suitable trade off between power dissipation and phase noise.

Additionally, synthesis of the digital phase detector using the ST 0.25µm library provides a range of area and maximum clock frequencies for various PVT (Process, Voltage and Temperature points). These indicate an area of some 30,000 µm**2, and a maximum operating frequency of over 200MHz.

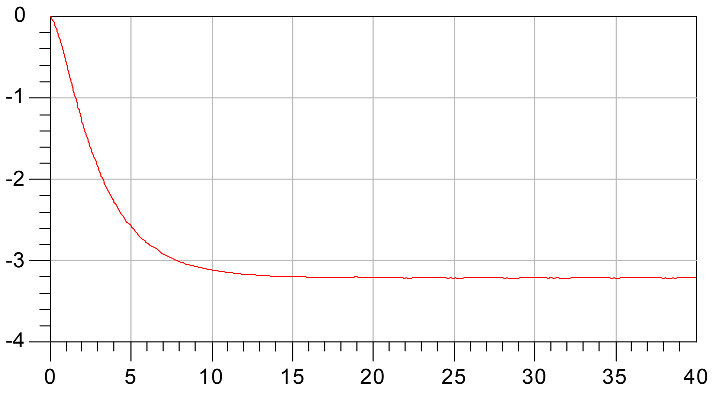

Figure 5: VCO control voltage v. time

3.4 Power Consumption of the Frequency Discriminator

With the digital [frequency discriminator] logic running at a reference clock rate of 100MHz it has been determined that the frequency discriminator block dissipates approximately 2mW with an operating voltage of 1.8V and 4mW with an operating voltage of 2.5V. These figures represent dynamic power, which is the power dissipated by a circuit when the circuit is active, that is whenever the voltage on a net varies because of some stimulus applied to the circuit[3]. These power estimates were obtained with Synopsys Power Compiler targeting a 0.25 µm standard cell library. The net transition counts required by the tool was obtained by means of simulation of a VHDL Register Transfer Level (RTL) description of the design with Synopsys VSS.

4. Analysis of vhdl Simulation Results

The frequency discriminator circuit described above was initially coded in vhdl and tested with Cadence simulation tools for a reference frequency of 100 MHz. A vhdl testbench was also coded to drive the frequency discriminator through its I/O and control ports.

Figure 6 shows the simulation results for a reference frequency of 100 MHz. The main points to note are:

- The exponential ramping (K3 loaded to half maximum value) of the VCO control voltage (v_lpf_out) that conforms to the behavioural model's predicted performance
- The pre-loading (K3 loaded to target value) of the output register (5 mS, 6 mS) before switching to a new target frequency makes acquisition of the target frequency almost instantaneous

This suggests that for mobile communications applications, a microcontroller within a handset would be able to pre-calibrate the synthesiser and store in its memory a set of K3 values appropriate to any desired channel frequency.

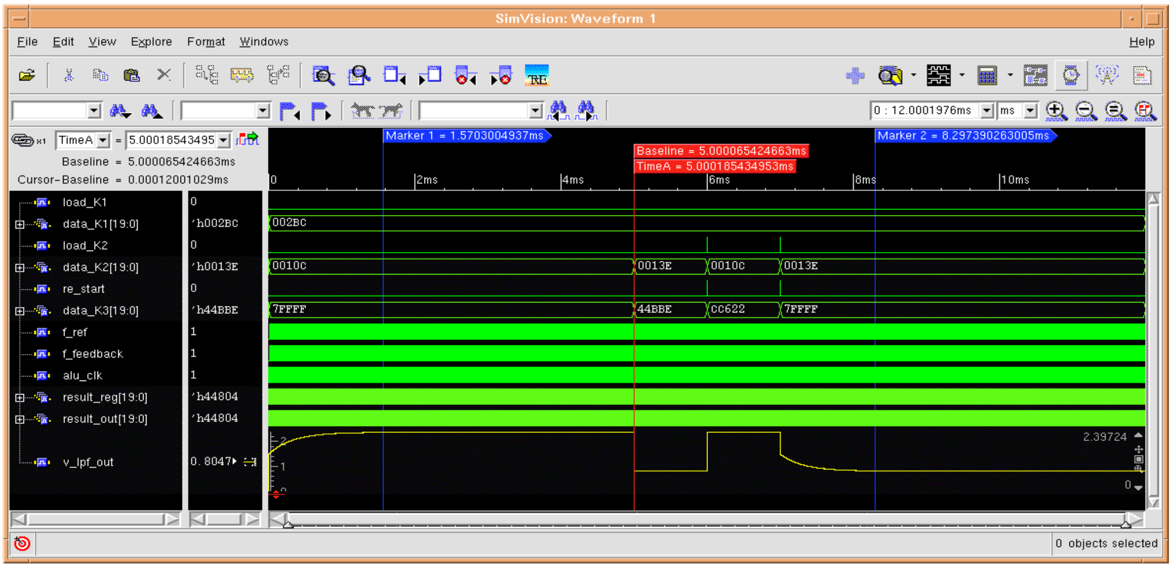

Figure 6: vdhd-coded synthesiser simulation (fref = 100MHz)

It is possible to the use the high-order bits of the synthesiser's output register to band switch the VCO. With this design, if the upper range of the current VCO band is exceeded while seeking to reach the target frequency, the effective voltage accumulating continues to ramp up in the output register; if the high order bits of the register are wired directly to the band select lines of the VCO, the synthesiser will automatically hunt through the bands until the target frequency is reached. No calibration or intervention by extra circuitry is required, which is a major advantage this design has over other synthesiser designs.

Table 1 shows a summary of simulation results for 7:5 ratio of reference frequency to feedback frequency. It can be seen that if large-value constants are loaded into the K1 and K2 (recalling that K2 = N2–N1) registers, then the time taken to reach a stable output frequency is reduced.

K1	K2	K3	V_max	V_min	Acquisition time
50	20	MAX/2	-	-	>10 mS
500	200	MAX/2	1.99918	1.99846	3.03 mS
5,000	2,000	MAX/2	2.00282	1.99567	202 uS
5,000	20,000	MAX/2	2.03644	1.96492	21 uS

Table 1:Acquisition time for 7 to5 ratio, alu_clk = 50 MHz

Table2 shows the effect of pre-loading the output register with the known final value. (This corresponds to the 'pre-calibrated' situation discussed earlier.) It can be seen that the acquisition time is reduced to nanoseconds, and in reality corresponds to the flushing time of the four-input average filter implemented within the test bench.

K1	K2	K3	V_max	V_min	Acquisition time
5	2	698638	1.99883	1.99882	120nS
50	20	698638	1.99881	1.99877	120nS
500	200	698638	1.99918	1.99846	120nS
5	3	768466	2.19861	2.19860	120nS

Table 2:Acquisition time with predicted value of K3 pre_loaded

5. Mixed-signal Simulation

A full mixed signal simulation has been performed using Verilog-ams. A simple single pole filter was used as the loop filter. Figure 7 shows results obtained for a 7:5 ratio of reference to feedback frequency.

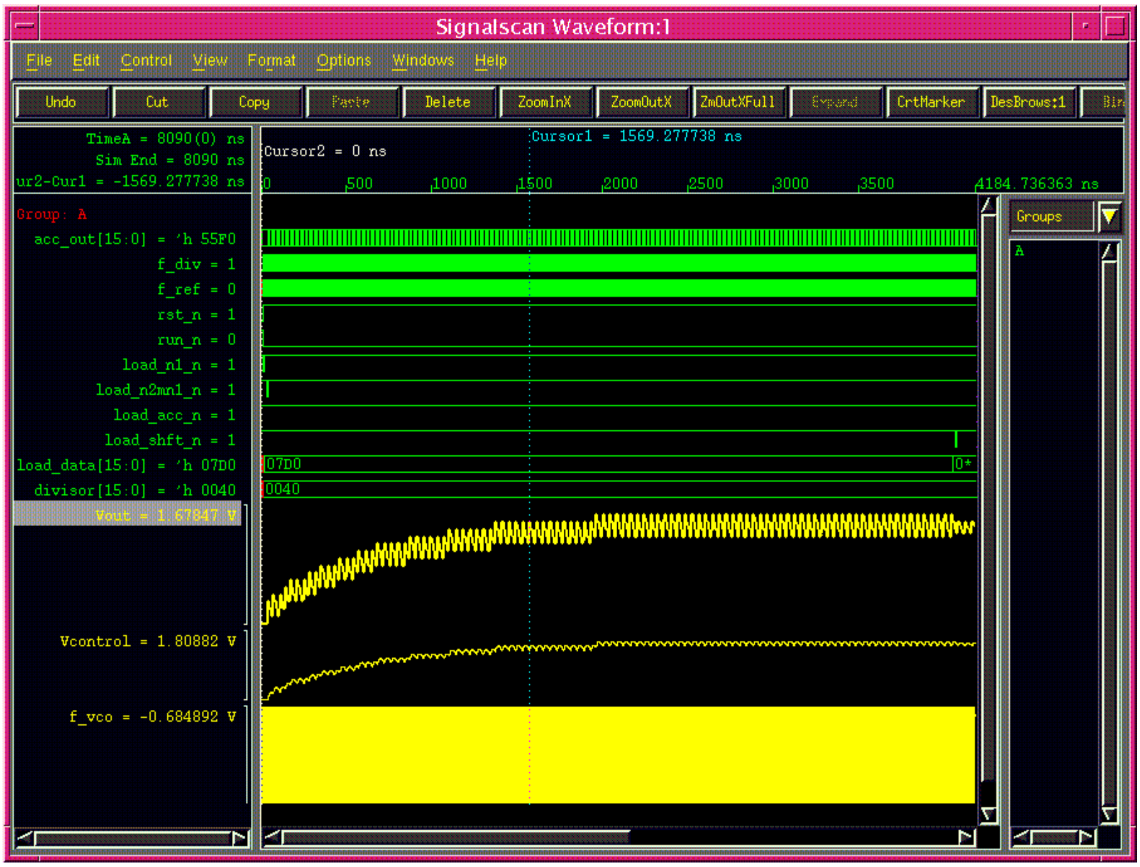

```
TimeA = 8090(0) ns
Sim End = 8090 ns
ur2-Cur1 = -1569.277738 ns
Cursor2 = 0 ns
Cursor1 = 1569.277738 ns
Group: A
acc_out[15:0] = 'h 55F0
f_div = 1
f_ref = 0
rst_n = 1
run_n = 0
load_n1_n = 1
load_n2mn1_n = 1
load_acc_n = 1
load_shft_n = 1
load_data[15:0] = 'h 07D0
divisor[15:0] = 'h 0040
Vout = 1.67847 V
Vcontrol = 1.80882 V
f_vco = -0.684892 V
```

Figure 7: Ratio 4:3 (4.8 GHz) – Start-up

5.1 Frequency Discriminator Output

For a frequency ratio such as 4:3 above, sampling the feedback frequency with the rising edges of the reference frequency clock will produce a low frequency output signal (3) repeating at intervals of four cycles of the reference signal (1) or three cycles of the feedback signal (2) as shown in Figure 8.

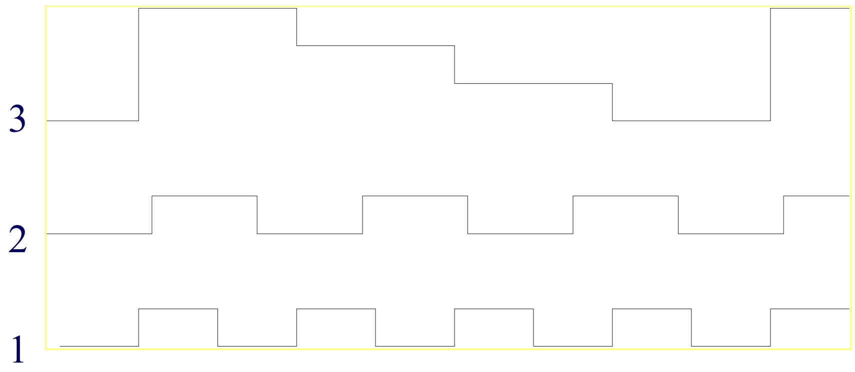

Figure 8: Frequency Ratio 4:3

These can be seen clearly in the Verilog-ams simulation (Figure 9). Such signals could prove problematic for larger integer ratios, e.g. 40:29.

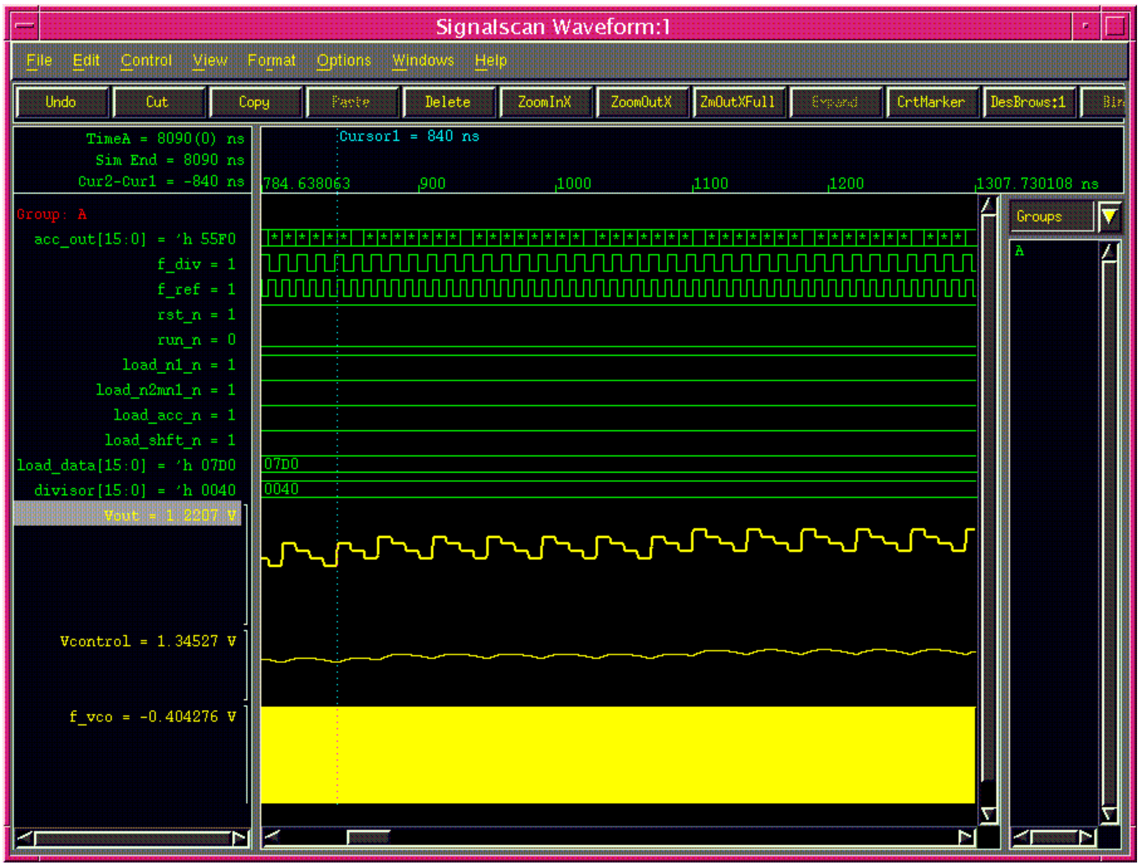

Figure 9: Ratio 4:3 (4.8 GHz) Frequency Discriminator Output Waveform

6. Conclusions

In this paper we have presented a wide-band hybrid frequency synthesiser. The bandwidth of the synthesiser is limited only by the data width of the digital circuit block and associated DAC inputs. The ability to generate signals over a wide range of frequencies as compared to a PLL-based synthesiser is considered an important feature. Another attraction of the design is the ability to directly band-switch the VCO without the need for extra hardware. The major issue to be addressed in the future is the elimination of unwanted low-frequency components present at the output of the frequency discriminator.

7. Acknowledgements

We wish to acknowledge the generous support of Agilent Technologies and their donation of ADS research without which the behavioural modelling of this synthesiser would not have been possible.

References

[1] R.Best. Phase locked loops: Design, Simulation, and Applications, McGraw-Hill Professional, 1999.

[2] P.V. Brennan, R. Walkington, A. Borkjak, Performance of PLL Synthesiser based on DDS feedback, IEEE Electronic Letters No. 19991512.

[3] R.J. Bosselaers:"Phase Locked Loop including an Arithmetic Unit": US Patent No 3913028, Filed US Patent Office, 1975.

[4] Synopsys Power Compiler Reference Manual, version 2000.05, May 2000.