UNIVERSITY OF SOUTHAMPTON

High frequency dividers in bulk and silicon-on-insulator CMOS technologies.

by

Ketan.Mistry

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Engineering and Applied Science

Department of Electronics and Computer Science

October 2004

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND APPLIED SCIENCE
DEPARTMENT OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

High frequency dividers in bulk and silicon-on-insulator CMOS technologies.

by Ketan.Mistry

One of the key components common in integrated receiver designs is a RF local oscillator. This particular block is implemented as a phase-locked loop frequency synthesiser in many published high frequency architectures. A sub-block of this loop is the frequency divider, which sits in a feedback loop between the phase detector and voltage controlled oscillator (VCO). It is this block, along with the VCO, that operate at very high frequencies, and will therefore be heavy consumers of power compared with low frequency blocks.

This thesis brings together the results of an investigation into frequency division, with a particular emphasis on high frequency dual-modulus dividers intended for fabrication in a CMOS process. The motivation behind this research is driven by the continuous demand for low power high speed circuits, as justified above. The static source-coupled logic forms the basis of the high frequency divider cells, and investigation into its behaviour and attempts to ascertain what restricts its high frequency characteristics. After a review of some published dividers, attention is turned to the phase selection architecture and a glitch-free control scheme that will play a key role in the design of two dual-modulus dividers. One such divider is the divide-by-64/65, which also demonstrates a new circuit topology in which stacking and current re-use exploit SOI CMOS technology. The second is a divide-by-16/17 circuit implemented in bulk sub-micron CMOS again with current-steering, but with the omission of stacking and thus makes for a useful comparison. This second divider also plays a crucial role in the design of a subsequent integer-N programmable divider (divisors 513-544), intended for use within a frequency synthesiser as part of a wireless IEEE802.11a complaint receiver IC, and it is detailed in the penultimate chapter. Another divider circuit developed in this work is a fixed divide-by-2 bulk CMOS cell capable of operating beyond 10GHz, intended for the very same wireless project. All divider circuits have been fabricated and successfully measured with results given in each of their individual discussions.

Contents

List of Figures

List of Tables

Acknowledgements

The work in this thesis has been supported by numerous individuals over the period of study. Therefore, I feel a moral obligation to mention as many of those people as possible. I sincerely apologises for any omissions.

First and foremost, I would like to thank Prof. Bill Redman-White for accepting me as his research student, as well as reading and re-reading the numerous editions of this thesis. I must also acknowledge his link with Philips Semiconductors as it proved invaluable during the latter half of my tenure.

Next, I would like to express my thanks to my two colleagues, James Benson and Dr. Nele D'Halleweyn. As well as assisting me on the challenges of SPICE modelling and SOI CMOS transistors, they also provided a more sociable and relaxed atmosphere in which to work.

Within the ECS department, all the research students (past and present) and support staff within the Microelectronics department are thanked for their assistance and/or friendship during my stay. I cannot recall any hostility by any individual that would have made for an uncomfortable work environment. Staff in the Mechanical workshop, Electronics Workshop and Computing support are mentioned, as their help was invaluable during the research. I should also thank EPSRC in providing the generous scholarship which eliminated the money woes associated with the basic day-to-day necessities.

It was close to the end of the second year, that a decision to send me to Philips Research Laboratories (Eindhoven, The Netherlands) was made. This brought me into contact with the Integrated Transceivers Group, where I worked on a CMOS wireless LAN project. I owe a great deal of thanks to those colleagues and friends, in particular, Dr. Pieter Hoojmans, Dr. Donnine Leenaerts, Nenad Pavlovic, Dr. Cicero Vaucher, Denis Jeurrisen and Peter De Vreede. The work environment and well-equipped lab (as well as encouragement to learn their operation) proved a valuable experience.

Lastly, I should also thank Prof. Orla Feely, for employing me in the Circuits and Systems group at the University College Dublin, whilst finishing this thesis. To employ me as a full-time member of staff, was a huge risk on her behalf, and hope she will not regret the time it has taken to close this thesis off.

Chapter 1

Introduction

For the next generation of mobile communications, IC designers are called upon to design in a higher frequency regime. A key component is the radio transceiver that provides the signal path interface for data to and from a resident signal processor. It is typically the smallest block in a communications chip but tends to consume the most power per unit area and requires the most design time. As the need for more channels to exist simultaneously becomes important (with each channel accommodating a larger bandwidth), the specifications and performance required become tighter and greater respectively. Chips are required to be manufactured at a low cost, and designs using standard CMOS processes always attract attention owing to their integratability. Alternative technologies such as bipolar and BiCMOS technologies have been known to deliver radio sections capable of operating at high speeds, and in the present day, bipolar processes have fewer mask stages, whilst BiCMOS processes have slightly more masks than very short channel CMOS processes. Where availability and integratibility are issues when deciding upon a technology (especially where 'logic' is needed), then plain CMOS would seems the best choice.

This chapter starts with a brief look at down-conversion schemes. As all of the design work has been performed in CMOS, a section follows describing the current CMOS technology, both in bulk and SOI. The last section puts the design of frequency dividers into context, explaining the roles of each sub-block in PLL-based CMOS frequency synthesisers, as well as supplementing each with examples of published designs.

1.1 Radio architectures and current implementations

Before starting this section, it should be emphasised that there is no intention to describe in detail radio architectures for IC implementation. Whilst frequency synthesisers

are present in wireless front-ends, they are not restricted to local oscillators, and can be used in clock recovery, and coherent frequency demodulation.

The aim of the front-end is to take a digital signal that has been converted to analogue and then up-convert this baseband signal to a higher band where it can be transmitted with enough power via an antenna. The methods by which to achieve this are not restricted to one particular up-conversion scheme, but characteristics associated with the standard usually dictate the choice of one arrangement over another. The Superheterodyne receiver, such as that seen in Figure 1.1, is the general structure for most common front-ends with the frequency planning determining the number of channels.

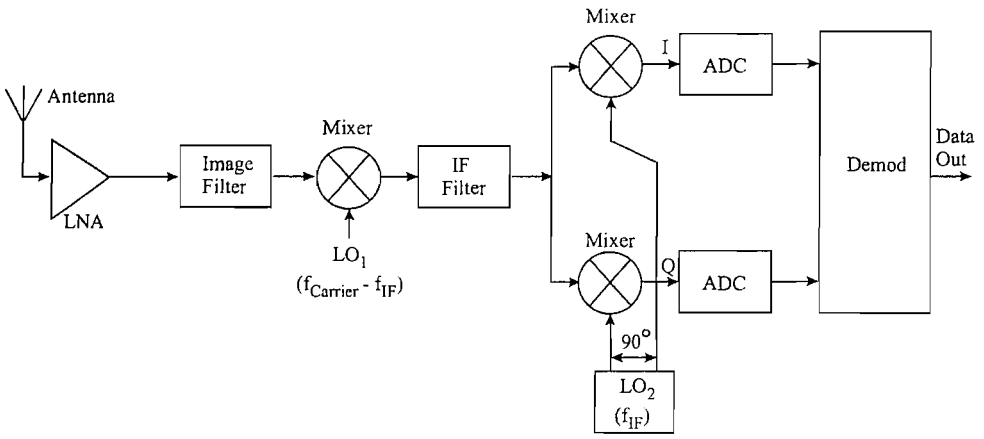

FIGURE 1.1: Block diagram of a Superheterodyne receiver, tailored for digital wireless transmission.

After a low noise amplifier and image reject filter, a mixer is usually encountered, driven by an RF input and a local oscillator (in the case of a receiver). During this mixing stage, the RF input is down-converted to an intermediate frequency(IF), whereupon a channel select filter will attenuate out-of-channel interferers. The trade-off has to be whether to choose a high IF or low IF. Assuming high side injection, where the local oscillator frequency is above the RF signal (for the purpose of a relaxed tuning range), a high IF will mean that the image reject filter should attenuate an image band 2 x IF higher and hence not downconvert it to the IF together with the desired channel. However, any signal close to the desired channel (perfectly acceptable as it is in-band) will need to be suppressed using a tight channel select filter operating at a high frequency and this is a major problem. By having a low IF, the image reject filter doesn't usually suppress any image band as much as it should and thus mixes down to the IF with a greater power (in this image band) than in the previous setup. The advantage is that a tighter channel select filter response can now be realised operating at a lower frequency. A high IF incidentally has need for a local oscillator (LO) with a lower fractional tuning range [49][11].

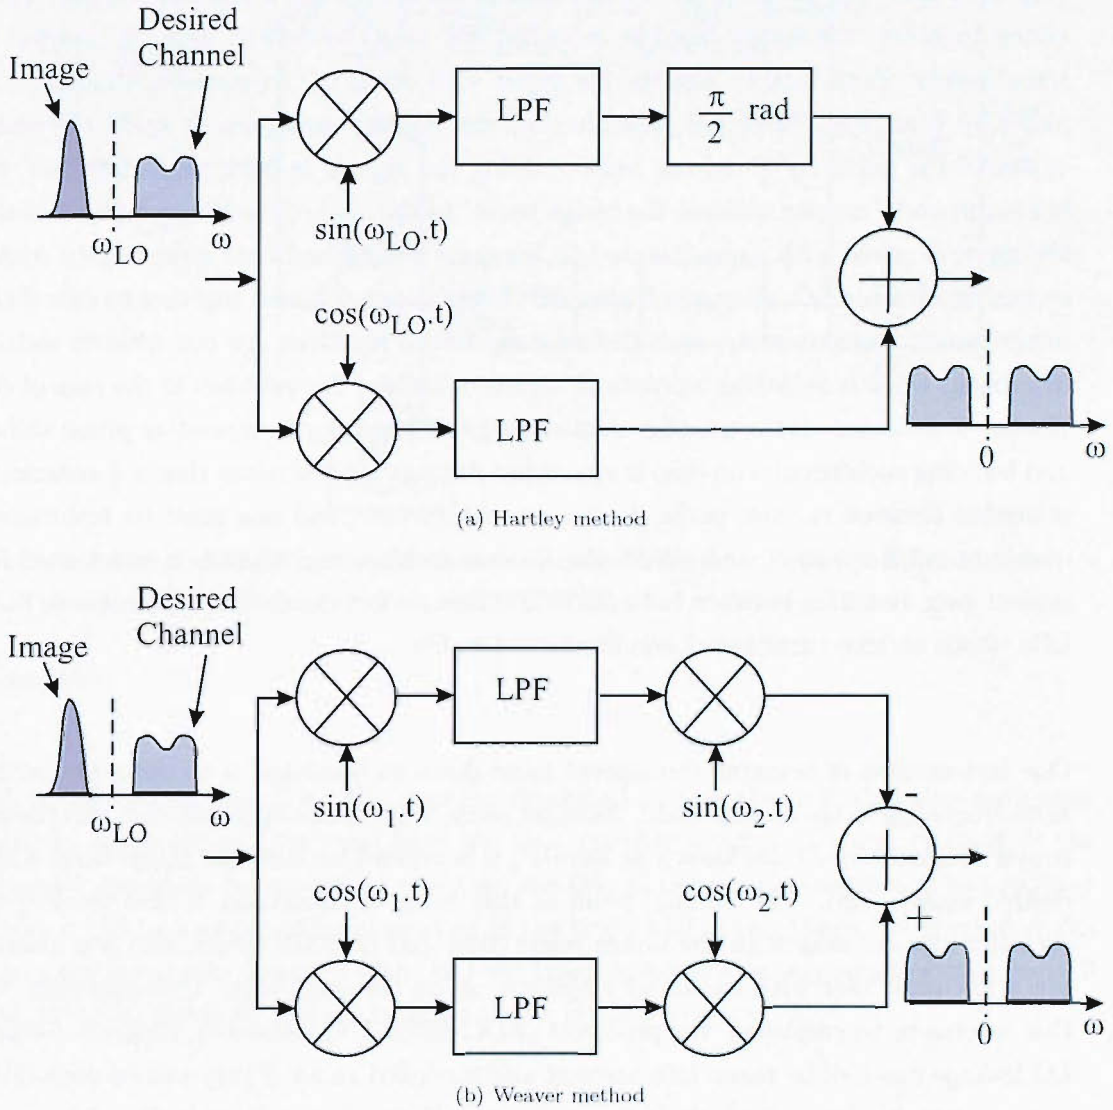

FIGURE 1.2: Two image-reject architectures.

By using a dual-conversion architecture, the trade-off above can be dealt with in two parts. By down converting to a high IF, the problem of image rejection can be tackled, whilst a second mixing stage will down-convert to a second and lower IF, where the channel select problem can be dealt with. Finally, a high-speed high precision data-converter can be used to produce a digital version after which advanced digital signal processing will use algorithms to perform I/Q demodulation.

One other solution to the image-rejection problem is to use a so-called image-reject architecture. Two common arrangements are the Weaver architecture and the Hartley architecture. The strategy with both is to down-convert the desired band into two paths

whereupon further processing will adjust the image component in one path. The aim here is to have the phase of the image band rotated through π radians so that, when added together, the image band is cancelled out and the desired band is boosted in signal power. Both look to mix the RF input with equal LO frequencies, that differ in phase by $\frac{\pi}{2}$ radians. After low pass filtering, the Hartley arrangement shifts the phase of one of the paths by $\frac{\pi}{2}$ radians before adding the signals in both paths together, resulting in an IF output without the image band. In the case of the Weaver architecture, the input is mixed with a quadrature LO, low-pass filtered and then mixed again with a second quadrature LO, after which the pair of signals are summed together to cancel the image band. Unfortunately, such theoretically trivial solutions are not without serious drawbacks when considering a practical implementation. The problem in the case of the Hartley architecture is the need for a constant-gain, frequency-independent phase shifter and building such circuits on chip is somewhat difficult. It is obvious that a $\frac{\pi}{2}$ coherency is needed between the two paths in such an architecture, and this must be maintained over a broad frequency band. With the Weaver architecture, there is a strict need for perfect gain matching between both paths and also perfect quadrature operation in both LOs. Both architectures have been illustrated in Fig. 1.2.

One last method of bringing the desired band down to baseband is to use a LO at the same frequency as the desired band. Straight away, it should be apparent that this direct-conversion technique (also known as zero-IF) is beneficial because the image band is the desired signal band. The 'selling' point of this mode of conversion is that the system architect can do away with the image reject filter and IF SAW filters, and can instead use a low-pass filter with baseband amplifiers all on the same chip. Unfortunately, for this scheme to be employed, the problems of DC offset, I/Q mismatch, flicker noise and LO leakage must all be taken into account and modelled to see if they can be controlled to an acceptable level. A block diagram of a zero-IF receiver is given in Fig. 1.3.

The zero-IF and low-IF methods have become more important as wireless standards move to higher frequency bands and tighter channel spacing. For both down-conversion schemes (which incidentally use quadrature mixing), the local oscillator is at the same frequency as (or close to) the centre of the desired channel or frequency band of interest. However, the VCO usually runs at twice or even four times the centre frequency of the band, depending on the method of quadrature generation. If one assumes that in both cases, the method of quadrature generation relies on the master-slave action of either a pair of latches or flip-flops, then a VCO running at twice the desired local oscillator frequency with a master-slave flip-flop (configured as a divide-by-2 circuit) will show less immunity to phase mismatch in its quadrature outputs (owing to its sensitivity of the mark-space ratio of the input), than a VCO running at four times the local oscillator frequency with a synchronous divide-by-4 circuit generating the required quadrature

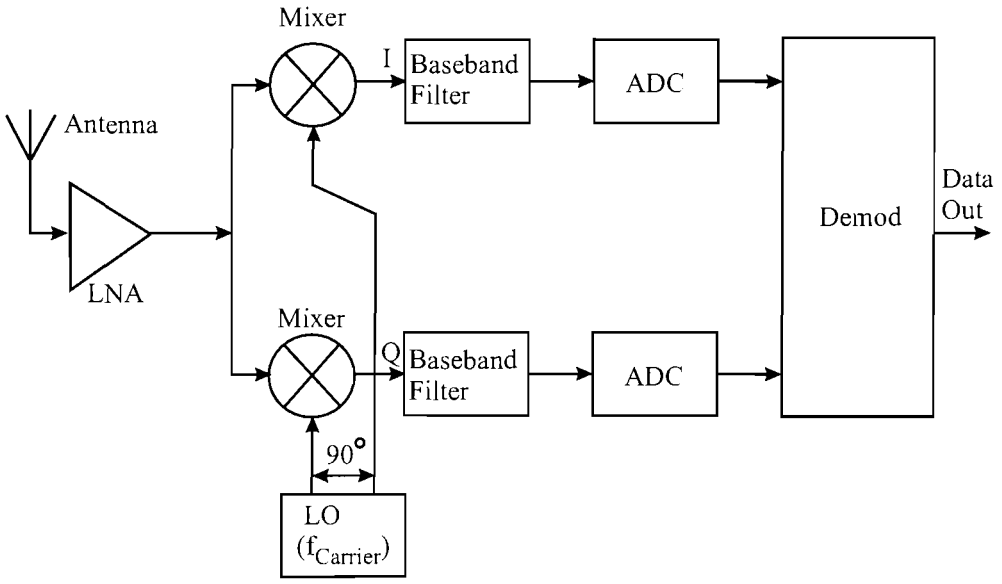

FIGURE 1.3: Block diagram of a direct conversion receiver, tailored for digital wireless transmission.

outputs.

A major driving point for the research described in this thesis is that any frequency divider investigated later must have the least possible number of stages running at the highest operating frequency. In doing so, the overall power consumption of any receiver design will be lowered. As will be clear in the latter half of this thesis, the dual-modulus dividers have been designed with the least possible loading on any signal source, as well as only one divide-by-2 stage running at the input frequency.

To end this section, a couple of examples have been given to show the viability of CMOS silicon technology to compete with more exotic SiGe technologies and bipolar processes in this area of analogue signal processing. One of the first complex gigahertz chips implemented the DCS-1800 standard in 0.25μm bulk CMOS [35]. Using the same phase select architecture (used in the design of the dividers in this thesis) for programmable division, a fully integrated transceiver (minus the antenna, antenna filter, PA, and baseband processor) has shown the capability of this 'digital' process for high performance analogue signal processing. Running on 2V, the circuit consumes 191mW in receive mode, whilst only taking 160mW during the transmit mode. The integrated frequency synthesiser holds a 16 modulus divider, a quadrature VCO and an active-passive combination filter to realise the fourth order type-II PLL. The VCO uses a centre-tapped 3 turn inductor with a Q-factor of 9.

More recently, designers are increasingly showing 5GHz radio frontends in RF CMOS

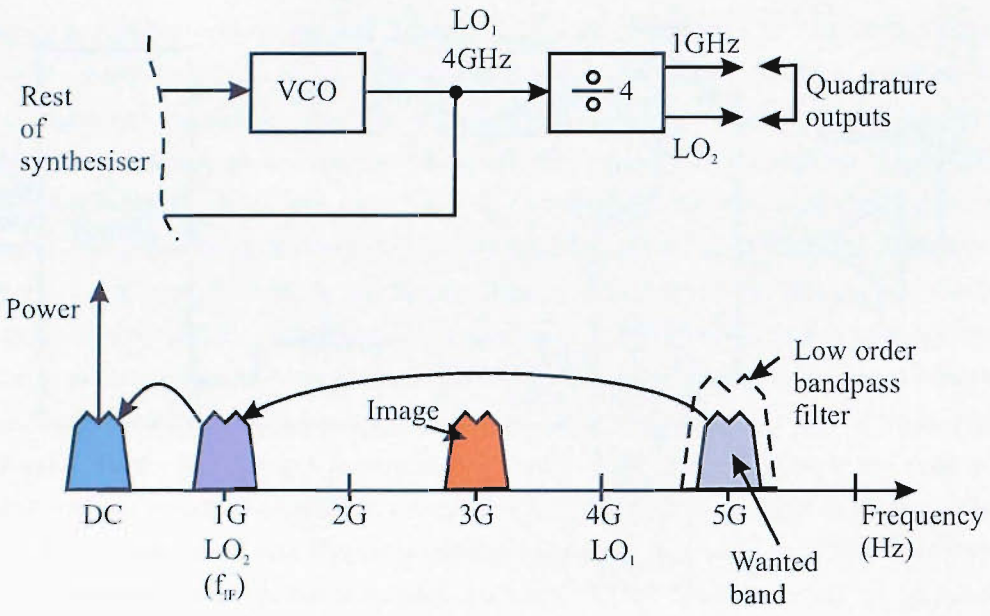

FIGURE 1.4: Frequency plan of the down-conversion method in an example 5GHz wireless LAN receiver [36].

processes with a particular example given in [36]. Concentrating on the front-end and not the digital baseband (which in this case is a separate chip), a dual-conversion architecture is employed in both the transmit and receive paths. In the case of the receive chain, a 4GHz LO down-converts the filtered 5GHz band to a 1GHz IF, before finally downconverting to baseband. This frequency plan (see Fig. 1.4)is chosen as there is only need for one frequency synthesiser designed to give a 4GHz output. The 1GHz local oscillator is derived by tapping the output of the 4GHz oscillator to drive a divide-by-4 prescaler, with quadrature LO. The transmitter works in the reverse manner with a power amplifier integrated on chip too, delivering an output power of 22dBm. The transmitter chain has been characterised to dissipate 790mW whilst the receiver consumes 250mW, with the synthesiser taking another 180mW (2.5V with 3.3V I/O in 0.25μm bulk CMOS).

1.2 Bulk CMOS technology

Over the past ten years, aggressive scaling of MOS devices and interconnects has delivered products with circuit elements running beyond 5GHz. Advancements in both processing and lithography has enabled scaling of their gate lengths to sub-100nm levels. Phase-shift mask technology, choices in reticle material and available light sources yield these feature sizes both in transistors and metallisation.

Table 1.1 shows the projections made in a 2001 ITRS road map, on the whereabouts

	2001	2003	2005	2007	2010	2013	2016
Technology node (nm)	130	100	80	65	45	32	22

TABLE 1.1: Timeline of technology nodes [1].

of CMOS technology in those years. The term 'Technology node' is almost synonymous with the 'DRAM half pitch' (which incidentally refers to the closest separation between two metal lines in adjacent DRAM cells.) The actual 'minimum' gate length of a MOSFET in every node is smaller than this value. However, as time has progressed, amendments to this table have resulted in the following roadmap, citing the technology of the past, present and future (Table 1.2.)

	1997	1999	2001	2003	2006
Technology node (nm)	250	180	150	130	100

TABLE 1.2: A revised table of the transistor gate lengths [32].

In terms of circuits produced, CMOS digital circuits reap the benefits of scaling. With the arrival of the 3GHz processor in production, ultra large scale integration demonstrates the complexity and packing density achieved with today's processing tools. With memory sections exploiting the small footprints for each transistor as well as clever layout, a myriad of functions can be retained on the same piece of silicon without the need to couple many dice externally. Even in mixed signal IC's, the CMOS digital section can easily be distinguished by the regularity and granularity of the regions unlike analogue sections such as those found in today's transceiver chips.

At present, 3.2GHz processors are in production (Intel Corp), fabricated using a 130nm process [26]. Advanced Micro Devices (AMD) has given a roadmap illustrating the technology of future processor families and architectures, with a migration from bulk 130nm CMOS process to SOI 130nm CMOS process [5]. This strategy is used to step away from the reliance on the progress of lithography, until it is available and mature. The SOI gives some performance gain in power and speed, though Intel Corp have contested any gains associated with the technology at such fine levels of geometry. Another key indicator of the adherence to the semiconductor roadmap, is the density in memory, particularly DRAM. Companies such as Hynix, Micron and Samsung have 1GBit DRAM chips (not modules) at the production stage, with some using 100nm technology on 300mm diameter wafers.

In recent times, the importance (and reward) in realising analogue functions in plain

CMOS optimised for digital applications has become greater and greater. Resisting the temptation to move to a slightly more expensive technology such as bipolar and BiCMOS can pay dividends in terms of power and cost when designing a mixed-signal IC. There are a number of products on the market offering analogue and mixed signal processing elements fabricated in CMOS technology ranging from filters and op-amps to data-converters and radio transceivers. Such products not only incorporate the basic MOS transistors, but also passives such as poly resistors, metal-metal capacitors and spiral inductors. Unfortunately, it is the requirement for such elements that results in the lag of analogue/mixed signal products because characterisation and extensive device (active and passive) modelling must be completed before any schematic can be simulated. Also, with high frequency design, the complex behaviour of the substrate and interconnects has frequently called for a number of fab runs with standalone elements and even design iterations. One last point to note is that sub-micron MOS devices tend to have poor signal performance such as high noise floors (main reason being the $\frac{1}{f}$ noise, whilst the thermal noise is comparable) and low current drives compared with bipolar technology. This is one reason for the lack of investment in high specification analogue MOS circuit design. Nevertheless, this has not stopped the emergence of high performance designs such as a 10GHz PLL [16] and a transceiver IC for a 5GHz wireless LAN [36] appearing in professional publications.

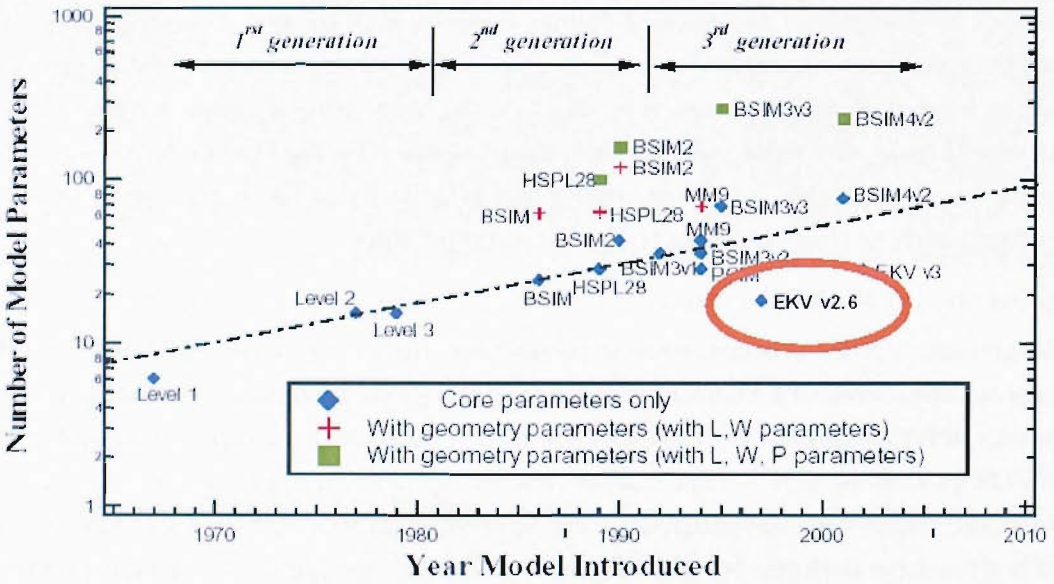

FIGURE 1.5: Graph showing the number of technology parameters vs the year a model or its revision was introduced [14].

As with any high performance IC, the emphasis placed on accurate, fast, robust active component models has never been greater. With each fabrication run costing many tens of thousands of US dollars (excluding the mask set), mistakes in peformance cannot be tolerated, and this is the reason behind activity within the industry to address the

needs of non-quasistatic effects for example, albeit a little late in time. Models such as BSIM [39], EKV [13] and MOS Model 9 [44] are regularly referenced in the circuit design and modelling communities, but with such fine geometry, the ability to predict even DC transconductance characteristics with models such as BSIM [39] can result in notable discreancies when choosing transistor dimensions considerably off those used for parameter extraction. A quick and dirty solution is to 'parameter bin,' but at the expense of a hideous number of parameters and then the model becomes distantly related to any physical model.

The graph shown in Fig. 1.5 captures the evolution of models, in particular those for MOSFETs. The so-called first generation models were the Level 1, 2 and 3 (SPICE) models. The legend in this graph makes reference to 'geometry parameters' and these are included in some models to reduce the error between a simulated value and an actual measurement, even at large channel widths. One such parameter is geometry dependent mobility degradation. However as the graph shows, there is a marked rise in the number of model deck parameters, as well as the device model becoming less physical.

The first model, occasionally referred to as the 'Shichman-Hodges' model, had the channel length modulation as the only geometry dependent effect. Other features were the gradual channel approximation and square-law for the saturated drain current. Unfortunately, no model for the sub-threshold current existed.

The second model makes up for the lack of sub-threshold conduction in the first, as well as accounting for velocity saturation, vertical field mobility degradation, and describing in more detail the depletion and threshold regions of their operation. The third model is more efficient mathematically and takes on a semi-empirical approach to modelling.

For the second generation of models, BSIM 1 [9] and 2 [47] were two examples at the time. Verison 1 was more of a digital model, being a mathematically tailored to break its operation into distinct regions, without much care at the transition point. Threshold voltage and mobility modelling were key to the model as well as a detailed sub-threshold current expression. In version 2, the addition of output conductance as well as modifications to the threshold voltage and mobility models made this version a more 'analogue' model.

The third and current generation, includes BSIM3, EKV and MOS model 9. The first incorporates a lot of smoothing functions for well-behaved model equations and their derivatives. MOS model 9 is a Philips Research model with compact, clean and robust modelling, again using smoothing functions, but has the advantage over BSIM 3 that there are few parameters in the model deck. The EKV model is geared towards

low-power low current analogue circuit design. Its developer claims to have used a new and clearer method for analytical modelling as well as drawing in some of the other 3rd generation model features. By moving the reference point from the source node to the substrate, such a model is useful where symmetrical device behaviour is called upon.

In the context of RF design, particularly for transceivers, it is important for the models to incorporate more of the parasitic effects observed with laying out the devices, and to model drain/source implant resistance and gate resistance in order to account for them at the circuit simulation stage, especially during the characterisation of phase noise performance of MOS-based oscillators. The MOS model 9 with an RF extension is one such model and has been used during the design of the bulk CMOS dividers reported in this thesis.

1.3 Silicon on Insulator CMOS technology (SOI)

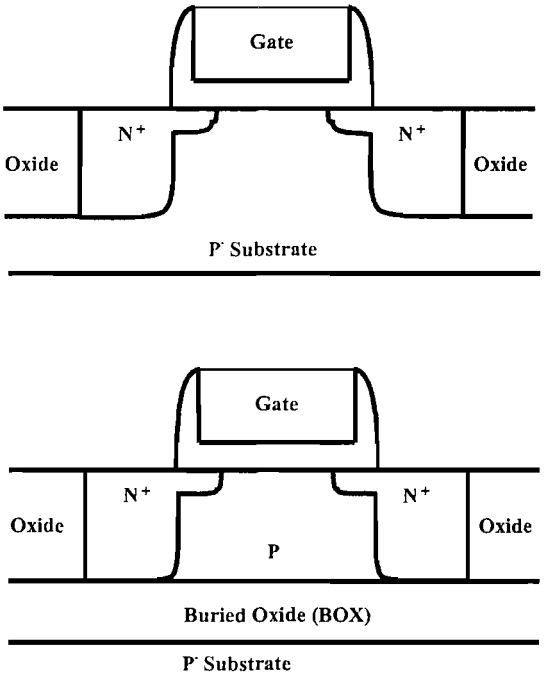

FIGURE 1.6: Cross sections of an NMOS transistor on Bulk (top) and SOI (bottom) wafers.

1.3.1 Devices

SOI technology has been used in many special applications such as radiation-hardened and high-voltage integrated circuits [12], owing much to the buried oxide. With considerable effort in process technology and wafer manufacturing, SOI is now starting to

become a serious option for applications demanding high speed and low power characteristics. One of the designs presented in this report is targetted for fabrication on SOI wafers. These wafers give rise to CMOS devices that no longer require wells for isolation between the different types due to the underlying oxide layer [48]. The lack of wells prevents 'latch-up' (improves short channel effects and soft error immunity) and can eliminate the need for well contacts, therefore resulting in a higher packing density as well as the bonus of fewer masks (less costlier mask sets)[1]. The oxide layer also serves mixed-mode circuits by minimising the noise from digital partitions that may affect any analogue partitions.

The isolation also serves to maintain the threshold voltage of individual devices between the power supply rails. Unlike bulk CMOS designs, placing say an NMOS device immediately on top of another NMOS device, keeps their threshold voltages the same (assuming the bodies are tied to their respective source terminals.)

As the electronics industry creates mobile devices, the need for low power consumption devices becomes more important. With fabrication on SOI wafers, the source and drain junctions extend all the way to the back oxide layer (Fig. 1.6), giving a reduced parasitic capacitance and therefore faster device operation. Lower power is partly attributed to the lower static leakage currents through the lower junction areas, as well as lower CV^2F terms (where C, V and F are the gate load capacitance, the difference between high and low logic levels on those capacitance terminals, and the frequency of operation, respectively.)

Saying this, the CMOS devices on bulk Si wafers are simlilar to those fabricated on SOI wafers, such as PMOS devices still requiring wells (assuming a p^- doped top substrate.) Ignoring body contacts, layout of transistors and their interconnects remain identical between both technologies.

1.4 Wafer fabrication

Before any device fabrication can take place, a wafer must be acquired. After slicing a cylindrical ingot, wafers for bulk CMOS processing are polished and a thin single crystal silicon layer is grown epitaxially. However, SOI requires a buried oxide layer within the wafer. This can be achieved by:

- implanting oxygen atoms into the silicon wafers at high energies (*SIMOX*),

[1] This is not true where MOS body tied devices are incorporated

- growing oxide layers on two wafers, implanting one with Hydrogen ions from the opposite wafer face to the oxide, bonding the wafers together, before using a *SmartCUT®* process to split the wafers at the Hydrogen implantation depth(as licensed by SOITEC in their *UNIBOND®* product family),
- two wafers with oxidised top layers are bonded, after which one side is polished down to the required silicon film thickness (*Etchback*).

Once a buried oxide is established, the top silicon layer must be polished and possibly annealed, before gate stacks, implant regions and interconnect layers can be formed.

1.4.1 Device configuration

SOI MOSFETs appear in three configurations: *partially depleted floating body* (PDFB), *partially depleted body tied* (PDBT) and *fully depleted* (FD) [34].

The partially depleted transistor (both floating body and body tied) is produced using thick silicon films (typically >150nm) where the depletion region of a strongly inverted MOSFET fails to reach the back oxide, eliminating any coupling between the back- and front gate oxides. This type of device has the advantage that the V_T is tailored through the choice of channel doping profile (see FD device). The main disadvantage lies with the PDFB configuration which is susceptible to *floating body* effects.

The FD devices usually have a thinner silicon film (or lower doping concentration) where the whole body region is depleted and the back gate influences the operation of the device. This configuration of device allows the threshold voltage to be reduced without unduly increasing any off-state leakage currents (crucial to low power devices). FD MOSFETs are considered to possess improved short channel behaviour whilst eliminating floating body behaviour, but reduced source/drain junction depth that delivers the improved short channel behaviour comes at the price of source/drain series resistance. Also, the requirement for the silicon film to be thin poses a problem in controlling the V_T of the device due to its sensitivity to process and thickness variations.

SOI technology also comes with other problems. For example, in a PDSOI device, the lack of well contacts implies the body region to be unconnected or 'floating,' where any charge injected into the body will cause fluctuations in the body potential. Also, the back oxide plays the role of a thermal insulator, resulting in localised heating (*self-heating*). Parameters such as mobility μ and V_T are affected, forcing device characteristics to deviate from their ambient temperature operation on a dynamic basis. Digital CMOS

circuits tend to have low power dissipation per device compared with analogue circuits which constantly dissipate power due to their biasing.

1.4.2 Body contacts

Unlike bulk CMOS designs, transistors fabricated on SOI wafers are not electrically connected under the field oxide that isolates adjacent transistors. Charge injected and removed from this 'body' region is usually transported via a body contact, though this does not have to be the case (see subsection 1.4.3 - '*Floating bodies*'). With 'Body tied' configurations, it is important to keep the channel-to-body contact resistance as small as possible, thus allowing fast transport of charge into and out of the isolated body. This is achieved by placing body contacts at one or both gate ends of a device. The design rules, used for the dual-modulus 64/65 divider, dictate how close and how many of these contacts should exist, as well as the maximum gate width for a specific gate length transistor which will explain why in that design, gate widths have been purposely capped to 7μm and 15μm for $< 0.7\mu$m and $\geq 0.7\mu$m gate lengths respectively. Where transistors have larger widths than these, the required transistors are broken into smaller transistors to make a 'finger' structure.

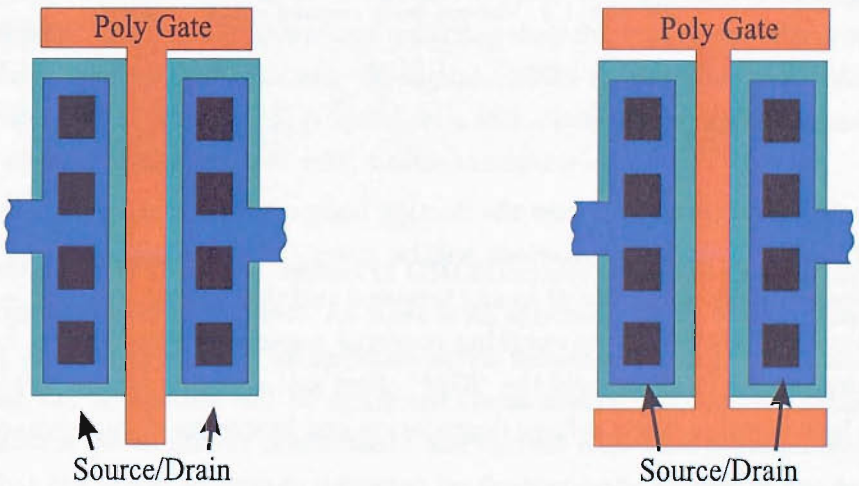

(a) T-Gate (b) H-Gate

FIGURE 1.7: Two examples of MOS transistor layouts.

As there are H-gate and T-gate layout arrangements (Fig. 1.7), there also exist various styles of body contacts. They can be:

- standalone - contact that exists outside the active region and is electrically con-

nected to a node by the first metal line (Fig. 1.8(a)),

- embedded - a highly doped region of an opposite type to the source implant regions, which is electrically connected by cutting into the source under a contact (Fig. 1.8(b)),
- gate end tie-downs - like the previous version but is instead placed at either or both of the gate ends, with the opposite type implant lying under the gate end (Fig. 1.8(c)).

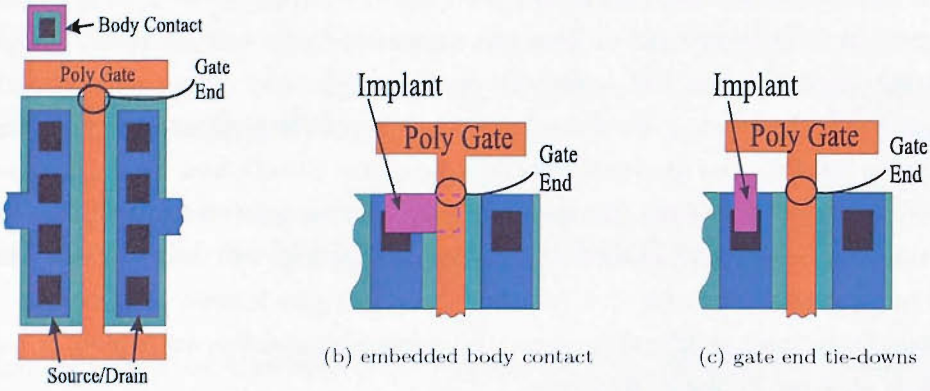

(a) standalone body contact

(b) embedded body contact

(c) gate end tie-downs

FIGURE 1.8: Various body contact configurations.

1.4.3 Floating bodies

As the dual modulus design uses the floating body configuration in its high speed stage, a brief look at a device's behaviour will be given. Floating body PDSOI devices have their body regions unconnected to any terminal and thus are allowed to float electrically. Unfortunately, there are two resulting concerns associated with floating bodies in partially depleted devices, namely the *'Kink' effect* and *parasitic bipolar effect*. Other side effects include breakdown voltage degradation and hysteresis during dynamic operation [42].

When the dc I-V characteristic of such a configuration is examined (see Fig. 1.9), a pronounced kink is observed for low gate biases whereby an increasing drain voltage would produce a larger than anticipated drain current. In the case of an nMOS transistor in its saturation mode, electrons will collide with the silicon lattice if there are high drain fields (impact ionisation). This generates many electron-hole (e-h) pairs, whose electrons travel into the drain implant and the holes migrate to the point of lowest potential. With floating bodies, these holes build up (though some do leak away) and raise the overall potential of this body. With reference to the 'body effect' or 'back bias' effect

observed in bulk CMOS, the drain conductance increases, which in turn causes more e-h pairs to be created and so on. There is, however, a point where the 'body-source' diode becomes forward biased, injecting electrons from the source into the body, and recombining with the stray holes. The I-V characteristic then behaves in an expected manner to that of short channel transistors (channel length modulation or drain induced barrier lowering). One major implication of having devices exhibiting a kink effect is that the output resistance of a device can drop which means the internal gain is also reduced.

Increasing the drain bias further would highlight the second effect, parasitic bipolar action. The silicon film under the gate stack can be thought of as a floating base terminal and the current into this base is again a result of impact ionisation. This parasitic device is turned on by the buildup of charge in the floating body. By using source/drain extensions, the effective collector/emitter areas are reduced, thus reducing the current gain of this parasitic device.

The characteristics described above prove useful when designing analogue circuits in a PDSOI CMOS process. However, this technology has been cited as a contender for high speed IC applications. Therefore, it is crucial to understand the effects of switching in actual circuits, and to note any key observations that may be attributed to floating-body action [42][33][46]. Authors of publications reporting their findings on the effects of floating bodies have concentrated primarily on digital CMOS topology circuits. Although our SOI divider design (Chapter 5) is based on a SCL circuit topology, it is possible to apply some of their findings, albeit with a little manipulation.

An immediate observation in the context of CMOS circuits is that their transient analyses show signals to operate rail-rail. As there is an inevitable delay between any rapid change on a device's input and its response at its output terminals, it is feasible to conclude that the transistors will be subjected (momentarily) to parasitic bipolar effects, and hence a rise in output conductance and current [42]. This previous study has concluded that the speed advantage delivered by floating-body devices is due to three characteristics:

- less source/drain junction capacitance (regardless of floating-body)
- dynamic 'kink' effect
- dynamic 'V_t' due to capacitive coupling.

When the body is tied to the source terminal, any speed advantage observed is attributed to the reduced junction capacitances at the drain and source implants, whilst

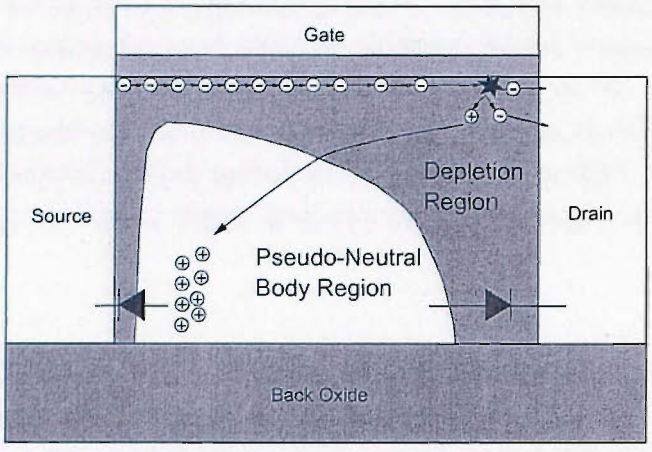

(a) Mechanism of the 'kink' effect,

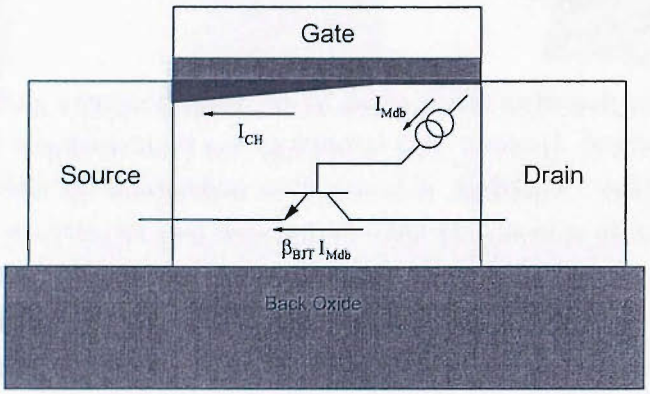

(b) schematic of the parasitic BJT,

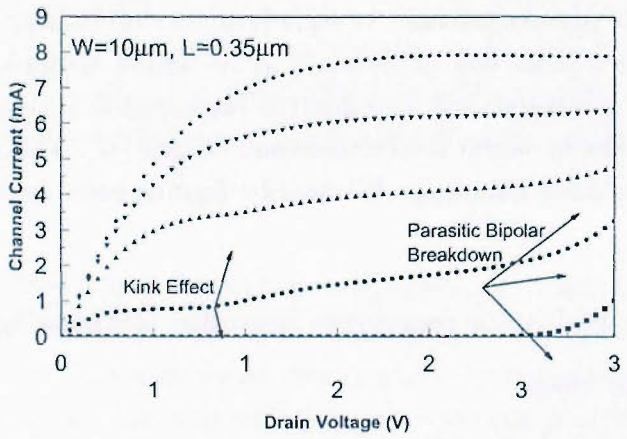

(c) I_{ds} versus V_{ds} (measured) during floating body operation.

FIGURE 1.9: Floating body effects.

the dynamic 'kink' effect, due to body charging, is a small contributor to any increase in speed. Dynamic kink behaviour is similar to its static counterpart, except the effect is now a function of switching time and its frequency (alternatively, the body region's carrier generation and recombination time). The authors responsible for the results [33] acknowledged their difficulty in gauging how much of an effect the dynamic 'kink' had on the speed and kept the initial DC body voltage constant throughout their entire simulation.

Capacitive coupling is responsible for the third mechanism and can be broken into two sub-groups: current overshoot and dynamic loading. The first of these improves the switching speed by increasing the forward V_{BS} of the drive transistor in a CMOS inverter (due to gate-body coupling), giving a higher switching current. The second case decreases the forward V_{BS} and/or increases the reverse V_{BS} of the load transistor, with the help of the drain-body coupling. When observing the transient behaviour, the body-source voltage follows the gate voltage due to the gate-body coupling. When the inversion layer forms, the body-source voltage will follow the drain voltage by means of the drain-body coupling.

As will be seen, there is a notable difference between the CMOS-style circuit and the SCL circuits, both in their circuit topology and the electrical behaviour. Nevertheless, the above account provides suggestions as to what is happening in the transistor bodies of those devices that are switching in the forthcoming frequency dividers. An effort is made in Chapter 5 to understand why the two highest speed divider circuits in an asynchronous chain may operate faster when their transistors have floating bodies.

1.4.4 Self-heating

The term above refers to the disadvantage of having a buried oxide for electrical isolation, whereby localised heat cannot dissipate as easily into the back side of the wafer. If a small device is operated with a low frequency signal, the silicon film can heat up if current density is high, degrading the I-V characteristics of the circuit. Self-heating is an important issue that must be accounted for when designing in sub-micron technology. To model a CMOS rail to rail LSI analogue circuit without this effect would result in gross discrepancies between silicon and simulation. In a high speed NMOS logic design, the problem is not as pronounced (assuming 50% or less duty cycle) because the device is off for a comparable amount of time as it is on (after a certain frequency, the designer should look at RMS current, as would be done in high frequency interconnect design). In the case of CMOS logic, power dissipation occurs during the transition between logic levels. Hence, a design operated at low frequency should dissipate less power than a circuit operating at a higher rate over the same length of measurement time.

1.4.5 Circuit simulator models

To go about designing circuits for ICs, circuit simulators become an essential tool. However, models are needed to describe the transistor behaviour when operating in certain conditions and configurations. For SOI, models such as BSIMSOI [38] and UFSOI [40] exist. However, specifically for the SOI dual-modulus divider, the 'Southampton Thermal Analogue' model or *STAG* [34], together with a Honeywell 0.8μm and 0.35μm parameter set, have been used for the design work. This model uses a *surface-potential* technique as opposed to traditional *threshold voltage* based techniques, relying on a single expression that is continuous for all regions of operation, making this very suitable for any analogue designer. Modelling a characteristic with more than one expression can lead to 'jumps' in the simulated performance and thus anomalous results.

As designers move further into the gigahertz range, the argument for better RF models becomes stronger. There is a dearth of robust compact RF models [30] for both analogue and digital design in any wafer technology. The importance of how designs operating in such a frequency range can affect system characteristics such as *bit error rate (BER)* and propagation path cannot be stressed enough. Models need to reflect the RF performance in foundry processes. Simulations of low frequency signal must be done in a high frequency environment. Interconnects are also ignored at schematic level, which is disastrous at such frequencies. Layout driven designs could possibly help achieve a closer estimate of the final realisation. Even though SOI aims to minimise it, substrate coupling is never eradicated and needs to be attended to, possibly with a layout driven approach.

1.5 Review of CMOS frequency synthesisers

In the majority of designs reported, a stable and accurate local oscillator is needed to down-convert a channel(s) to a lower intermediate frequency or even baseband, as well as up-converting baseband data to a higher intermediate or even transmit frequency. The emphasis on 'stable' and 'accurate' stems from the need to select the whole of a desired channel. One way of achieving such a function is through the use of frequency synthesisers. They must meet certain specifications such as low phase noise and low power consumption. The stringency of these requirements is dependent on the type of application, for example, a single chip satellite tuner will call for a low phase noise oscillator compared with wireless LAN oscillators. However, power consumption is not a design priority in satellite tuner designs as it is usually mains operated, unlike portable wireless LAN devices.

A frequency synthesiser can be either incoherent, coherent direct, coherent indirect,

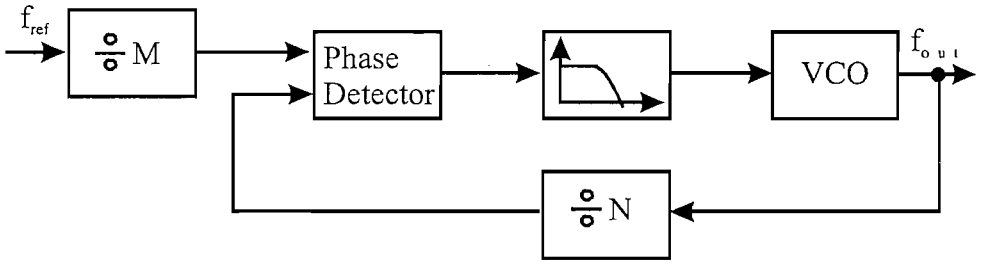

FIGURE 1.10: Block diagram of a basic frequency synthesiser based on a PLL.

coherent direct digital [50]. For the rest of this section, only the coherent indirect synthesis and more specifically, the phase-locked loop with a programmable divider will be described. This type of loop contains a phase detector/comparator (generates a slowly varying ac signal when not in frequency lock, and then a steady dc output whose amplitude and polarity reflect the phase difference between its two inputs), a loop filter (that removes higher order products resulting from the phase detector), a VCO (generating the desired frequency) and a programmable frequency divider (used to obtain a frequency close to the reference frequency). A charge pump is usually inserted between the phase detector and loop filter in order to deliver charge to the loop filter, lessening the burden on the phase detector. Another reason for its inclusion is that the static phase settles to zero, giving the characteristics of a 'Type-II' PLL (PLL with two integrators in the loop - the VCO being one), without having to actually add one [18]. In the context of high speed applications, the loop is configured as a frequency multiplier, so a very stable reference exists at low frequency, but the VCO is driven to run at a much higher frequency. The multiplication results because of the frequency divider in the loop and the output frequency is given by the following simple expression:

$$f_{out} = N.\frac{f_{ref}}{M}. \tag{1.1}$$

where N is the division ratio within the loop, M refers to any prescaler present before the input to the phase detector, f_{ref} is the frequency of the source at the input and f_{out} is the frequency of the output after the VCO. Six key characteristics exist for frequency synthesisers:

- phase noise
- frequency stability
- output frequency range
- frequency increment
- switching time
- spurious output

The specifications for these requirements determine the cost, power consumption, and size of an IC synthesiser. Phase noise is a figure of merit that demands attention, especially in applications that require tight channel spacing. Having the power from the carrier leak into adjacent channels can be detrimental to the performance of the transceiver. As the channel spacing requirement is relaxed, the 'skirts' either side of the carrier may be more prominent depending upon the application. Figure 1.11(a) shows the prinicpal behind 'reciprocal mixing', where signal power from the unwanted adjacent band can interfere with the desired band. Figure 1.11(b) highlights the effect of phase noise 'masking' an adjacent signal in a desired group of frequencies.

1.5.1 Phase detector

The phase detector, as its name implies, quantifies the difference in phase between its two inputs. With signals of sine wave characteristic, an analogue multiplier, such as the Gilbert cell (figure 1.12), is a very common implementation. After multiplication, a DC term and double-frequency term emerge, with the latter being removed by a subsequent low pass filter. Such a cell gives a zero average output when the input phase difference is $\frac{\pi}{2}$ radians (regardless of the sense of phase difference.) This aspect forms part of the specification of the loop arrangement. Also, the lock range (where the loop is frequency locked, but not necessarily phase locked) is between 0 rad and π radians phase difference. In the case where one or both inputs are square waves, a commutating multiplier can be the answer. The important difference between this and the last phase detector is that the square wave at the input(s) has spectral components other than the fundamental. This implies that there is a possibility the loop will lock on to one of the harmonics. This can be useful in producing super-harmonic locked oscillators (see later) to produce a 'divided' output. However, when undesirable, the VCO tuning range must be restricted to prevent such locking. The spectrum of a square wave does fall off with a $\frac{1}{f}$ behaviour and thus becomes harder to lock to higher harmonics. Saying this, it must be attended to in practical loops based on such a detector (by means of aggressive filtering).

In many transceiver designs, digital logic blocks are frequently used for such functions. The exclusive-OR gate (talked about in the previous paragraph as a Gilbert cell, but driven hard with large signals) is such an example. When driven with square wave signals, the output is related to the input by an XOR relationship, with the only difference being the DC level on the output.

The above multiplier-based detectors are in fact quadrature detectors (they indicate a phase difference of $\frac{\pi}{2}$radians) and there are instances where a need for zero phase difference locking exists. This can be achieved using a sequential phase detector, providing a

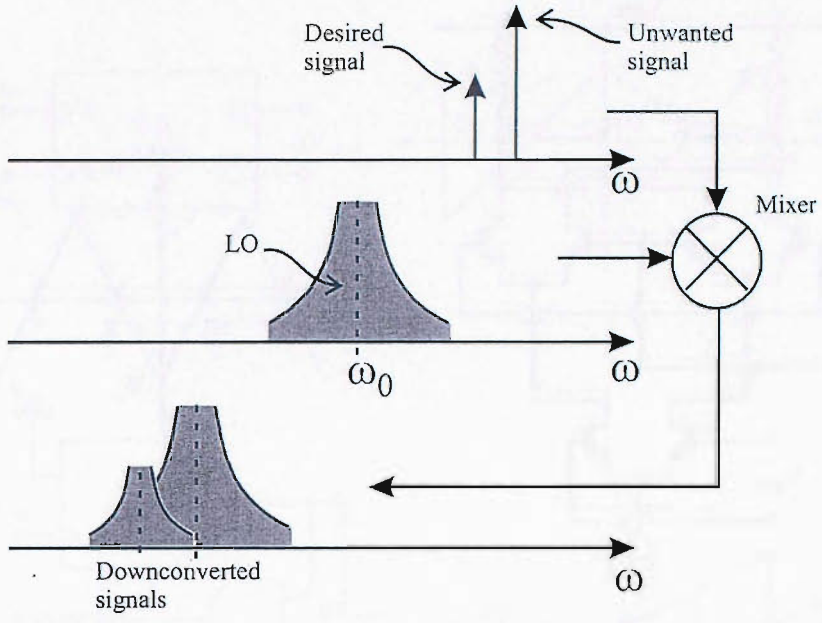

(a) Reciprocal mixing

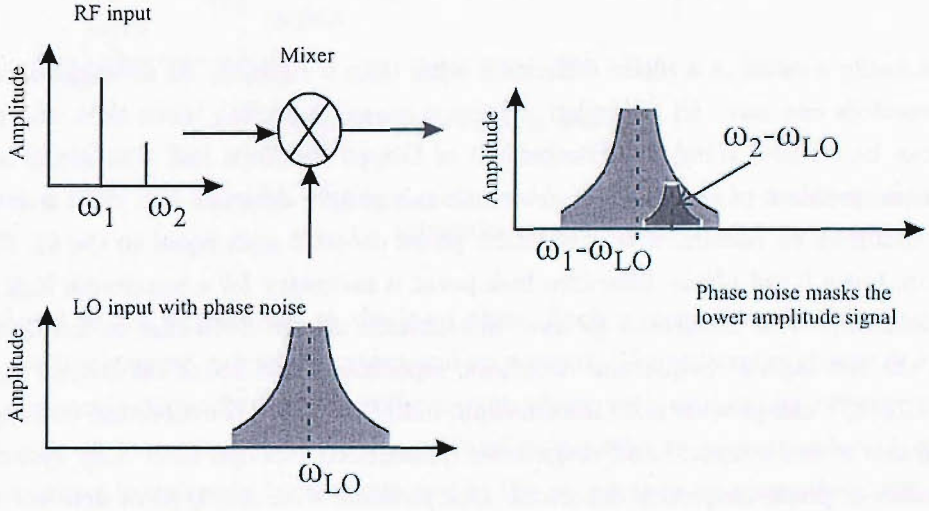

(b) Phase noise can mask the adjacent signal when needed

FIGURE 1.11: Two problems of phase noise on the local oscillator during down-coversion.

0 rad (or even π radians) phase difference condition in lock, with some having outputs that are proportional to differences greater than 2π radians. The drawbacks associated with such detectors is their sensitivity to transitions and add a sampling operation into the loop [49]. In its simplest form, the SR flip-flop (Fig. 1.13(a)) can implement such a function with the output showing a 'sawtooth' characteristic as a function of the input phase difference over many multiples of 2π radians cycles. One problem (depending on the topology,) is that speed difference can result in static phase error with the average

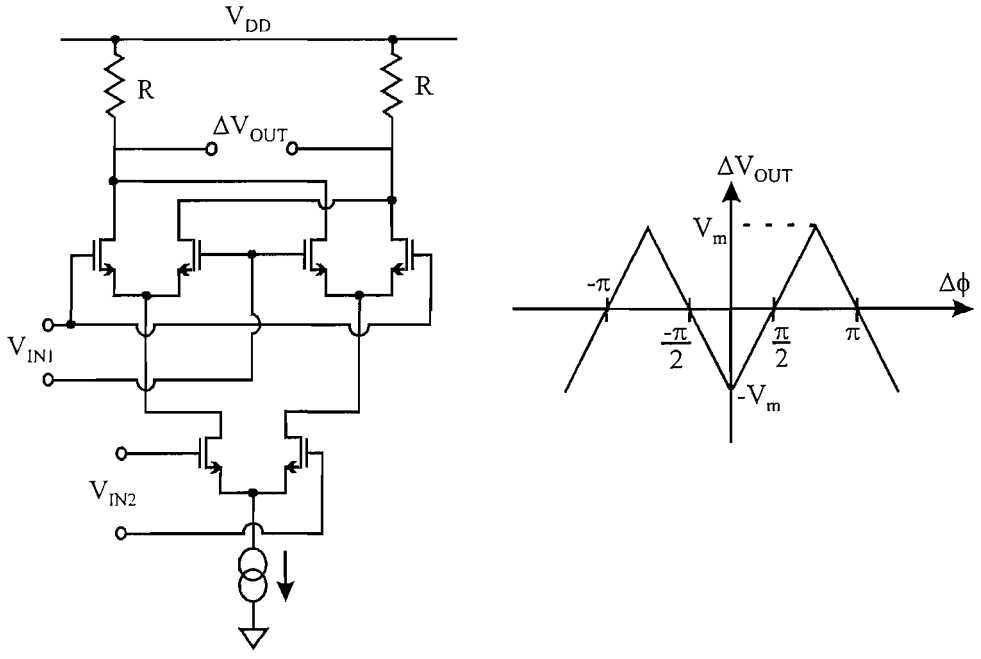

FIGURE 1.12: Example of an analogue multiplier, with the specific schematic of a Gilbert cell.

output being a result of a phase difference other than π radians. As mentioned, sequential detectors can have an extended detection range, spanning more than one period. This can be realised using an arrangement of D-type flip-flops and this circuit also removes the problem of static phase error. Such a simple detector can yield a detection range spanning 4π radians with a constant phase detector gain equal to the SR flip-flop solution, but a 0 rad phase difference lock point is necessary for a maximum lock range.

Occasionally, it is important to have information on the difference in frequency between the two inputs. Sequential detectors, especially those based on D-type flip-flops (Fig. 1.13(b)), can provide such information, indicating both the direction (through signalling one of two outputs) and magnitude (manifested through their duty cycles), and are known as phase-frequency detectors. One problem with this type of detector is that near the zero phase error mark, the gain may differ from what is actually expected as the up-down signal pair are 'impulses' around that lock point. Also, there is a sensitivity to missing edges, and the repercussions are to correct this error. The former problem can be solved with the aid of a DC offset (thus a static phase error), greatly suppressing the nonlinearities.

1.5.2 Loop filter

Following the phase detector is a loop filter, which serves to remove the double-frequency product term. Frequently, a charge-pump sits between these blocks and delivers a charge

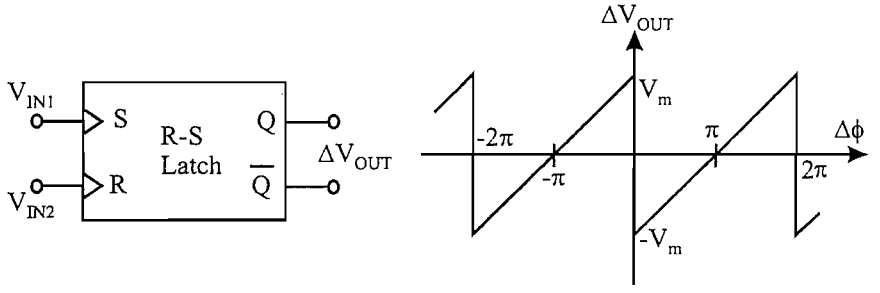

(a) RS Latch based phase detector

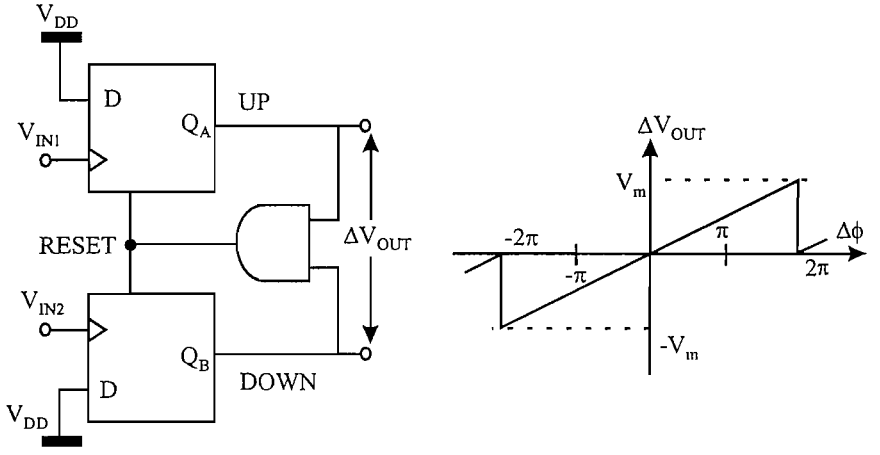

(b) D-type phase frequency detector

FIGURE 1.13: Two examples of the phase detectors and their respective transfer characteristics.

proportional to the phase error to the loop filter. Such a component has three output states: positive current, negative current and no current. The amount of charge delivered can be minuscule depending on the difference in phase with an impulse often resulting close to lock. A loop filter converts this train of pulses into an average value. A simple implementation has already been mentioned in the paragraphs on phase detectors using D-type flip-flops, and it is here that a charge-pump is often integrated with the phase detector.

The loop filter is the component that sets the order of the loop and is thus a key factor associated with the loop dynamics. Technically, there is no need for an explicit filter. In such a case, the PLL would be classed as a first order loop with a low-pass transfer characteristic (taking the output before the entering the VCO.) After linearising the loop, as pictured in Fig. 1.15, it is found that the loop bandwidth is dependent on the phase detector gain and the so-called VCO 'gain', which has dimensions radians per second per volt (and maybe some other gain stages around the loop.) Unfortunately, the output of a multiplier phase detector contains sum- and difference frequency components, both of

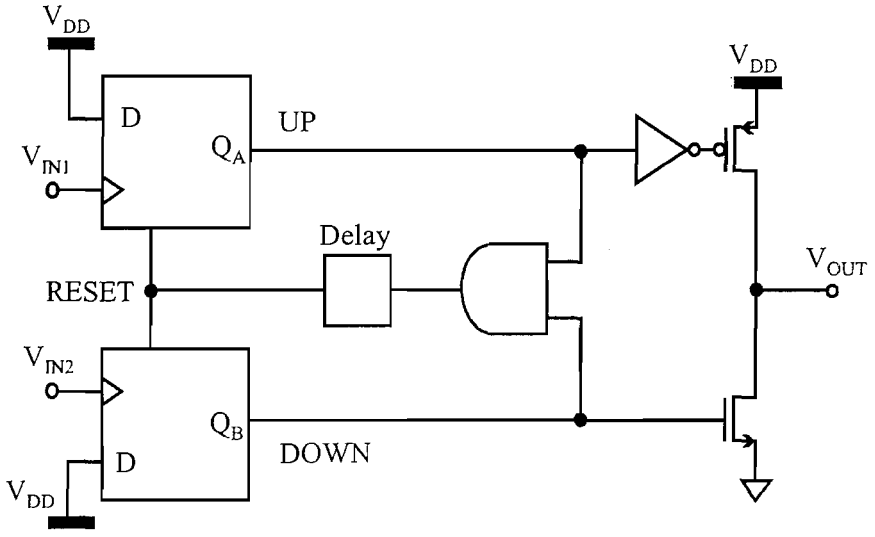

FIGURE 1.14: Charge Pump attached to a phase frequency detector.

which are fed directly to the input of the VCO and thus run the risk of locking to a higher frequency if within the VCO tuning range. Another problem is the out-of-band signals present on the input reference frequency, that are also passed to the VCO, without any attenuation.

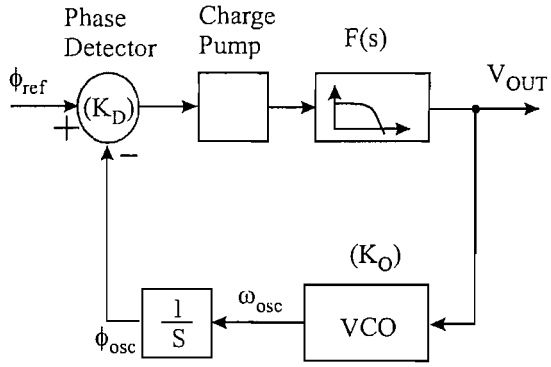

FIGURE 1.15: Linearised model of a PLL.

Taking a closer look at various configurations (with the help of Fig. 1.15,) no loop filter would yield the following transfer function:

$$\frac{V_{OUT}}{\omega_{ref}} = \frac{K_O \cdot K_D}{s + K_O \cdot K_D} \cdot \frac{1}{K_O}. \tag{1.2}$$

The closed-loop transfer function is shown in Fig. 1.16, where K_V is in fact $K_O \cdot K_D$. This configuration is unconditionally stable, with there being a $\frac{\pi}{2}$ rad phase shift in the loop, for frequencies greater than $K_V \cdot rad \cdot s^{-1}$.

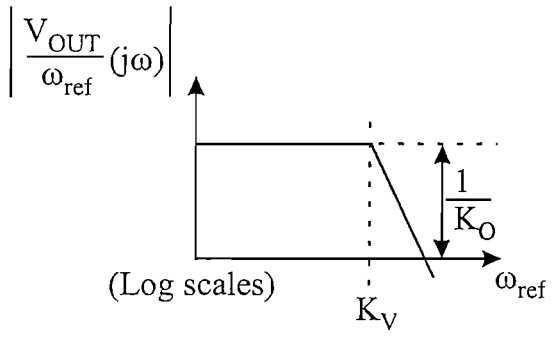

FIGURE 1.16: Transfer function of PLL with no loop filter.

This can be improved by adding a simple single-pole RC filter, thus generating a second order PLL. When designing the loop, it is desirable to have a low cutoff frequency associated with this block, but has to be capped to a value so that there isn't excessive peaking in the frequency response. One problem (depending on the application) is that the loop bandwidth is dependent on the loop gain and so too is the locking range. The transfer function for this second order PLL is:

$$\frac{V_{OUT}}{\omega_{ref}} = \frac{1}{1 + \frac{s}{K_V} + \frac{s^2}{\omega_1 \cdot K_V}} \cdot \frac{1}{K_O}, \tag{1.3}$$

where K_V is the same as before, ω_1 is the product R·C (components of the filter.) The transfer characteristics of this arrangement is given in Fig. 1.17.

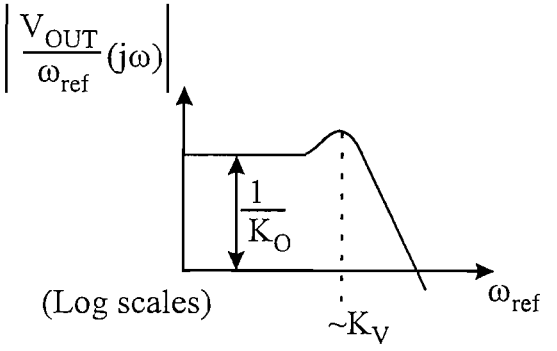

FIGURE 1.17: Transfer function of PLL with 1st order loop filter.

The 3dB bandwidth of this structure is $\sqrt{2} \cdot K_V rad \cdot s^{-1}$, and therefore ω_1 is set as low as possible (to reject out of band signals) before peaking becomes a problem. It is usual to have a maximally flat low pass configuration with the complex conjugate poles placed at $\frac{\pi}{4}$ rad from the negative real axis.

By adding a 'zero' to the transfer function of the filter (Fig. 1.18(a)), the pole frequency can be lowered without the risk of under-damping. This filter has a pole at $\omega_1 =$

$\frac{1}{C\cdot(R_1+R_2)}$ and a zero at $\omega_Z = \frac{1}{C\cdot R_2}$. The phase margin, with such a modification, improves so long as there is enough margin for the phase response to recover beyond $\frac{\pi}{4}$ rad (see Fig. 1.19.) With this solution, the loop bandwidth and loop gain are now independent of one another, thus allowing flexibilty in design.

In modern PLL design, it is usual to have a charge-pump before the loop filter and this removes the resistor, R_1 (Fig. 1.18(a)) from the loop filter (the input voltage source and R_1, are replaced with a current source and its internal resistance). Another point is that a shunt capacitor is placed across the remaining RC filter, resulting in a second order filter and a third order PLL. Its sole purpose is to remove spurs present on the VCO input line and its associated pole is usually placed some distance away from the dominant pole, so as to keep the steady-state response the same as the 2^{nd} order PLL.

1.5.3 Voltage controlled oscillator

The next and probably most challenging block is the VCO as, by default, it operates at the highest frequency within the loop. This function can be implemented in a number of ways with each having their merits and flaws. The examples include LC-oscillators and ring oscillators (both shown in Fig. 1.20). The specifications of the PLL/frequency synthesiser dictate the choice of topology. Ring oscillators [10] tend to be easier to design, and can be analysed with a cascade of equations and a rudimentry feedback path. However, in the newly emerging wireless applications, the call for cheaper, yet high performance radio chips translates into the need for oscillators with high specification spectral purity running at high operating frequencies. LC-tank oscillators [10] have been studied and characterised to deliver a low phase noise output at high frequencies; this explains their attention in recent publications [10]. Modern monolithic processing has

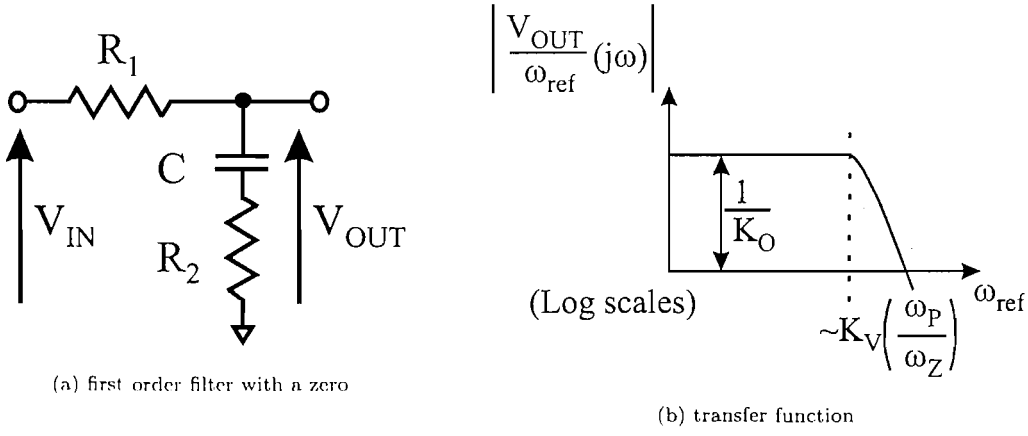

(a) first order filter with a zero

(b) transfer function

FIGURE 1.18: Loop filter and the corresponding characteristic of the configured PLL.

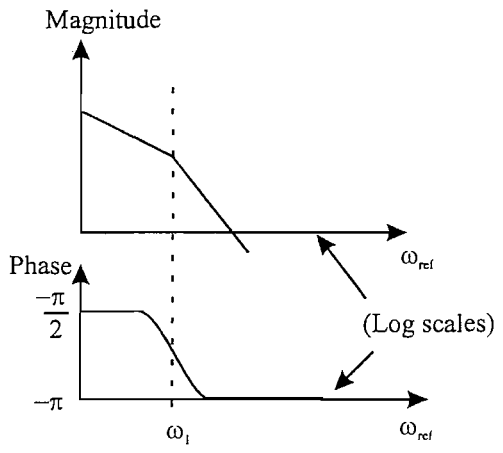

(a) bode plot of PLL with a single pole filter

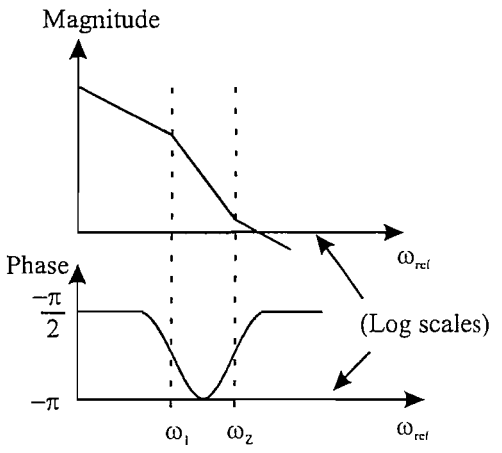

(b) bode plot of PLL with a single pole and zero filter

FIGURE 1.19: Bode plots of the PLL.

enabled passive structures such as spiral and 'horseshoe' inductors to be fabricated on the same die as MOS and bipolar circuitry [49]. By analysing the process steps, a little imaginative layout can generate useful (occasionally parasitic) devices such as varicaps, also found in VCOs. Although possessing poor Q-factors compared with discrete resonators, monolithic integration is certainly possible.

Ring oscillators based on CMOS-style inverters are easy to design and implement as their behaviour has been characterised and described using large-signal functions. The oscillation frequency can be set by modulating the tail current in the case of an inverter based on a differential pair (static logic) or by modulating the power supply of the inverter chain in the case of inverters based on a classical CMOS inverter (dynamic logic.) As for its performance, they tend to have a wide tuning range, low silicon area and their behaviour across process, power supply and temperature is modelled well owing to good MOS models. Power supply and substrate noise tend to be an issue, but can be minimised by employing a differential inverter topology. However, the output swing remains a concern with such a circuit.

The LC-tank based oscillators, if a perfect resonator (ie with no losses) were allowed, would have energy flowing between the electrical and magnetic storage elements at the rate of the calculated resonant frequency:

$$\omega_o = \frac{1}{\sqrt{L \cdot C}}. \tag{1.4}$$

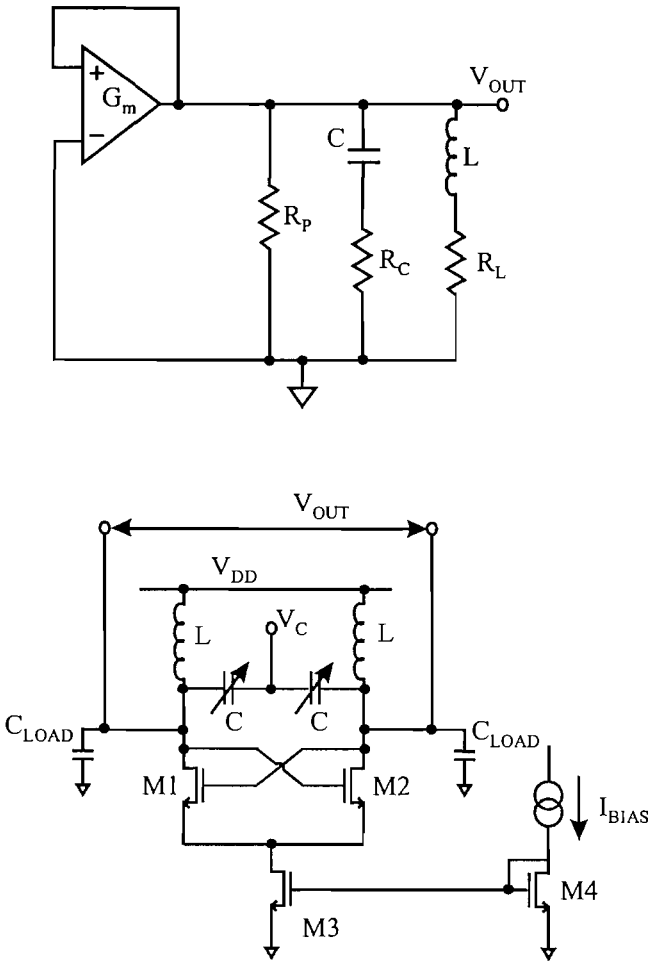

(a) general LC-tank oscillator with a transistor implementation below it,

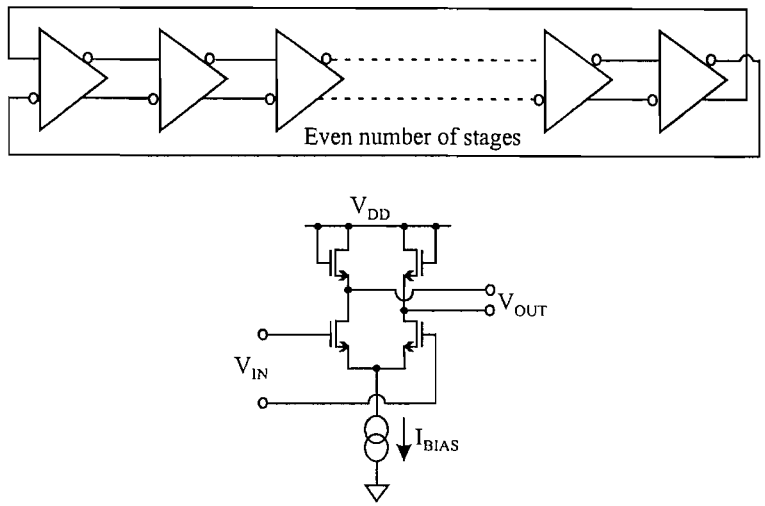

(b) general ring oscillator with a transistor level view inside one of the inverters

FIGURE 1.20: Two examples of voltage controlled oscillators.

Such a network should yield an infinite Q factor, with the definition given by:

$$Q = \frac{(stored\ magnetic\ energy) + (stored\ electrical\ energy)}{energy\ dissipated\ per\ cycle} \cdot 2 \cdot \pi. \qquad (1.5)$$

Unfortunately, reality shows the resonator to possess a finite Q factor. In the context of LC-tanks, the value depends on a number of issues such as materials, frequency of operation and physical design of the various components in the tank. In a monolithic silicon implementation, the degradation in Q factor is an aggregate of various loss mechanisms from skin effect in the conductor of the inductor to eddy currents in the silicon directly below a planar inductor design. Inductors fabricated in foundry silicon processes have Q factors of 5-10 at 2GHz [20] versus Q factors in the 100's for sub-100MHz discrete components. Such 'losses' are modelled electrically in the tank as resistive, with the series resistance of the inductor being captured in another resistor, in parallel with the parallel LC network. The losses are usually compensated for by matching it with an active negative transconductor (effectively a negative resistor.)

LC-tank oscillators tend to give a larger output swing and can operate at higher frequencies than ring oscillators, but have the downsides that they consume a great deal of area, require very accurate models (which can imply large complex models for the passives and stresses the need for better models) and has a narrow tuning range (due to the varactors), though switching out discrete capacitors can alleviate this.

The Q-factor degradation poses a problem for receiver designs as it translates into a spreading of the carrier power into the neighbouring sidebands which is captured in the performance metric 'phase noise.' From a channel selection point of view, the density of channels within an allocated band depends on this noise power at an offset to the resonant frequency of the local oscillator.

This is a good point to promote one of the main themes of this thesis, namely the reduction of power consumption in the high frequency analogue sections of a receiver. By loading the VCO with the mixer and frequency divider, the loaded Q-factor of the resonant tank will degrade. Such loading provides additional mechanisms by which energy in the tank is lost and therefore, more power must be consumed by the negative-g_M stage. This is another reason why the dual-modulus frequency dividers discussed in later chapters only have one divide-by-2 flip-flop loading the VCO outputs, whilst running at the highest frequency.

To be able to quantify and control this phase noise in oscillator design relies on knowing how its output is modified and identifying any distinct regions in a representation of

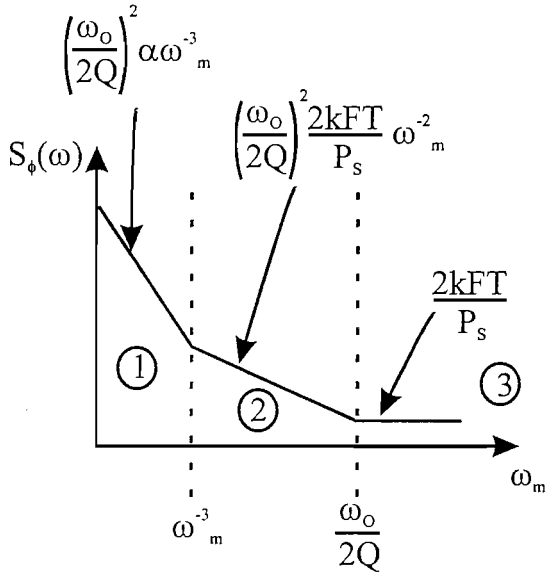

FIGURE 1.21: Power spectral density of the phase noise of a feedback oscillator (assume log scales for both axes). Region 1 is attributed to the flicker-noise of the active devices, whilst the characteristic in region 2 is due to the Q-factor of the oscillator. The last region quantifies the system noise floor.

its output. One of the first models to talk about phase noise is based on a feedback oscillator model [17], and through a heuristic approach, the power spectral density of the phase noise $S_\Phi(\omega_m)$ has been described as:

$$S_\Phi(\omega_m) = \frac{2 \cdot F \cdot k \cdot T}{P_S} \cdot \left(1 + \left(\frac{\omega_O}{2 \cdot Q \cdot \omega_m}\right)^2\right) \cdot \left(1 + \frac{\omega_{\frac{1}{\omega^3}}}{|\omega_m|}\right), \tag{1.6}$$

where ω_m is the offset from the centre frequency of the resonator, F is a device noise factor (empirical), k is Boltzman's constant, T temperature in Kelvin, P_S is the absolute power of the resonant frequency ω_m, $\omega_{\frac{1}{\omega^3}}$ is the corner frequency distinguishing the $\frac{1}{\omega^3}$ and $\frac{1}{\omega^2}$ regions, and lastly Q is the load Q-factor. With there being an empirical parameter in the equation (obtained by fitting parameters to device data, and applying also to the device flicker noise corner frequency) this is not a reliable model to explain the shaped noise. Saying this, the white noise and $\frac{1}{\omega^2}$ regions are predicted well, though the $\frac{1}{\omega^3}$ corner frequency should not be assumed to be equal to the device flicker noise corner frequency [4]. A representative plot of the asymptotes are given in Fig. 1.21.

It seems appropriate to break this equation down into the various regions and relate them to various metrics of the general oscillator. The large noise power seen rolling off at a rate inversely proportional to the cube of the offset frequency (region 1) is attributed to the device noise of the active elements. As one moves further away from the carrier (after what was thought to be the device flicker noise corner [17]), the poor Q-factor of the oscillator becomes an issue, shaping the noise power with a roll-off that is inversely

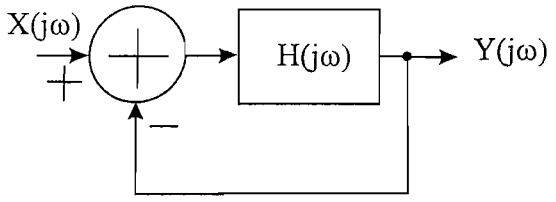

FIGURE 1.22: General representation of an oscillator in a closed-loop feedback arrangement.

proportional to the square of the offset frequency. This approximation continues until half the feedback bandwidth $\left(\frac{\omega_O}{2 \cdot Q \cdot \omega_m}\right)$ [17] is reached, after which the combined system noise floor becomes dominant. This is not only due to the oscillator components, but also any amplifiers attached to its output, and even measurement equipment. PLL designers tend to design for large output swings to give a low phase noise output [10]. There can be instances where poor matching between transceiver stages can yield a noise floor that swamps the lower offset regions of the spectrum. Likewise, it is possible to have excellent Q-factors in a design such that the composite spectrum shows no $\frac{1}{\omega^2}$ region.

LC-tank based oscillators are known to be less noisy than inductorless versions [10][3]. However, ring oscillators have played an important role in microprocessors and due to their differential operation and ease of operation, such oscillators can also play an important role in transceiver design. It is important to characterise the phase noise associated with this type of oscillator and though no explcit tank exists, a Q-factor for this style can still be defined. The derivation is purposely omitted and the result is shown below ('open-loop' Q is how much the closed-loop system opposes a change in frequency):

$$Q = \frac{\omega_O}{2} \cdot \sqrt{\left(\frac{dA}{d\omega}\right)^2 + \left(\frac{d\Phi}{d\omega}\right)^2}, \tag{1.7}$$

where A and Φ are the magnitude and phase responses of a linear transfer function, $H(j\omega)$ (see Fig. 1.22.) The noise power spectral density is then shown to be shaped by the following transfer function:

$$\left|\frac{Y}{X}[j(\omega_O + \omega_m)]\right|^2 = \frac{1}{4 \cdot Q^2} \cdot \left(\frac{\omega_O}{\omega_m}\right)^2. \tag{1.8}$$

Three types of noise have been identified for this oscillator and shown to have their power spectral densities shaped by that transfer function [10]. Additive noise is simply a direct addition of noise (power) components to the output and hence the total noise at an offset frequency is the sum of those components each generated by a noise source at a frequency near the oscillation frequency. HF multiplicative noise arises because nonlinearities associated with the oscillator can cause an injected component to 'fold' around the carrier, therefore doubling the noise power at an offset from the carrier tone (twice the power calculated for the additive noise component.) LF multiplicative noise

is where noise power close to DC is upconverted to and around the carrier frequency. In a static current ring VCO, the frequency of oscillation is changed by changing the tail current over each cell, and any noise on this input line is frequency modulated to around the carrier, again being shaped by the transfer function of the oscillator.

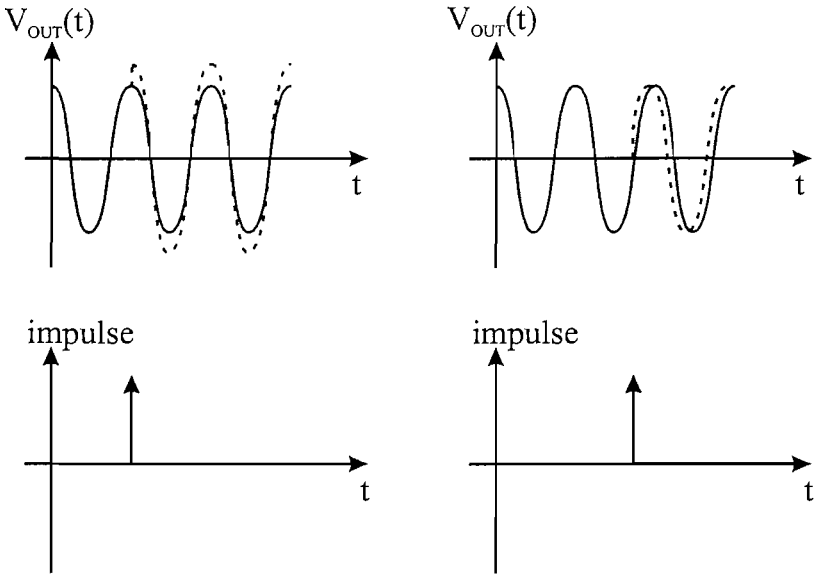

FIGURE 1.23: Concept of the 'impulse sensitivity function' used in the analytical model of phase noise, showing the principle of time-variance with an impulse injected at two different instances of a simple periodic signal.

In recent times, a study has been performed on high performance silicon oscillators, hoping to understand and model phase noise without the reliance on an empirical model. One important model, again general to all types of oscillators, moves away from linear time-invariance theory, as it is found that the phase noise response is dependent on the instance in an output cycle when a perturbation occurs [4]. If disturbed during the crest of the output waveform, the oscillator should settle back with very little effect on the phase. However, if disturbed in the vicinity of a zero-crossing, then the effect is more pronounced. These points are captured in Fig. 1.23. An 'impulse sensitivity' function is defined for the different types of noise (as well as being specific to the style of oscillator) and the advantage is then the removal of empirical parameters giving rise to an analytical expression. The study goes on to show that noise near integer multiples of the oscillation frequency is the major contributor to the low offset frequency noise seen around the carrier tone. Asymmetry in the output waveform can affect the upconversion of the low frequency noise, thus allowing the $\frac{1}{\omega^3}$ corner frequency to be lower than device flicker noise corner frequency.

The observant reader will note the great deal of attention paid to the noise concerning the zero-crossing instants and not the noise resulting in amplitude variations. In modern

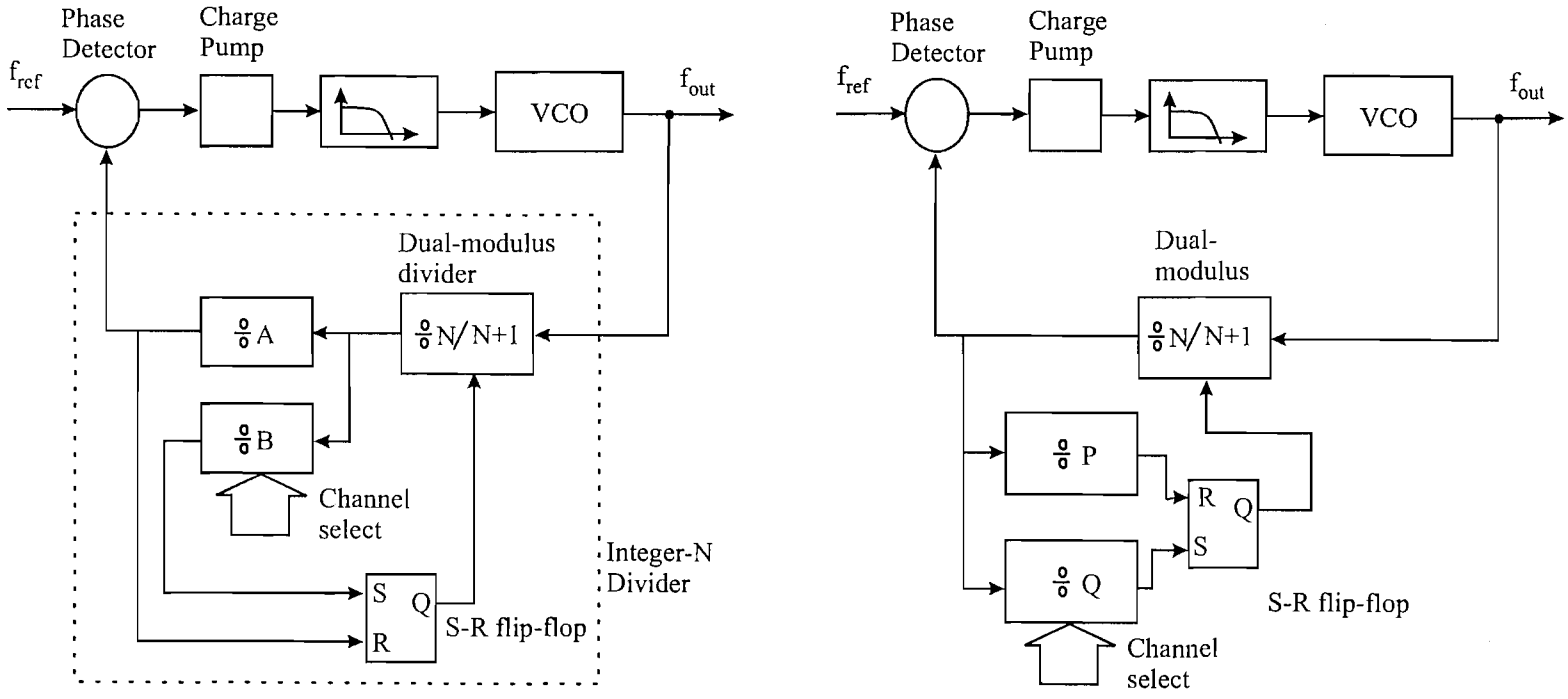

(a) integer-N, (b) fractional-N.

FIGURE 1.24: Block diagrams of two different PLL-based frequency synthesisers.

day receiver designs running in GHz bands, any such noticable amplitude modulation is most likely to be of a lower frequency and should thus be attenuated by the loop bandwidth of a PLL-based frequency synthesiser. The variation will also be limited by the low power supply encountered on common IC implementations. If the variation is large, then the local oscillator waveform will suffer 'clipping', but the resulting distortion should occur at multiples of the fundamental tone (desired VCO output frequency), again attenuated by the loop bandwidth of the PLL. Another reason for the lack of concern is that in most phase modulation/demodulation schemes, a switching mixer is employed. This type of mixer is only concerned with the 'instant' at which a local oscillator input crosses its mean value and hence any amplitude noise is not significant.

1.5.4 Synthesisers

So far, only a PLL has been described. However, by adding a divider (programmable or fixed-ratio) between the VCO and phase detector, the properties of this loop can be changed. For example, by having a fixed divider in the loop, the PLL becomes a frequency multiplier where the output frequency is the reference frequency multiplied by the division ratio. In the case of the programmable divider, one has the output frequency varying in multiples of the reference input arriving on the phase detector input. This last category can be broken up into two PLL types: integer-N and fractional-N synthesisers. Both are illustrated in Fig. 1.24.

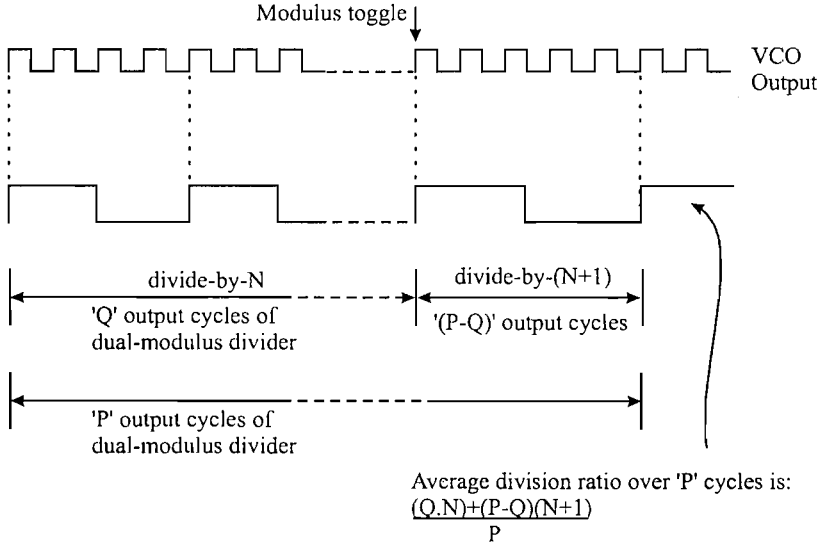

FIGURE 1.25: The principle behind fractional-N division, based on a dual-modulus divider.

The integer-N PLL is where the division ratio is an integer and hence the output fre-

quency is synthesised as an integer product of the reference frequency, not necessarily starting at zero. The fractional-N synthesiser has the same blocks as the integer-N synthesiser, but the divider now has fractional moduli, between one integer modulus and a consecutive integer modulus. A very simple implementation exists and in fact is the same for both. Depending on where the output is taken, one can have either. By having the output of the dual-modulus divider acting as the output of the programmable divider yields a fractional-N synthesiser. With the help of Fig. 1.25, it can be seen that the number of output cycles is counted and compared against a programmed value, and when equal, the division modulus of the dual-modulus divider is toggled for the remaining cycles of the fixed modulus counter. The average value is usually a fraction between N and N+1. In the case of the integer-N synthesiser, the output is the most significant bit of the fixed modulus counter. So now, the output truly is a multiple of the number of input cycles. More elaborate sigma-delta loops (Fig. 1.26)are often found owing to their averaging and noise-shaping characteristics.

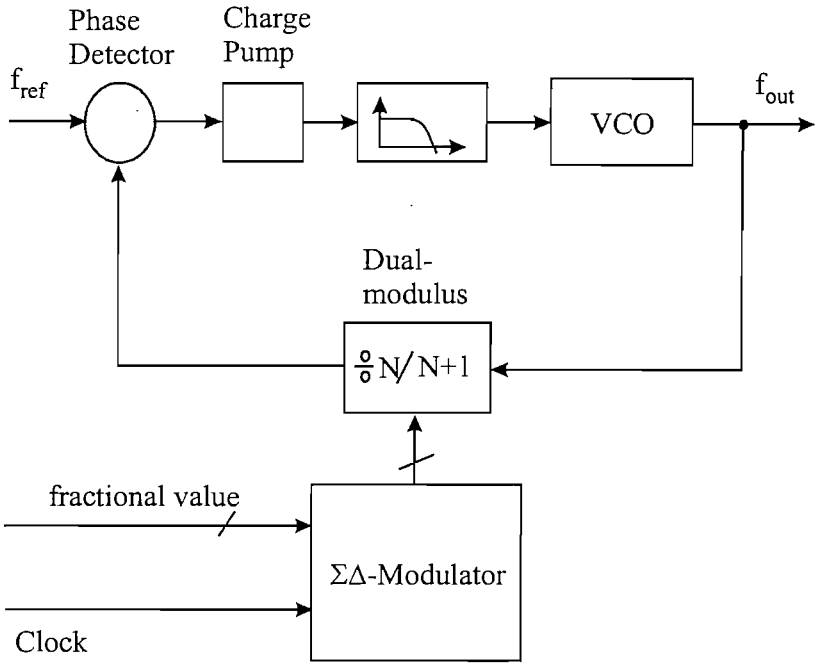

FIGURE 1.26: Block diagram of a fractional-N synthesiser, with a delta-sigma modulator controlling the dual-modulus divider, so as to randomise the error before shaping the noise out of the band of interest.

The design and implementation of fractional-N synthesisers is pursued by those requiring a local oscillator with a low phase noise output and a fast settling time. As designs target higher frequency bands and demand a relatively high frequency resolution (to squeeze more channels into a band) so the integer-N architecture becomes unworkable, because a lower reference frequency (relating to a finer resolution) would dictate a narrower loop bandwidth thus giving a synthesiser with a poor settling time. It is usual

to have the reference frequency of the PLL equal to the channel spacing of specification (usually done using a standard crystal oscillator which is subsequently divided down using a prescaler). When multiplying this frequency, there is a $20\log_{10}(N)$ degradation in phase noise performance at the output (phase noise of the VCO experiences a high-pass filtering action with the dominant corner frequency equal to the loop filter bandwidth). So by lowering the division ratio, N, the result is an improvement in phase noise (though other contributors may swamp any performance gain due to this parameter). The problem with dividers in terms of phase noise is that if the phase detector is not the major source of noise then any input variations will appear on the output multiplied by the division ratio N. Therefore, the output spectrum of the phase noise at low offset frequencies will be a factor of N^2 higher in power than the phase fluctuations at the input [31][2].

With the ability to choose fractional division ratios (between two or more integers that aren't necessarily consecutive,) the design can use a higher reference frequency. This allows the loop bandwidth of the PLL to be increased, thus changing the dynamics to give a faster settling time. Increasing the loop bandwidth can indirectly lead to a lower phase noise degradation in the context of passive loop filters, as it is possible to have lower value resistors (assuming a RC based filter). With narrow loop bandwidth PLLs, the time constant of the dominant pole must be large, thus implying a large value of resistor and resistor noise to go with it. However, a problem cited concerns the phase detector. For every decade increase in reference frequency, the phase noise on the output tends to degrade by 10dB, hence overall gain in performance is approximately $10\log_{10}(N)$.

One of the major disadvantages in using fractional-N synthesisers is the spurious content on the output. Reference spurs occur on the output owing to mismatches in the charge pump negative and positive current branches, as well as non-ideal phase frequency detection and imperfections in the compensation circuitry. The result is small pulses emitted at a rate equal to the reference frequency multiplied by the division resolution. In the case of integer-N division, this resolution is simply 1. For the fractional-N divider, this is equal to the minimum fractional step. The ramifications are that it modulates the VCO control input, and is hence upconverted around the carrier tone. The power of this tone may be comparable, if not larger than, the 'wanted' tone leading to an error in the modulation/demodulation of the channel. Loop bandwidths, as a rule of thumb, are usually set 10 times lower than the reference frequency, to give approximately 20dB suppression (first order passive filter assumed) to this spur before it reaches the VCO control input. In the case of fractional-N synthesisers, this spur is even harder to suppress as it is at a fraction of the reference frequency. This has prompted the design of higher order loops in such synthesisers (by means of higher order loop filters) [8]. There is also a reluctance amongst designers to go down the route of fractional-N synthesis, as

it usually involves the need to design a good delta-sigma modulator as well as the PLL itself.

Practical examples of synthesisers have become more numerous, as previous investment in studying these structures is now aiding circuit designers involved in radio frequency design and planning. One such design is based on a 5GHz wireless LAN standard [20] [23](Fig. 1.27.) This chip is a frequency synthesiser fabricated in 0.24μm bulk CMOS and runs with 1.5V and 2.0V power supplies for the analogue and digital sections respectively. The fabricated circuit consumes 25mW, of which ~16% is accounted for by the VCO and first divide stage. A key block in this circuit, attributable to this low consumption is an injection-locked frequency divider. A superharmonic injection-locked divider [21] works by having the incident frequency as a harmonic of the oscillation frequency. If one considers a feedback loop that is unconditionally unstable with zero input and has its own free-running oscillation frequency, then by having a nonlinear block in such a configuration, harmonics are generated, and appropriate filtering removes all except the desired harmonic. This harmonic is fed back to the input and compared with the incident frequency (Fig. 1.28(b)). This idea is implemented with an appearance similar to a LC-tank based oscillator with more spiral inductors used to realise the tank (Fig. 1.28(a).) The difference is the current to this oscillator is modulated by the VCO output and it is found the common source connection is a good place to inject the incident signal. Subsequently, a pulse swallow frequency divider is used to achieve

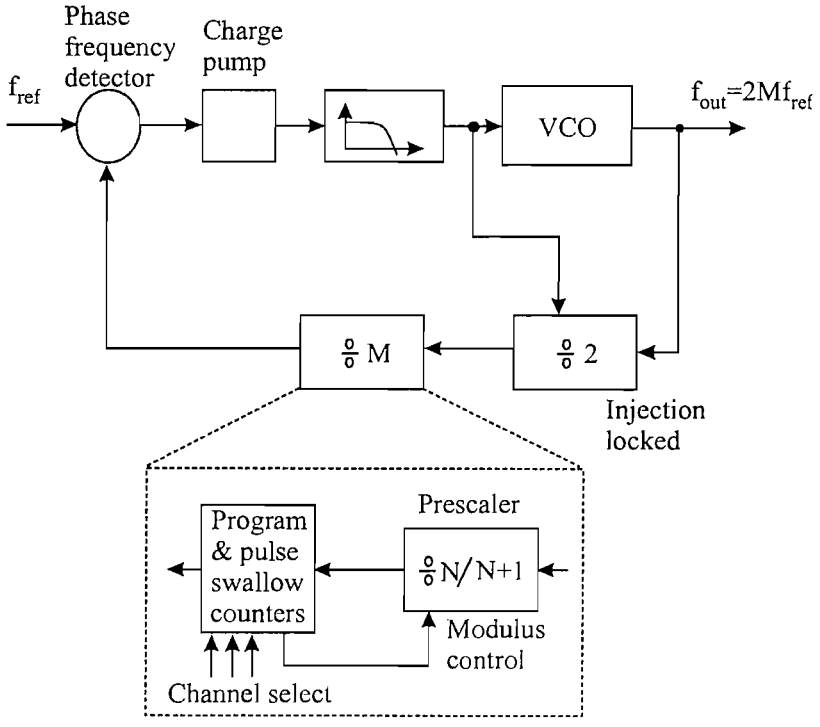

FIGURE 1.27: Block diagram of a 5GHz frequency synthesiser with an injection-locked frequency divide-by-2 [20].

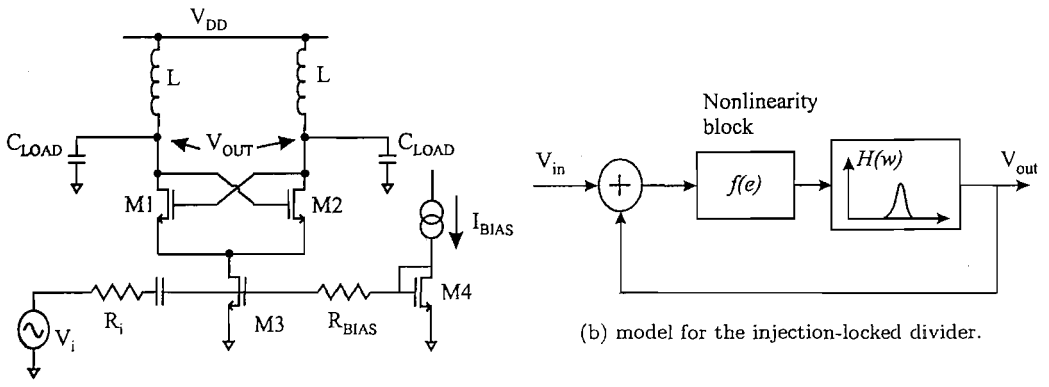

(a) schematic of the divider,

(b) model for the injection-locked divider.

FIGURE 1.28: Superharmonic injection-locked divider.

the programmability, by using a divide-by-22/23 unit to achieve a programme ratio of 220-227 inclusive (similar to the arrangement in Chapter 8, except that a fixed counter of 10 and a 3-bit comparator are used). Incidentally, the dual modulus divide-by-22/23 is built upon a cascade of divide-by-2/3's. The synthesiser runs on a 11MHz reference, and has a LO spacing of 22MHz.

One last example of how far standard bulk CMOS can be pushed is with the demonstration of a 10GHz PLL [16](Fig. 1.29). Using a 0.18μm bulk CMOS process, a third-order PLL has been realised. In this example, a standard LC-tank based oscillator is designed with PMOS energy restorers and together with a varicap in parallel with a fixed capacitance delivers a tuning range of around 700MHz. However, in order to drive further stages (dual-modulus divider and quadrature generator) a buffer is used, with low Q inductors in their loads. With division ratioes of 1024 and 1028, the PLL draws 34mA from a 1.8V supply.

1.6 Summary

This introductory chapter is written with the intention of preparing the reader with some of the terminology to be encountered as well as giving some justification for how the later designs have been tailored. Examples are given demonstrating the applicability of a digital CMOS process to a high perfomance complex analogue problem. The chapter introduces the reader to the radio receiver architectures commonly used in modern IC implementations. This helps the reader appreciate the enthusiasm for low power circuits and understand the context in which terms such as 'channel selection' arise. Though the circuits designed have never deviated from frequency division, there has always

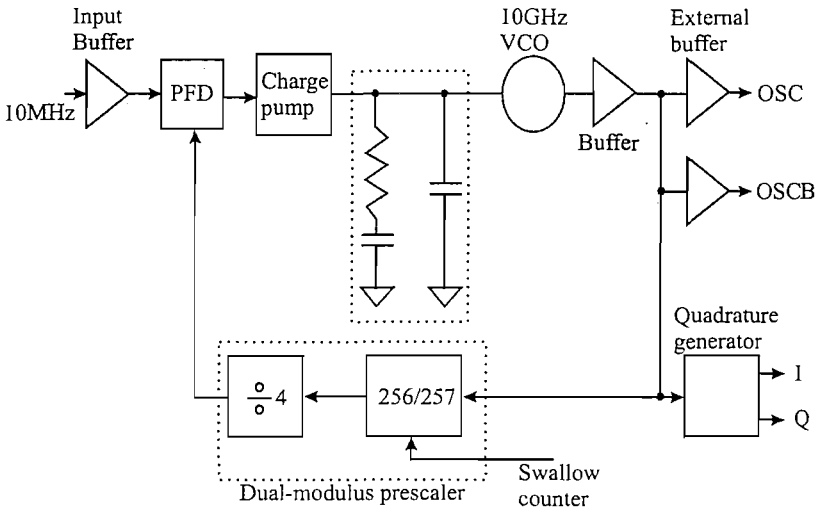

FIGURE 1.29: Block diagram of a 10GHz frequency synthesiser [16].

been an application in mind, so as to put performance values attained into perspective, thus allowing a comparison against other published designs. The application has been frequency synthesisers for local oscillator generators in radio transceiver design. This brief section justifies the role of the frequency synthesiser (and eventually the frequency divider) but deliberately stays away from the rigorous maths showing modulation/demodulation, as well as showing how effects such as phase noise can mar the transceiver's performance.

A look at bulk CMOS technology is then given, showing its current state, followed by the same treatment for SOI CMOS technology. In the case of SOI CMOS technology, the various operating modes are also brought to light, as is the individual transistor's structure, that have allowed the concept of the stacking to be realised in Chapter 5. An explicit section exists on the importance of good transistor models that enable designers to produce circuits in sub-micron geometries.

By looking into the various blocks of the synthesiser, the reader can begin to appreciate some of the decisions made that have shaped the divider designs to be seen. By explaining the operation and necessity of each PLL component, it can be shown how each interacts with the others. From the discussion, it can be seen that a lower division modulus is desirable, with designers taking the steps to implement fractional moduli so as to get the same frequency granularity, but with a higher reference frequency. It is also imperative not to load the VCO with too much input capacitance as the output frequency is a function of the capacitance on the output. This is clearly seen in a later chapter where a bulk CMOS divide-by-2 loads a VCO and careful interaction between the designer of each must be present to target the desired range of operating frequencies. This, however, is first seen and realised in the SOI divider where the decision

to go with that circuit topology was largely due to the initial cascading of divide-by-2 stages, before arriving at the rest of the dual-modulus divider. Other divider's (some of which will be shown in Chapter 3) have been tried but were thought to load the VCO too much or consume too much power. The decision to go for an integer-N divider in the bulk CMOS programmable divider was made based on the fear of fractional-spurs. Rather than investing more time in designing a high order delta-sigma modulator, it was felt that the integer-N divider would suffice, and the finished design together with measurements justify this course of action. Although it results in a synthesiser with a higher phase noise at low offset frequencies, the wireless system specification was generous enough to dictate a high phase noise at 10MHz (channel spacing.) This may not be the case in another application, however such as digital satellite tuner designs, where a more stringent phase noise specification exists.

1.7 Outline of the thesis and extent of originality

This thesis has been compiled into a number of chapters, describing various designs as well as looking at aspects of source-coupled logic and other divider circuits.

As the majority of the work has been in SCL, Chapter 2 presents a review of this form of logic. Example circuits, simulations, and predictions accompany discussions on the DC and AC characteristics of the basic inverter.

Chapter 3 will then review some of published divider circuits. Examples of high performance fixed-, dual- and multi-modulus dividers are given, looking into their circuit topology and their architecture.

One of the key parts of the work conducted is captured in Chapter 4. Without making explicit reference to the process, a description of the modified phase select architecture is given. The chapter concentrates on the suggested phase-selector and looks at how its elegant implementation allows broadband error-free operation. Whilst the phase selector architecture is attributed to work by Dr. Jan Craninckx and Prof. Michael Steyaert [28], the method of controlling the selected phase with an 8-state controller so as to eliminate erroneous outputs at high frequencies is the work of this dissertation's author, as was the circuit topology of the current steering phase-selector.

The SOI dual-modulus divider is described in some detail in Chapter 5. The concept of stacking and current re-use is strongly promoted, and should convince a reader of the

feasibility of a higher power supply. Schematics, layout, simulation and measurements all accompany the clever use of PDSOI MOS transistors. The idea behind stacking SOI circuit blocks was based on initial ideas of Craig Easson and Prof. William Redman-White. It has been developed for divider circuits by this author to the stage mentioned in this chapter. The integration of the current steering phase selector is also the work of this author. The complete design and chip testing was performed at the University of Southampton.

The second phase of the circuit research was undertaken using bulk CMOS technology, and this is made clear in the subsequent three chapters. The bulk dividers, which were fabricated by Philips NV, for the 5 GHz wireless LAN project are illustrated in these chapters. With the chosen receiver architecture to be direct-conversion, the quadrature generator of Chapter 6 shows the schematics through to layout with measurements of the 10GHz (and higher) divide-by-2 cell. The author has developed this circuit from a standard circuit topology to that with optimised performance in the desired frequency band, targetted for a process that is predominantly digital. Succeeding this block in the actual frequency synthesiser comes the programmable divider. Specific to the wireless LAN project is a divide-by-513 to 544 block, based on a divide-by 16/17 dual-modulus divider. The dual-modulus divider is again based on the current steering phase-selector circuit developed for the SOI process, but without any stacking. The complete dual-modulus divider has again been developed and completed by the author, with some assistance on its inclusion in the multi-modulus divider by engineers Nenad Pavlovic and Dr. Domine Leenaerts. All these dividers have been fabricated in 0.18μm bulk CMOS technology, with measurements (performed by the same Philips engineers) showing satisfactory operation.

After the conclusion, an Appendix has been added to describe the design of the high frequency testboard used to characterise the SOI divider die.

Before ending this chapter, Figs. 1.30(a) and 1.30(b) show the fabricated dividers of this work compared against other published designs [24][15][37][25][6][28][45][43][7][41][22][29][19][27]. In both cases, the maximum input frequency for a particular technology node shows how the designs discussed in this thesis fair against a sample of other designs. Although the design in Chapter 5 doesn't compare favourably against state-of-the-art designs, its elegance and novelty are sadly overshadowed by other circuit cells necessary in its testing.

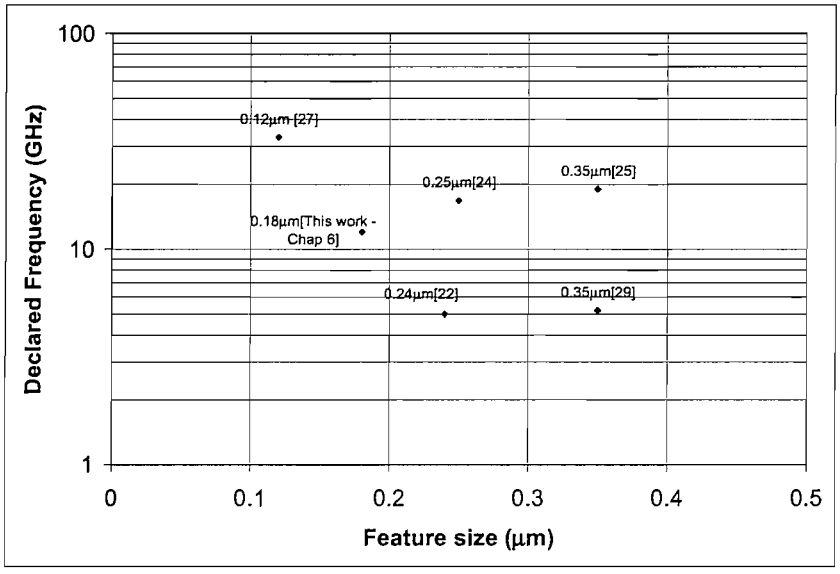

(a) comparison of fixed divide-by-2 circuits,

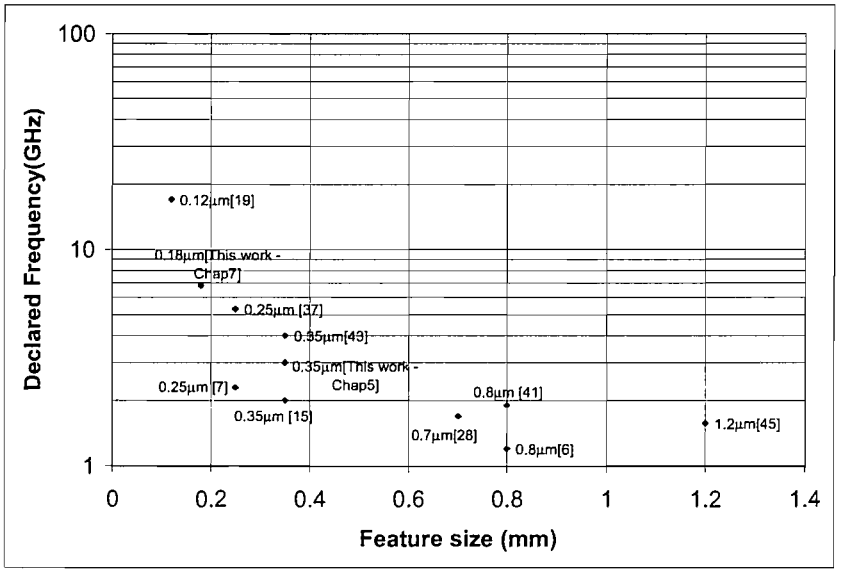

(b) comparison of multi-modulus dividers

FIGURE 1.30: Comparison of maximum input frequency against other published designs.

Bibliography

[1] International technology roadmap for semiconductors. Technical report, Semiconductor Industry Association, 2001.

[2] A.Hajimiri. Noise in phase-locked loops (invited). In *IEEE Southwest Symposium on Mixed Signal Design*, pages 1–6, Austin, USA, 2001.

[3] A.Hajimiri, S.Limotyrakis, and T.H. Lee. Jitter and phase noise in ring oscillators. *IEEE Journal of Solid-State Circuits*, 34(6):790–804, June 1999.

[4] A.Hajimiri and T.H.Lee. A general theory of phase noise in electrical oscillators. *IEEE Journal of Solid-State Circuits*, 33(2):179–194, February 1998.

[5] Advanced Micro Devices (AMD). Amd products preview (2003). Http://www.amd.com/.

[6] B.Chang, J.Park, and W.Kim. A 1.2GHz CMOS dual-modulus prescaler using new dynamic D-Type flip-flops. *IEEE Journal of Solid-State Circuits*, 31(5):749–752, May 1996.

[7] B.Chi and C.Chi. New implementation of phase-switching technique and its applications to GHz dual-modulus prescalers. *IEE Proc. Circuits, Devices and Systems*, 150(5):429–433, Oct 2003.

[8] B.Daly. Comparing integer-N and fractional-N synthesisers. *Microwaves & RF*, pages 93–133, September 2001.

[9] B.J.Sheu, D.L.Scharfetter, P.-K.Ko, and M.-C.Jeng. BSIM:berkeley short-channel IGFET model for MOS transistors. *IEEE Journal of Solid-State Circuits*, 22(8): 558–565, August 1987.

[10] B.Razavi. A study of phase noise in CMOS oscillators. *IEEE Journal of Solid-State Circuits*, 31(3):331–343, March 1996.

[11] B.Razavi. Design considerations for direct-conversion receivers. *IEEE Trans. on Circuits and Systems-II*, 44(6):428–435, June 1997.

[12] C.-T.Chuang, P.-F.Lu, and C.J.Anderson. SOI for Digital CMOS VLSI: Design considerations and advances. *IEEE Journal of Solid-State Circuits*, 86(4):689–720, Apr 1998.

[13] C.C.Enz, F.Krummenacher, and E.A.Vittoz. An analytical MOS transistor model valid in all regions of operation and dedicated to low-voltage and low-current applications. *Analog Integrated Circuits and Signal Processing Journal*, 8:83–114, July 1995.

[14] C.Lallement, F.Pêncheux, and W.Grabinski. High level description of thermodynamical effects in the EKV 2.6 most model - Special Session. In *MIXDES*, Wroclaw, Poland, June 2002.

[15] C.S.Vaucher et al. A family of low-power truly modular programmable dividers in standard 0.35μm CMOS technology. *IEEE Journal of Solid-State Circuits*, **35**(7): 1039–1045, Jul 2000.

[16] D.-J.Yang and K.O. A monolithic CMOS 10.4GHz phase locked loop. *IEEE VLSI Symposium Proceedings*, pages 36–37, 2002.

[17] D.B.Leeson. A simple model of feedback oscillator noise spectrum. *Proceedings of the IEEE*, 54(2):329–330, February 1966.

[18] F.M.Gardner. Charge pump phase-locked loops. *IEEE Trans. Communications*, COM-28(11):1849–1858, November 1980.

[19] H-D.Wohlmuth et al. A 17GHz dual-modulus prescaler in 120nm cmos. In *IEEE Radio Frequency Integrated Circuits Symposium*, pages 479–482, Philadelphia, USA, June 2003.

[20] H.Rategh et al. A CMOS frequency synthesiser with an injection-locked frequency divider for a 5GHz wireless LAN receiver. *IEEE Journal of Solid-State Circuits*, 35 (5):780–787, May 2000.

[21] H.Rategh and T.H.Lee. Superharmonic injection-locked frequency dividers. *IEEE Journal of Solid-State Circuits*, 34(6):813–821, June 1999.

[22] H.R.Rategh, H.Samavati, and T.H.Lee. A 5GHz, 1mW CMOS voltage controlled differential injection locked frequency divider. In *IEEE Proc. Custom Integrated Circuits Conference*, pages 517–520, San Diego, May 1999.

[23] H.Samavati et al. A 5GHz CMOS wireless LAN receiver front end. *IEEE Journal of Solid-State Circuits*, 35(5):765–772, May 2000.

[24] H.Wang. A 1.8V 3mW 16.8GHz frequency divider in 0.25μm CMOS. In *IEEE International Solid-State Circuits Conference*, page TP12.1, 2000.

[25] H.Wu and A.Hajimiri. A 19GHz 0.5mW 0.35μm CMOS frequency divider with shunt-peaking locking-range enhancement. In *IEEE International Solid-State Circuits Conference*, page 26.3, 2001.

[26] Intel. Intel products preview (2003). Http://www.intel.com/.

[27] J-O.Plouchart et al. A power-efficient 33 GHz 2:1 static frequency divider in 0.12-μm SOI CMOS. In *IEEE Radio Frequency Integrated Circuits Symposium*, pages 329–332, Philadelphia, USA, June 2003.

[28] J.Craninckx and M.Steyaert. A 1.75GHz/3V dual modulus divide-by-128/129 prescaler in 0.7μm CMOS. *IEEE Journal of Solid-State Circuits*, 31(7):890–897, Jul 1996.

[29] J.M.C.Wong, V.S.L.Cheung, and H.C.Luong. A 1V 2.5-mW 5.2-GHz frequency divider in a 0.35-μm CMOS process. *IEEE Journal of Solid-State Circuits*, 38(10): 1643–1647, Oct 2003.

[30] L.Geppert. RF Bridges to the Network - Special report. *IEEE Spectrum*, 38(1): 69–71, Jan 2001.

[31] M.Chrzanowska-Jeske. Fractional/integer-N PLL basics - technical brief SWRA029. Technical report, Texas Instruments - Wireless Communication Business Unit, 1999.

[32] M.Chrzanowska-Jeske. Effects of scaling interconnects - lecture. Technical report, Portland State University, 2003.

[33] M.M.Pelella and J.G.Fossum. On the performance advantage of PD/SOI CMOS with floating bodies. *IEEE Transactions on Electron Devices*, 49(1):96–104, January 2002.

[34] M.S.L.Lee. *Compact Modelling of Partially Depleted Silicon-on Insulator MOSFETs for Analogue Circuit Simulation.* Ph.D. thesis, University of Southampton, 1997.

[35] M.Steyaert et al. A 2V cellular transceiver front-end. *IEEE Journal of Solid-State Circuits*, 35(10):1895–1907, December 2000.

[36] M.Zargari et al. A 5GHz CMOS transceiver for IEEE802.11a wireless LAN systems. *IEEE Journal of Solid-State Circuits*, 37(12):1688–1694, December 2002.

[37] N.Krishnapura and P.R.Kinget. A 5.3GHz programmble divider for HiPerLAN in 0.25μm CMOS. *IEEE Journal of Solid-State Circuits*, 35(7):1019–1024, Jul 2000.

[38] University of California at Berkeley. BSIMSOI. Model for Silicon-on Insulator MOSFETs for Analogue Circuit Simulation. Website www.berkeley.edu.

[39] Department of Electrical Engineering and Computing Science. 94720 BSIM3. version 3.0 manual. Technical report, University of California, Berkeley, 1996.

[40] University of Florida. UFSOI. Model for Silicon-on Insulator MOSFETs for Analogue Circuit Simulation. Website www.uf.edu.

[41] P.Larsson. High speed architecture for a programmable frequency divider and a dual modulus prescaler. *IEEE Journal of Solid-State Circuits*, 31(5):744–748, May 1996.

[42] P.Lu et al. Floating-body effects in partially depleted SOI CMOS circuits. *IEEE Journal of Solid-State Circuits*, 32(8), Aug 1997.

[43] R.Ahola and K.Halonen. A 4GHz CMOS multiple modulus prescaler. In *IEEE International Conf. Electronics, Circuits, and Systems (ICECS),*, pages 323–326, Lisboa, Portugal, 1998.

[44] D.B.M. Klaassen R.M.D.A. Velghe and F.M.Klaassen. MOS model 9 - Unclassified report. NL-UR 003/94. Technical report, Philips Research Laboratories, Eindhoven, The Netherlands, 1995.

[45] R.Rogenmoser, Q.Huang, and F.Piazza. 1.57GHz asynchronous and 1.4GHz dual modulus 1.2 μm CMOS prescalers. In *IEEE Proc. Custom Integrated Circuits Conference*, pages 16.3.1–16.3.4., San Diego, May 1994.

[46] S.Krishnan and J.G.Fossum. Grasping SOI floating-body effects. *IEEE Circuits and Devices magazine*, 14(4):32–37, July 1998.

[47] S.M.Gowda, B.J.Sheu, and J.S.Cable. An accurate MOS transistor model for submicron VLSI circuits - BSIM_plus. *Proceedings IEEE Custom Integrated Circuits Conference*, pages 23.2.1–23.2.4, May 1991.

[48] S.M.Sze. *Semiconductor Devices, Physics and Technology.* John Wiley and Sons, 1985.

[49] T.H.Lee. *The Design of CMOS Radio-Frequency Integrated Circuits.* Cambridge University Press, 1999.

[50] U.L.Rohde. *Digital PLL frequency synthesizers - Theory and design.* Prentice Hall, 1983.

Chapter 2

Source coupled logic (SCL) theory and examples

2.1 Introduction

Throughout the succeeding design chapters, it will become apparent that the fabricated circuits are predominantly designed using a basic source coupled logic (SCL) topology. This chapter serves to describe such logic signalling and the methods to implement logic functions to handle such signals. A SCL inverter circuit forms the basic cell used to analyse the DC and transient behaviour associated with this differential logic signalling, as is often done when performing the equivalent analysis for CMOS topology circuits. Whilst the section on DC behaviour is fairly straightforward, the transient analysis of the circuit's behaviour relies on simple equations generated by other authors, collecting data from simulations for various permutations of circuit and device parameters. With each set of data, an effort is made to link the results with the small-signal behaviour of the inverter, so that any trends can be verified. Before the chapter ends, a look at some other SCL logic functions is presented, as well as the dependency of the switching speed on the common-mode input of a multi-input gate. The chapter ends with a look at the classic SCL topology master-slave flip-flop.

It must be reiterated that the study in this chapter refers to 'CMOS' as a process technology, accounting for the fact that it is a *complementary* MOSFET process. It can also refer to the schematic arrangement of PMOS and NMOS transistors giving a rail-rail logic swing, whereby the output node of any logic function is pulled 'up' to the V_{DD} power rail when a particular combination of logic 0's is present on the gate terminals of the PMOS transistors, and the output node is pulled 'down' to digital ground when a particular combination of logic 1's is present on the gate terminals of the NMOS transistors. This definition is not central to this chapter.

2.2 Source coupled logic

This logic style has been used in its sister form, emitter coupled logic (ECL), for many years where high speed logic circuits are required. Compared with CMOS or TTL topologies, the SCL style permits differential signalling as well as allowing high speed charging and discharging of output nodes. In the case of our divide-by-2 cell, a CMOS rail-rail topology divider would employ both p- and n-channel MOSFETS, where the p-type FETs present a considrable load on the previous circuit (VCO or maybe another divider). Although CMOS rail-rail dividers have a lower static power consumption compared with those used in these designs, SCL based designs offer a greater degree of control on switching and current consumption, as well as operating at higher frequencies.

Figures 2.1 (a) and (b), show the difference between a SCL and CMOS rail-rail combinatorial gates where the SCL version sits on a current source. The output is a pair of complementary outputs whose voltage difference is set by the tail current and one of the load resistor values. Immediately, the benefits of SOI emerge, because modern bulk CMOS technology insists on the n-wells and substrate being tied to the highest and lowest potentials respectively. The well known back-bias effect manifests itself here, forcing different threshold voltages as the number of MOS devices stacked on top of a like MOS device increases. The buried oxide found in SOI designs electrically isolates devices on the same die, thus allowing the relative V_T to remain constant, paving the way for many devices to be stacked. However, stacking is ill-advised when considering cascading such gates. If the output of gate (n-1) is to feed the bottom-most input of a 3 input NAND gate, level translation of some sort is required to give the right biasing, and this can degrade the performance if the voltage translation is large. This incidentally, is achieved using either a coupling capacitor between stages and setting the bias with a resistive divider on the input of the subsequent stage, or by using a source follower. The former can add significant load capacitance to ground on the output of the logic gate, through a parasitic capacitance to ground, whilst the latter is single-ended and consumes more power (in bulk CMOS circuits, the body-source transconductor in the source follower arrangement is active, reducing the transistor's current drive capability and thus slows down when charging any load capacitance on its output node). Also, as the number of inputs on the SCL gate rises, the delay in response to changes on the lower differential pair increases, and this is due to the extra load that must be switched through each switching level above the input in question.

Below is an analysis of the electrical characteristics of the SCL gate structure. This has been broken down into the reasonably well understood DC behaviour, as well as the less well analysed AC and transient behaviour.

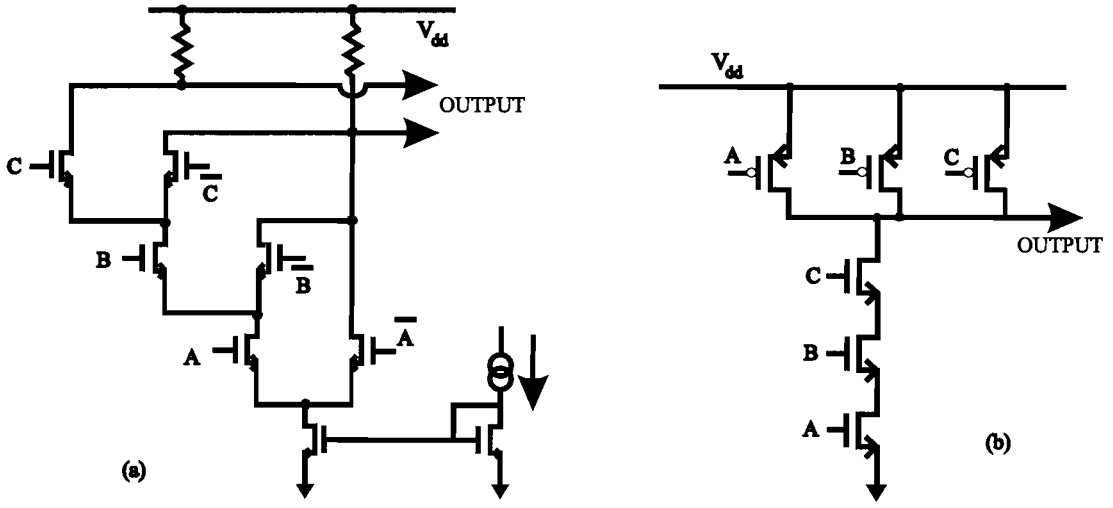

FIGURE 2.1: Schematic diagram a SCL NAND gate (a) and a CMOS NAND gate (b).

2.2.1 DC behaviour

The properties and operation of a SCL gate can be illustrated with the help of an inverter. The schematic in figure 2.2 shows such an inverter with a capacitive load that represents another stage.

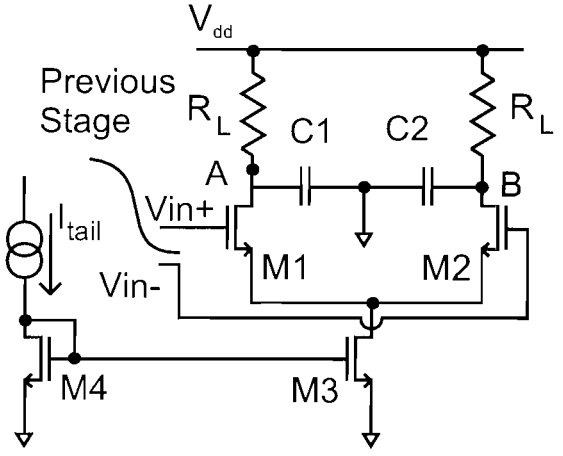

FIGURE 2.2: Schematic diagram a SCL Inverter.

When the inputs are fully switched, the outputs are either V_{dd} (High) or $V_{dd} - I_{TAIL} \cdot R_L$, both with respect to $0V$. From the simple saturation drain current expression for a single NMOS transistor:

$$I_{DS} = \frac{\mu_N \cdot C_{OX} \cdot W}{2 \cdot n \cdot L}(V_{GS} - V_T)^2, \tag{2.1}$$

we see that the minimum differential input voltage (ΔV_{in}) required to switch the outputs is given below:

$$I_{TAIL} \cdot R_L > \Delta V_{in} = \sqrt{\frac{2 \cdot n \cdot I_{TAIL} \cdot L}{\mu_N \cdot C_{ox} \cdot W}}. \tag{2.2}$$

where W and L are the gate width and length of each MOS device in the long tail pair, respectively. C_{ox} is the gate capacitance per unit area, n is the subthreshold slope, whilst μ_N is the mobility in an n-channel. So, by increasing the aspect ratio of the device (and the g_m), there is an increase in the DC gain and hence the differential inverter can switch earlier. Unfortunately, high sensitivity can be a penalty too, where erroneous operation can hamper decisions further along, even though signals propagate using differential signalling.

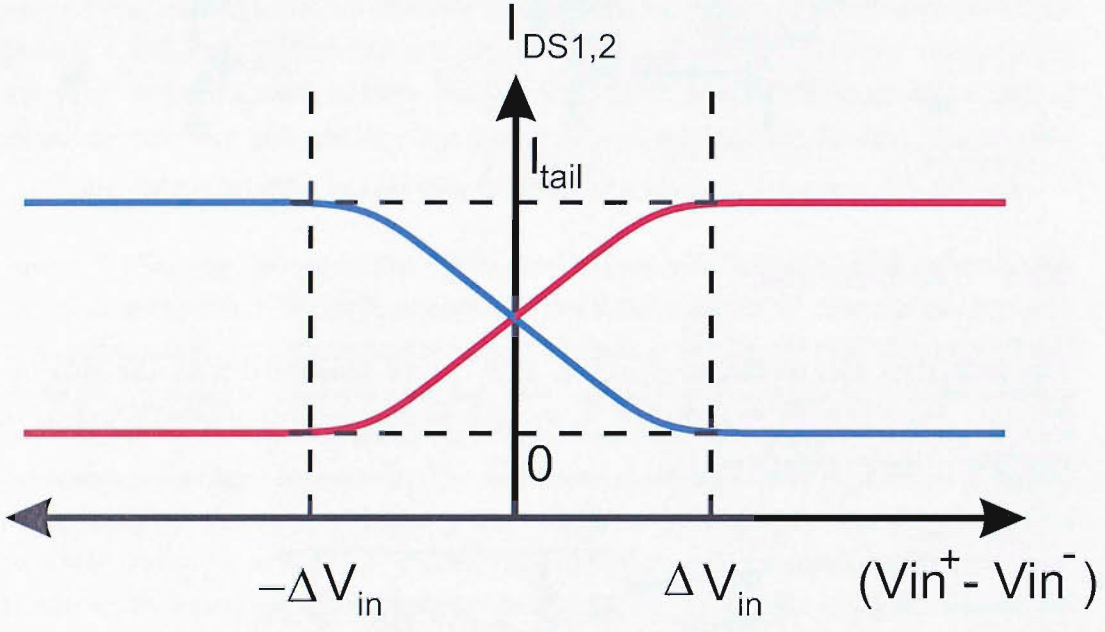

FIGURE 2.3: DC characteristcs of an SCL inverter.

An interesting point to note is that the output swing must be limited so that the MOS devices do not operate in triode region, leading to a higher capacitance seen on its gate terminal. When the active devices M1 or M2 have enough drain-source bias to allow them to operate as current sources, the maximum current is sunk through one MOS transistor or the other (ideally) and the load capacitances (C1 and C2) are charged at the maximum rate. Where there is a high output swing, there will be less voltage dropped across either M1 or M2 (determined by the transisitor whose gate terminal is at V_{DD}), and assuming the same gate bias, the device that is ON may be in triode region, pulling less current to charge the output capacitive loads, and thus a lower rate of change in output voltage. Devices operating in the triode region also tend to exhibit a higher capacitance on their gate terminals, which would present a problem for a previous stage if it were to operate at high frequencies.

In the case of a ring oscillator [2], identical SCL-inverters are DC coupled where the common-mode output voltage can exist in a certain range. If this voltage is too high, then the result is that our current source experiences a high drain bias along with the

second order effects such as higher drain conductance that accompany it. The result is a current source that behaves less than ideal and the differential pair's common-mode rejection deteriorates with the likelihood of a false metastable state. A low common-mode input voltage would force the current source to operate with a lower internal impedance, giving exactly the same problem as before.

2.2.2 Transient behaviour

A transient analysis of the circuit is capable of giving a time domain response at a discrete excitation frequency. Any large signal distortion can be monitored in this analysis, but to find the top speed of the circuit is usually an iterative task. Whilst small-signal AC analysis sweeps through a range of frequencies, linearising the divider does not make sense as the upper sections of the circuit never operate linearly (for the the majority of time relative to an output period,) with true operation dictating one half be energised and the other lying dormant. Another problem with the AC method being applied directly to a frequency divider is that it is hard to tell whether it is functioning, as the output is half the frequency of the input. By taking a FFT of the transient analysis (one period), will lead to a frequency domain representation of the output, but this does

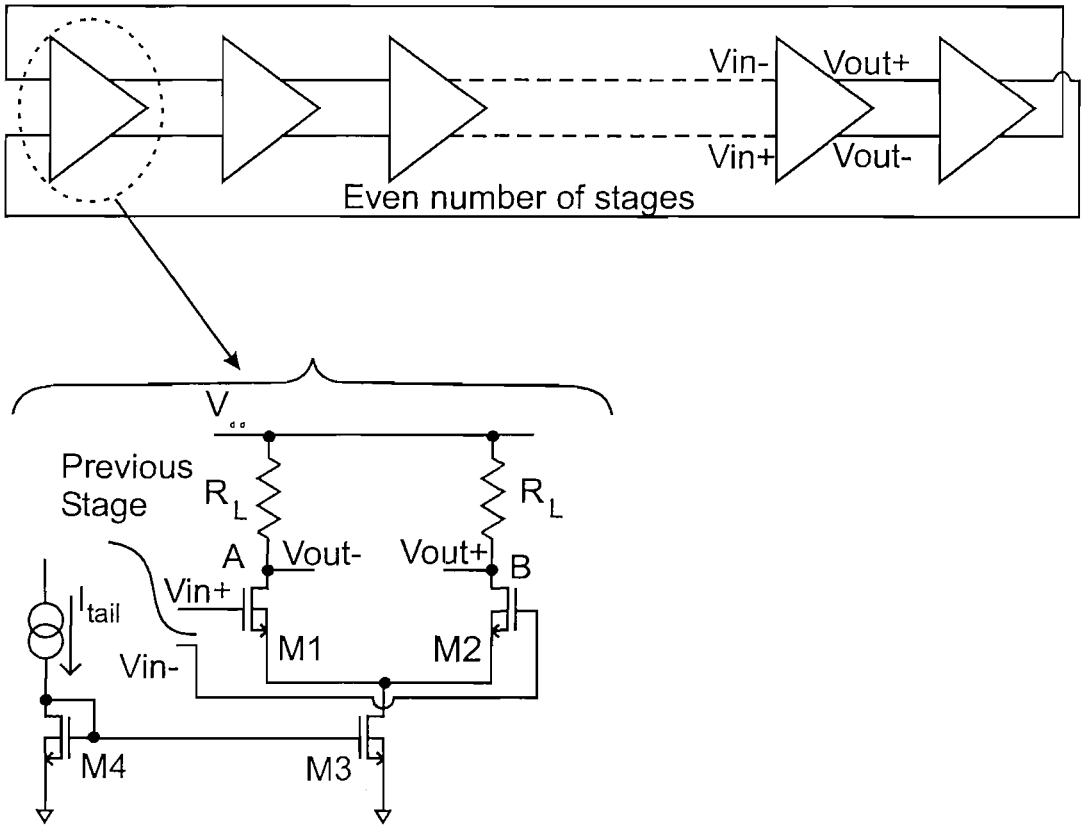

FIGURE 2.4: Schematic of the ring oscillator circuit with an even number of stages, each with a SCL inverter.

not solve the problem of finding the top speed with respect to certain design parameters.

In order to understand how and where extra speed may be obtained in our divider, a ring oscillator has also been examined with a binary number of stages (Fig. 2.4). By having a ring oscillator, one can find the time it takes for a gate to respond to an input. Each stage consists of a differential inverter sitting on a MOS transistor current source (body-tied NMOSFET, W=20μm, L=0.5μm, M3 in the diagram.) The differential inverters are appropriately cascaded to give a ring oscillator, toggling without the need for a level shifter between two successive stages. Four circuits (4, 8, 16, 32 stages) have been simulated (transient) simultaneously with a long enough run to allow them to reach their steady state operation. A number of simulations have been performed with various permutations and combinations of the parameters associated with each differential stage (supply voltage, bias current, load resistance, channel length, channel width.)

In the last few years, investigations on these oscillators have been reported [4] [5] [8], where the emphasis has been on their design in sub-micron technologies. Based on the aforementioned work, a simple expression has been formulated to predict the trend in operating characteristics, as well as predicting the actual time delay. Equation (2.3) calculates the oscillation frequency of the ring oscillator:

$$f_{osc\_smr} = \frac{1}{2 \cdot N \cdot \tau_{PD}} \tag{2.3}$$

where N denotes the number of stages. The propagation delay of each stage is given by τ_{PD} and is expressed as:

$$\tau_{PD} = R_L \cdot (C_{gs} + C_M + C_{gb} + C_{bd}) \cdot ln(2) \tag{2.4}$$

with R_L being the load resistance, C_{gs} being the 'gate-source' capacitance, C_M being the 'Miller' capacitance across the 'gate-drain' terminals, C_{gb} being the 'gate-body' capacitance and C_{bd} accounting for the 'body-drain' capacitance. These terms are modelled with expressions of varying complexity. This first order model is derived with the help of Fig. 2.5. No small-signal transconductor (and hence no high frequency zero [5]) is included in the expression owing to it being a simple a first order equation. The factor '$ln(2)$' refers to the step input applied, which would be around 0.8 should a ramp input be applied (value based on the ramp time being twice the propagation delay of the inverter). These two numerical factors deserve a little more explanation to help any reader appreciate their origin. The following analysis is done with respect to the oscillation frequency of a ring oscillator.

The factor relating to a step input at the input of the RC integrating circuit (as shown

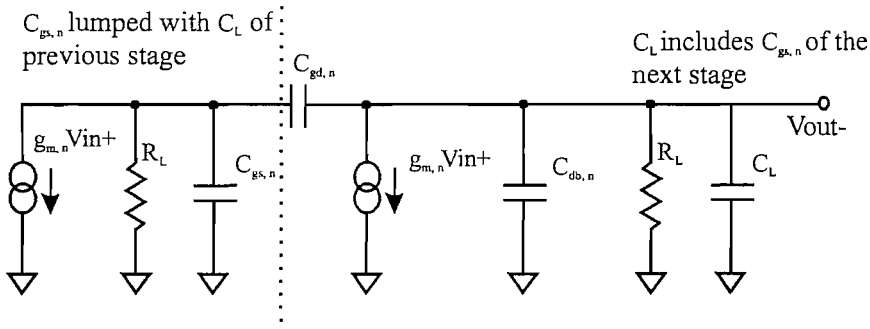

FIGURE 2.5: Schematic showing the half circuit of a SCL inverter, such as the one shown in Fig. 2.4.Note that the analysis concentrates on the network to the right of the dashed line.

in figure 2.6(a) with 'R' referring to R_L), starts with the integral:

$$V_{OUT}(t) = \frac{1}{C} \cdot \int i(t)\mathrm{d}t \tag{2.5}$$

or in its differential form:

$$i(t) = \frac{\mathrm{d}V_{OUT}(t)}{\mathrm{d}t} \cdot C = \frac{V_{IN}(t) - V_{OUT}(t)}{R} \tag{2.6}$$

and after separating the variables:

$$R \cdot C \cdot \int_0^{V_{OUT}} \frac{1}{\mathrm{d}V_{IN}(t) - \mathrm{d}V_{OUT}(t)} \mathrm{d}V_{OUT} = \int_0^t \mathrm{d}t \tag{2.7}$$

$$-ln(V_{IN}(t) - V_{OUT}(t)) + ln(V_{IN}(t)) = \frac{t}{R \cdot C}. \tag{2.8}$$

By assuming the following conditions:

- at t<0, $V_{IN}(t) = 0$,
- and t≥0, $V_{IN}(t) = V$

and that we are interested in the time taken,τ, for the output of our RC network, V_{OUT}, to reach half the final value, V, we get:

$$R \cdot C \cdot ln\left(\frac{V}{V - V_{OUT}}\right) = t \tag{2.9}$$

$$R \cdot C \cdot ln(2) = \tau. \tag{2.10}$$

With the help of figure 2.3, it should be clear that we have presumed the differential outputs of the 'observed' gain stage to have toggled at the 'crossover' point, and hence would be interested in the time taken to reach that point.

The second choice of factor denotes the use of a slow ramp input in the model of the RC

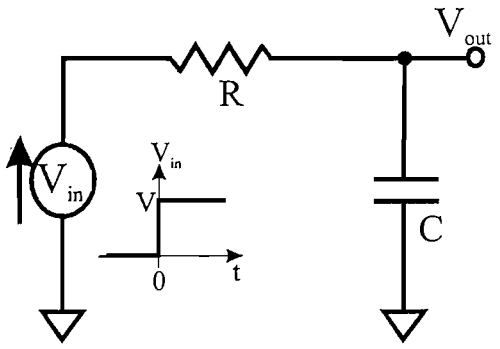

(a) A simple RC circuit with a step input.

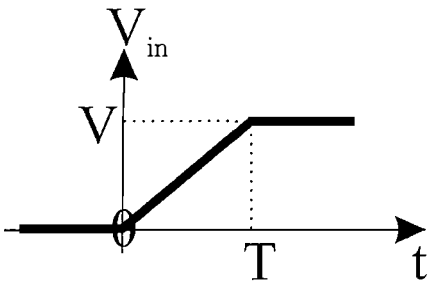

(b) Specification for the ramp input.

FIGURE 2.6: Two different input conditions.

network on any output of a gain stage. Naturally, one must be specific on the ramp's specification, namely its rate of change as shown in Fig. 2.6(b). If the rate is too quick, then the time delay tends towards the value observed for the step input. If the rate is too slow, then no useful delay can be measured. The factor '0.8', arises from the assumption that the input stops ramping after a time equal to twice the propagation delay of each inverter. This value is chosen to determine the time taken between the instant when the inverter's input toggles through half the logic swing, and the moment the output toggles through half the logic swing in the opposite direction [4]. This delay calculated should not include the RC delay on the output node of each inverter stage.

The capacitance representing the load on the output of the inverter is the 'gate-source' capacitance of one of the differential pair transistors (in saturation) on the following stage and is given by:

$$C_{gs} = \frac{2}{3} \cdot W \cdot L \cdot C_{ox} \tag{2.11}$$

where C_{ox} is the gate oxide capacitance.

The Miller capacitance is represented as the following:

$$C_M = C_{gdo} \cdot W \cdot \left(1 + cos\left(\frac{\pi}{N}\right)\right) \tag{2.12}$$

where C_{gdo} is the 'gate-drain overlap' capacitance. The factor, $(1 + cos(\frac{\pi}{N}))$, describes the voltages at the input and output of each inverter moving in opposite directions, and as N increases, this capacitance tends to a value twice the overlap model parameter [8]. This is in fact a crude approximation to the inverter chain's true time-domain behaviour with the most inaccuracy when there are fewer inverter stages. When N is large, the Miller capacitance is twice the gate-drain overlap capacitance of the active device in the half circuit. When N tends to a smaller integer, the phase difference between the

input and output nodes of each inverter stage will become less than π rads. Therefore, the voltage across the gate-drain overlap capacitance will reduce too, becoming lower than 2 times the input voltage as the number of stages drops to 1. The term $cos(\frac{\pi}{N})$ models this reduction in voltage, by reducing the value of the Miller capacitance instead. However, one must be careful not to confuse the $cos(\frac{\pi}{N})$ with the low frequency gain of the inverter stage.

A very crude representation of the $gate - body$ capacitance in saturation is given as follows:

$$C_{gb} = \frac{1}{5} \cdot W \cdot L \cdot C_{ox} \tag{2.13}$$

and is similar to the gate-source capacitance. However the difference is that when in operation, the saturated region tends to detach itself from the drain implant and a small capacitance exists between the gate and the body region of the transistor. This region consists of the gate oxide followed by a depletion region in the body. The fractional term is found to be between $\frac{1}{10}$ and $\frac{1}{5}$, though this is through a 'curve-fitting' approach as opposed to any physical origin. Therefore, without any citation or analytical proof, it cannot be relied upon and naturally needs to be studied.

The last capacitance is taken from [1] and is very dependent on the transistor structure and process. The 'body-drain' capacitance is present because of the depletion region between the heavily doped drain implant and the low doped body. In the case of the NMOS transistor, the drain voltage is usually higher than the body potential, though it is less so in body-tied PDSOI technology. For the purpose of this work, the following expression is used (as defined in the compact model [6]):

$$C_{bd} = \frac{C_J \cdot W \cdot t_{film}}{(1 - FC)^{1+m_{grad}}} \left(1 + FC(1 + m_{grad}) + \frac{V_{dd} - I_{tail} \cdot R_L - (V_{dd} - 1.3)}{PB} \cdot m_{grad} \right). \tag{2.14}$$

where t_{film} is the silicon film thickness, FC is the 'junction forward bias coefficient' (default set to 0.5), PB is the junction potential (default set to 0.8V), V_{dd} is the supply voltage of the circuit, and I_{tail} is the tail current in each stage. The m_{grad} term is the grading coefficient of the junction. Two values exist: 0.33 for a graded junction and 0.5 for an abrupt junction. In PDSOI technology, the drain/source implants reach the buried oxide and hence no depletion region exists here. However, the sidewall of the drain implant has a graded junction and hence a value of 0.33 is used. In bulk CMOS technologies, the abrupt junction would also have to be accounted for in the expression owing to the lack of a buried oxide. It should also be stated that the value of 1.3 in the term '$(V_{dd} - 1.3)$', refers to the common source voltage of the differential pair (with respect to circuit ground), and is specific to the chosen V_{dd}. The last term in the

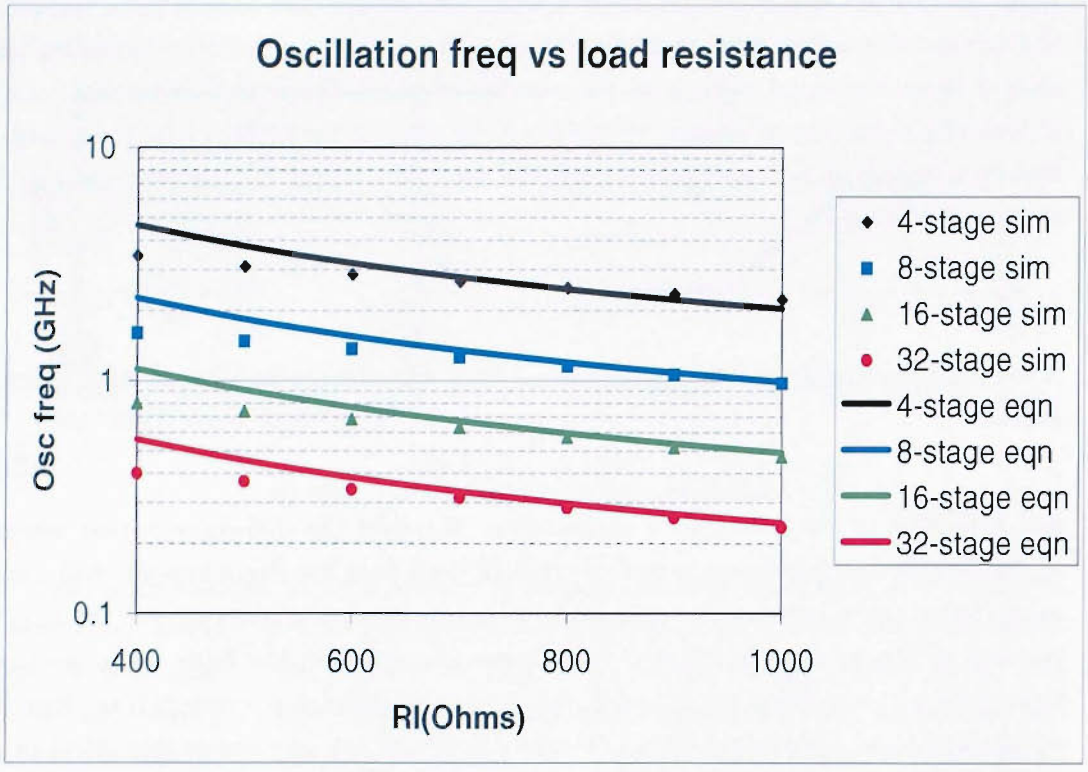

FIGURE 2.7: Simulated oscillation frequency vs load resistances, for various number of stages. (Ib=1mA, W=40μm, L=0.35μm, V_{dd}=3.3V)

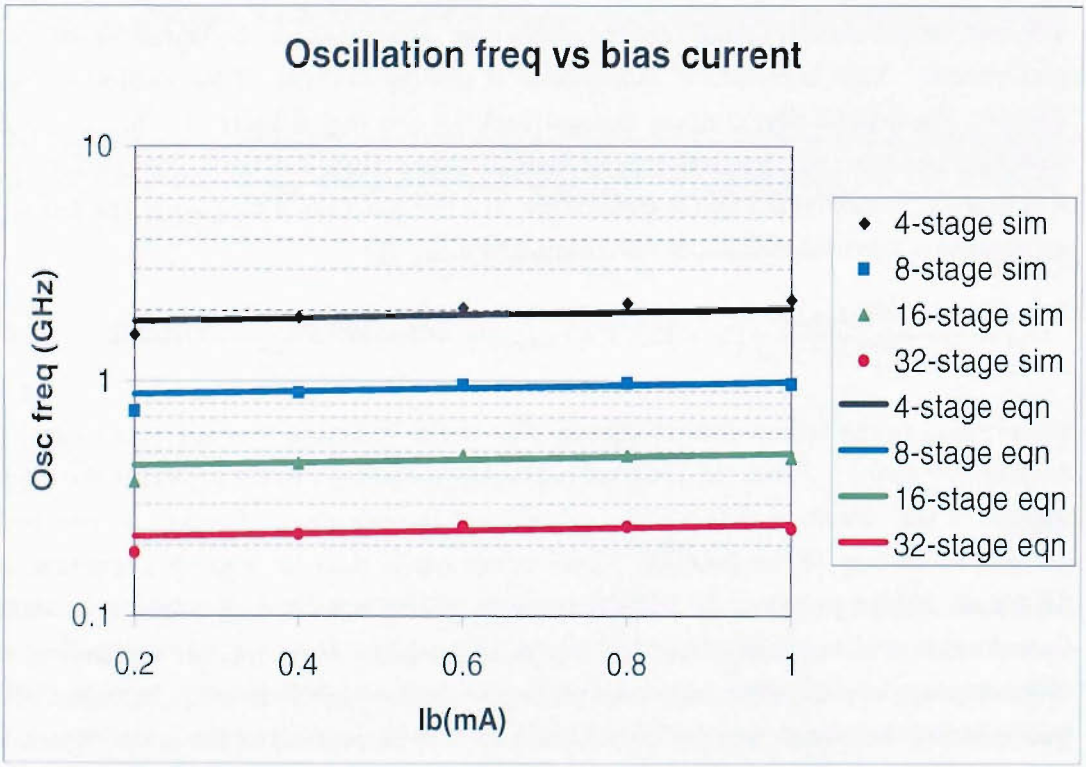

FIGURE 2.8: Simulated oscillation frequency vs bias currents, for various number of stages. (Rl=1kΩ, W=40μm, L=0.35μm, V_{dd}=3.3V)

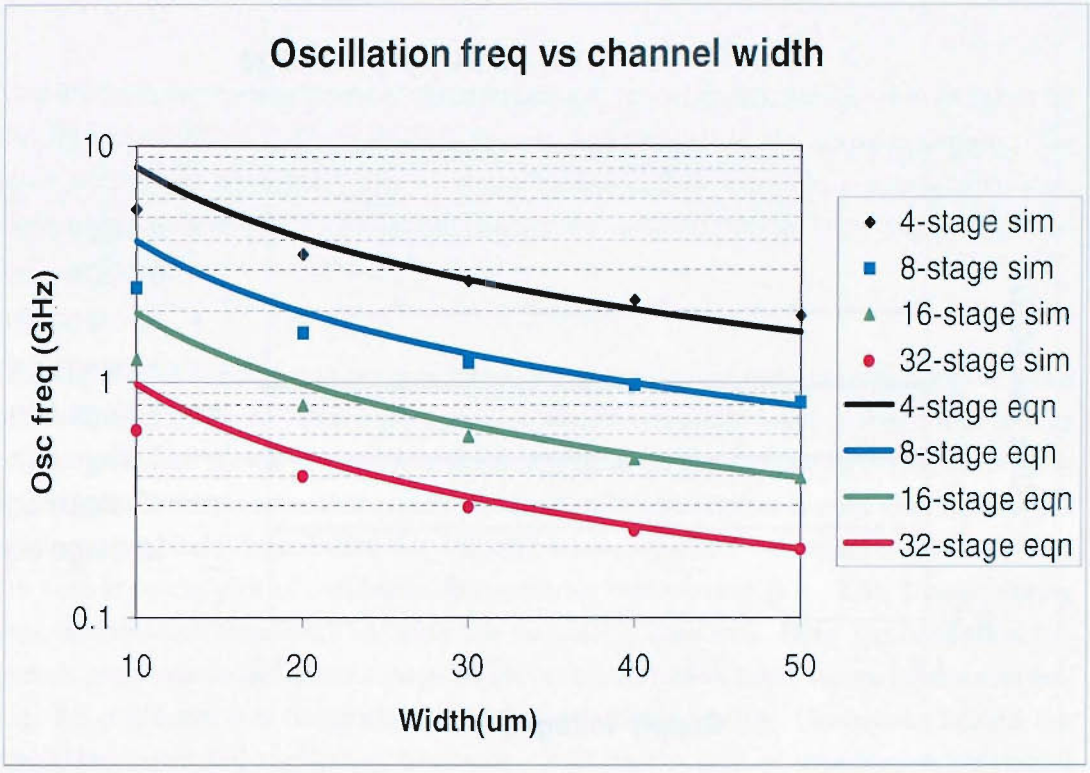

FIGURE 2.9: Simulated oscillation frequency vs differential pair channel widths, for various number of stages. (Ib=1mA, Rl=1kΩ, L=0.35μm, V_{dd}=3.3V)

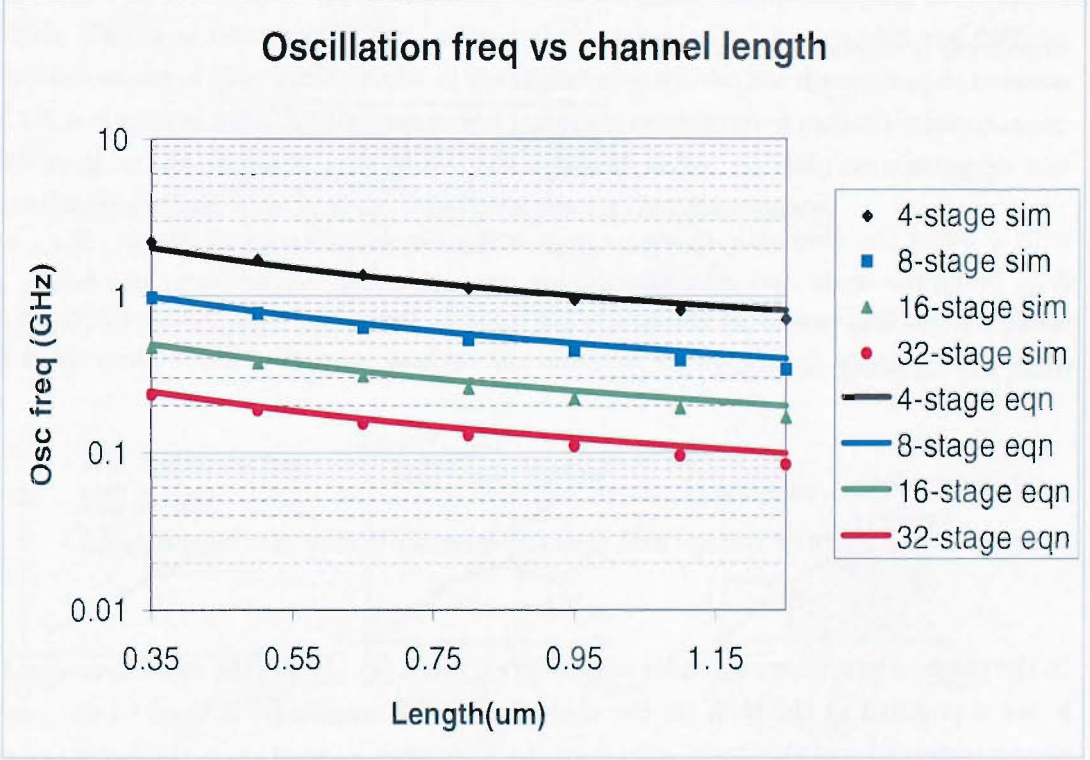

FIGURE 2.10: Simulated oscillation frequency vs differential pair channel widths, for various number of stages.(Ib=1mA, Rl=1kΩ, W=40μm, V_{dd}=3.3V)

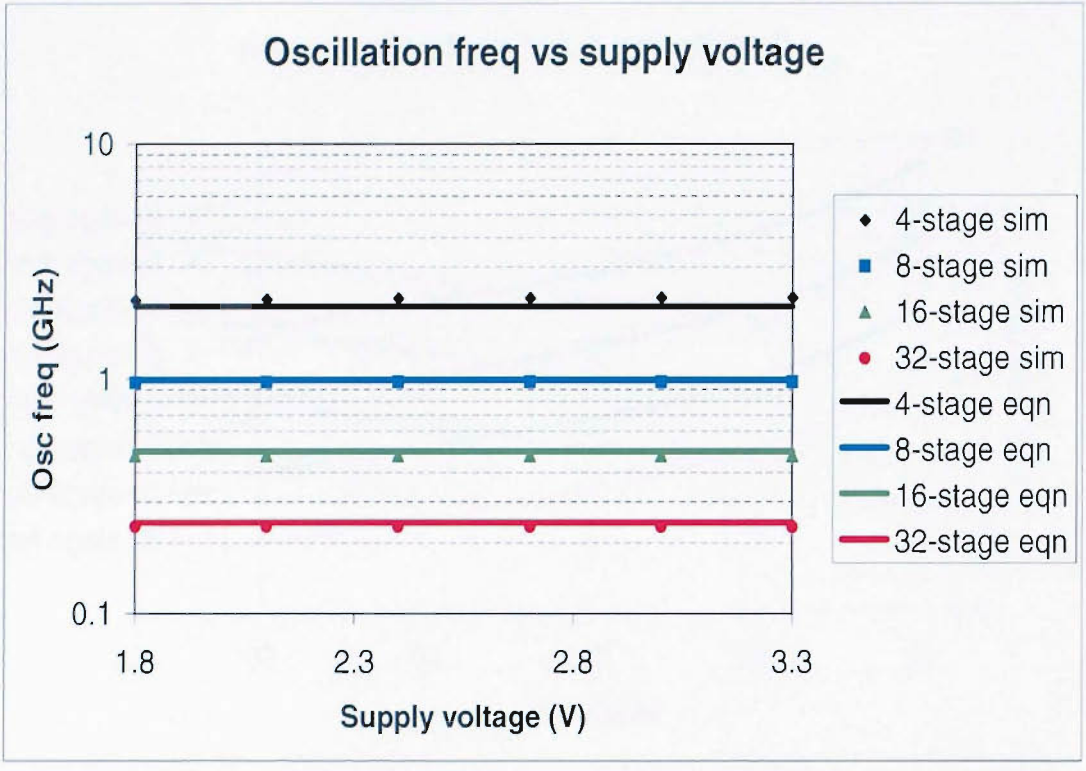

FIGURE 2.11: Simulated oscillation frequency vs supply voltage, for various number of stages. (Ib=1mA, Rl=1kΩ, W=40μm, L=0.35μm)

expression is expanded as follows [1]:

$$C_J = \sqrt{\frac{q \cdot \epsilon_o \cdot \epsilon_{rsi} \cdot (N_{plus} \cdot N_{sub})}{2 \cdot \phi_{bi} \cdot (N_{plus} + N_{sub})}}, \tag{2.15}$$

with q being the electronic charge, $\epsilon_o \cdot \epsilon_{rsi}$ being the permittivity of silicon, N_{plus} and N_{sub} being the drain and substrate doping concentrations, respectively, and finally ϕ_{bi} being the 'built-in potential' across the p-n junction. This last variable can be evaluated using the following equation [1]:

$$\phi_{bi} = \phi_t \cdot ln\left(\frac{(N_{plus} \cdot N_{sub})}{{n_i}^2}\right), \tag{2.16}$$

where ϕ_t is the thermal voltage and n_i is the intrinsic carrier concentration.

In the publications where a similar expression is given [4], [5], [8], the dependence on V_{db} is not a problem as the bulk (in the case of a NMOS transistor) is fixed to the lowest circuit potential and the drain voltage of the transistor in question is the output of the inverter. For body-tied SOI, the body region of the transistor is shorted to the source terminal of the transistor, requiring the user of this expression to be a little intuitive on the value to choose.

In order to show the usefulness of these equations, model parameters have been taken for the *Honeywell SSEC 0.35μm PDSOI process*, and applied to the above equations. The same parameters have been used with the STAG CMOS model in a circuit simulator, allowing a comparison of oscillation frequencies against various large signal electrical and physical parameters (Fig. 2.7 - 2.11).

Looking at all five plots, one can conclude that, the evaluated expression follows the trend of the simulations well. There are instances where aymptotic behaviour is witnessed, as well as predicted curves intersecting simulated data. In the age of short-channel devices, it is even more impressive that the zero-field mobility parameter is used with no account of the vertical field degradation nor velocity saturation. The only instance where this is not true is in the plot of oscillation frequency vs bias current (Fig. 2.8), though saying this, the equation does seem to follow the simulation data well. Here, the oscillation frequency responds to the linear change in bias current, but at lower values of bias current, Fig. 2.8 concludes it is obviously not modelled by the equation. One reason behind the lower-than-predicted oscillation frequency could be the lack of capacitance associated with the MOS current source transistor and is not accounted for in the equation.

Another concern is the over-estimation of the frequency when stepping the channel width. The error between the simulation and the equation is between 50% and 60% for very low values of gate width, whilst at the higher gate widths, the discrepancy is between $< 1\%$ and around 14%. A thorough model including neighbouring parasitic capacitances, dominant at low values of gate width, has not been given. As well, accounting for any mobility variations may go some way to narrowing this discrepancy.

It is useful to analyse the behaviour of the inverters and its corresponding ring oscillator at a schematic level, as it helps confirm the observed trends as well as guide a designer

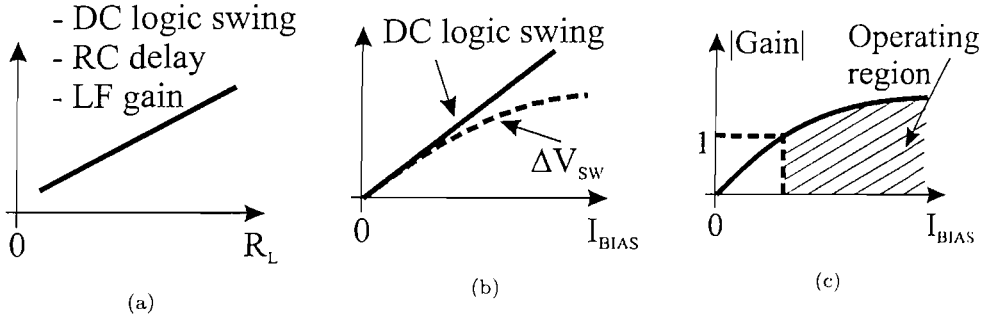

FIGURE 2.12: Expected trends for the logic swing, low-frequency gain and RC delay, against load resistance and bias current.

to a more optimum design.

Figure 2.7 shows the oscillation frequency against load resistance. As R_L increases, the small-signal gain, RC (lumped) time constant on the output of every inverter and the DC logic swing, all rise linearly. The small-signal transconductance of each stage, switching threshold range and the input capacitance seen on the input of every stage all remain steady. If it is assumed that no MOS device in the inverter stages enters the triode region, then the delay (associated with the time constant) is proportional to R_L, as in Fig. 2.12(a).

Setting R_L back to its default value, a trend is observed for the bias current I_{bias}. With increasing I_{bias}, we would only expect the DC logic level to rise linearly, with increases in ΔV_{sw} and g_m (and thus gain), both rising proportionally to $\sqrt{I_{bias}}$ (Figs. 2.12(b) and 2.12(c)). Again, the lumped capacitance on each inverter stage remains the same (strictly speakng, the Miller capacitance would rise, though this is not shown in our equation). Though not as conclusive, the logic swing rises at a faster rate than the threshold switching range and the small increase in oscillation frequency seen in Fig. 2.8

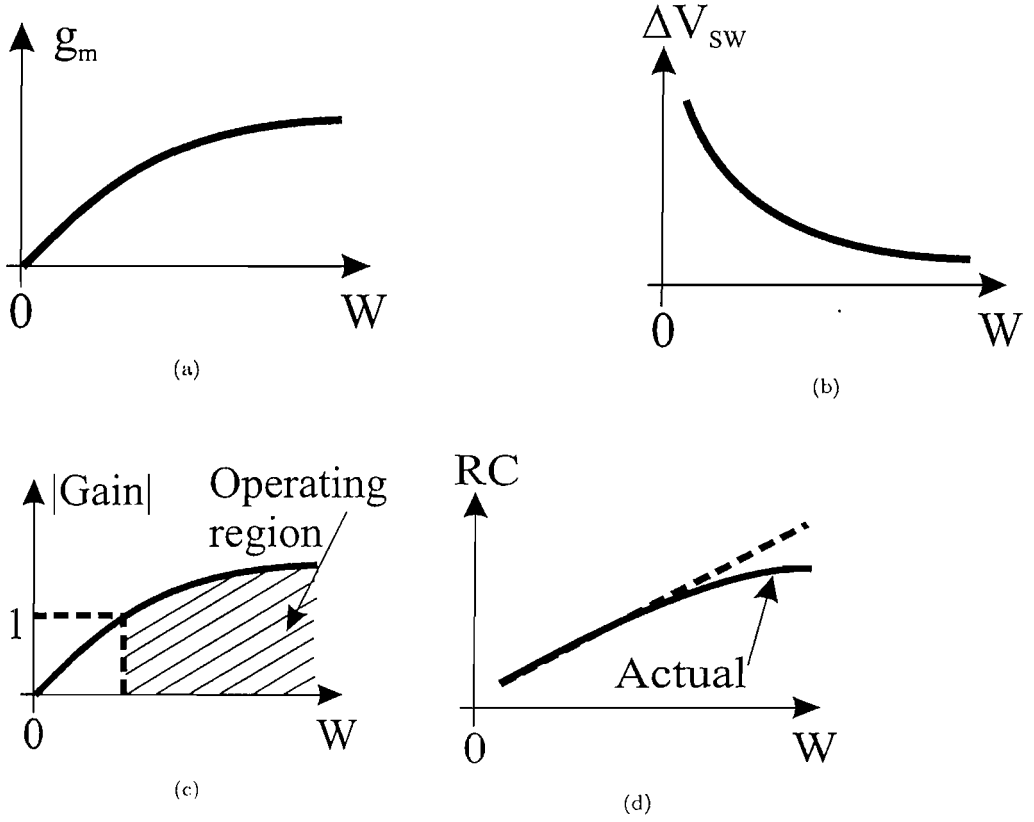

FIGURE 2.13: Expected trends for transconductance, switching range, low-frequency gain and RC delay, as a function of channel width.

can be attributed to the larger current drive of the gain stage.

The importance in distinguishing between the two lateral dimensions associated with the channel is underscored in Figs. 2.9 and 2.10. As the channel width is increased, the small-signal transconductance of the inverter (and thus its small-signal gain) will increase at a rate proportional to $\sqrt{W}$ (Figs. 2.13(a) and 2.13(c)), whilst the range of switching thresholds would decrease at a rate proportional to $\sqrt{\frac{1}{W}}$ (Fig. 2.13(b)). The lumped capacitance on the input of every inverter stage, would rises linearly too. This latter mechanism is dominant, as seen in Fig. 2.9 (albeit the inverse of Fig. 2.13(d)), but the increased drive resulting from the increase proportional to $\sqrt{W}$ in small-signal transconductance reduces the delay in transitioning between logic levels. This goes someway to explaining the discrepancy between the circuit simulation and evaluation of the analytical expression, equation (2.4), therefore justifying the need to include a g_m term in the expression.

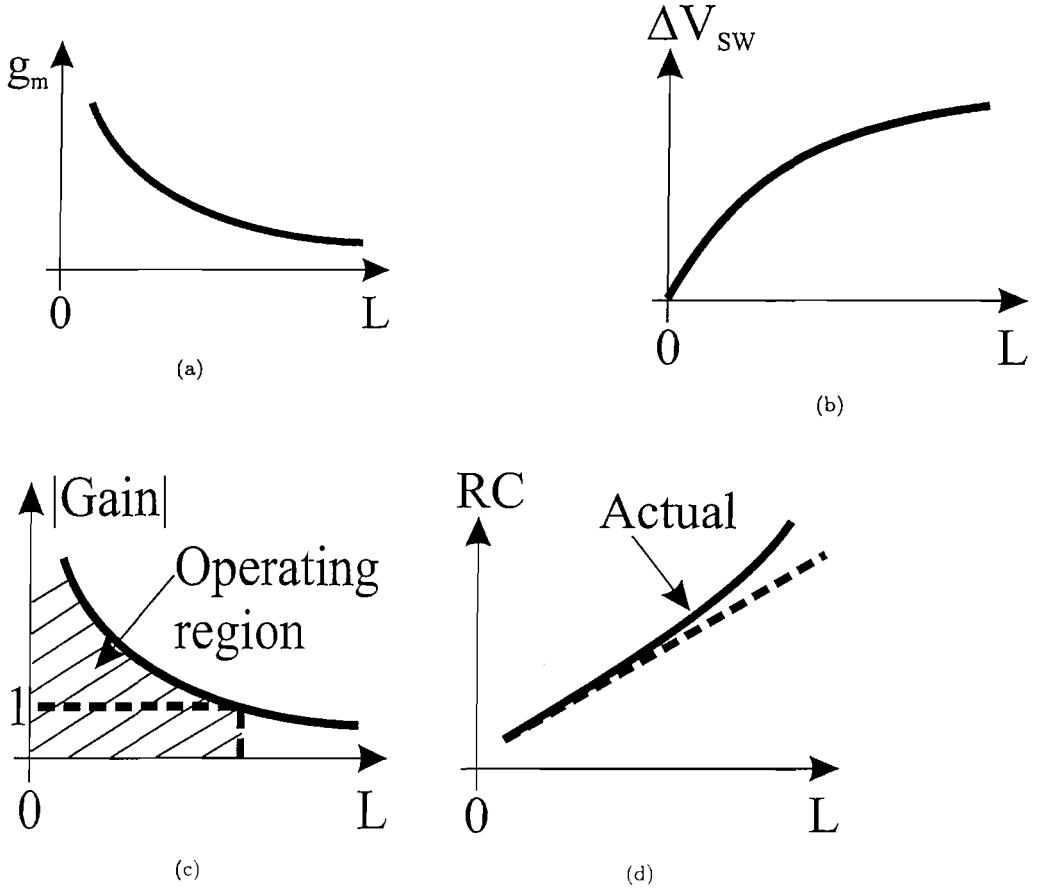

FIGURE 2.14: Expected trends for transconductance, switching range, low-frequency gain and RC delay, as a function of channel length.

The scaling with the second dimension, namely the channel length, has a few subtle differences compared with the channel width. By increasing the length, the small-signal transconductance will decrease at a rate proportional to $\sqrt{\frac{1}{L}}$ (Fig. 2.14(a)), whilst the switching threshold range increases at a rate proportional to $\sqrt{L}$ (Fig. 2.14(b)). Again, any increase in channel length would give a linear increase in every inverter's input capacitance, and thus the delay per stage. As in the case of channel width, a discrepancy exists between the evaluated analytical expression, equation (2.4), and the circuit simulation; specifically, an overestimate of the ring's oscillation frequency at large values of channel length. Although the major influence on the delay is the capacitance, the reduction in small-signal transconductance of the differential pair reduces its current drive, leading to a lower oscillation frequency than predicted by the analytical expression (Fig. 2.14(d)).

The remaining evaluation shows the oscillation frequency plotted against the common power supply of the entire ring oscillator. Fig. 2.11 showsthe results from an evaluation of the analytical expression and circuit simulation with varying power supply voltage. A close observation indicates that at high toggle frequencies, the expression fails to predict a higher oscillation frequency with increasing supply voltage. The only reference to the parameter V_{dd}, is in equation (2.14). As the figure demonstrates, it plays a superficial role, with negligible effect on the plot. Unlike CMOS topology inverters, once the tail current source below the differential pair has a voltage exceeding its compliance any variation in the supply voltage should not, theoretically (excluding channel-length modulation), affect the biasing of the devices (provided that the variation does not push any device in to triode region). Realistically though, increasing the bias across the drain-source terminals of a MOSFET device will result in deviation from the perfect DC model of a MOSFET, especially in short-channel devices. Increased supply voltage translates to an increased average bias across the active devices, inducing effects such as impact ionisation. Coupled with the effect of channel length modulation, an increased drain-source bias is likely to result in a higher DC drain current, justifying why the simulation predicts a higher toggle frequency than the analytical equation.

Remaining with the graphs, an important point to note is that it is not possible to extrapolate the equations at either end without it losing its credibility. The ring oscillator is a feedback circuit, and for it to sustain controlled oscillation, there must be an adequate loop gain (and at least 2π rad phase difference between an arbitrary starting point and the return signal). Although the equations may predict a faster oscillator, the actual result will be a decaying oscillation as the small signal gain, $g_m \cdot R$, falls to unity and below.

As well as adequate loop gain, every stage in a high frequency oscillator design should

have a low RC product on its output node. The capacitive load lumps together the output-input capacitances on an output node between two successive inverter stages. There must also be adequate current drive by which to charge and discharge those capacitances. Unfortunately, this implies a larger current through the load resistor, leading to a larger output swing and thus conflicting with the bias of the differential pairs in each stage. Of course, in SCL logic designs, one has the option to reduce the load resistance, but this will conflict with the small-signal gain of each stage, thus affecting its switching speed. The other option is to raise the power supply to accomodate the larger output swing, but then there is the question of increased power consumption and technology limitations. The means to achieve a higher current drive can result in increased capacitance on each inverter stage's input node, reducing the expected gain in oscillation frequency.

Naturally, from these graphs, this VCO (actually current-controlled oscillator) has had its oscillation frequency predicted to a certain degree, depending on the circuit parameter found to be crucial in the design. Of course, at sub-micron level, this equation is no substitute for transistor circuit simulation but does show the trend as well as giving a starting point for its optimisation.

2.3 Example circuits

Following on from the SCL inverter ring oscillators, this section reports on some preliminary simulations of a stacked NAND-gate-based ring oscillator. In this study, the stacking capabilities of SOI transistors are highlighted but there is no discussion (analytically) of its speed problem. Another motivation for these simulations is to do with the 'zero-detect' set of NAND gates in an initial choice of dual-modulus architecture [7].

This implementation consists of a string of flip-flops in an asynchronous configuration. An *end-of-count* NAND gate detector signals the end of a binary count, triggering a flip-flop to block the next incoming cycle (see Fig. 3.8). This process repeats itself thereafter unless a different integer modulus is required. One important point about this divider is that the pulse counter does not signal when all 0's or all 1's are detected when counting up or down respectively (in the case of a divide by 32\33 counting up, the *end-of-count* detector would trigger after the transition from 11111 to 00000, requiring careful design of critical paths). Instead, the design triggers once all 1's or 0's have been detected when counting up or down respectively. This allows the input to the *end-of-count* to accumulate, gradually easing the layout of the final divider.

Unfortunately, a problem with this design in high frequency applications is caused by the stacking nature of the differential pairs, with actual results given in section 3.2.

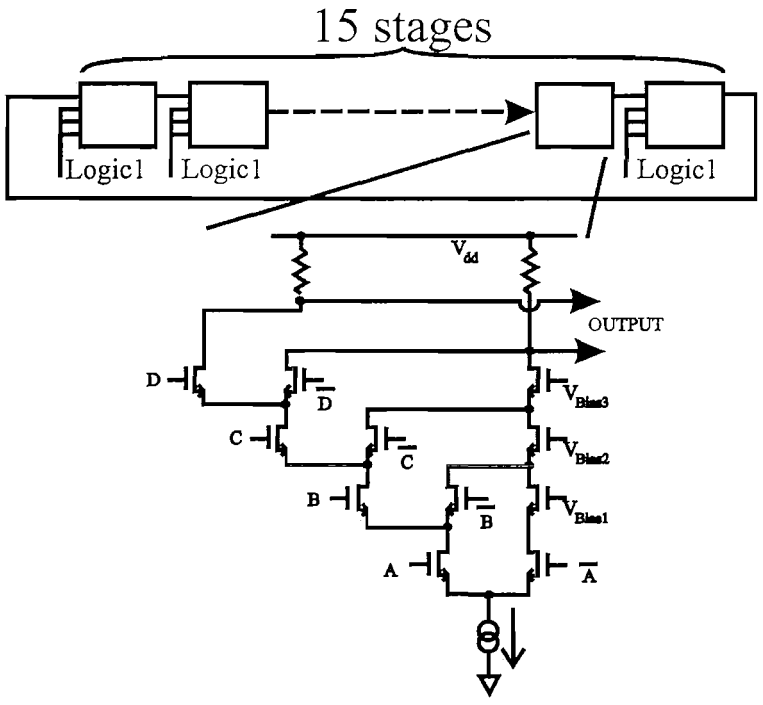

FIGURE 2.15: Ring oscillator test with NAND gate elements. Either A, B, and/or C are fixed or B, C, and/or D are fixed with results given in Fig. 2.16(a) and 2.16(b).

Together with the help of the discussion on stacking divider stages in section 5.1.2, it is noted that in a stacked version of [7], the output of the first divide-by-2 (and also the second highest speed signal) exists as the first output from the stack. All outputs are DC coupled to the NMOS NAND gate which also consists of stacked long tailed pairs. This first output must switch the current between the entire stack of long tailed pairs and cascodes, and this is found to restrict the overall speed of the entire divider. To illustrate this, simulations have been carried out on ring oscillators (Fig. 2.15). The first ring oscillator is built around SCL based inverters, each drawing 1mA from a 3.3V supply. The device dimensions in the differential pair are $\frac{20\mu m}{0.35\mu m}$. The second ring oscillator has unit cells composed of 4 input NAND gates using SCL stacking. Each long tailed pair also has the dimensions $\frac{20\mu m}{0.35\mu m}$. The current consumed in each stack is 1mA drawn from a 3.3V supply. For this implementation, three of the inputs to every NAND gate are anchored to logic 'high'. Both oscillators give a differential output of $\pm$300mV (load resistor of 600Ω) centred around 3V. Looking at Figs. 2.16(a) and 2.16(b), it can be seen that toggling is a lot quicker if the input from a previous stage arrives on the top-most differential pair rather than arriving on the lowest pair. The concept of stacking used within this architecture together with simulation results is presented in section 3.2.

An observation that must be cleared up is that the SCL-inverter ring oscillator appears to be running half as fast as the SCL-NAND ring oscillator (upper toggled inputs.) The explanation for this apparent anomaly is that each inverter is DC coupled to the

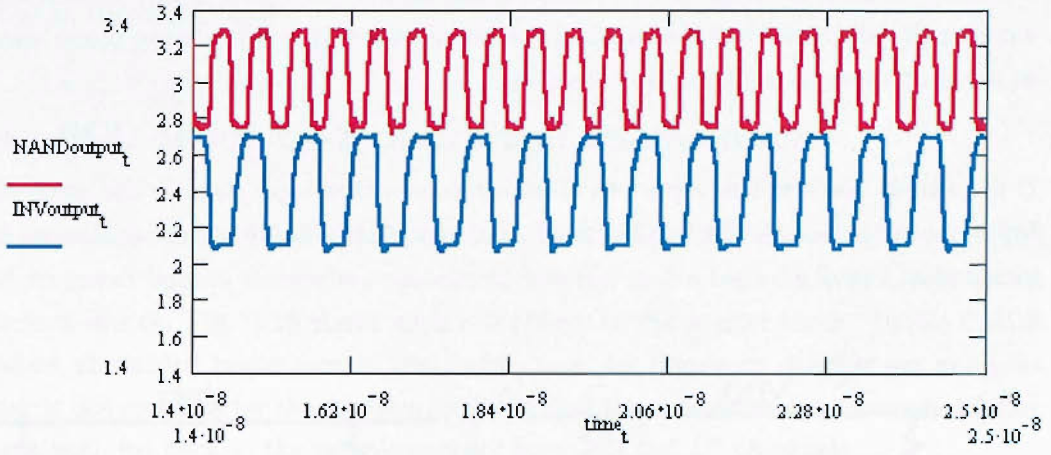

(a) Single-ended output of NAND gate with toggling input set on the highest set of inputs (blue trace refers to a toggling input on the 'D' inputs)

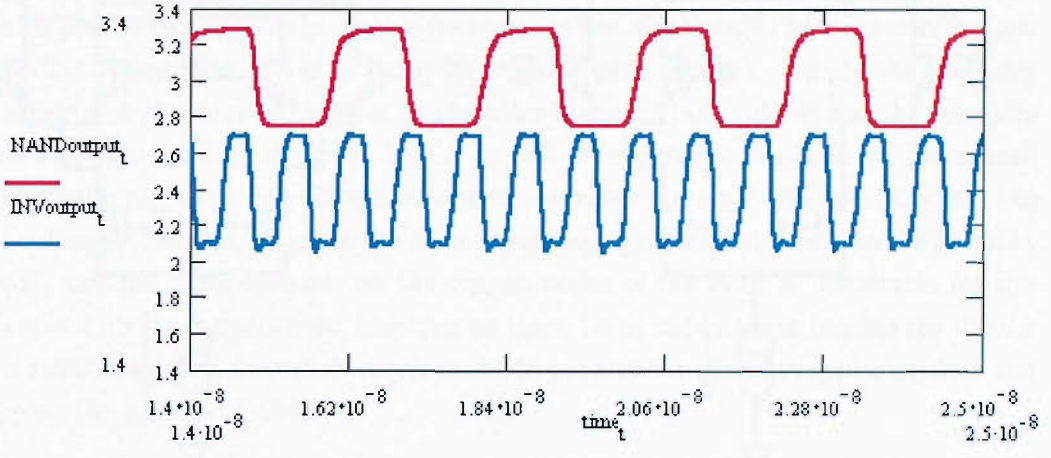

(b) Single-ended ouput of NAND gate with toggling input set on the lowest set of inputs (blue trace refers to a toggling input on the 'A' inputs)

FIGURE 2.16: Simulations of 15 stage ring oscillators. The red trace corresponds to the output from the ring oscillator based on a 4 input SCL NAND gate. The blue traces refer to a ring oscillator based on SCL inverters.

next, whilst the NAND gates are coupled by means of 'perfect' level-shifters. With the admission of using voltage-controlled voltage sources as level-shifters, each stage would be buffered from the next, hence giving an improvement in performance. The inputs arriving on the terminals of the NAND gate are internally level shifted as follows: inputs 'A' and '$\overline{A}$' shifted down 1.2V; inputs 'B' and '$\overline{B}$' shifted down 600mV; inputs 'C' and '$\overline{C}$' shifted down 300mV, and inputs 'D' and '$\overline{D}$' are not shifted at all. The reason behind these level shifters was to simply save time rather than design appropriate source-followers.

Simulations show the SOI divider based on the architecture of reference [7] to produce incorrect division at frequencies of 4.5GHz and above. This poor speed performance in a conventional divider has thus pushed the investigation into considering phase selecting as an alternative to dual-modulus division.

It is a timely place in this report to illustrate some other logic gates in the same circuit topology. As well as the D-type flip-flops, NOR-gates and NXOR-gates are also available, pictured in Fig. 2.17.

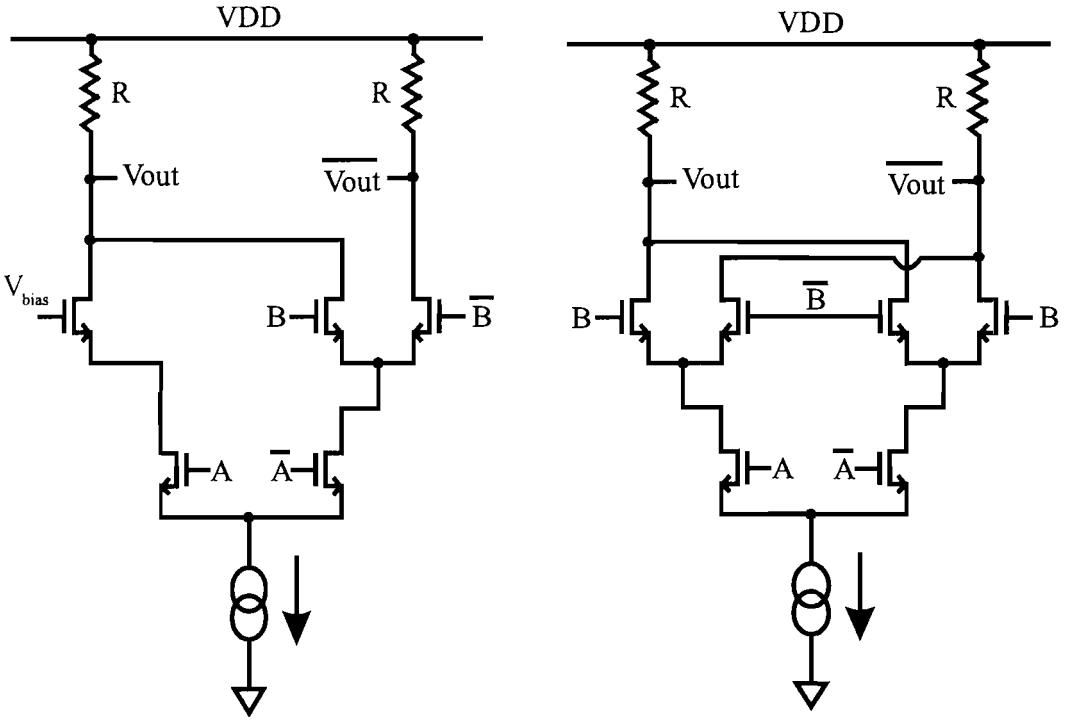

(a) NOR gate.

(b) Exclusive-NOR.

FIGURE 2.17: Other examples of SCL gates.

In the case of the NOR-gate, a 'cascode' transistor has to be placed above the transistor controlled by signal 'A', in order to balance the number of levels of gates below each load resistor. It also prevents 'stressing' the 'A' transistor with a high field across its drain-source terminals, therefore preventing impact ionisation and imbalance in the differential output. Its DC voltage would inevitably be generated by some stable bias network. As in the case of the NAND-gate, the inputs have different common-mode levels and this forces a slight complication in the designer's synthesis of a combinatorial logic solution.

The EXNOR gate is equivalent to a current steering Gilbert Cell mixer. Fig. 2.17(b)

shows a differential circuit where the inputs are applied on different levels of the circuit just as in the NAND gate.

2.4 SCL divide-by-2 circuit and its operation

The master-slave latches in each D-type are built around the above logic style. The level triggered latches themselves contain an inverter and a bistable, which both sit on a current source. Fig. 2.18 shows such a topology in the master block. Unlike CMOS dividers, the added bonus here is that, when used for frequency division, an extra inverter is not required for the positive feedback and the complementary outputs (Q and $\overline{Q}$) are both fed back to the complementary input ($\overline{D}$ and D) terminals.

The operation of the flip-flop is as follows. When the clock is high (according to the diagram), the differential pair senses its inputs and translates them to its output ports with respect to V_{dd}. When the clock signal toggles low, the bistable becomes enabled and stores the output from the inverter in its cross coupled network. During the switching of current from inverter to bistable, the inverter begins to power down and the common-mode rejection ratio of the latch begins to fall. Without the requirement for a small differential input switching voltage, the latch can enter metastability, resulting in a loss of the division pattern. In order for the outputs to be transfered across to the bistable, a small amount of capacitance on the output nodes of the latch is favourable for this high speed divider's operation. Insisting on there being capacitance implies the divider has a minimum input frequency requirement to prevent the divide stage losing data and stopping the process of division.

The dilemma here is that increasing the capacitance to accommodate the low frequencies comes with some penalty as its upper frequency limit decreases too. However, in useful applications such as short range wireless communications, the prescaler is expected to operate in a narrow band (few hundred MHz) and hence optimising this block to operate at the high end of its capacity is usually far more desirable.

The master-slave arrangement modifies the functionality from a level-triggered device to an edge-triggered flip-flop. The inverter of the master and bistable of the slave come to life on one clock edge, whilst the bistable of the master and the inverter of the slave energise on the opposite clock edge. The binary frequency division is attained by passing like outputs to like inputs between the master and slave, and then cross coupling the outputs of the slave to the inputs of the master latch (Fig. 2.18). A slave output typically lags by a period of half an input clock period behind the equivalent master output and, as the latches give out complementary outputs, the overall result is four phases, separated by half an input clock period ($\frac{\pi}{2}$ rad). This will become useful in the

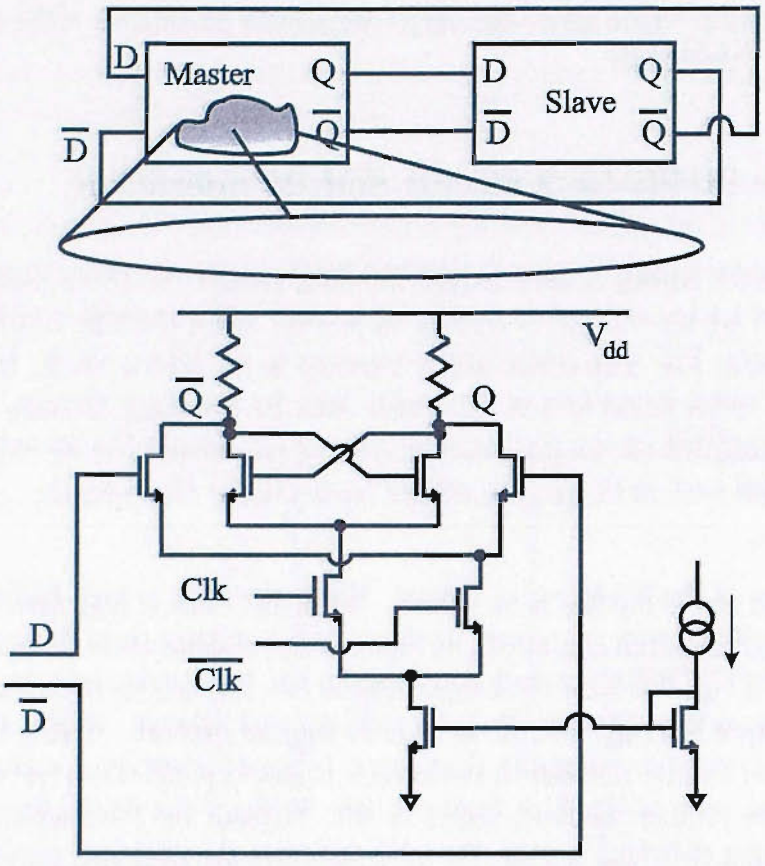

FIGURE 2.18: Schematic diagram a Master-Slave flip-flop.

phase selector section.

When cascading dividers, the differential gains of the inverter and bistable are important parameters that need attention if the divider is to sustain accurate division without any metastable states. The gains of both must be greater than unity, with the inverter remaining stable up to the divider's maximum operating frequency. For a set logic swing (eg, outputs generate a difference no greater than 600mV in magnitude), increasing the current drive by increasing the bias current will in fact reduce the gain (the load resistance value has to come down by the same factor.) Incidentally, a larger output swing may force differential pairs into the triode region, whilst a lower swing reduces the top operating frequency (due to a lower differential transconductance and lower gain). For example, increasing the current by a factor of two would decrease the gain by 3dB (again, this assumes the output swing is maintained and that the load resistance is scaled accordingly). Increasing the aspect ratio (by increasing the width) would indeed help the current drive of the pair, but the high frequency performance would be degraded owing to the extra loading on previous stages. Having too high a swing would result in one of the differential pairs moving into the triode region (as V_{ds} would become less than that gate overdrive).

2.5 Summary and note on designing SCL dividers

In this chapter, attention has been paid to the DC and AC characteristics of SCL combinatorial logic where it is possible to extract a trend and go about designing some of the logic circuits. Unfortunately, analysing the divider circuit is a little more involved where the problem is not simply the RC delay of a cascade.

A simple model for a *long* inverter chain resulting in the delay per stage, is presented. It achieves the figure by lumping various capacitances (intrinsic and extrinsic) into one capacitance on the node between every inverter output and the subsequent inverter's input. When observing actual transient simulations of ring oscillators, a selection of transistor operating modes and terminal capacitances are expected. One can say this delay is approximately constant with increasing I_{BIAS} if R_L and channel width are scaled such that the DC logic swing is proportional to ΔV_{SW}. Thorough transistor models would show that the parasitic capacitance dominates over the intrinsic capacitance, at small channel widths and bias currents, resulting in a loss of toggle frequency.

A point to note is that SCL dividers lack an analysis and design methodology based on performance specification parameters. Designers of high speed dividers have longed for a 'recipe' allowing them to choose the dimensions of the transistors or simply provide a starting point for subsequent iterations to arrive at (or get close to) the desired operating performance.

The problem stems from the non-linear nature of the circuit. The required scenarios are:

- the highest operating frequency given the output swing and common-mode voltage and technology,
- the lowest power consumption given the desired operating frequency, output swing and common-mode voltage and technology,
- a certain output swing, given the operating frequency, power supply and technology.

In traditional linear design, a circuit operating point is found about which a small-signal analysis can be performed using both 'hand' equations and a simulator. Assuming no distortion occurs, the gain and phase responses of the circuits can be ascertained with

the help of a simulator. With our non-linear SCL divider, the input is usually a large signal operating around a bias point and by definition, should remain stable for a period of time before toggling between states. The case for linear small-signal analysis is also flawed owing to the intended difference between input and output frequencies.

One solution is to use a brute force method of simulating the divider for a particular set of technology/model parameters with different permutations and combinations of large signal parameters. From this, trends can be found and the designer can converge on the desired performance at the expense of a multi-dimensional matrix and large simulation times.

An alternative is to use linear periodically varying analyses [3]. Firstly, a periodic analysis computes the periodic operating point with the application of the large clock signal about which the circuit is linearised and the small-signal information is applied. A linear time-varying analysis computes the response. To describe the complexity of the

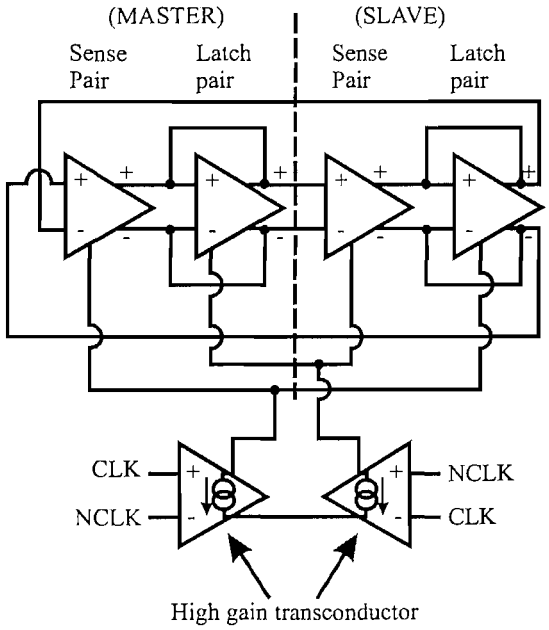

FIGURE 2.19: Symbolic view of the SCL divide-by-2 circuit.

task, Fig. 2.19 should help illustrate the large signal dynamic behaviour. The input signal drives the transconductors situated at the bottom of the diagram, toggling pairs of amplifiers above them. Although the 'latch' amplifiers have an explicit positive feedback and the 'sense' amplifiers have none, a negative feedback manifests itself too at high frequencies owing to Miller capacitance between the gate and drain terminals of the transistors. A linear small-signal AC analysis (such as one in SPICE simulators) is meaningless as the amplifiers in the upper half of the diagram energise and de-energise at a rate equal to half the input clock frequency. Hence the biasing and operating points

of these amplifiers change too. Although the outputs from this divider show a large AC signal toggling around a common-mode voltage, it is the master-slave action of the flip-flop that prohibits such a mode of analysis. In some of the forthcoming divider circuits of this dissertation, an attempt was made to use periodic steady-state analysis, but without success. In the design of the SOI divider, our circuit simulator did not possess such a tool and therefore was unfeasible. In the case of the divider circuits fabricated in bulk CMOS technology, all designs failed to converge whilst starting the tool, and owing to time pressures, were not investigated any further.

Bibliography

[1] J.Benson. STAG v2.7 model user guide. http://www.micro.ecs.soton.ac.uk/stag, 2002.

[2] Ketan.Mistry. 9 month report. Technical report, Microelectronics, ECS, University of Southampton, July 2000.

[3] K.Kundert. Introduction to RF simulation and its application. *IEEE Journal of Solid-State Circuits*, 34(9):1298–1319, September 1999.

[4] M.Alioto and G.Palumbo. Oscillation frequency in CML and ESCL ring oscillators. *IEEE Transactions on Circuits and Sytems-I*, 48(2):210–214, February 2001.

[5] M.Alioto and G.Palumbo. Design strategies for source coupled logic gates. *IEEE Transactions on Circuits and Sytems-I*, 50(5):640–654, May 2003.

[6] M.S.L.Lee. *Compact Modelling of Partially Depleted Silicon-on Insulator MOSFETs for Analogue Circuit Simulation.* Ph.D. thesis, University of Southampton, 1997.

[7] P.Larsson. High speed architecture for a programmable frequency divider and a dual modulus prescaler. *IEEE Journal of Solid-State Circuits*, 31(5):744–748, May 1996.

[8] S.Docking and M.Sachdev. A method to derive an equation for the oscillation frequency of a ring oscillator. *IEEE Transactions on Circuits and Sytems-I*, 50(2): 259–264, February 2003.

Chapter 3

Review of published frequency dividers

In this chapter, we review the design of high speed divider circuits, and go on to examine various dual-modulus prescaler architectures. Because of specific interest, a discussion of phase selection dual-modulus dividers has been delayed until Chapter 4.

3.1 Fixed modulus prescalers

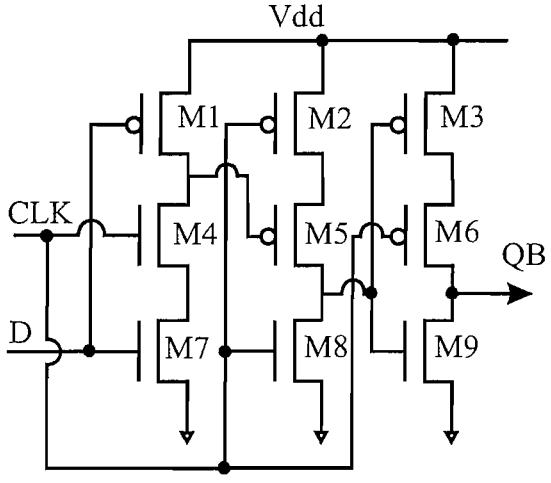

FIGURE 3.1: Schematic of a True Single Phase Clocked flip-flop [10].

In the design of high speed asynchronous dividers, evaluating single divide-by-2 stages can provide a useful estimate of the highest frequency attainable. Although loading must be emulated, they can provide useful starting points around which to build a circuit.

Classic static CMOS topology dividers have not been published as much as SCL based

dividers in the context of high frequency synthesisers. However, those published tend to be based around True Single Phase Clocked (TSPC) flip-flops [10] shown in Fig. 3.1. Reference [16] talks about a design based on modified TSPC flip-flops. These divide stages are used in the synchronous divider for the first high speed stage, but it has been pointed out [16] that its low frequency operation is questionable, with incorrect toggling due to lack of storage on certain nodes. Adding dummy loads can alleviate this, but its high frequency operation is slightly compromised.

Another modified version of the TSPC flip-flop is reported in [5], where it is used in the a divide-by-4\5 block. Unfortunately, it has been highlighted that this flip-flop also fails to function correctly at some frequencies below the maximum toggle frequency owing to its edge response characteristics [11]. Disasterous glitches have shown its unsuitability, but when run fast enough (at approximately 600MHz) with a capacitive load, the divider does operate satisfactorily.

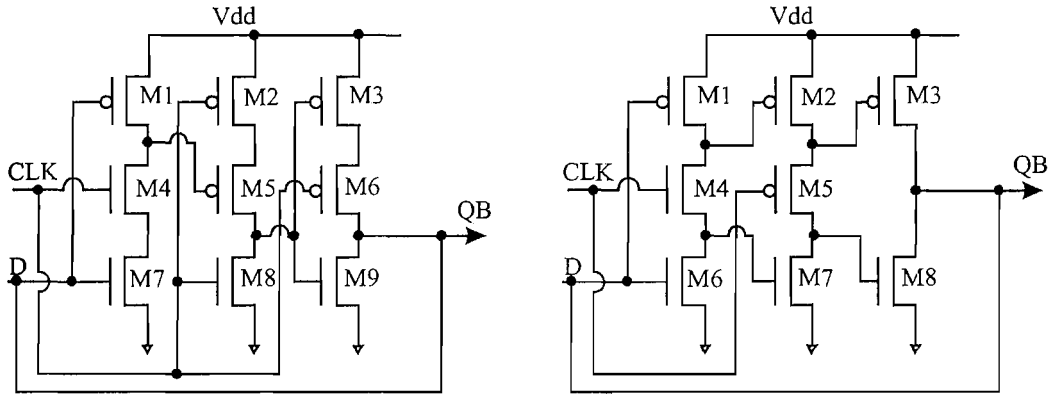

(a) TSPC divider (b) modified divider

FIGURE 3.2: TSPC circuits configured as dividers.

The TSPC flip-flop can be reduced to 8 MOS devices compared to 9 [2] (see Fig. 3.2), whose operation will follow dynamic logic principles. The reduction in devices effectively reduces the load on the clock input by 25% and in turn, reduces the power consumption. The simulated performance of [2] shows division by 2 at 2.2GHz in a 5V supply rail for a design fabricated in 0.8μm CMOS process.

As input frequencies have risen, dividers have migrated over to using source-coupled logic (SCL). Technically, logic elements sit on current sources, copying the bipolar designs. Published designs have shown flip-flops where the current source transistors are removed yet the divider still functions, owing to the large input swing. The authors of publication [3] talk about the use of a divide-by-2 stage based on a SCL topology that excludes the

clock and current source transistors under the long-tailed pairs. Instead, PMOS load transistors are clocked with complementary clocks, derived using a pass gate and inverter (to equalise delay). With this block fabricated using a 0.1μm CMOS process and 'ring' shaped MOS layouts (Fig. 3.3(a)), 13.4GHz operation is achieved when the circuit in Fig. 3.3(b) runs with a 2.6V supply, whilst consuming 28mW. According to the authors, the ring shaped MOS layouts generate circuits operating 25% faster than if a typical 'textbook' device were to be used, due to the lower drain capacitance.

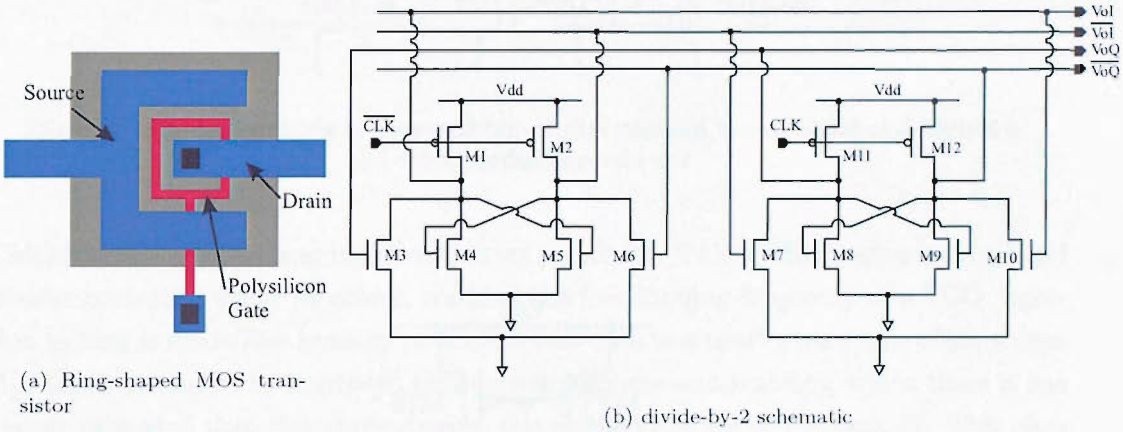

(a) Ring-shaped MOS transistor

(b) divide-by-2 schematic

FIGURE 3.3: Very high speed frequency divider [3].

Where power conservation is desired, designs can couple a mixture of SCL dividers with CMOS dynamic logic dividers. A dual modulus divider [9] has been presented using SCL based dividers for the high speed section and CMOS flip-flops [17] for the lower speed partition. It has been found that the SCL dividers would toggle faster without the current sources (clock transistors sit on the ground net as shown in Fig. 3.4), although the penalty here is that these clock transistors are required to be driven harder. This is not a problem when tailoring the design for operation within a frequency synthesiser, because its input is driven by a VCO geared towards low phase noise operation. The output voltage in the dividers is generated by a folded NMOS diode fed by PMOS current sources. After the first divide-by-2 stage, a second SCL divider is used, but this time it sits on a current source so as to accomodate low voltage differential signalling.

Another form of ratioed logic employs pseudo-NMOS implementations of the SCL divider with [14] cited as an example. These authors acknowledge the limitations of TSPC flip-flops because of their large input capacitance (parallel connections of NMOS and PMOS transistors) as well as the poor transconductance of the PMOS device. The design shown in Fig. 3.5 uses no long tailed pairs within the divide stages and their results (from a 0.25μm CMOS fabricated design) have demonstrated functionality at 5.5GHz on a 1.8V supply.

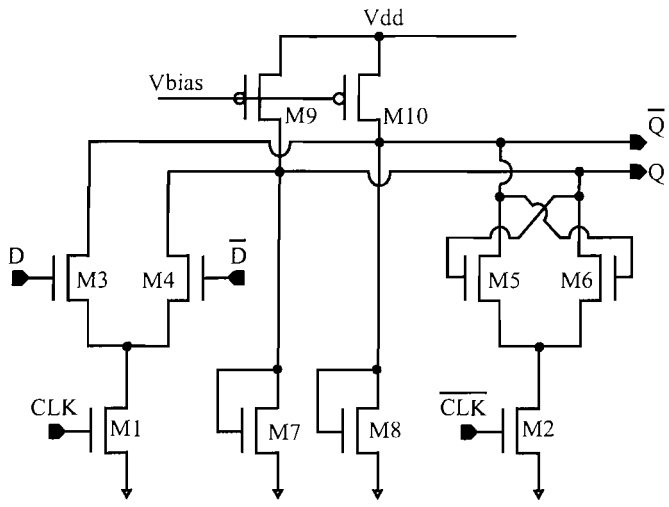

FIGURE 3.4: Schematic of half the D-type flip-flop used in the first implementation of the phase selector divider [9].

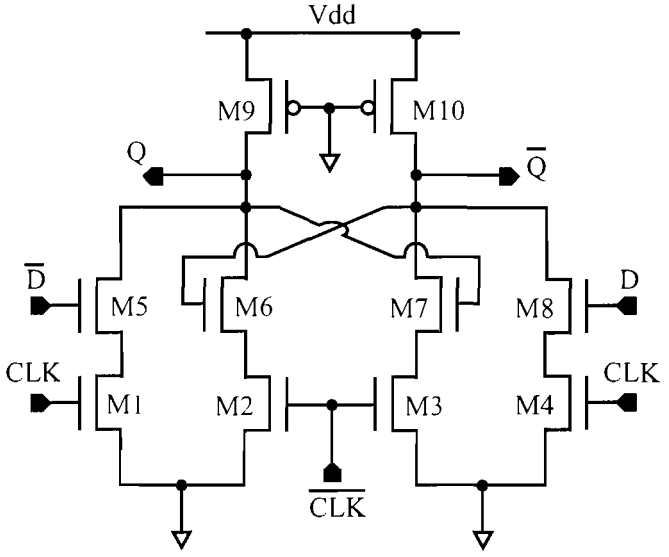

FIGURE 3.5: Pseudo NMOS schematic of half the D-type flip-flop used in a modified version of the phase selection scheme [14].

A recent presentation has resurrected the idea of clocking loads in SCL dividers, thus changing the properties of the divide stage [7]. A 16.8GHz divider fabricated in a 0.25μm CMOS process operating within a 1.8V supply voltage illustrates a similar divide topology to that observed in [9], but the speed enhancement is attributed to the load devices (see Fig. 3.6). PMOS loads are clocked in such a way that they present a low impedance when toggling but exhibit a high imedance when latching.

The last divide-by-2 reviewed accomplishes 19GHz operation when fabricated in a 0.35μm

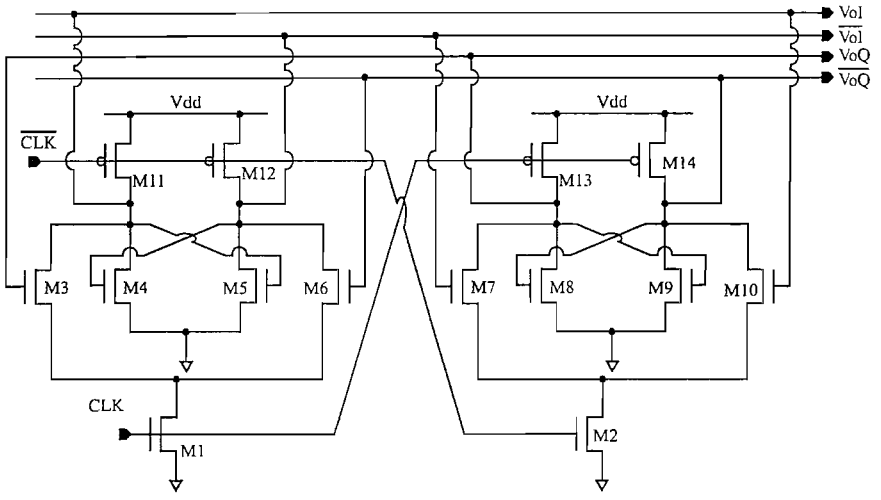

FIGURE 3.6: Schematic of a high speed divide-by-2 circuit, where the active loads are clocked too.

CMOS technology by using injection locking principles. This is effectively a narrow band divider controlled using varactors, tracking the free running frequency of a VCO. Injection locking is where free running oscillators lock onto to a nearby harmonic of the source [1]. This can lead to an improved phase noise response and stability, whilst there is less energy disspated than the above design, due to the Q factor of the tank [8]. This class of frequency dividers, although limited in functionality, does serve its purpose in PLL based frequency synthesisers. Stable reference sources tend to exist in a small frequency range and, where finer resolutions are required, these fixed ratio prescalers can help achieve that goal.

When constraining the modulus to be single radix, the divider tends to be feedforward, easing the architecture design. Asynchronous arrangements are favourable in this application because the global feedback of the synchronous counter tends to load the driver with parallel clock inputs. With the asynchronous divider, there is an added bonus of the total power increasing logarithmically as a function of the number of divide stages (with an N stage divider, and additional stage N+1 should consume half the power of stage N as its fundamental input frequency is half that of stage N's input, assuming it presents half the capacitive load to the output of stage N, and so on.) If there is a 3 stage asynchronous divider (composed of divide-by-2 cells,) and if the first divider has an input capacitance of C, bias current I and input frequency f_{in}, then the second stage should have an input capacitance of $\frac{1}{2}C$, its bias current should be $\frac{1}{2}I$ and its input frequency should be $\frac{1}{2}f_{in}$. The third stage will have an input frequency equal to $\frac{1}{4}I$, an input capacitance of $\frac{1}{4}C$ and a bias current of $\frac{1}{4}I$. Hence, the total current of this divide-by-8 will be $\left(1+\frac{1}{2}+\frac{1}{4}\right)\cdot I = 1\frac{3}{4}\cdot I$. This is true in theory, and ignores the fact that the gate width of a transistor cannot be divided down to unrealistic dimensions, nor can the current be scaled down as aggressively to infintesimal values.

Reference [12] gives two examples of such a divide arrangement. Its flagship design is a divide-by-128, whilst results of a divide-by-16 are also presented. Measured results for the divide-by-128 show the 12GHz divider consuming 60mW using a mixture of 2.4V and 2V power supply rails. The high speed flip-flop consumes 32mW from a 2.4V supply whilst the remaining 6 stages collectively consume 28mW in a 2V environment (combination of TSPC and SCL based dividers). The divide-by-16 is composed entirely of SCL based dividers and has functioned with input frequencies in the range of 9-14GHz, although no power consumption values are given.

An interesting point to note in both designs is that their input sections are terminated with 1000Ω on-chip resistors in order to generate the necessary input signal (-3dBm into a 1000Ω termination corresponds to a $1V_{peak}$ input). The use of aluminium bondwires for the test arrangement in a way determines the choice of termination resistance.

SCL based designs must also tackle the issue of interstage coupling to ensure correct biasing of devices. Source followers can be implemented but designs in bulk CMOS suffer from body effect, and hence attenuation of the AC signals, below the load capacitance - transconductance pole. Above this pole, voltage division is observed owing to the presence of reactive elements. AC coupling with a 'local' voltage divider biasing can work but the authors [12] decided against it because parasitic capacitance (associated with the on-chip capacitor) to the substrate would yield another reactive voltage divider, thus attenuating signals. Instead, two consecutive SCL stages are directly coupled with the current in the first stage increased to lower its common-mode output voltage. Lowering the power supply of the first stage and increasing the load resistance are both detrimental to the high frequency performance of the divider.

3.2 Dual modulus dividers

Single loop synthesisers can have division ratioes between 2 and several thousand. According to analysis in [18], the settling time of the synthesiser increases as the frequency resolution decreases, which is a poor characteristic in designs that hop carriers more than a thousand times per second (the resolution to step size ratio in a single loop is equivalent to the reference frequency). With the division ratio N taking on fractional values, the output frequency can be changed in fractional increments of the reference. It is not possible to design a straightforward divide-by-N block where $N \notin \mathbb{Z}$ ($\mathbb{Z}$ is an integer). If the ratio is toggled between two integer ratios, the average ratio appears to be fractional.

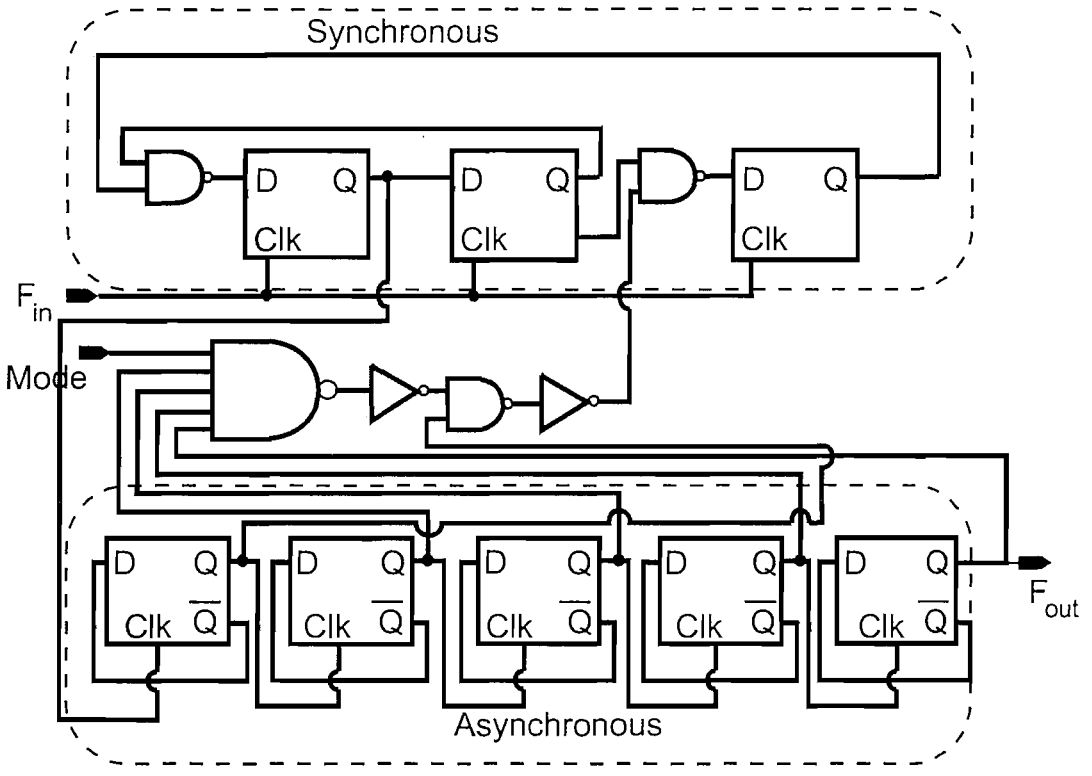

FIGURE 3.7: Schematic of a 128/129 dual modulus divider.

Reported dual modulus designs tend to use two consecutive ratioes. By using a division ratio of N+1 for one output cycle (output of the overall divider) from a chosen number of cycles, M,

$$L_{M-i} = L_M, L_{M-1}, L_{M-2}, \cdots, L_1, L_0 \tag{3.1}$$

(when M-i=0, L_i=N+1 else L_i=N, where M$\in \mathbb{Z}$), the average ratio is :

$$\frac{\sum_{i=1}^{M} L_i}{M} = \frac{\overbrace{N+N+N+\cdots+N+(N+1)}^{M\ terms}}{M} = \left(N+\frac{1}{M}\right) \tag{3.2}$$

and therefore the output from the synthesiser is $f_{out}=\left(N+\frac{1}{M}\right)\cdot f_{ref}$. In order to implement fractional synthesis, an accumulator or counter is required to decide when to toggle between the ratioes. This style of programmable divider gives good resolution whilst keeping the settling time low.

One of the earliest implementations of a dual-modulus divider, which is referenced frequently, is shown in Fig. 3.7 [17]. The output of the VCO is immediately followed by a synchronous counter. This is optimised for speed and not power, with three flip-flops employed in the high-speed synchronous divider, giving a divide-by-4/5 capability. Additional flip-flops cascaded asynchronously together with the combinatorial logic present

in the feedback path give the resulting 128/129 divider. Their results show the 1.2μm CMOS design to function at 1.4GHz when in divide-by-128/129 mode, drawing 6.9mA from a 5V supply. Another divide-by-128/129 design is presented in [2], but their modified TSPC flip-flops have yielded a dual modulus design capable of 1.22GHz operation in a 0.8μm CMOS technology, even though the power consumption is reduced by some 5mW.

The authors of publication [13] report a design based on a similar arrangement to the previous design, except the dual moduli are 15/16. Here, a 2 bit Johnson counter (synchronous) precedes a 2 bit asynchronous divider. Logic in the global feedback enables the synchronous counter to divide by 3/4, and coupled with the TSPC based asynchronous divide-by-4, the overall result is a divide-by-15/16 unit. 1.4GHz operation is achieved whilst consuming less than 13.5mW on a 5V supply.

Moving away from that architecture, [15] introduces the idea of clock preprocessing.

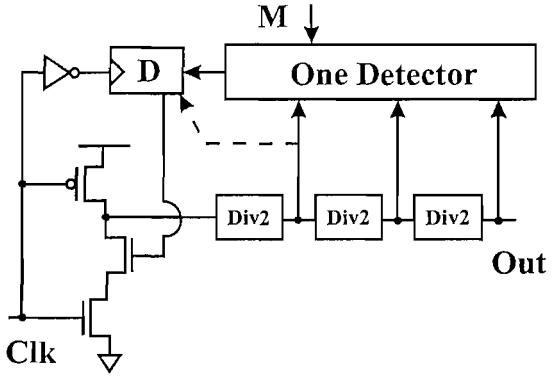

FIGURE 3.8: Block diagram of a divide by 8/9 divider.

This implementation incorporates asynchronous division throughout and gates the clock at the very beginning of the circuit. This gating is controlled by the state of the actual divide chain. The block diagram in Fig. 3.8 shows a dual-modulus divider incorporating a 'One detector', which is simply a multi-input NAND/AND gate. This controls a standalone D-type flip-flop configured to trigger on a complementary edge to the asynchronous divider. It must be stated that this is a narrow band design because the inverter before the single flip-flop not only provides another phase, but also gives a delay so that a clock edge may trigger the ripple divider before being blocked. At low speeds, the blocking can occur faster than would be liked, causing a state of deadlock. It is this strategy that gives a successful operation. Also, at high speeds, the output of the first divide-by-2 stage is fed directly to logic embedded within the single D type flip-flop as it is identified as being a critical path. The dual-modulus divider as a whole is able to function correctly up to higher speeds because the detector is 'primed' and 'decided', without having to wait for the first divide-by-2 to toggle. The divider exploits asynchronous behaviour by scaling subsequent divide stages accordingly in order to reduce

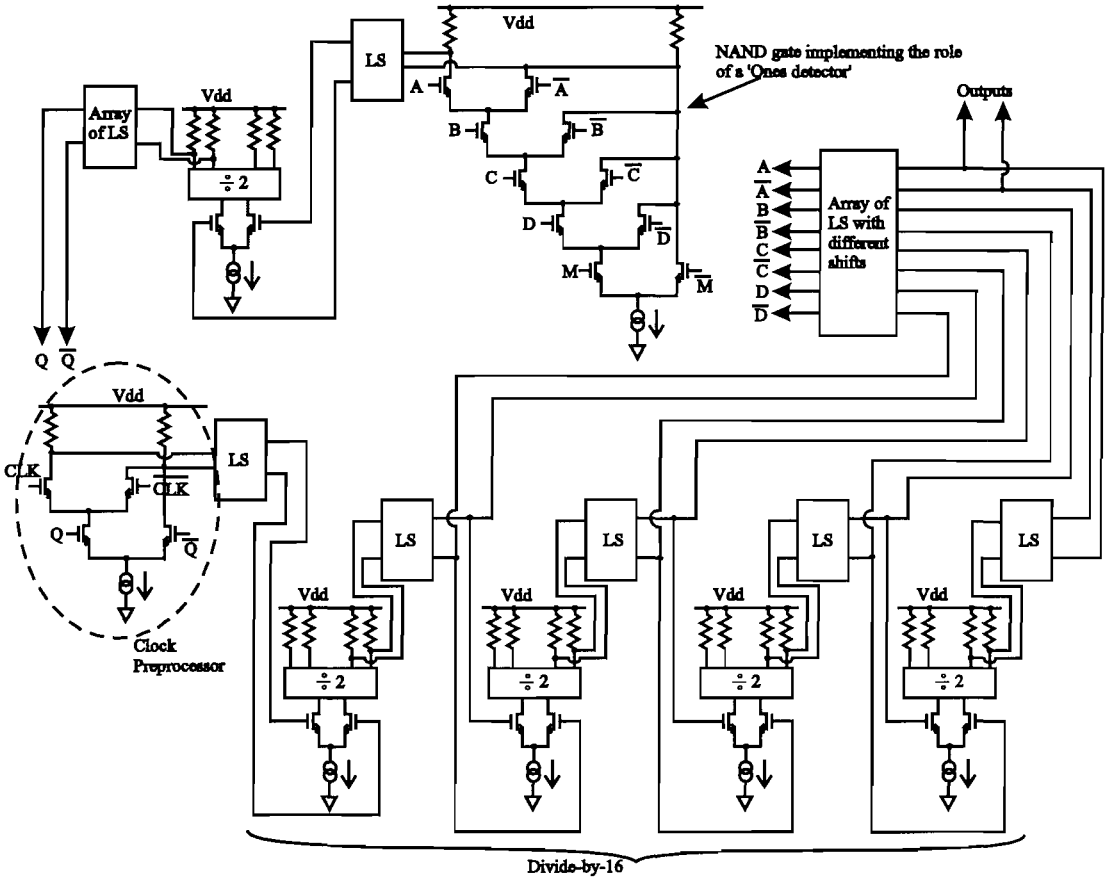

FIGURE 3.9: Schematic of a divide-by-16/17 dual modulus divider based on an asynchronous architecture [15].

power. Measurements have demonstrated this configuration to operate successfully with a 1.90 GHz input frequency on a 5V power rail when fabricated in a 5V bulk CMOS 0.8μm process.

This architecture was our initial choice for a dual-modulus divider implementation in SOI, but simulations showed its performance to be inferior when compared with the dominant design of this thesis. The problem was traced to the 'ones detector' and its propagation delay marred the performance of the rest of the circuit. The maximum speed simulated is shown to be approximately 4GHz after which errors begin to arise in the non-binary division.

The results in Fig. 3.10 refer to the design seen in Fig. 3.9 where the divide-by-2 stages were never cascaded one of top of the other. Simple voltage-mode cascading, followed by 'up' level-shifters to a NAND gate, showed the unoptimised design to function with a maximum 4GHz input. At 5GHz, an erroneous divide-by-19 results, with the NAND gate to blame. This is a heavy consumer of power, even though it fits within a 3.3V power supply, but makes a useful comparison for current-reuse designs.

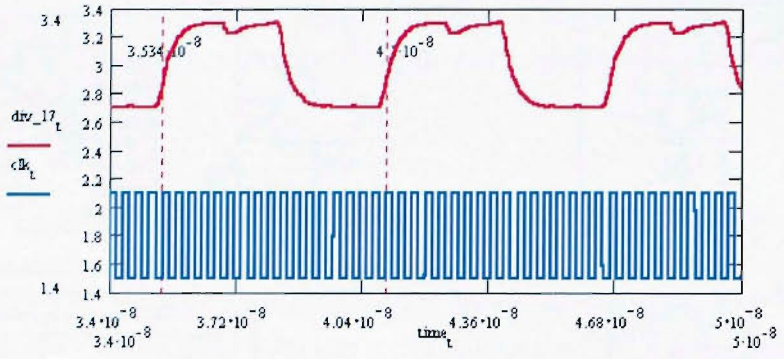

(a) Divide by 17 with a 3GHz input signal

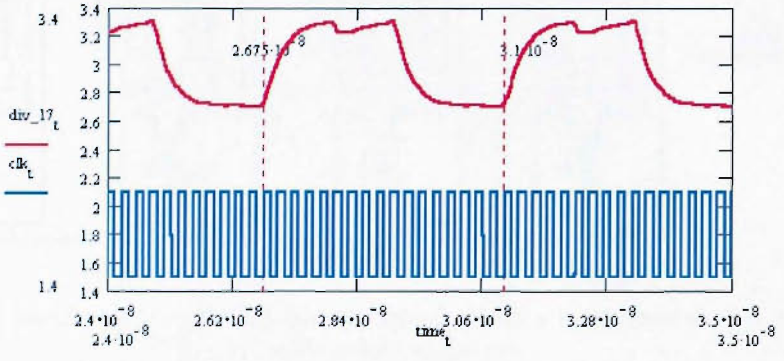

(b) Divide by 17 with a 4GHz input signal

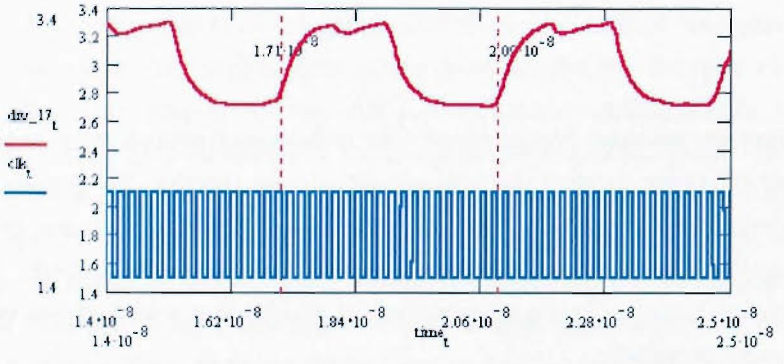

(c) Divide by 17 with a 5GHz input signal

FIGURE 3.10: Simulation results of the 16/17 dual modulus divider (based on the Larsson arrangement [15]) operating in divide-by-17 mode. The red trace corresponds to the divided output, whilst the blue trace refers to the input clock.

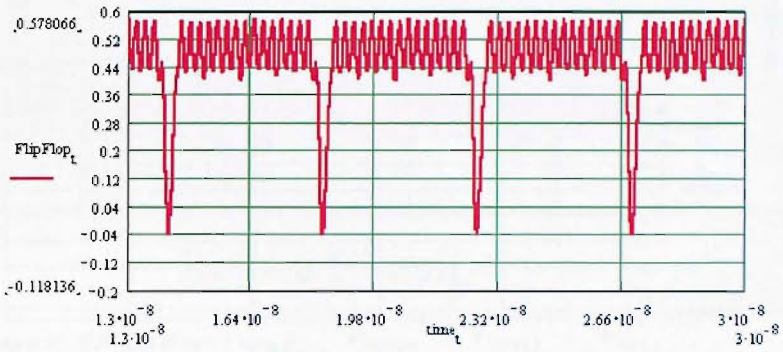

(a) the gating signal from the flip-flop, with the divider running at 4GHz

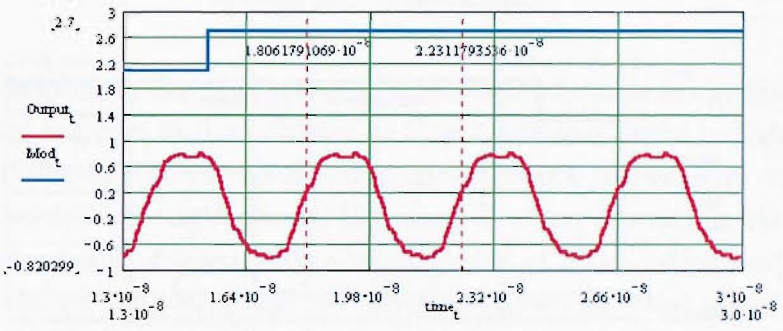

(b) the output signal, with the divider running at 4GHz

FIGURE 3.11: Simulation results of the 16/17 dual modulus divider (based on the Larsson arrangement [15]), operating in divide-by-17 mode. This topology uses a stacked divider with current-mode coupling.

Next comes the modified version of [15], whose divide chain is composed of a divide stack (current-mode cascading.) This stack contains 4 divide-by-2 stages, with cascodes between each. Figure 3.11 gives the results of a simulation with the divider running at 4GHz. No modulus control is implemented here, and so divide-by-17 is hard programmed.

By taking the marker values in Fig. 3.11(b), one can easily verify 17 input clocks cycles for one output cycle. In the waveform of Fig. 3.11(a), it must be stated that this trace is in fact a differential output, which raises concern at the lack of swing on this signal. The effect of this is made clear in Fig. 3.12, where an erroneous division ratio of 16 emerges. At 4.5GHz, the flip-flop (seen in Fig. 3.8) cannot signal fast enough to gate the clock input, hence providing no dual modulus functionality at all. The output of the flip-flop is seen to be less than half of the differential swing (calculated by the current through a load resistor in the flip-flop arrangement). Incidentally, the results are based on a divider that uses current re-use with the NAND gate, by feeding the output currents

directly in to the bottom of the gate. The subject of NAND gates has been touched upon in Section 2.3.

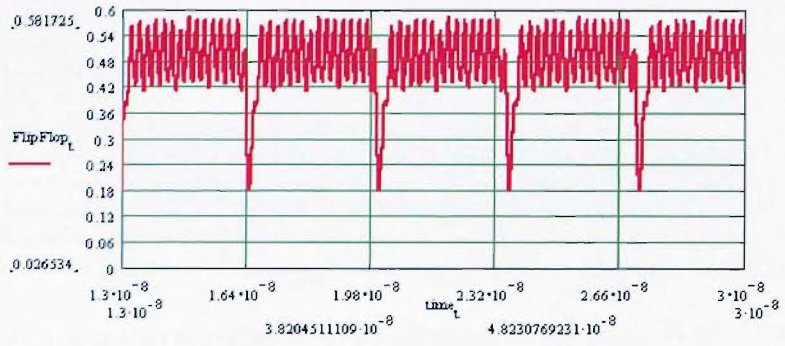

(a) the gating signal from the flip-flop, with the divider running at 4.5GHz

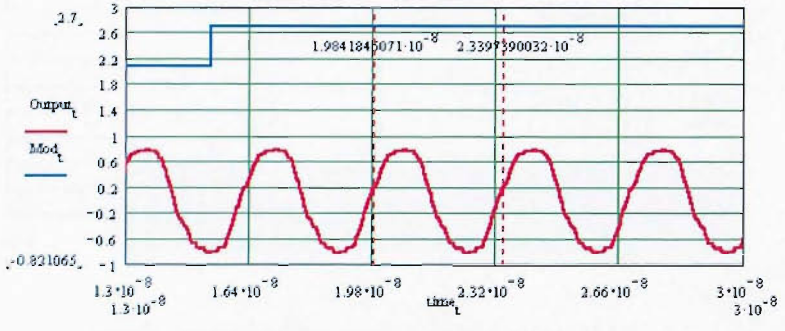

(b) the output signal, with the divider running at 4.5GHz

FIGURE 3.12: Simulation results of the 16/17 dual modulus divider (based on the Larsson arrangement [15]), operating in divide-by-17 mode. This topology uses a stacked divider with current-mode coupling.

This divider design cannot be compared against published designs owing to it being abandoned after unconvincing results. The simulations of the stacked NAND gate showed the fast toggling input to lie at the bottom-most input, and this prohibits any effort to optimise this design for speed and power. There is also the problem of detecting an all 1's condition if a synchronous divide-by-8 is used instead.

3.3 Programmable dividers

This last class of frequency divider takes an input frequency and divides the value by an integer. In integer-N PLLs, the higher division ratios permit more channels to be

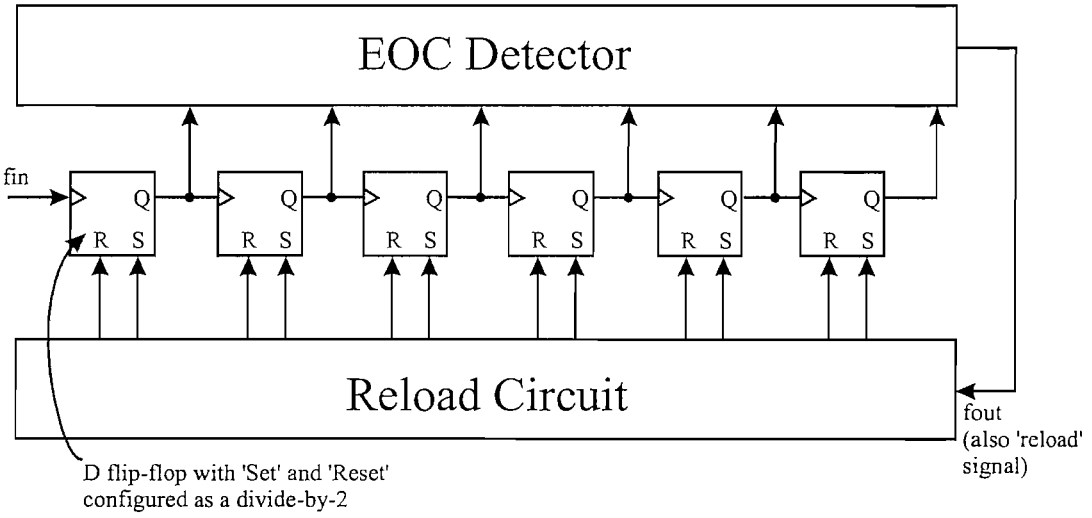

FIGURE 3.13: Block diagram of a conventional programmable divider, using an *end-of-count* (EOC) detector [6].

selected. An interesting point to note about these dividers for such an application is that the divisor need not start at 0, as division by 2 or 3 isn't very useful in high frequency regimes. For instance, Hi-Performance Local Area Network (HiPerLAN) requires only four consecutive division ratios, though the range of values depends on the frequency step size. Programmble dividers with a large number of divide ratios tend not to be used because of their complexity and low speed characteristics (dual modulus dividers offer a better solution in terms of speed, power and area).

An example showing complexity is given in reference [6](see Fig. 3.13). This programmable divider relies on and *end-of-count* (EOC) scheme, where a loaded modulus is counted down in binary, before starting again. A complete period denotes the passing of a number of input cycles (equal to the load modulus). The design mentioned can be programmed to divide between 2 and 63 inclusive with a top speed of 723MHz within a 5V supply. High speed operation is achieved with the use of an enhanced end-of-count detector. Fabricated using a 0.8μm CMOS process, the configuration consumes 17.12mW of power.

Although the aforementioned has a better power/frequency performance, the design presented by [15] manages 1.75GHz using the same minimum feature size (5V operation). By amalgamating the dual modulus architecture presented earlier with an asynchronous divide-by-N, a programmable divider can be achieved. It highlights that, by counting down, a loaded value can be reduced to zero gradually, and only quick detection of the LSB portion of the counter is needed. The paper also points out the difficulty in loading an asynchronous divider, because each stage is clocked by a previous stage and must wait for the correct edge to ripple through. Hence, a modified asynchronous divider with

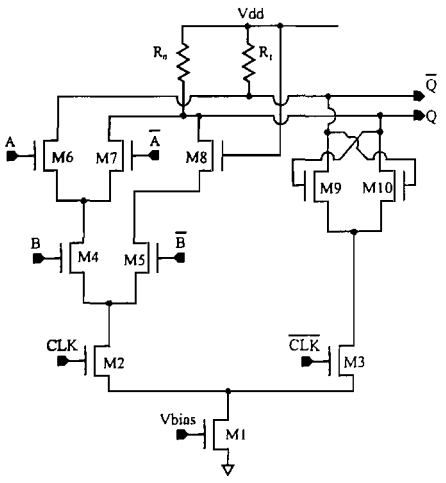

(a) D-latch with an embedded NAND gate

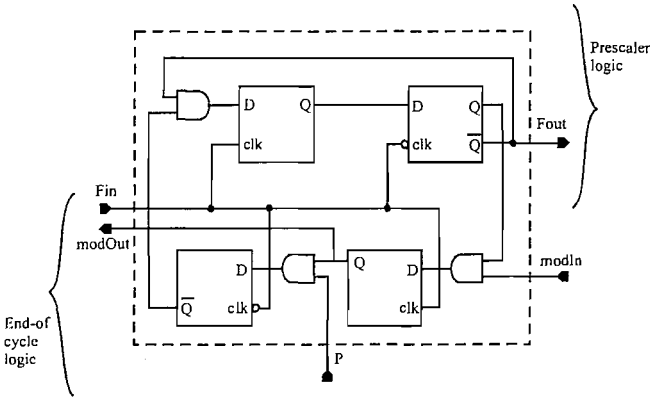

(b) inside each divide-by-2/3 module

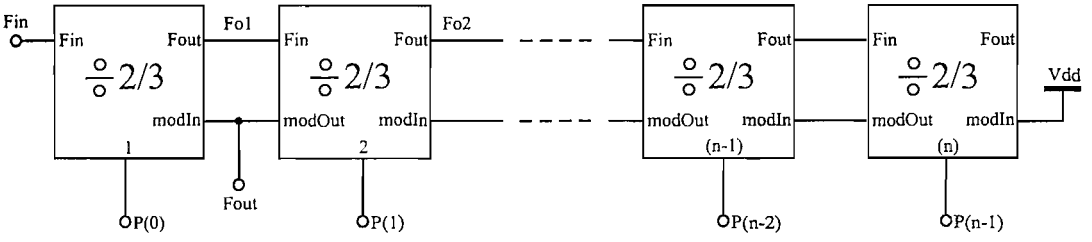

(c) asynchronous cascade of divide-by-2/3's giving a programmable divider

FIGURE 3.14: Design based on a modular approach to programmable division.

synchronous loading is used instead, permitting faster loading.

A novel and clever programmable divider is presented in [4] where the aim is modularity. Divide-by-2/3 cells are created using SCL based synchronous dividers. Each cell, shown in Fig. 3.14(b), employs 4 D-type latches as well as some combinatorial logic. SCL design techniques are exploited, whereby logic functions are embedded within the flip-flop (Fig. 3.14(a)). These cells are cascaded as in Fig. 3.14(c), according to the number of moduli required, though the actual moduli values are not as flexible. The design of 17 bit and 18 bit dividers were characterised in a 2.2V supply, although no absolute speed values are made available.

The last programable divider has an architecture that is similar to the phase selector design mentioned previously [9]. This particular divider [14] is designed for HiPerLAN applications running at 5.3GHz. The division ratios are limited between 220 and 224 inclusive as a requirement of HiPerLAN. Without the phase select circuitry, the fixed division modulus amounts to 216 (3 stages of divide-by-2 followed by 3 stages of divide-

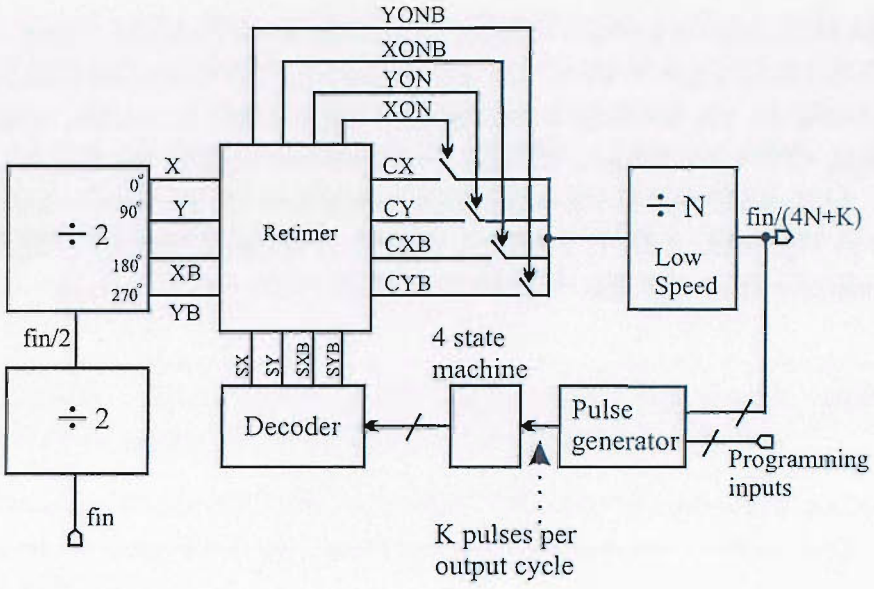

FIGURE 3.15: Block diagram of a 4 moduli programmable divider [14].

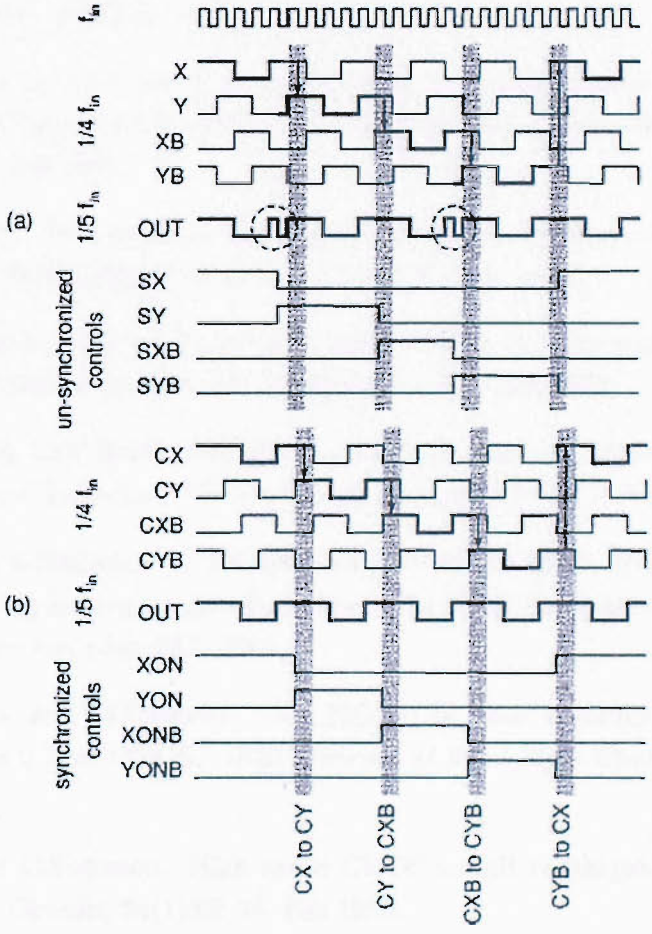

FIGURE 3.16: Waveforms showing glitches in the output when a control signal falls outside its 'window', and the remedy of the retimer block [14].

by-3). The phase selector is placed after the second high speed divide-by-2 stage, with the remaining divide-by-2 and divide-by-3 stages following it. The design uses pseudo-NMOS logic to implement the divide-by-2 functionality, and is found to operate satisfactorily in a 0.25μm CMOS technology. Glitches are eliminated through the use of a retimer (see Fig. 3.15), which synchronises the control signals to the respective clock lines as captured in Fig. 3.16. 5.3GHz frequency division is managed in a 1.8V supply, with 5.5GHz achieved at the expense of a 22% increase in power supply voltage.

Bibliography

[1] A.Bateman. Radio design course - transmitter and receiver architectures. Http://www.avren.com/Courses/TX_RX_Architectures_plain.htm.

[2] B.Chang, J.Park, and W.Kim. A 1.2GHz CMOS dual-modulus prescaler using new dynamic D-Type flip-flops. *IEEE Journal of Solid-State Circuits*, 31(5):749–752, May 1996.

[3] B.Razavi, K.F.Lee, and R.H.Yan. Design of high-speed, low-power frequency dividers and phase locked loops in deep submicron CMOS. *IEEE Journal of Solid-State Circuits*, 30(2):101–109, Feb 1995.

[4] C.S.Vaucher et al. A family of low-power truly modular programmable dividers in standard 0.35μm CMOS technology. *IEEE Journal of Solid-State Circuits*, 35(7): 1039–1045, Jul 2000.

[5] C.Yang et al. New dynamic flip-flop for high-speed dual-modulus prescaler. *IEEE Journal of Solid-State Circuits*, 31(5):744–748, May 1996.

[6] H.Chang and J.Wu. A 723MHz 17.2mW CMOS programmable counter. *IEEE Journal of Solid-State Circuits*, 33(10):1572–1575, Oct 1998.

[7] H.Wang. A 1.8V 3mW 16.8GHz frequency divider in 0.25μm CMOS. In *IEEE International Solid-State Circuits Conference*, page TP12.1, 2000.

[8] H.Wu and A.Hajimiri. A 19GHz 0.5mW 0.35μm CMOS frequency divider with shunt-peaking locking-range enhancement. In *IEEE International Solid-State Circuits Conference*, page 26.3, 2001.

[9] J.Craninckx and M.Steyaert. A 1.75GHz/3V dual modulus divide-by-128/129 prescaler in 0.7μm CMOS. *IEEE Journal of Solid-State Circuits*, 31(7):890–897, Jul 1996.

[10] J.Yuan and C.Svensson. High speed CMOS circuit technique. *IEEE Journal of Solid-State Circuits*, 24(1):62–70, Feb 1989.

[11] K.Sung and L.Kim. Comments on "new dynamic flip-flop for high-speed dual-modulus prescaler". *IEEE Journal of Solid-State Circuits*, 35(6):919–920, Jun 2000.

[12] B.De Muer and M.Steyaert. A 12GHz CMOS divide by 128 in 0.25μm CMOS. In *European Solid-State Circuits Conference*, pages 220–223, Stockholm, Sweden, 2000.

[13] N.Foroudi and T.A.Kwasniewski. CMOS high speed dual modulus frequency divider for RF frequency synthesis. *IEEE Journal of Solid-State Circuits*, 30(2):93–100, Feb 1995.

[14] N.Krishnapura and P.R.Kinget. A 5.3GHz programmble divider for HiPerLAN in 0.25μm CMOS. *IEEE Journal of Solid-State Circuits*, 35(7):1019–1024, Jul 2000.

[15] P.Larsson. High speed architecture for a programmable frequency divider and a dual modulus prescaler. *IEEE Journal of Solid-State Circuits*, 31(5):744–748, May 1996.

[16] R.Rogenmoser, N.Felber, Q.Huang, and W.Fichtner. 1.16GHz dual-modulus 1.2 μm CMOS prescaler. In *IEEE Proc. Custom Integrated Circuits Conference*, pages 27.6.1–27.6.4, May 1993.

[17] R.Rogenmoser, Q.Huang, and F.Piazza. 1.57GHz asynchronous and 1.4GHz dual modulus 1.2 μm CMOS prescalers. In *IEEE Proc. Custom Integrated Circuits Conference*, pages 16.3.1–16.3.4., San Diego, May 1994.

[18] U.L.Rohde. *Digital PLL frequency synthesizers - Theory and design.* Prentice Hall, 1983.

Chapter 4

Design and implementation of a new dual modulus divider architecture

The theme of the research has been to investigate and deliver ways to reduce the number of blocks running at the maximum signal frequency within the analogue RF section of a receiver. In doing so, one is able to reduce the power consumption of the whole implementation, and reduce the loading on the previous stage.

This chapter will examine the architecture of a particular dual-modulus divider[1] which addresses the issue of reducing the number of blocks operating at the highest signal frequency. Following this will be a few sections on how this original concept has been developed further, namely the design of an elegant phase-selection circuit and a glitch-free controller.

4.1 Review of a phase selecting dual-modulus prescaler

This chapter presents the development of a new dual-modulus divider architecture. The starting point for the evolution of the designs is the phase selector approach [1], with a block diagram given in Fig. 4.1. Here, we see the high frequency signal (delivered by a signal generator or maybe a VCO) arriving at its inputs and being divided by 4 using two divide-by-2 cells coupled in an asynchronous fashion. Each divide-by-2 circuit consists of a master-slave pair of D-type latches, resulting in D-type flip-flop whose outputs are fed back (in a negative sense) to the input of the flip-flop. Specific to the SCL divider, four outputs are present and there exists a $\frac{\pi}{2}$rad phase difference between certain pairs of outputs. This characteristic is justified by the master-slave action (giv-

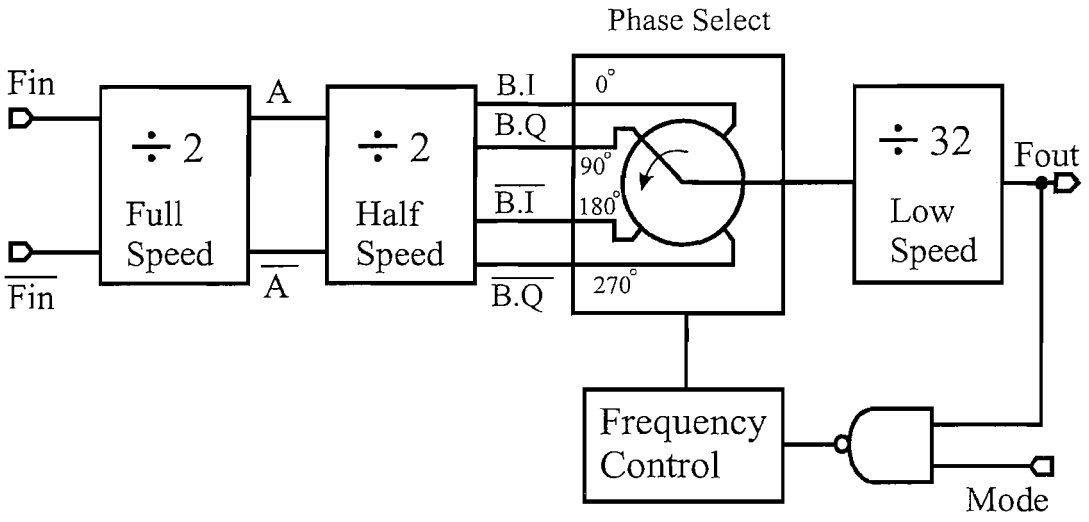

FIGURE 4.1: Block diagram of 128/129 frequency divider published by Craninckx [1].

ing rise to the $\frac{\pi}{2}$rad phase difference between the output of one latch and the output of the other) and the differential operation (resulting in the remaining πrad and $\frac{3\cdot\pi}{2}$rad phases). Referring back to the figure, division in a fixed binary mode uses outputs that are πrad apart to couple a divider to a subsequent divider. After the first division by 2, a $\frac{\pi}{2}$rad phase difference between two particular outputs is equivalent to half a cycle of the input frequency. After the second division by 2 (i.e. aggregate division is 4) a $\frac{\pi}{2}$rad phase difference between two particular outputs on this second divider now corresponds to one whole cycle of the input frequency. A phase-selector can be placed at this point in order to drive the remaining divider stages with an input that can be lengthened by one cycle of the 'full' divider. The publication describes a divide-by-128/129 frequency divider that requires at least another 5 divide stages in order to achieve the division-by-128. The second consecutive modulus is obtained by periodically switching between the four outputs of the second cascaded divide-by-2 cell. Each consecutive pair of outputs differs by $\frac{\pi}{2}$rad, which corresponds to one complete cycle of the input. Switching between these phases is identical to losing an input cycle. By omitting one cycle and hence the transition that accompanies it (both a rising and falling edge are lost), each of the divider stages after the phase selector holds its state for the period of one clock cycle (applied at the input to the whole dual-modulus divider.) Therefore, the period of an output cycle is equal to 1 'skipped' input clock cycle as well as 128 'divided' input cycles.

(NB From this point onwards, the notation used to refer to the 4 input phases will be based on that used in quadrature modulation. Hence $B{\cdot}I \Rightarrow$ 0rad phase (used as a datum for the rest), $B{\cdot}Q \Rightarrow \frac{\pi}{2}$rad, $\overline{B{\cdot}I} \Rightarrow \pi$rad, $\overline{B{\cdot}Q} \Rightarrow \frac{3\pi}{2}$rad. The 'B' prefix refers to the second divide stage at which these outputs are taken. As the first two dividers are connected asynchronously, the clocking of the dividers ripples through the pair. This means that the rising edge (in a rising edge triggered flip-flop) of the input clock cannot

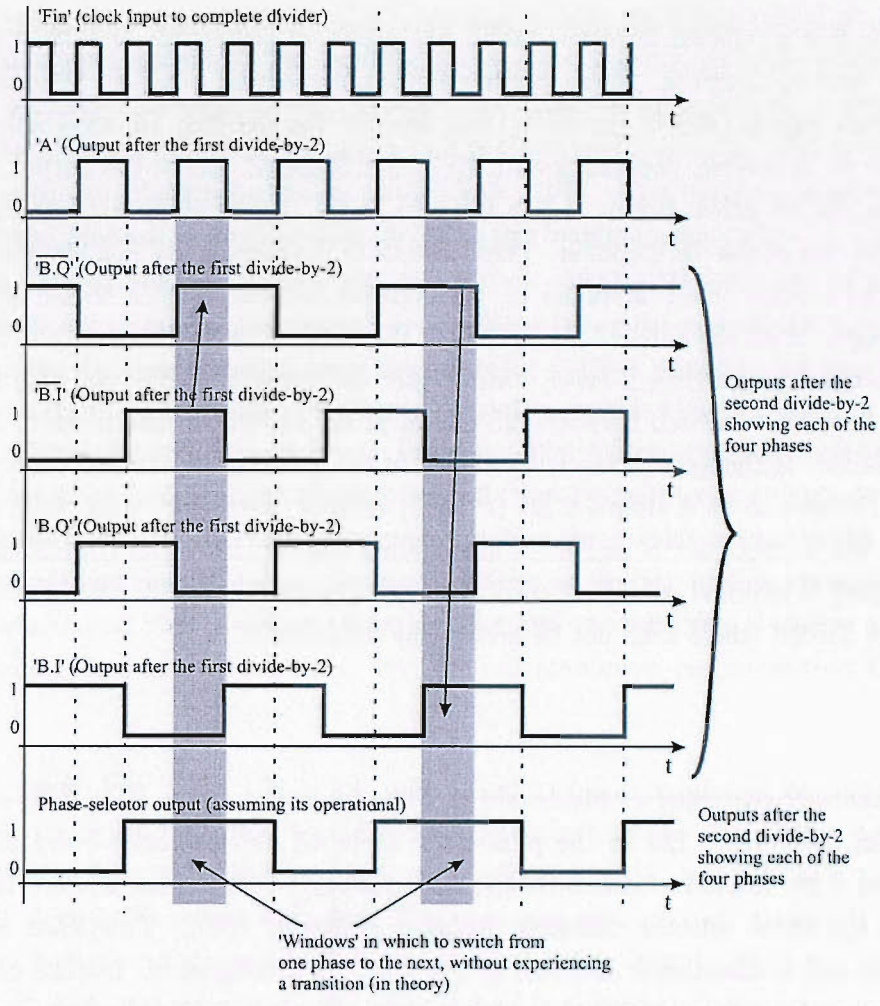

FIGURE 4.2: Waveforms of the 4 phases associated with the master-slave flip-flop, and their relation to the input clock.

be assumed to be in synchronicity with the rising edge of the $B{\cdot}I$ signal. Nevertheless, the outputs are $\frac{\pi}{2}$rad apart.)

The phase selector's role in the divider is to channel one of the four inputs to its output. As described, the frequency of each input is the same, the only difference being the $\frac{\pi}{2}$rad phase relationship between certain pairs of inputs. Whilst the divider operates in the (N+1) mode, where (N) is a power of 2, the phase selector should output a different signal that lags by $\frac{\pi}{2}$rad from the current input. This can happen at any stage during the 'full' division cycle. So for example, with a divide-by-128/129 divider, 32 clock cycles are still sent to the divide-by-32 divider following the phase selector when dividing by 129. The only difference between the divide-by-128 and divide-by-129 is that one of those cycles clocking the lower speed divide-by-32 is in fact one input cycle longer than the others (input to the actual divider). The topology by which the 'selected' input is carried to the output of the phase selector has been accomplished in more than one

way. The first published implementation has a pair of switchable amplifiers on each of the 'I' and 'Q' outputs. These amplifiers not only convert between differential and CMOS logic signals (and at the same time, amplify the signals), but they also output the signal or its inverse, depending on a logic control signal. Hence the output of these amplifiers can be either $B{\cdot}I$ or $\overline{B{\cdot}I}$ in the case of the 'I' amplifier, and either $B{\cdot}Q$ or $\overline{B{\cdot}Q}$ in the case of the 'Q' amplifier. These signals then progress to a combinatorial logic circuit where either the 'I' amplifier or 'Q' amplifier outputs are sent to the remaining divide stages. With the outputs being CMOS, the subsequent divide-by-2 cells will have a CMOS topology yielding a lower static power consumption compared with an SCL version. In order to switch between the phases at the inputs, a state machine of some sort is needed. Although not explicitly shown in the publication, the frequency control block is thought to be a simple 2-bit (4 state) counter, with some additional logic to generate the correct control signals. These signals can be rail-to-rail CMOS; hence no level shifting is required, though according to the schematics, signals and their complements are needed (these need not be necessarily differential).

The sequence of switching as illustrated in Fig. 4.2 is $B{\cdot}I$, $B{\cdot}Q$, $\overline{B{\cdot}I}$, $\overline{B{\cdot}Q}$, ensuring continuous division by 129 as the phases are switched between once every 32 cycles of the last 5 bit counter. Unfortunately, with a 4 state FSM, there is more than one input to the phase selector changing during a change of state. This leads to 'race' conditions and is disasterous in terms of the divider's functionality. Instead of having one operating boundary condition at high frequencies, there exist two, with the second being at a lower frequency. The upper operating frequency limit is set by the technology, whereas the lower limit is set by the 'race' situation described. The result is a circuit with a 'bandpass' characteristic (using filter terminology.) The 'C0' signal controlling whether to select the 'I' or 'Q' pairs was found to have transitions running slightly ahead of the 'C1' and 'C2' switchable amplifier control signals (see Fig. 4.3). However, the ripple nature of the asynchronous counter thought to be used in this controller leads to the delay in transitions. The result is a serious glitch that can momentarily change the selection of the phase to be sent to the output of the phase selector, causing the subsequent dividers to toggle (change state). One of the suggestions is to use a clocked D-type flip-flop in order to allow signals to settle before the state is passed to the phase selector inputs. The drawback is justifying the delay as well as the area and power consumption of such a cell. Their solution was to insert a buffer/inverter with a delay to lower the rise and fall transition times of the phase selector control inputs. With this delay, the slower (in terms of period) control signals are given a chance to catch up with the delayed control signal. Unfortunately, this is not really a solution and instead shifts the problem somewhere else. For this divider to work in true broadband fashion, there needs to be a way of changing one input at a time.

From a practical point of view, a dual-modulus divider isn't very useful on its own in today's wireless LAN chips. The need for a multi-modulus divider arises from the desire to seek a channel amongst several others. The dual modulus divider is used with an accumulator and comparator, resulting in a programmable divider. Unfortunately, the NAND gate can block the state machine clock input. The blocking is asynchronous, yielding the problem of prematurely clocking the frequency control unit. This in turn causes a different and incorrect input phase to be selected and the overall division in the chain is anything but the desired ratio. What is needed is some form of 'memory' which blocks the state machine clock yet holds the current inputs to the phase selector. So, in Craninckx's design [1], the frequency control unit holds the inputs to the phase selector when triggered to operate as a divide-by-128. When triggered to divide-by-129, the frequency control resumes where it left off, thus preventing any unwanted glitches. The drawback here is the extra power consumption due to this latch, though its circuit design should be that it doesn't load the output of the last divide stage of the true divider chain, and that its clock transistor has just enough drive to follow the output of the actual dual modulus divider. By this last statement, we mean that the divider

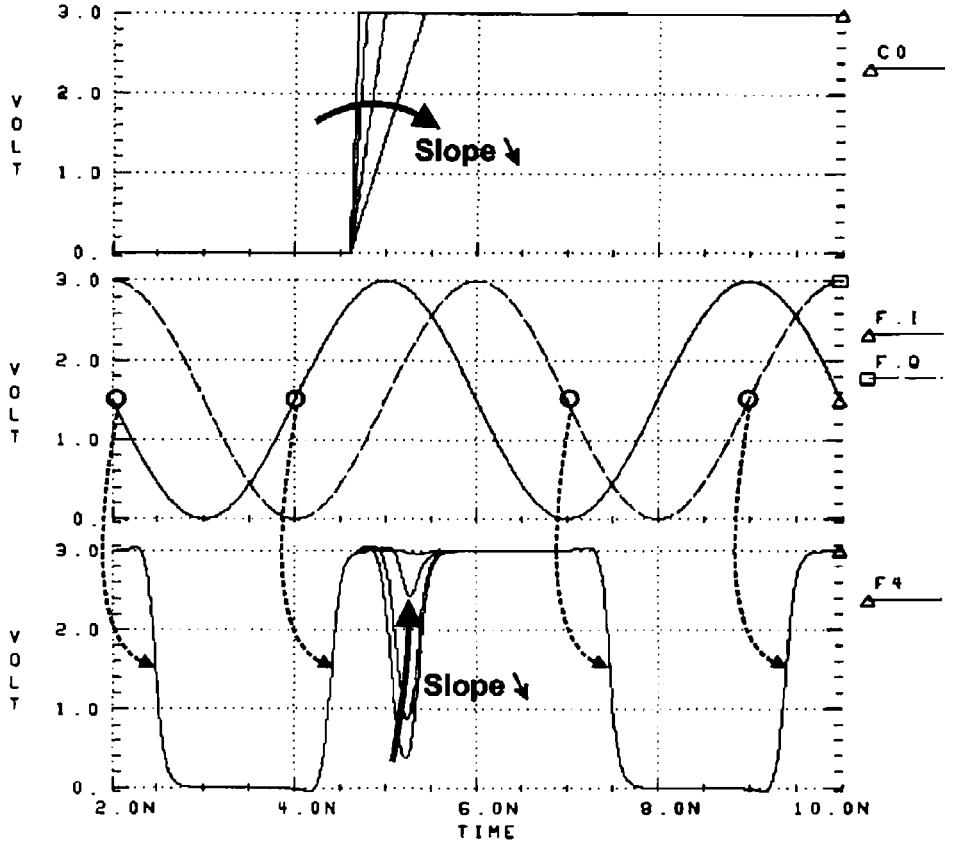

FIGURE 4.3: Plots showing the erroneous operation associated with the slope of the phase selector control transitions [1]. The thick arrows in the first and third plot, demonstrate erroneous frequency division, as the transition speed in one control input increases, thereby shifting the switching instant of the phase selector.

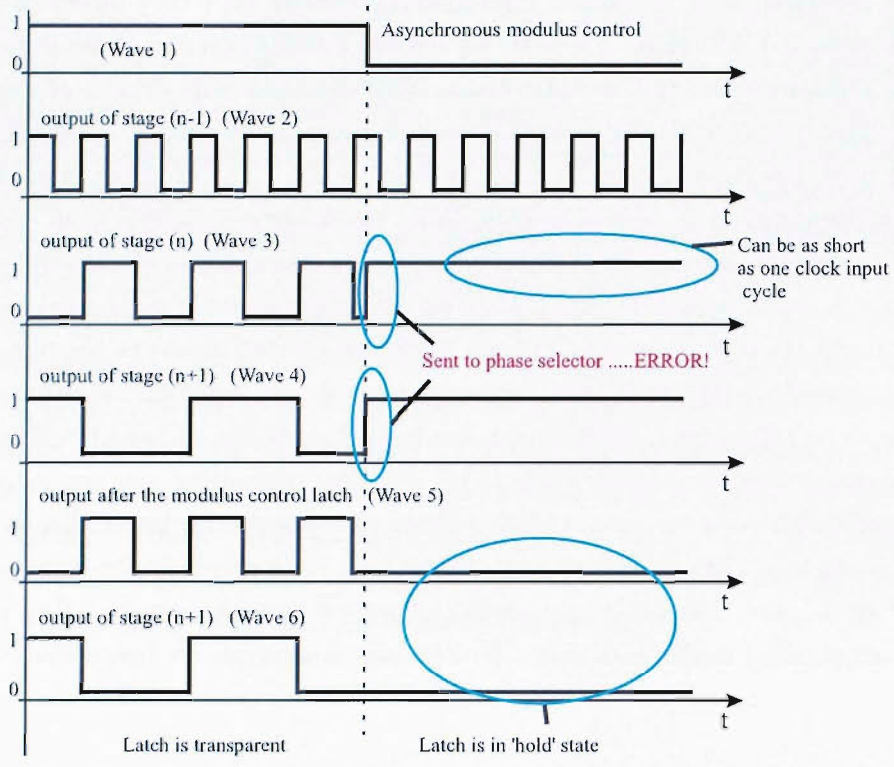

FIGURE 4.4: Time domain plots illustrating the problem with an embedded NAND gate in our implementation.

may be called upon to divide-by-128 for the least significant bit of a slow accumulator to toggle and toggle back again before dividing by 129 for the remainder of the overall cycle, as pictured in Fig. 4.4. This problem, together with its solution, manifested itself in the bulk programmable divider described in chapter 8.

For their fabricated design, 2.5mA was consumed by the first divide stage and 1.5mA by the second. The CMOS divide-by-32 sinks 1.5mA whilst the phase selection circuit sinks 2.5mA, all within a 3V power supply. From a performance point of view their divider runs at 1.75GHz, although a higher toggle frequency can be achieved by increasing the supply voltage. The author explicitly states the benefits of this design are that: 1) the clock generator (usually a VCO) is loaded with one D-type flip-flop instead of several in the case of a synchronous divider appearing at the front of the dual modulus divider [3]; and 2) that there is only one sub-block in the divider running at the highest frequency and therefore the whole divider runs with reduced power consumption. The high frequency feedback loop present in synchronous counters is also eliminated. Bearing this in mind, together with the ability to reduce the power consumption per divide stage from stage to stage in a binary fashion, this makes for an attractive circuit topology for low power transceivers.

A variant of this circuit corrects the flaw in the above description, using a 'retimer'

circuit [2], with the block diagram pictured in Fig. 3.15. As a precursor, the fact that a 'window' exists in which the inputs to the phase selector must change is described in detail, together with a description where no retimer is included (figure 3.16.) The circuit architecture differs from the previous one in that a pulse generator is used to give the multi-modulus feature of the circuit. The thinking behind this is to have the phase selector change phase more than once in one output cycle, hence more than one high frequency input cycle is swallowed per output cycle. The inputs to the divider select which modulus (between 220 and 224 inclusive) to select. Ultimately, a 4-state FSM is still employed but 'sychronising' of the phase selector's inputs is implemented in order to solve the glitch problem of the Craninckx implementation. Here, the problem of switching between the phases earlier than it should be done is clearly illustrated, together with the fatal consequences. This design solves that particular problem by delaying the action of the phase selector inputs to coincide with a window in which the current and next phase outputs have the same output level, thus preventing a glitch in the lines after the phase selector. This divider uses a pseudo-NMOS logic style for the circuit topology, and is clocked to run at 5.5GHz input with a 2.2V supply pulling 26.8mA whilst being driven by a 300mV_{peak} single-ended input. The circuit is fabricated in a 0.25μm CMOS technology and is discussed in chapter 3.

4.2 Circuit topology of new phase selector

The following is based on the Craninckx phase selector and has exactly the same operations, where it may invert either of the quadratures (master or slave outputs of the previous divide-by-2 stage) before selecting which phase to pass. However, a new current-steering circuit topology is presented, allowing the entire function to be performed in a current-mode fashion.

The circuit illustrated in Fig. 4.5 is the key to the dual-modulus designs in the rest of this text. A closer look at this circuit will show that it is in fact a '4-to-1 line selector' where a combination of the inputs (denoted by C0, C1, C2 and their complements) determine which phase to 'steer' through to the current-mode inputs of the divider succeeding the phase selector. Its operation is exactly the same as the original implementation of the Craninckx design, where either the $B{\cdot}I$ or $\overline{B{\cdot}I}$ phases, and $B{\cdot}Q$ or $\overline{B{\cdot}Q}$ are to be passed to the next stage of the phase selector. Here, another set of control signals determines whether to pass forward the chosen 'I' signal or the chosen 'Q' signal. The whole task is merely current steering but accomplishes the function in a far neater way. Being current-mode means that when coupling cells together in a signal path, the signals do not have to change between being a current and a voltage. As will be seen in later chapters, this omission is valuable to the designer and is a strength of the SOI dual

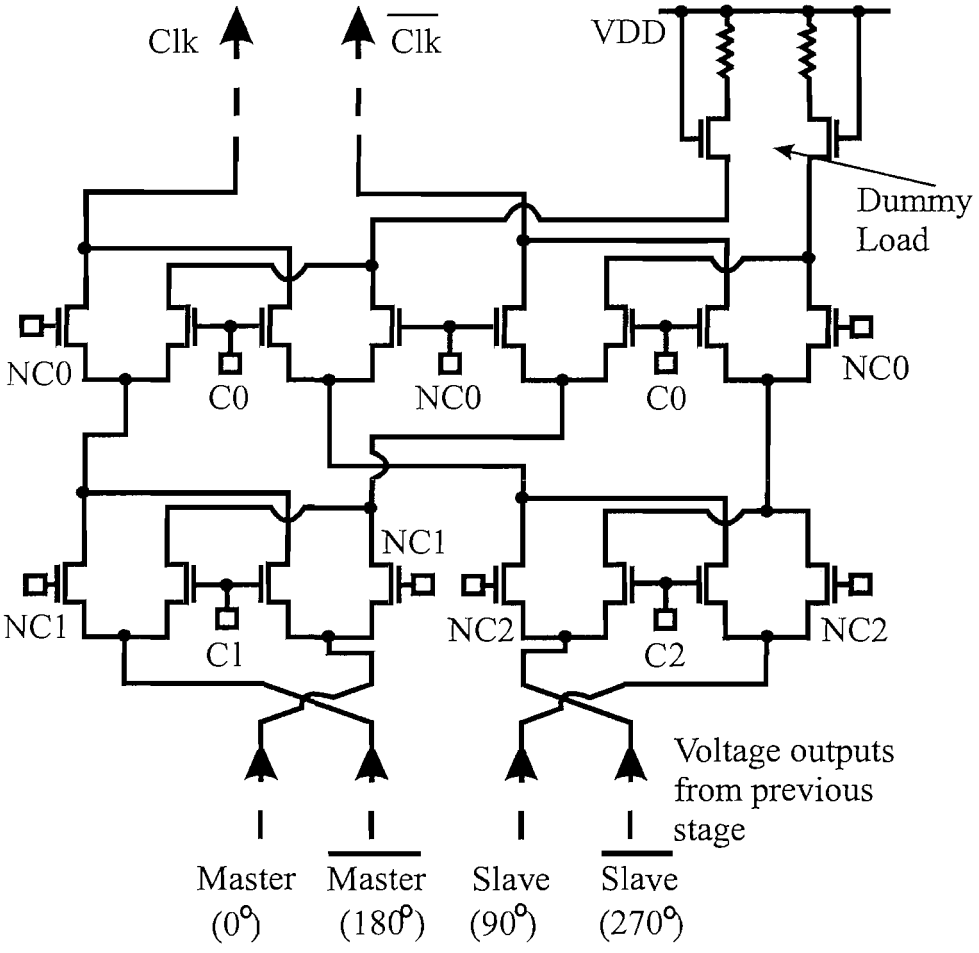

FIGURE 4.5: Schematic of the new phase selector.

modulus divider. Unlike the Craninckx implementation, the phase selection remains faithfully differential throughout its task. This allows smaller signal swings and hence quicker transition times.

$C2$	$\overline{C2}$	$C1$	$\overline{C1}$	$C0$	$\overline{C0}$		Output .
1	0	0	1	0	1		$B{\cdot}I$ (S0)
0	1	0	1	0	1		$B{\cdot}I$ (S1)
0	1	0	1	1	0		$B{\cdot}Q$ (S2)
0	1	1	0	1	0		$B{\cdot}Q$ (S3)
0	1	1	0	0	1		$\overline{B{\cdot}I}$ (S4)
1	0	1	0	0	1		$\overline{B{\cdot}I}$ (S5)
1	0	1	0	1	0		$\overline{B{\cdot}Q}$ (S6)
1	0	0	1	1	0		$\overline{B{\cdot}Q}$ (S7)

TABLE 4.1: Table showing the combinations for the various phases.

Table 4.1 shows all the combinations of the inputs and phase to which they correspond. At a first glance, it is obvious to say there is a level of redundancy with unused combinations of inputs. However, all combinations (henceforth referred to as 'states') must be used in order to remove the output glitch that is inherent with this topology (discussed in the following section). The table has been arranged in such a way that , only one input (including its complement) toggles between states. For example, between state S4 and S5, the {C2, C1, C0} combination changes from $\{0,1,0\}$ to $\{1,1,0\}$ and the output phase remains the same $\{\overline{B}{\cdot}I\}$, but between state S5 and S6, the combination changes from $\{1,1,0\}$ to $\{1,1,1\}$ and the output phase changes $\{\overline{B}{\cdot}Q\}$.

One problem with stacking both stages of the phase-select is that the bias voltages for each level differ. It is for this reason that good transistor models are needed with the ability to predict accurate DC behaviour, as well as providing some insight into the AC response of a short channel MOS device (assuming an active level shifter is used.) With the SOI divider (to be described in chapter 5), the level shift was accomplished with a cascade of two source followers per output at one point, because of the high power supply. In the case of the bulk dual-modulus divider (see chapter 7), the amount of level shift is less, but it is still needed.

When designing with this topology, the dimensions of the transistors should be kept minimal so as to avoid capacitively loading the slow speed, low current consuming dividers at the tail end of the divider chain. Sadly, as will be seen in the later chapters, the binary reduction in current is not seen to the extent it should be in the case of the SOI implementation, but a more serious attempt is seen in the bulk dual modulus divider. Each conducting device in this structure should be operating in saturation mode, with as little V_{ds} as possible.

Looking at the topology, it appears as a set of quadrature mixers with their outputs connected to a '2-to-1 line selector.' However, the design is more easily explained when thought of as a line selector. Neither of the inputs to this cell is a single-tone with many harmonics, regardless of the difference in phase.

4.3 Operation and discussion of the 8-state finite state machine

After the phase selector block, the remaining dividers forming part of the division ratio are added onto the end. Thereafter, a state machine of some sort is attached in order to

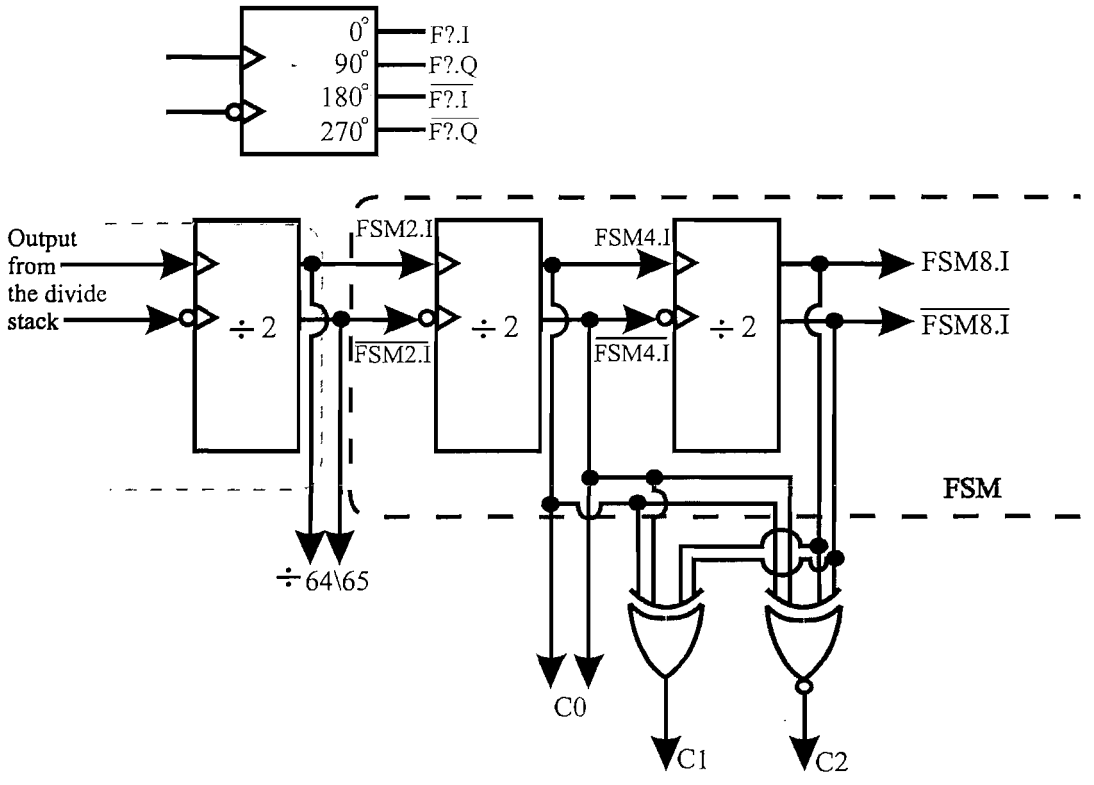

FIGURE 4.6: Schematic for the back end of the dual modulus divider, where the FSM has only 4-states.

achieve the dual modulus division (not the case in the Krishnapura design where a pulse generator precedes the state machine, in order to achieve the multi-modulus division [2]). In the previous designs, a 4-state FSM is employed and toggles through each of the consecutive phases. Unfortunately, asymmetry in the circuit can give rise to some phase transitions occuring too early, whilst others happen well inside their valid time 'window' [2].

To be able to control the new phase selector with a 4-state FSM, 4 sets of the inputs to the phase selector (out of a possibility of 8) must be chosen, by picking one of the two possibilities for each of the four quadrature phases. For the example in Fig. 4.6, the states labelled S0, S3, S4, and S7 are used with a 'Moore' finite state machine cycling through each of the four states at a rate equal to the output frequency of the dual-modulus divider.

As there are only 4 states, the FSM becomes extremely simple, with combinatorial logic translating the output of the 4 state counter. The asynchronous counter will count in binary from 0-3 inclusive and jump back to 0 before repeating the cycle again and again.

After Karnaugh mapping, the resulting Boolean equations are as follows (with refernce to Fig. 4.7):

$$C2 = \overline{D1}.\overline{D0} + D1.D0, \tag{4.1}$$

$$C1 = \overline{D1}.D0 + D1.\overline{D0}, \tag{4.2}$$

$$C0 = \overline{D0}. \tag{4.3}$$

Equations (4.1) and (4.2) are clearly XNOR and XOR expressions respectively, but owing to the choice of circuit topology for the D-type flip-flop, such a function can be omitted as the sequence is generated on one of the quadratures from the second divide stage.

Unfortunately, two problems exist. The first is if a major latency exist between the actual output of the divider and the input to the phase selector which translates to a lack of synchronicity. This gives rise to a glitch and hence incorrect division, as the extra rise and fall transitions from the phase selector will trigger an unwanted toggle in the subsequent divider(s). As illustrated in [2], this problem is not consistent in every consecutive phase jump, with some toggling occuring prematurely and others occuring in the 'window.' This can arise due to the mismatch amongst 'like' devices, low-current operation, and unbalanced parasitic RC constants. The other problem is a race between C0 and the C1 or C2 pairs of signals which results in erroneous division (pictured in Fig. 4.8). In this scenario, the C0 signal usually toggles before either C1 and/or C2 as it is generated earlier in the chain and our state machine is a ripple counter. If the current in subsequent stages is reduced by half, this problem becomes apparent at a

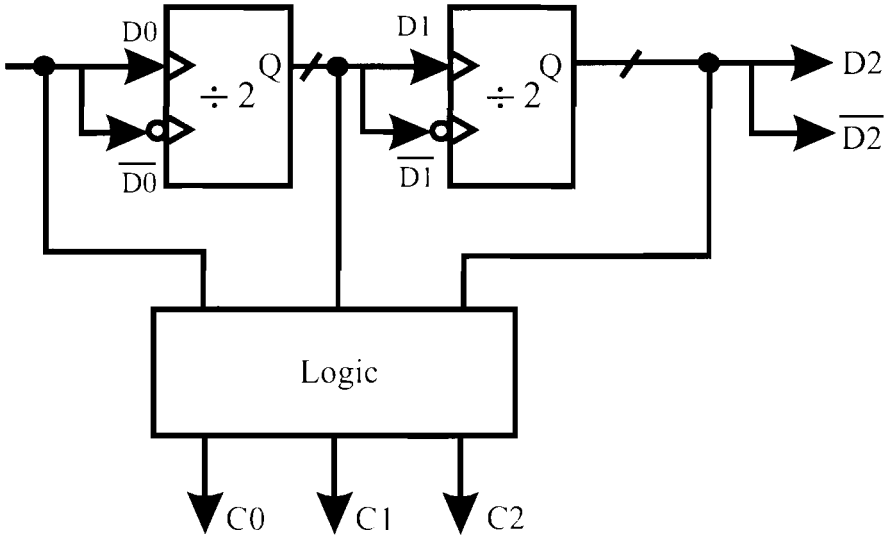

FIGURE 4.7: Schematic showing the FSM as an asynchronous counter with additonal logic giving the desired phase selector input combinations.

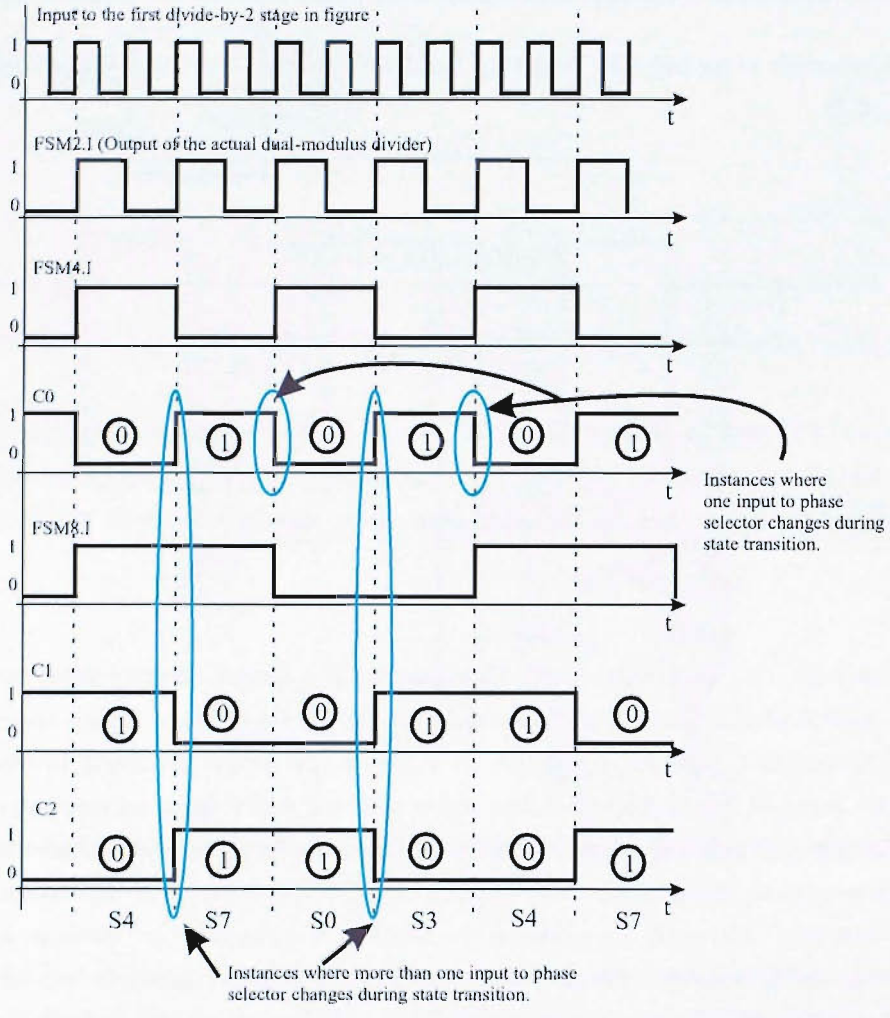

FIGURE 4.8: Waveforms of the 4-state FSM, highlighting the hazardous glitch problem.

lower frequency.

Fig. 4.9 shows waveforms from two different instants in a divide-by-33 stretch of a simulation run with the 4-state FSM. (The divide-by-32 is not shown because errors in its operation would arise in the divide stack, which will be shown later.) Looking at the particular time intervals framed by the markers, Fig. 4.9(a) gives a division ratio of approximately 33, whereas Fig. 4.9(b) produces a ratio of 29, showing inconsistency. The upper figure shows C0 changing with C1 and C2 remaining static. The lower figure shows all three control signals toggling, on the negative edge of an output cycle producing a 'race' condition, and clearly giving erroneous results. Although this result is shown for a 4GHz input, the same was observed for a 3GHz input too. Looking at other input combinations yields the same error, and this has motivated the design of an alternative state machine.

When implementing with a 2-bit synchronous counter (counting from 0-3 inclusive), the

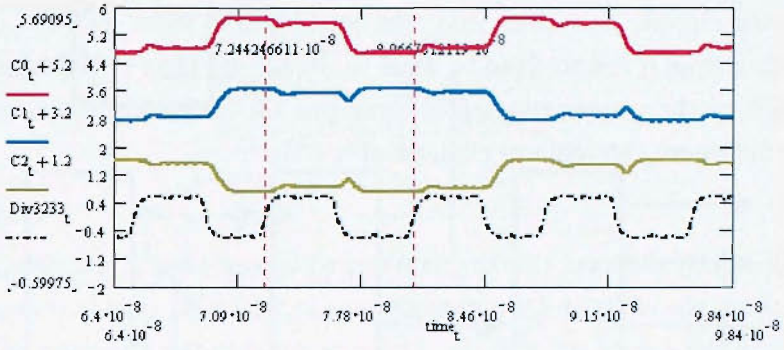

(a)

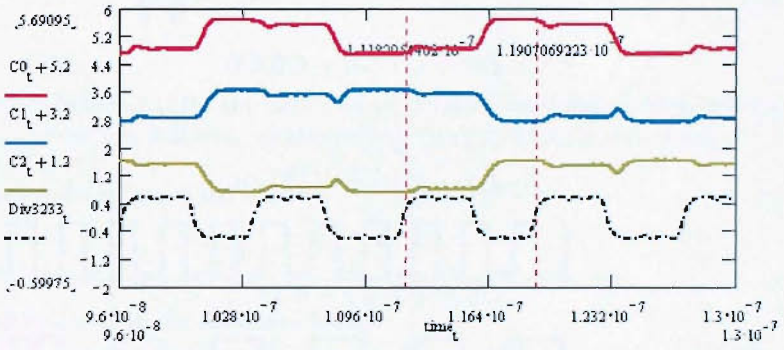

(b)

FIGURE 4.9: Simulation results of the 32/33 dual modulus divider operating in divide-by-33 mode with a 4GHz input. Note that the complication of having more than one phase select input changing one time manifests itself in the form of a division ratio error.

resulting Boolean expressions are:

$$D0 = \overline{Q0}. \tag{4.4}$$

$$D1 = \overline{Q1}.Q0 + Q1.\overline{Q0}. \tag{4.5}$$

However, a problem cited is the loading of two D-type flip-flops on the output of the previous divide-by-2 stage. By classic Boolean logic, the XOR/XNOR should exist on the outputs to generate C0, C1 and C2 as well as the feedback to the 'D' input of the second flip-flop. Unfortunately, all outputs cannot be generated by choosing the appropriate phase of the SCL flip-flops, and there still exists a race between all three of the signals implying the layout of their interconnects can be an important issue with regards to balance in the capacitive load.

To get around the problem of more than one signal changing at one time, all 8 combinations must be toggled through. Looking at the original table associated with our phase selector circuit, it is clear that two combinations exist for each quadrature and when toggling from 0 rad to $\frac{\pi}{2}$rad to πrad to $\frac{3\pi}{2}$rad and then round again. When going through each of the consecutive states, only one bit (including the complement in the context of differential signalling) changes at a time.

However, a slightly different thinking is required when using all 8 states. Although there are more states, the effective (or apparent) rate at which the phases change, must remain the same as in the 4 state solution, i.e. at a rate equal to the frequency of the output. As before, a synchronous counter is considered where it generates the desired combinations without additional glue logic after the output. After Karnaugh mapping, the Boolean equations for the logic are:

$$D2 = C1.\overline{C0} + C2.C0, \tag{4.6}$$

$$D1 = C1.\overline{C0} + \overline{C2}.C0, \tag{4.7}$$

$$D0 = \overline{C2}.\overline{C1} + C2.C1. \tag{4.8}$$

Unfortunately, this state machine must operate at a speed twice that of the actual divider output. Hence, it can be driven by the second from last divide stage in the asynchronous chain. Another problem witnessed by observation is that combinatorial logic is present at the output, leading to the problems of extra power consumption and conversion between different logic styles. Nevertheless, the solution exists and should be kept as a contingency.

The last solution is the one implemented in the dual-modulus dividers presented in the remainder of this thesis. Here, a 2-bit asynchronous counter generates 8 states, and although it possesses an inherent 'ripple' characteristic, its implementation is an elegant one. Again, it is the master-slave action of the flip-flops that has achieved an 8-state machine, with the desired effective phase transition frequency.

Looking at Table 4.2, the configuration of this state machine is generated without the need for Karnaugh mapping or Boolean logic, as observations should demonstrate that no logic is required. In this table, the first three columns refer to the ripple counter output, whereas the next three columns exhibit the wanted outputs for the phase selector circuit.

Each line in this table is one half a cycle of the input to the FSM, which happens to

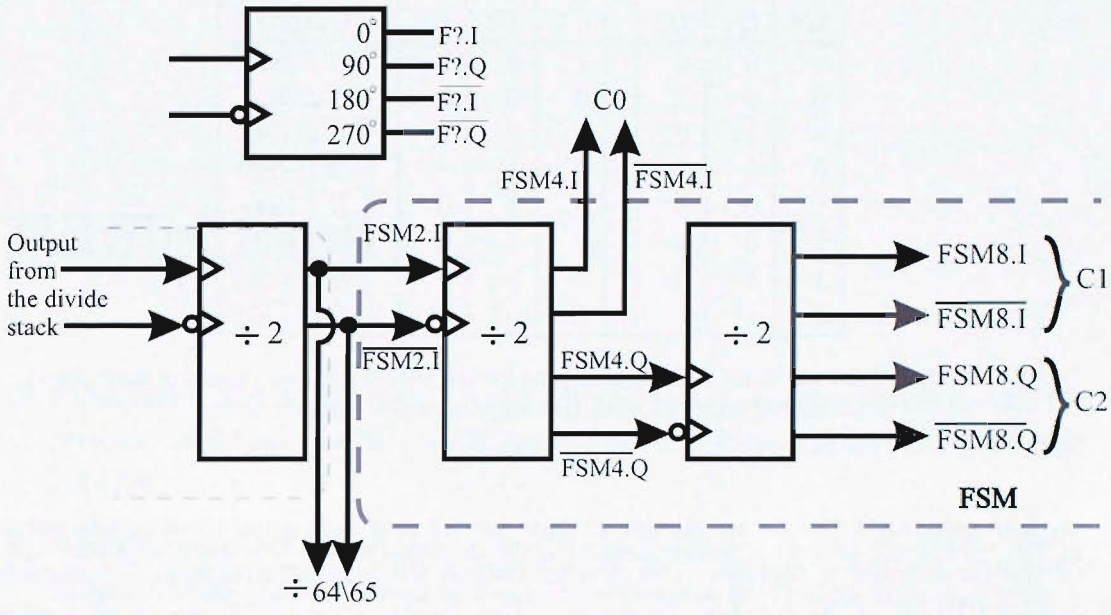

FIGURE 4.10: Schematic for the back end of the dual modulus divider, where the FSM now has 8 states, implementing Gray code style switching.

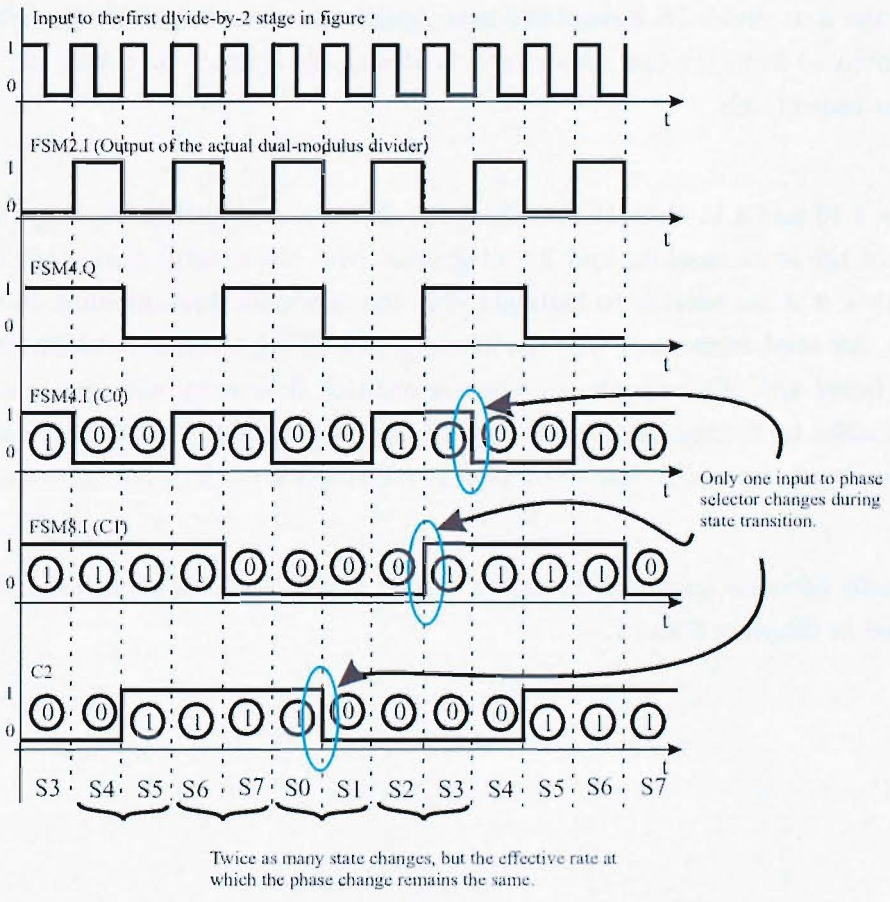

FIGURE 4.11: Waveforms of the 8-state FSM, showing the 'Gray' code style switching.

be the in-phase output of the final divider stage (excluding any state machine on the output). Hence this input is the same as that of the 4-state FSM but, looking at the

$Q2$	$Q1$	$Q0$		$C2$	$C1$	$C0$		Output.
0	0	0		1	0	0		$B{\cdot}I$ (S0)
0	0	1		0	0	0		$B{\cdot}I$ (S1)
0	1	0		0	0	1		$B{\cdot}Q$ (S2)
0	1	1		0	1	1		$B{\cdot}Q$ (S3)
1	0	0		0	1	0		$\overline{B{\cdot}I}$ (S4)
1	0	1		1	1	0		$\overline{B{\cdot}I}$ (S5)
1	1	0		1	1	1		$\overline{B{\cdot}Q}$ (S6)
1	1	1		1	0	1		$\overline{B{\cdot}Q}$ (S7)

TABLE 4.2: Table showing the combinations for the various phases (ignoring their complements) together with the outputs of the ripple counter.

pattern generated, the actual change of state occurs at a rate equal to twice the output of the dual-modulus divider. The desired output C0 is generated from the in-phase output ('I') Q1 of the counter. To generate the outputs C1 and C2, one must consider their transitions relative to the output of Q1. It should be apparent that C1 changes on a phase $\frac{\pi}{2}$rad lagging behind the output C0. Hence, the quadrature output ('Q') drives the next divide-by-2 stage in the state machine. The required signals C1 and C2 are generated from the last divide-by-2 in this ripple counter by taking its 'I' and 'Q' outputs respectively.

Figures 4.10 and 4.11 show the result of the derivation, explicitly showing the configuration of the state machine and its integration with the general dual-modulus divider. With this, it is noteworthy to highlight that the minimum dual-modulus divide ratio is 8/9 for this implementation, whereas for the 4 state FSM, the minimum divide ratio can fall to being 4/5. This simple yet effective solution does away with any combinatorial logic thanks to the quadrature outputs of the SCL divide-by-2 stages. The penalty, however, is an increase in the static power consumption of the whole divider.

Its benefit becomes apparent through a correct higher speed of operation. This will be included in chapters 5 and 7.

Bibliography

[1] J.Craninckx and M.Steyaert. A 1.75GHz/3V dual modulus divide-by-128/129 prescaler in 0.7μm CMOS. *IEEE Journal of Solid-State Circuits*, 31(7):890–897, Jul 1996.

[2] N.Krishnapura and P.R.Kinget. A 5.3GHz programmble divider for HiPerLAN in 0.25μm CMOS. *IEEE Journal of Solid-State Circuits*, 35(7):1019–1024, Jul 2000.

[3] R.Rogenmoser, Q.Huang, and F.Piazza. 1.57GHz asynchronous and 1.4GHz dual modulus 1.2 μm CMOS prescalers. In *IEEE Proc. Custom Integrated Circuits Conference*, pages 16.3.1–16.3.4., San Diego, May 1994.

Chapter 5

Dual-modulus 64/65 divider in 0.35μm SOI CMOS

This chapter concentrates on the design of a dual-modulus SOI divider (modulo 64/65). The goals in this part of the study were to investigate new circuit architectures made possible in SOI, and whether any power and/or speed advantages could be obtained. The specification for this divider has been open and hence the choice of division ratio is arbitrary, but its high-frequency operation is in the interest of targetting this design for a modern communications application. Therefore, it has been designed to run as fast as possible with realistic inputs and load terminations.

One of the highlights of this design is the novel method of stacking current steering divide-by-2 stages (SOI CMOS) asynchronously, introducing a 'true' division of bias current as one moves further down the division chain. State-of-the-art silicon frequency divider designs have not yet demonstrated such aggressive reduction in bias current from stage to stage. Another strong feature of this chapter's work is the incorporation of the current steering phase selector and accompanying glitch-free controller in SOI CMOS technology. The stacking is not limited to divide-by-2 stages, and the dual-modulus divider to be presented has a synchronous divide-by-8 stage too. The concept of stacking divide-by-2 stages exceeding the designated power supply for the technology without destroying any of the devices, is yet another example of advancing the state-of-the-art as the choice of voltage swing is power supply independent. It is hoped the reader will appreciate how this style does away with level-shifters by adopting this current-mode cascade style. Towards the end of the chapter, a successful fabrication together with measurements shows that the preceding work to have been taken from concept to realisation.

5.1 Circuit design

5.1.1 Phase selection prescaler architecture

Many of the ideas discussed in the first four chapters are brought together in this section, culminating in an interesting and novel topology. It is hoped the reader will intuitively understand how the circuit evolves from a single divide-by-2 sub-block, to that given in the title of this chapter. As the structure, design and operation of the standard SCL divide-by-2 have already been discussed, the text below will simply move one level higher in the schematic hierarchy when referering to it, although there will be times when a closer look inside the instances is called for.

To achieve division by 64 asynchronously, the number of cascaded divide-by-2 cells is:

$$log_2(64) = 6. \tag{5.1}$$

Obviously, this does not take into account any additional dividers required by the control logic, to be presented later. The fact that there are two consecutive division moduli simply means that a high frequency input cycle is lost somewhere. A block diagram of the divider has been captured in Fig. 5.1, and we will show in the following subsections the required blocks to use current steering as well as the development of new circuit topologies.

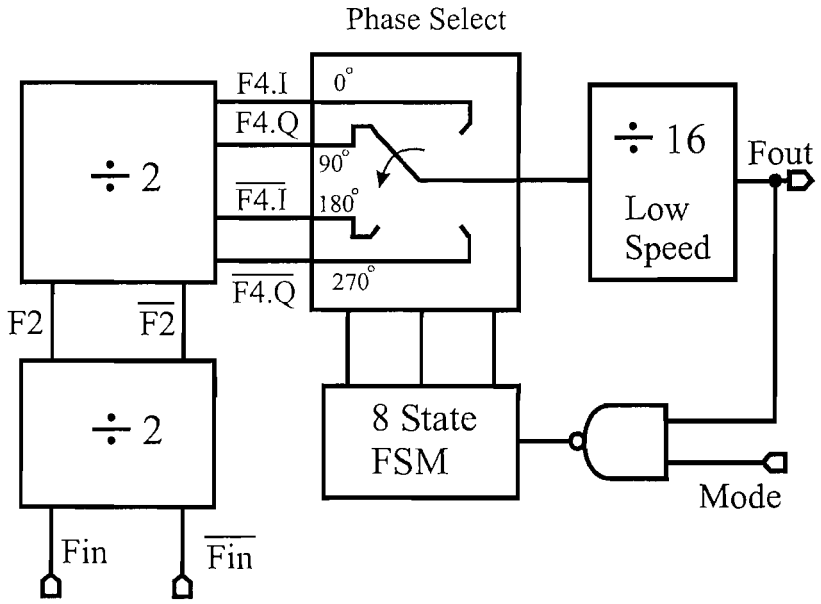

FIGURE 5.1: A block diagram of this chapter's dual modulus divide-by-64/65 divider.

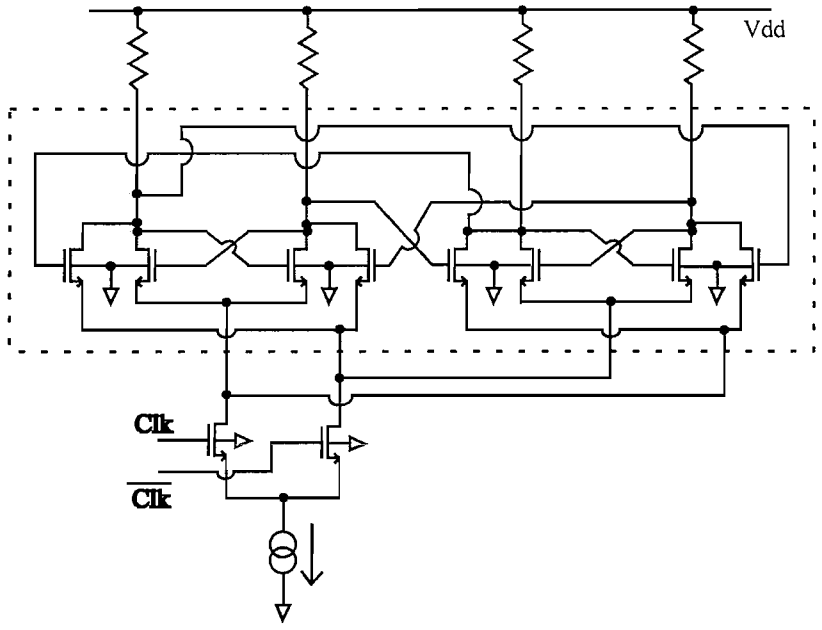

(a) Bulk.

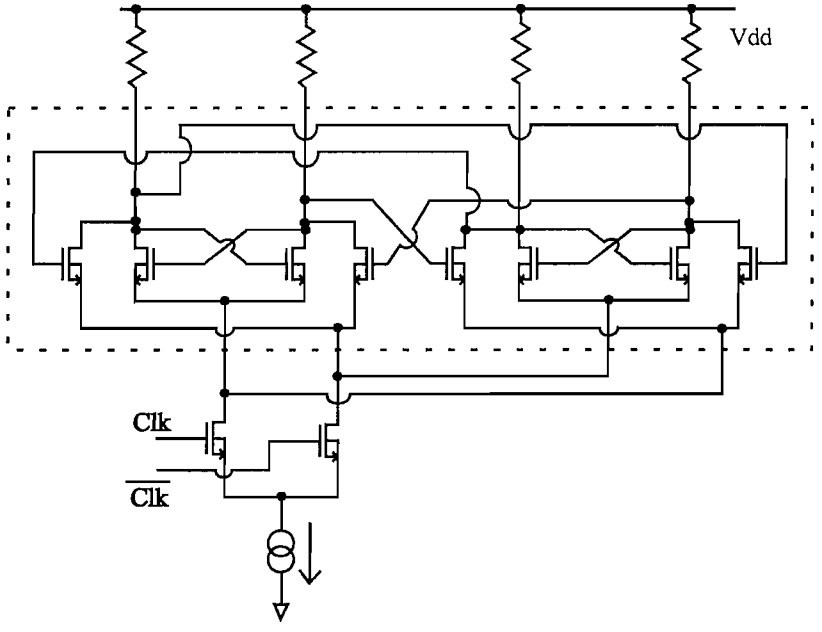

(b) SOI.

FIGURE 5.2: SCL divide-by-2 in bulk and SOI CMOS.

5.1.2 Stacked current steering logic in SOI

A flexibility exists with PDSOI that is not present with bulk CMOS and this is shown in Fig. 5.2, though the difference is subtle. In bulk CMOS processing, and assuming a normal p^- doped substrate, all bodies in which the channels of NMOS devices invert

must be fixed to an electrical value (usually the circuit ground). As an NMOS transistor is stacked on top of another NMOS transistor, the back-gate effect causes the threshold voltage of the upper transistor to rise, thus requiring a larger voltage drop between its gate-source terminals in order to achieve the same level of drive. In classic submicron circuit design, where the power supplies reduce with smaller geometry process generations, this is clearly prohibitive. One suggestion might be to use N-well and then introduce a p^- well within it, after which both wells would be tied to the source terminal of the cascoded transistor. The problem here is that additional masks are needed and isolation relies on a reverse bias between the N-well and bulk substrate, as well as a packing density penalty. With the SOI MOS process, the back oxide and shallow trench isolation together generate silicon 'islands' in which the MOSFET exists. There is no requirement to fix the potential of the silicon film above the buried oxide region to the same value (for every NMOS or PMOS transistor) and hence transistors stacked in a cascode fashion theoretically possess the same effective threshold voltages (assuming the bodies are tied to their respective source terminals, as illustrated in Fig. 5.2(b)). As long as the field across any gate oxide and buried oxide (though much larger) doesn't exceed the breakdown voltages of the oxides, the power supply to the circuit may increase beyond that rated for the technology in conventional use. This is certainly useful for any designer as the effect of back-gate bias can be more or less eliminated from the design equations.

Within the basic divide-by-2 cell, there are 3 levels of NMOS transistors stacked upon one another followed by a load (maybe another MOSFET or a passive element). This is another design problem faced by IC designers (using bulk CMOS processes) as newer geometries become available, and they must design circuits to operate within an ever decreasing power supply. With the 3 levels of MOS transistors, the headroom for large signal swing is severely limited and is naturally a problem for circuits designed to exhibit high signal-to-noise ratios. In zero IF architectures, bipolar technology is preferred as it tends to be 'cleaner' in terms of noise compared with CMOS devices. The bipolar transistors superior flicker noise performance warrants its use in the local oscillator section of the receiver. Carrier trapping and releasing at the interface states and thermal noise due to gate and channel resistance are both significant in CMOS. If CMOS was insisted, then you would have to increase the current and area of the responsible devices (those responsible for generating the in-band noise) as well as increasing the power supply of the circuit, and the more the technology scales to smaller dimensions, the worse it gets (bipolar devices have a higher g_m for a given cut-off frequency) [1]. With PDSOI, extra headroom can be afforded, within reason, by simply using a higher supply voltage. A cause for concern when designing mixed signal ICs might be the multiple power supplies on the same die and the need to generate stable voltage references on-chip or power them externally via area-expensive bondpads.

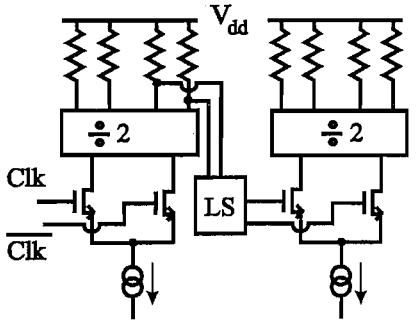

(a) Voltage mode cascading

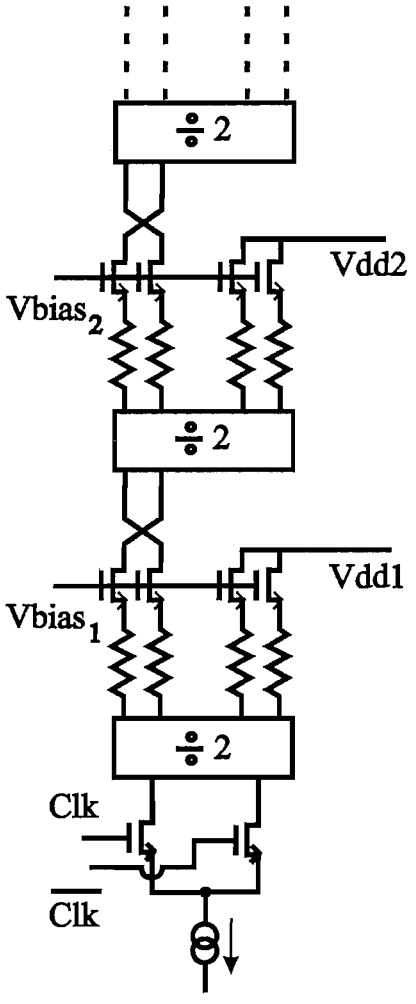

(b) Stacking with current mode cascading

FIGURE 5.3: Diagrams showing methods available in SOI design for cascade divide-by-2 stages asynchronously.

This SOI divider attempts to exploit this benefit of local isolation between transistors. The advantage of the proposed topology is due to the elimination of high frequency level shifting between successive SCL divide-by-2 cells when cascaded asynchronously. Figure 5.3 illustrates the difference between the conventional voltage-mode cascading (available to both PDSOI and bulk CMOS technologies) and the current-mode cascading employed here (only available to PDSOI CMOS technology). The level shifters are not a mandatory requirement, even in low supply voltage CMOS technology. With enough DC current and load resistance through each latch, the common-mode output voltage of the divider can be pulled considerably lower than might be expected intuitively (only at high frequencies though). This allows direct coupling of stages, as seen in chapter 6, encountering less risk than would be in a bipolar design, where collector-base saturation is a concern. However, keeping the signal swing low and the output common-mode voltage high does have the advantage of lower power consumption per stage, thus calling

for level shifting of some sort, as drawn in Fig. 5.3(a). One method is to use a resistive divider, which theoretically adds no capacitance on its inputs or outputs. This type of level-shifter is relatively intolerant of mismatch, as the ratio of the two resistor is more important. However, the resistor must be large in value so that it does not look like a low impedance to the driving stage, drawing current, which in turn raises the issue of noise. There is also the issue of attenuation of the signal and this would be by the same factor as the reduction in bias. Another method is to AC couple the output of one divider capacitively with the input of another, resetting the bias voltage to the desired input common-mode voltage using some form of high impedance divider. Unfortunately, capacitance on-chip has the drawback of parasitic capacitance to substrate as well as large area consumption when used with a resistive divider, to ensure the low frequency pole is sufficiently far from the toggle frequency on that node, resulting in a narrower band circuit (NOTE: the issue of capacitance to substrate is less severe in the case of SOI technology [11] as there is field oxide above buried oxide, hence giving speed gains of 10% or more over the equivalent bulk device. In some cases, speed gains of 20% have been witnessed at the 0.22μm technology node [7]). On-chip capacitors tend to be lossy; without modelling and characterisation of the capacitors at high frequencies, it is unwise to use them in a first attempt of a complex circuit. A third method uses a pair of source followers which are not differential and account for a notable portion of the power budget at high frequencies.

Figure 5.3(b) shows the SOI method of stacking [12], where the current from a lower divide-by-2 cell drives and clocks the subsequent upper divider cell. In the traditional divide-by-2 cell, the current is steered through one of a pair of loads, giving rise to a voltage drop across it. Assuming bias translation (where necessary), the AC voltage is applied to the 'clock' transconductor, giving rise to an AC current switching between different elements in both the master and slave latches. The stacking method shown in the figure removes this redundancy (applicable to SOI technoogy) and at the same time lowers the nominal power supply per divider stage, as the active current source has been reduced. Assuming no leakage in the devices, the current consumption of stage N is exactly half that of stage N-1. Figure 5.4 shows the transistor level abstraction of the stacked divider. Being a PDSOI body-tied design, explicit body-to-source connections have been shown on the schematic. The transistor arrangement within the dotted box can be considered a macro cell that is rigid, except for the varying transistor dimensions. Two points needing to be addressed in this figure are the cascode transistors between stages and the cross-over of the current line. The cascodes are there to mimic both a voltage source and the current source, depending on which terminal one observes. For the lower divider, looking into the source terminal of the cascode transistors (assuming saturation), one should expect a low impedance; hence variation in current should have little effect on the voltage at that node. Looking into the drain terminal of those very same transistors should yield a higher impedance (assuming saturation), its behaviour

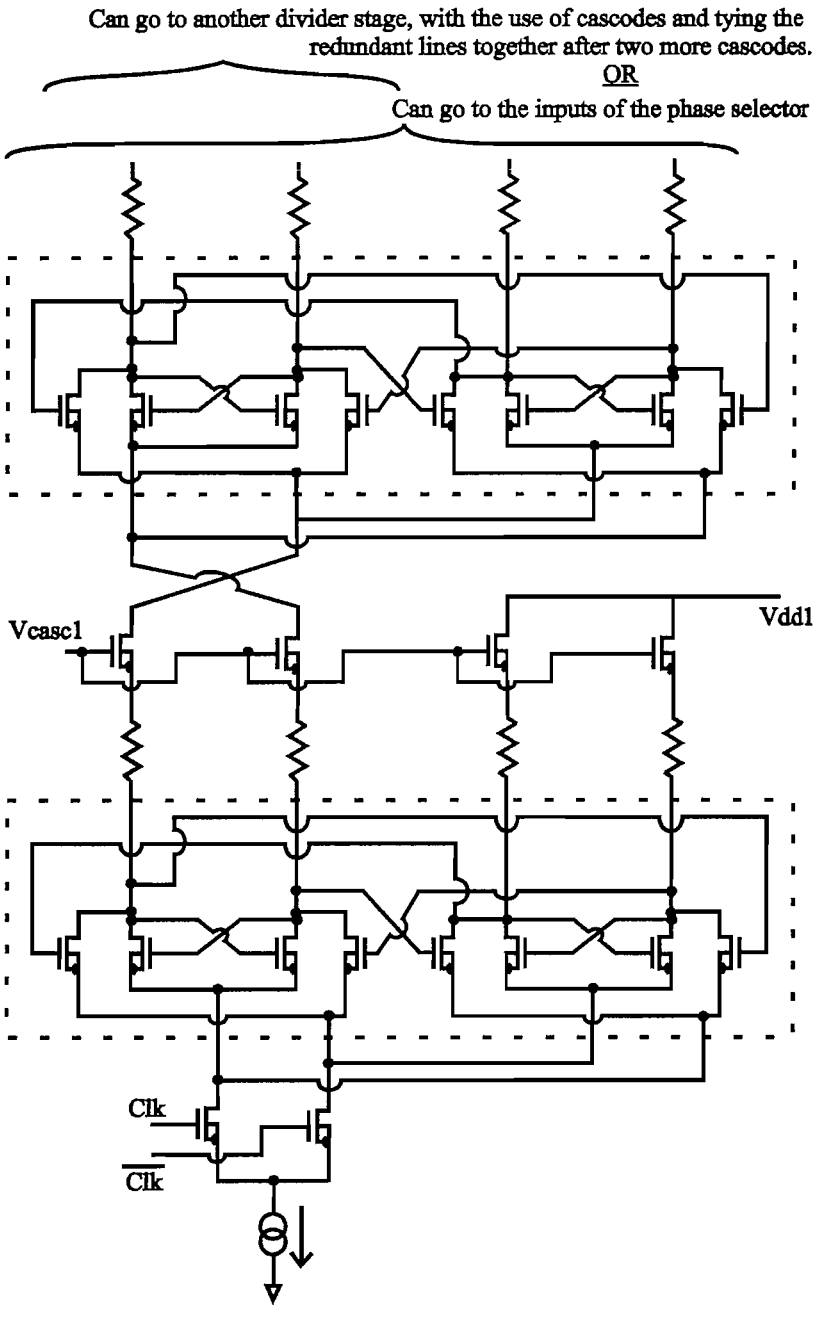

FIGURE 5.4: A closer look at the transistor implementation of stacking with current mode cascading.

likened to that of a current source. The crossing over of the output lines is due to the inversion that takes place between the output of divider N and the input to divider N+1 and this ensures all the dividers toggle on the same edge. When the output is high, no current flows, giving rise to a single-ended voltage that is closer to V_{dd} than its differential neighbour. This voltage couples to the same transistor as before, forcing it to draw current, and this would be inconsistent with the SOI divider if there were no crossing.

After two divide-by-2 cells have been stacked one on top of the other, a pair of outputs need to be capped with cascodes and tied to some fixed potential (labelled as Vdd1); there are unused in the switching process. It should be obvious that half the current is being lost here as static DC and, with stacking in mind, it would be useful for other blocks further in the chain. Fortunately, as will be described later, the current into this node will be supplied by a later divide-by-2 cell providing valuable current reduction (see Fig. 5.11).

5.1.3 New current steering phase selector in stacked SOI logic

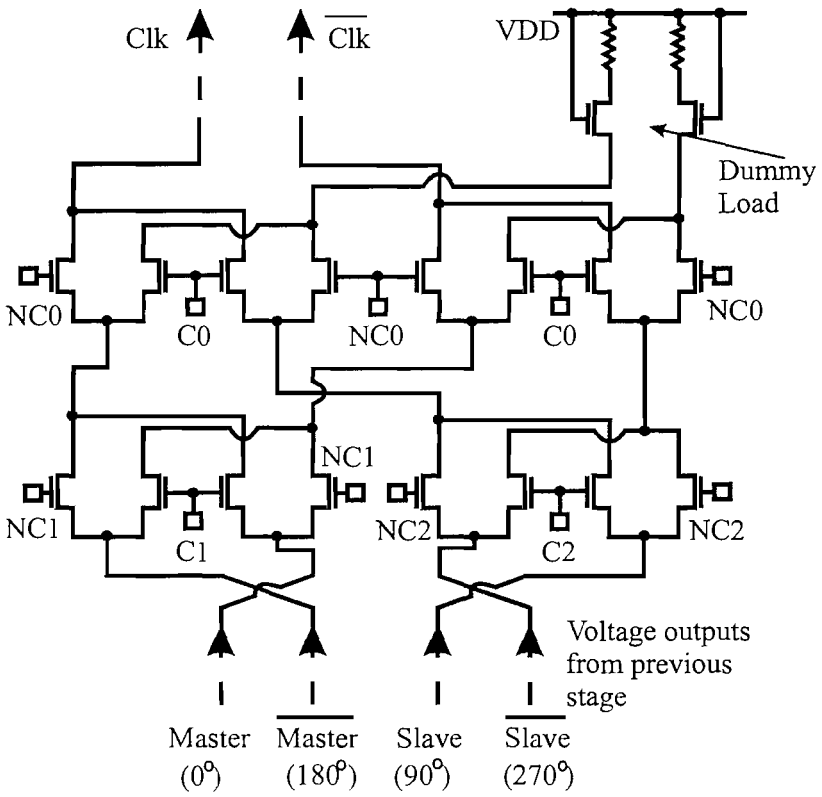

FIGURE 5.5: Schematic of the new phase selector.

As is the focus of this design, the next block to be stacked is the phase selector shown in Fig. 5.5. Unlike with the first divide-by-2, all four current switching outputs are used and hence no current reuse by another cell is possible. Differential pairs of signals are selected and inverted (if necessary) before emerging at the top of the block with the unused pair of signals simply terminated into a dummy load. As will be described further into this section, the off-stack dividers (not cascaded like those shown in Fig. 5.3(b)) hang off the same power supply as the main stack and so the outputs controlling the phase-selector inputs must be level shifted down. As the phase-selector has two switching levels, two different arrangements of level shifters are needed. When the dual-

modulus divider is set to run as a divide-by-64 unit, the outputs of the FSM do not change and hence inputs to the phase selector remain static. When chosen to divide-by-65, the toggling state signals of the FSM arrive at the inputs to the phase selector and cause it to change phase once in every 64 input cycles.

5.1.4 Fixed divide-by-16

From here onwards, the design as in Fig. 5.1, proceeds as a chain of dividers connected in asynchronous fashion. To recall, the first divide-by-2 stage in the stack sits on a 2mA current sink. The second stage hence runs on 1mA and, after the phase selector, 500μA is left for the upper stages. To maintain the same logic swing and gain per stage, this binary division of current forces the load resistances of subsequent stages to be increased at an exponential rate, and this is not practical when considering the layout of the devices (parasitic capacitance does not scale down with decreasing channel widths). If one considers the current reduction if all six divide-by-2 stages were stacked asynchronously within the power supply, one would find that only 31μA would be available to drive each D-latch in the last stage, needing 20kΩ load resistors to give the same 600mV swing used throughout the design. The problem was also compounded by the higher supply and hence larger level-shift further on down the chain. Unfortunately, simulations showed the divider to give very slow transition times and this would affect the changeover (i.e. metastability) further along the chain. An example consisting of a stacked divide-by-128 circuit is illustrated in Fig. 5.6, with Figs. 5.7 and 5.8 showing the outputs at each of the stages up the stack. In particular, Fig. 5.8(d) clearly shows a marked reduction in rise time.

This was solved using a synchronous divide-by-8 at the output, ie on top of the phase selector, with the last division-by-2 supplied by its own current source. A schematic of this arrangement can be found in Fig. 5.9. The synchronous divider is made up of 4 D-type flip-flops driven by the same clock frequency and the 500μA of current is divided by the 8 D-latches, resulting in 62μA of current though each. The immediate benefit is that the outputs do not ripple through this section of the divider and hence the transitions occur much faster. The load resistance in this section is lower, with the final divider also benefiting from a lower power supply.

It is interesting to note that although the signal path travels from the bottom of the stack to the top, conventional current still travels from the top supply rail to ground. It seems that current is being lost as more divide stages are encountered, but current is actually conserved because it accumulates from an off-stack divider and a dummy load on top of the phase-selector.

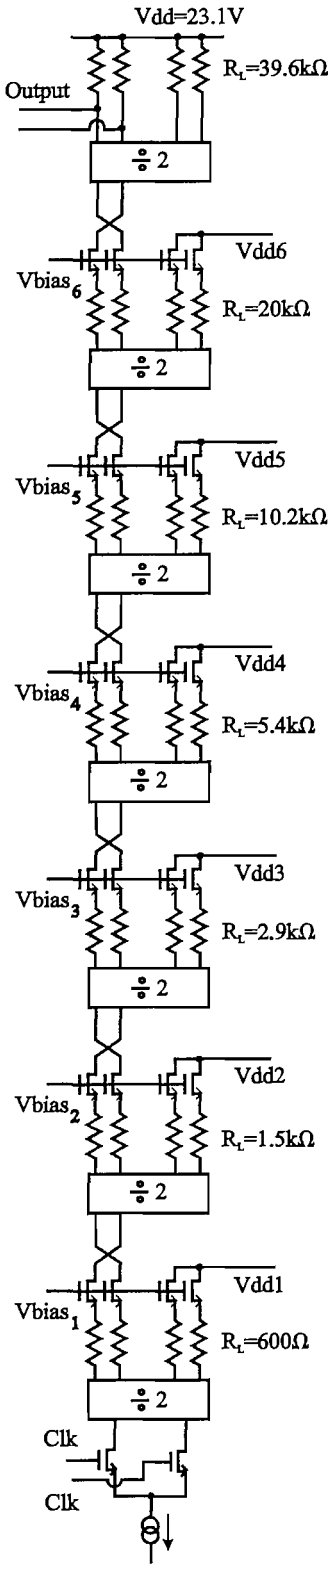

FIGURE 5.6: Schematic of an asynchronous divide-by-128 circuit employing current-mode cascading of divide-by-2 stages throughout.

5.1.5 Modulus control

As seen in the block diagram of Fig. 5.1, modulus control is achieved by means of 'gating' the toggling of the FSM. This diagram is a little misleading because the NAND gate should block the clock and not force the output of the second off-stack divider to take up a particular logic state. Another problem that exists here is that the operation of the NAND gate causes erroneous division when used in a programmable divider loop, if care is not taken to match the trigger edge type of the divider to the 'inhibit' output of the NAND gate. This effect is not obvious when used as a dual-modulus divider, as it is simply a transient that settles into the correct division rather quickly. However, recalling the truth table of a NAND gate, the output fails to toggle if another input is at logic '0', and remains at logic '1'. This means that in our SCL divider, having the NAND

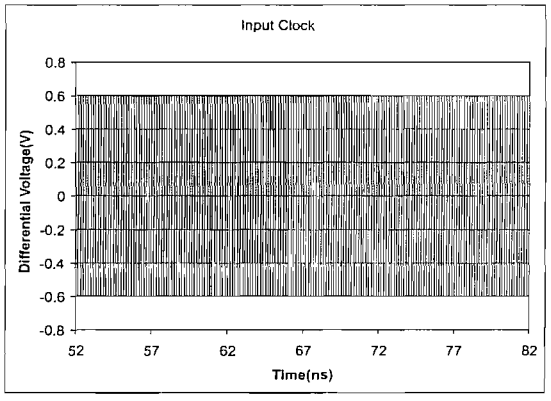

(a) clock input

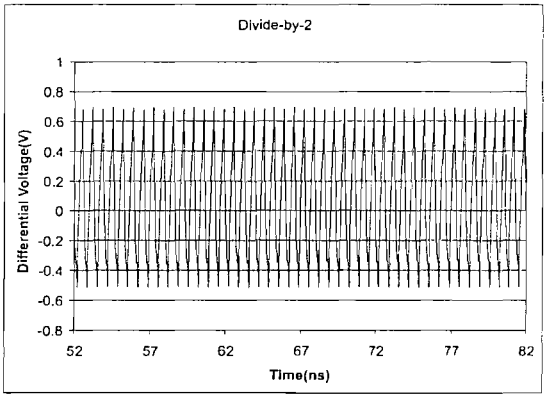

(b) divide-by-2 output (R_L=600Ω)

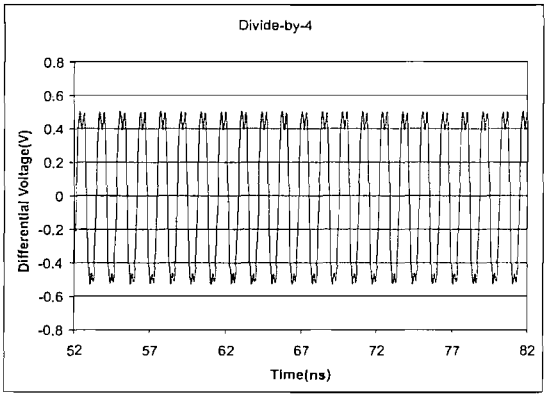

(c) divide-by-4 output (R_L=1.5kΩ)

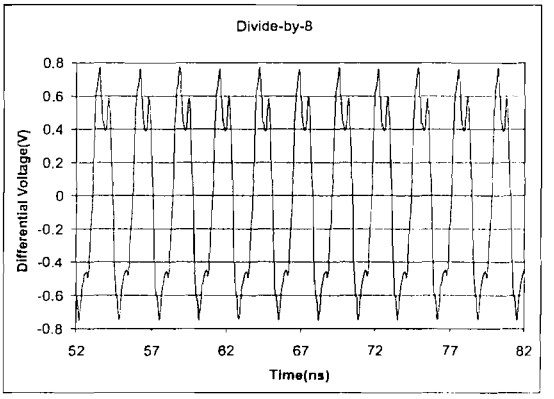

(d) divide-by-8 output (R_L=2.9kΩ)

FIGURE 5.7: Plots of a divide-by-128 (employing current reuse), as an example showing the limitations of excessive current division in a divider stack. Simulated inside a 23.1V supply, with the clock arriving on 20μm wide NMOS transistors.

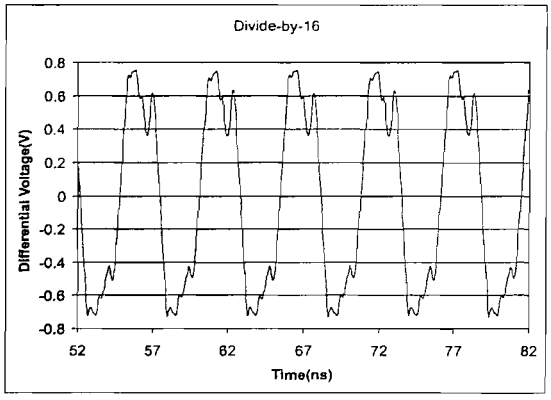

(a) divide-by-16 output (R_L=5.4kΩ)

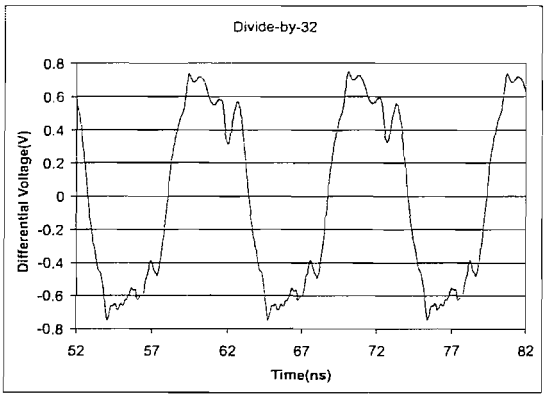

(b) divide-by-32 output (R_L=10.2kΩ)

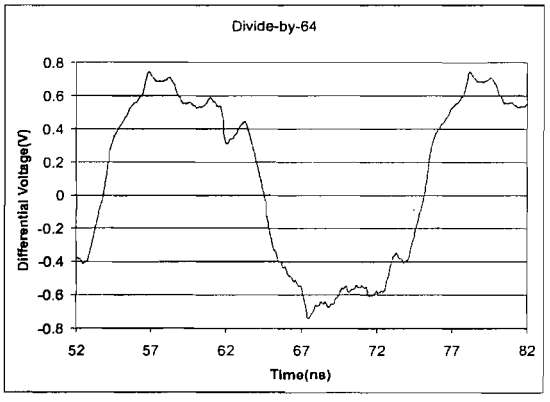

(c) divide-by-64 output (R_L=20kΩ)

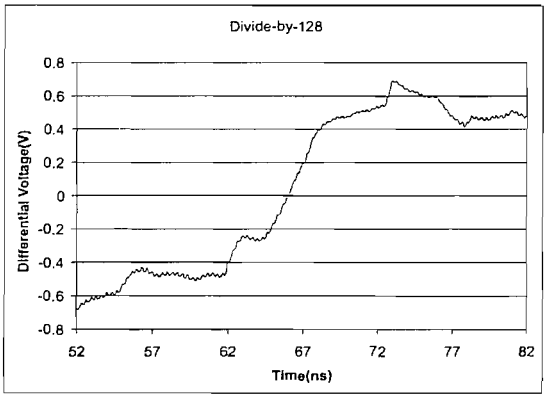

(d) divide-by-128 output (R_L=39.6kΩ)

FIGURE 5.8: Plots of a divide-by-128 (employing current reuse), as an example showing the limitations of excessive current division in a divider stack. Simulated inside a 23.1V supply, with the clock arriving on 20μm wide NMOS transistors.

gate embedded into the output stage of one of the FSM dividers has a direct effect on the output of the divider (waves 3 and 4 in Fig. 4.4). In fact, what should happen is that further clock edges to these last two FSM stages would be inhibited with the state of the FSM held, preventing any edges clocking the remaining stages. This requirement has been noted and correctly implemented in the programmable divider designed in bulk CMOS to be described later. Incidentally, the NAND gate performing the modulus control is a differential SCL type. Therefore, a CMOS-to-SCL level translator and CMOS buffers are placed between the corresponding bond pad and the NAND gate.

5.1.6 Power supply tolerance conditions

With the stack dividing-by-32, the last divide-by-2 in the dual-modulus divider plus two more divide-by-2 stages for the FSM remain to be added to the divider. In the final

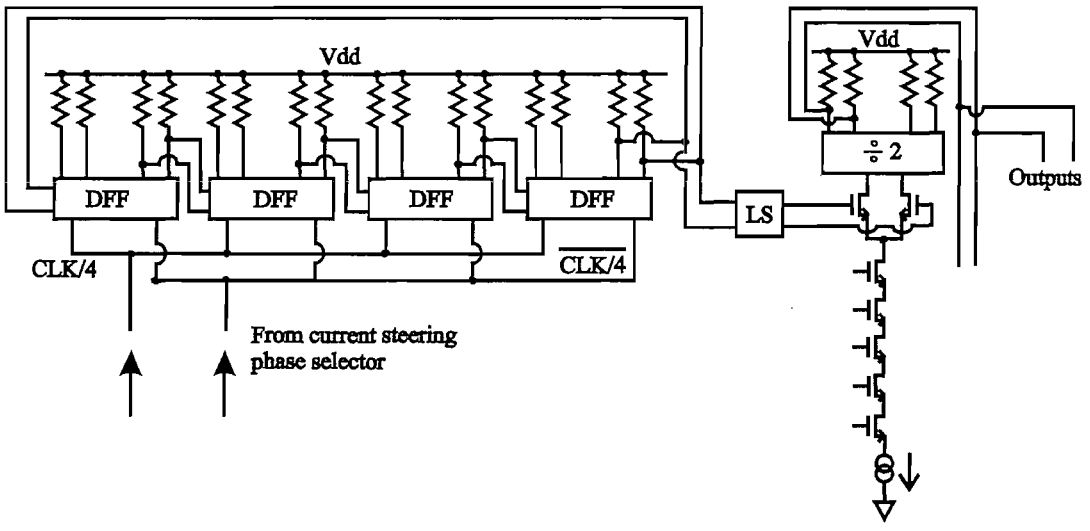

FIGURE 5.9: Schematic of the synchronous divide-by-8 stage sitting above the phase selector with its outputs voltage-mode coupled to a further divide-by-2 stage sitting on its own current source.

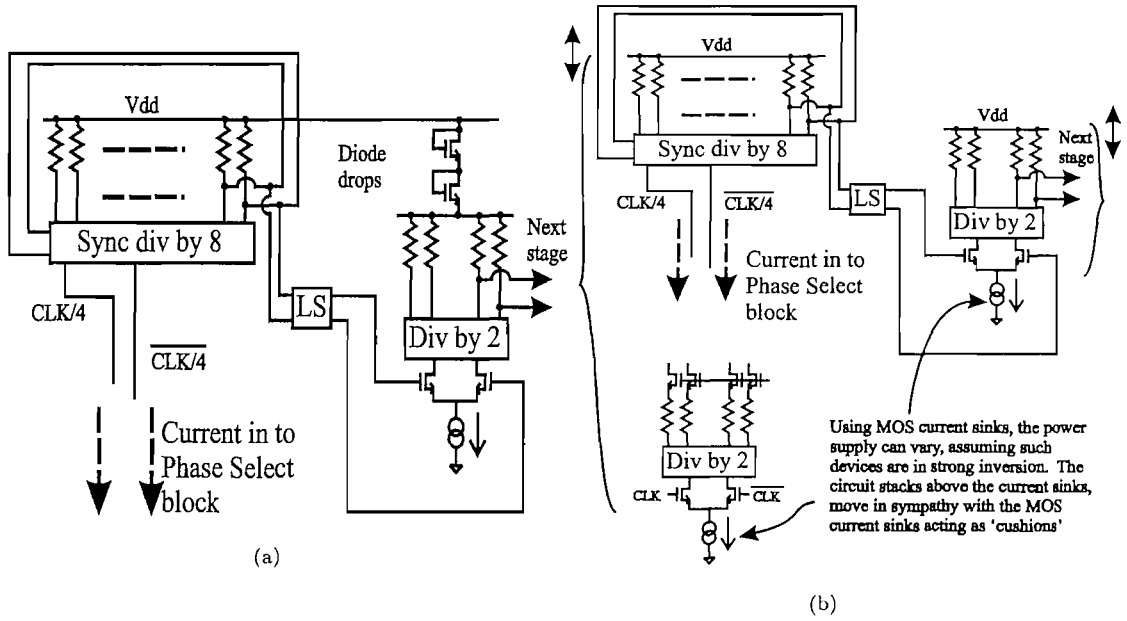

FIGURE 5.10: Difference between multiple power supplies and 'cushioning' effects of the MOS current sinks.

design, the three off-stack dividers are directly 'hung' from the 6.8V supply. The choice to have one power supply on chip as opposed to one for the stack and another for the off-stack elements requires many pads This appears to contradict the need for a 2.5V power supply as well as 1.7V, 3.3V and 5.1V pads but these are due to the fact that the 2.5V supply is for the low frequency output and the modulus control, whilst the latter three are for the protection against ESD. Two alternatives are available here (see Fig. 5.10): either the divider can 'hang' from the 6.8V supply and be supported by a

number of cascodes so as to ensure no gate oxide is stressed or the power supply can be dropped (using MOS diodes) to a level such that the divider outputs are compatible with the phase selector inputs. The idea of the MOS diodes works because the current sunk into the SCL dividers is relatively static (unlike a CMOS divider.) Sadly, though different power supplies generated on chip may be feasible, it is not desirable in terms of testability. To support the dividers in our design, a number of cascode transistors sit under the running divider macro (including the 'clock' transistors). Source followers (also supported with cascodes) running at relatively high currents are used to level shift between the output of one stage and the clock input of a subsequent stage. They are also required to level-shift between the output of the FSM dividers and the inputs to the phase selector. For a future re-design, a reduction in current consumption owing to these cells (or even an alternative) should be pursued. The idea of using a classical CMOS topology in the latter stages would remove the need for level shifting and hence lower the power consumption, but the requirement for quadrature outputs, even in the latter stages of the divider, has steered the design towards the current topology. Variation on the 6.8V power rail can effectively be 'cushioned' as long as the current sources at the bottom of the stack are sufficiently in saturation.

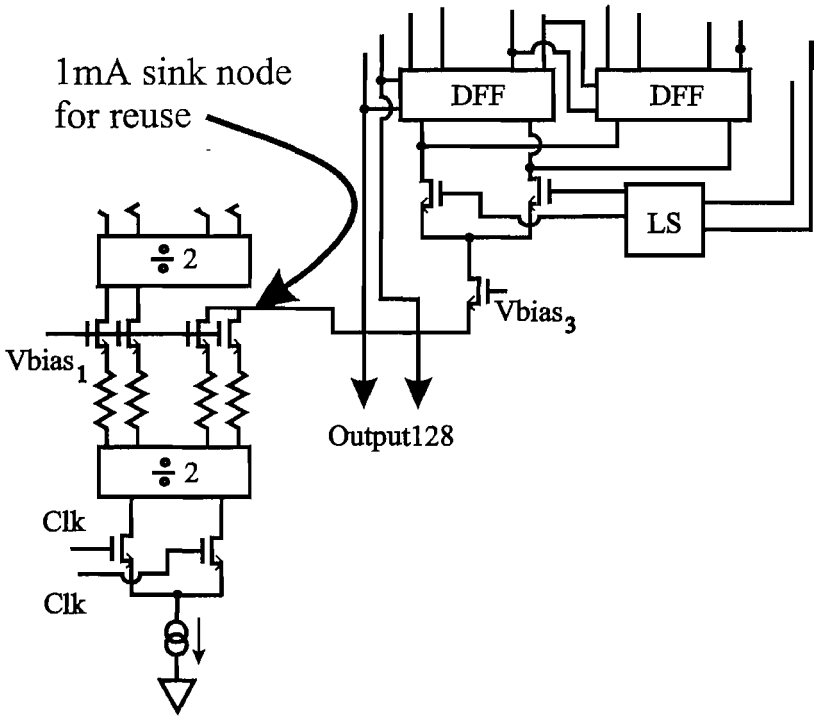

FIGURE 5.11: Example of current re-use for an off-stack divider [12].

5.1.7 Current re-use outside the divide path between separate signal paths

One of the key features of this SOI design is the reuse of current, and the second off-stack divider (used for the FSM) capitalises on that facility in a very straightforward manner. As pictured in Fig. 5.11, current from an off-stack divide stage is sunk into the common point of one of the latches (master or slave) in one half of the first divide-by-2 stage of the divider stack. This particular divider uses 1mA of current and is sized accordingly. However, with clever optimisation it should be possible to have all three off-stack dividers running off the main stack. The author admits to a 'conservative' use of current in the slower stages of this divider and this was decided upon to increase the chances of a functional design after fabrication. However, at present, 1mA, 1mA and 2mA are consumed by the first, second and third off-stack dividers, respectively (in order of decreasing speed.) By optimising the aggregate current of the last three divide stages to be less than 1mA, the core of the dual-modulus divider could be made to function on 2mA. Thereafter, attention should be focussed on the inter-stage coupling, namely the source followers.

In order to bias the gate terminals of the cascodes, a simple current-controlled bias generator is included on chip. Owing to an excellent DC model within STAG (referring to the strong-inversion mode), faith can be put in the simulation of such a network and optimising its total current consumption. Due to time pressures, this was not exploited to the degree it could have been and hence the current consumption is rather high. Diode connected MOS transistors either sit on a NMOS device or hang from a PMOS current-source transistor biased in a current mirror fashion. With the ability to short the body to the source terminal of every individual transistor, the bias (and hence the stress) across the gate-source terminals can be kept to a minimum whilst stacking higher. A chain of diode-connected MOS transistors also works to generate the on-chip input common-mode voltage. Two series-connected 10kΩ polysilicon resistors sit between the bias generation node and the RF ground of the 50Ω terminations to provide the input common-mode bias. A centre-tap exists between the 10kΩ resistors and is brought out to a pad where this voltage can be sensed and/or overriden.

5.1.8 Source followers

A clever aspect of the divider design is the stacking of SCL divide-by-2 stages, accomplished by exploiting PDSOI CMOS technology. However, this has not stopped the design from employing source followers for the role of shifting the common-mode bias of some lower frequency signals to a different (lower) bias. The divider stages that are

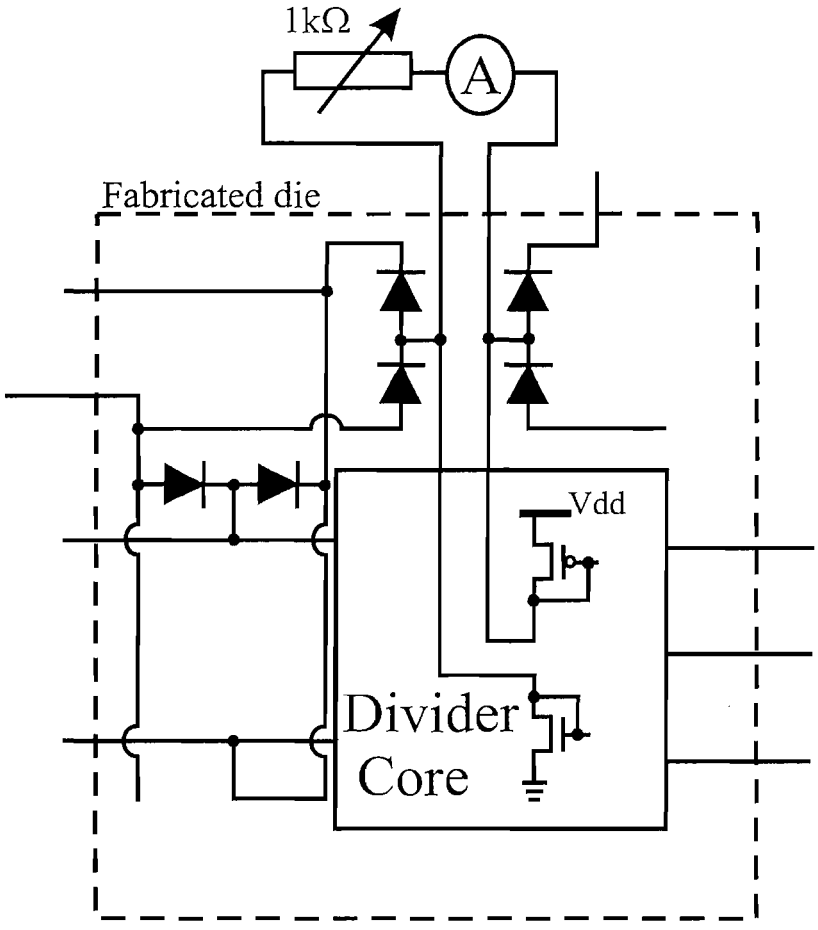

FIGURE 5.12: 1kΩ variable resistor used for setting the master current bias.

not on the main stack still require a form of voltage-mode coupling and source followers were implemented. As these off-stack dividers are close to the high power supply rail, cascodes are used to carry the difference in voltage gradually between the output of the common-drain stage and the circuit ground. These cascodes naturally serve to give the MOS current source a higher impedance looking into its terminal at the common-drain output node.

With the help of Figs. 5.13(a) and 5.13(b), it is clear that with a perfect current source as a load, the same current must be pulled from the common-drain NMOS device (ignoring leakage and assuming a subsequent circuit does not have a low input impedance). Therefore, the *gate-source* voltage will be constant and the output will follow the gate voltage, with a drop in common-mode voltage dependent upon the bias current. The large-signal behaviour can be expressed as follows:

$$V_{IN} - V_{OUT} = V_{GS} \tag{5.2}$$

and, if it is assumed that the current source is implemented as a common-base connected

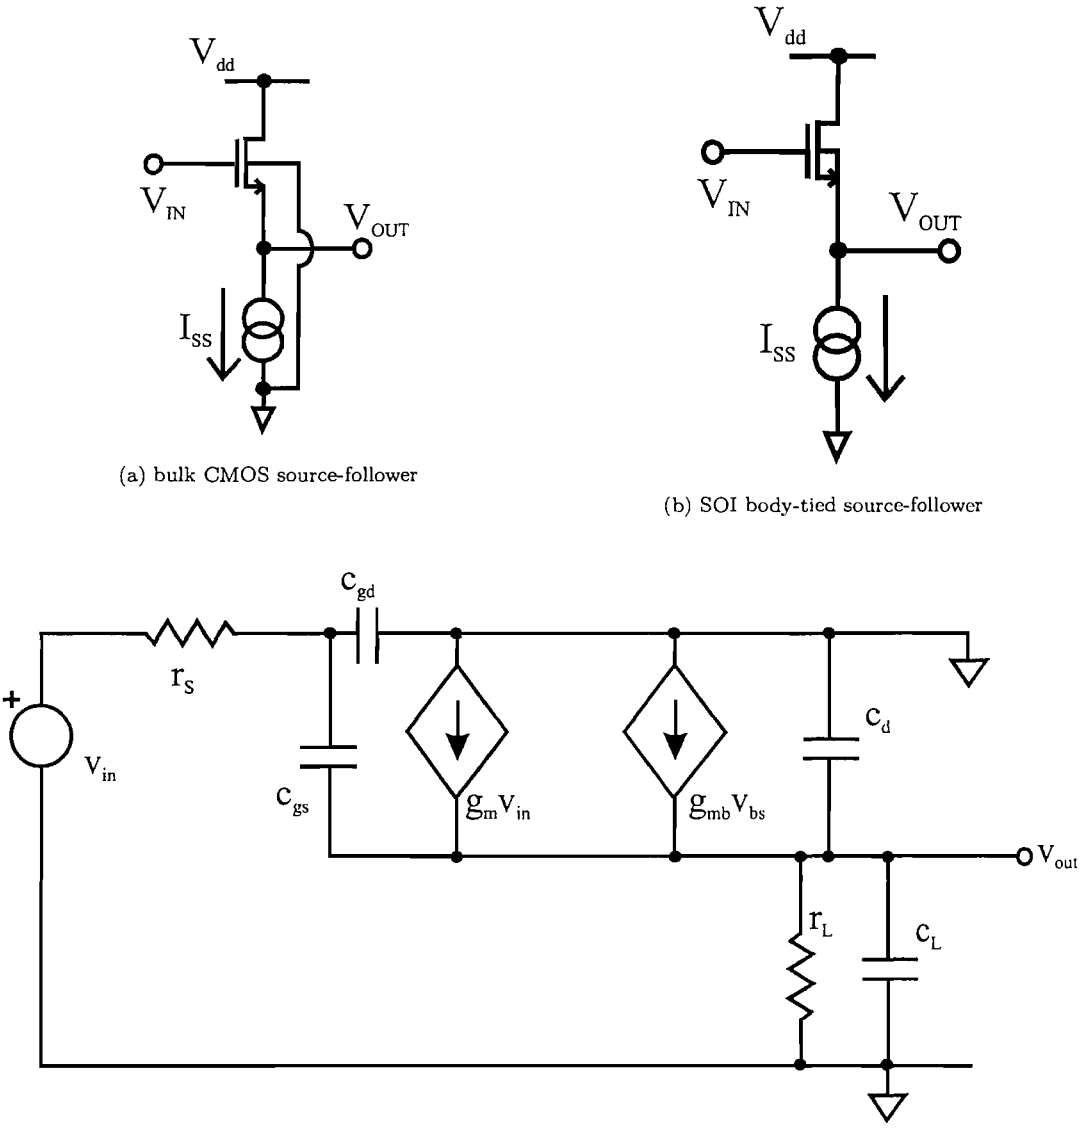

(a) bulk CMOS source-follower

(b) SOI body-tied source-follower

(c) small-signal representation of the source-follower

FIGURE 5.13: Schematics showing the difference between a bulk and SOI source follower.

NMOS transistor, the expression remains true as long as the current source compliance does not go below its $(V_{GS} - V_T)$. One can add a little more detail to equation (5.2), resulting in the following:

$$V_{OUT} = V_{IN} - \underbrace{V_{to} - \gamma\left(\sqrt{2{\cdot}|\phi_f| - V_{BS}} - \sqrt{2 \cdot |\phi_f|}\right)}_{V_t} - \underbrace{\sqrt{\frac{I_{SS}{\cdot}2{\cdot}L{\cdot}n}{W{\cdot}\mu_n{\cdot}C_{OX}}}}_{Min.\ V_{DS}\ for\ saturation}, \quad (5.3)$$

where V_{BS} is the voltage between the body and source terminals of the device, V_{to} refers to the 'zero V_{BS}' threshold voltage, n refers to the sub-threshold slope of the common-

drain connected MOS transistor, $2{\cdot}|\phi_f|$ is the surface potential of the channel at the source terminal side, and γ is the body factor given by:

$$\gamma = \frac{\sqrt{2{\cdot}q{\cdot}N_{sub}{\cdot}\varepsilon_{si}}}{C_{OX}}. \tag{5.4}$$

Here, q represents the electronic charge, N_{sub} is the substrate doping concentration, ε_{si} is the dielectric permittivity of silicon, and C_{OX} is the capacitance per unit gate area. These equations show that, by having a positive V_{BS}, the NMOS device will be harder to bias into saturation mode. Hence, in the bulk implementation of Fig. 5.13(a), V_{BS} is a non-zero value, whereas in the body-tied PDSOI follower of Fig. 5.13(b), the V_{BS} term is zero, forcing equation (5.3) to the simpler form of equation (5.2).

Another circuit performance metric associated with the source follower is its pass-band gain. Although one must be aware of this circuit's bandwidth, a more crucial aspect is the effect of this body-source voltage on the small-signal passband gain, a_v. Figure 5.13(c) helps one to generate the expression for a_v by summing the currents at the output node, V_{out}:

$$a_v = \frac{g_m}{g_m + g_{mb}}, \tag{5.5}$$

where g_m and g_{mb} refer to the drain-source and body transconductances, respectively. In its derivation, it is assumed that the common-drain amplifier's current source has negligible input capacitance to ground and a very large input resistance. In the case of the bulk CMOS source follower, the body transconductor is turned on and the amplifier's small-signal gain is less than 1. In contrast, a body-tied SOI follower design will have V_{BS} equal to zero, and hence the body transconductor in this case would be turned off, leading to a 'near' unity value for the small signal gain, ignoring second-order effects. With a number of cascode devices employed in each source follower in our dual-modulus divider design, the requirement that the current source have a high internal impedance is not so far fetched.

5.1.9 Floating-body transistors in high-speed divide-by-2 stages

One important design feature is that the transistors (minus the clock, current source and cascodes) in the first two divide-by-2 stages are implemented as floating body transistors (the body terminal is left electrically unconnected). This was found to give 1GHz improvement in the simulation, almost certainly due to the reduced capacitive loading, though simulation of such effects in circuits with widely varying time constants is prohibitively long and a thorough analysis of behind this is lacking. Incidentally, this action was motivated by results seen in digital rail-to-rail CMOS circuits [10].

In chapter 1, a summary was given to determine the mechanisms behind an increase in switching speed, citing investigations [16][13][18] on the behaviour of CMOS circuits. To our knowledge, however, there has been no discussion on the floating body effects in static logic and, more specifically, SCL circuits. Hence, this issue is raised again, this time in the context of our divider circuits. Figure 5.14 shows half of the lower most (highest speed) divide-by-2 stage of our main divider stack. After the second level of the MOS transistors (bottom-up), the sense and latch pairs have floating body PDSOI devices. Diodes representing the intrinsic junction between the source/drain implant and body regions have been marked explicitly on the diagram to aid our explanation. The body-tied devices can be identified as having a shorting link between the body- and the source terminals. From a large-signal point of view, the floating-body differential pairs in the diagram steer the maximum current available from one transistor to the other in a symmetrical fashion (both in the direction of input voltage and output current).

When the circuit is correctly biased and switching a large current, not all of the charge

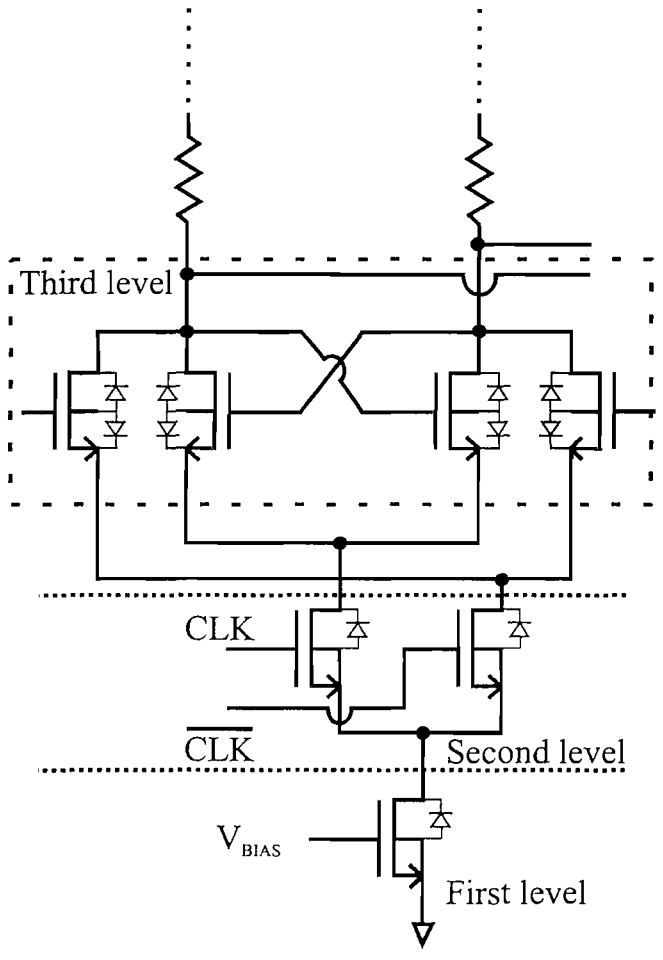

FIGURE 5.14: Schematic of one half of the highest speed divide-by-2 stage, with intrinsic diodes representing the junctions at the interface between implant and body regions.

is removed from the bodies of any sense or latch transistor when turning off. This is true if the circuit is operating at high speed, where the time-constants associated with removing charge are larger than the switching interval. This would be supported if it were possible to conclude that there is a difference in time-constants between removing the charge and introducing the charge by injection or drift/diffusion mechanisms (possibly with the help of the 'hysteresis' effect associated with floating body devices [5]). Hence, the body potential would rise with respect to circuit ground to an equilibrium value. By looking at the basic SCL divide-by-2 circuit, it is intuitive to say that the equilibrium value is set by the fixed bias voltages and the current source. Looking at the switching action in a little more detail, any rise in the body node would increase V_{BS}, decreasing the associated transistor's value of V_t and therefore switching more current. At the same time, the body-drain diode would have less reverse bias across its terminals and therefore less depletion capacitance. This is unlike the action seen in the body-tied transistor, at the lower levels, with the body-node being anchored to the source. As one ascends our main divider stack (resulting in lower frequency outputs), the action may become more prominent and help in its optimization by reducing the aspect ratios of the transistors, and therefore their drain capacitances too, in each divide-by-2 stage.

With this circuit being a static current form of logic, a preset amount of current is switched through any one device in a differential pair, unlike dynamic CMOS logic. In addition, the SCL divider is current-mode logic, where the emphasis is placed on switching that preset magnitude of current as fast as possible, as opposed to charging and discharging certain voltage nodes. Hence, the effect of dynamic V_t may not be applicable in our circuit style, and may show that SCL circuits will not benefit as greatly from floating bodies when implemented in PDSOI, unlike rail-rail CMOS style circuits. Of course, the above is based on supposition and would inevitably need to be rigorously investigated (experimentally) to prove or disprove our theory.

5.1.10 Output buffers

Thinking ahead to the measurement phase, it is important not to disrupt the operation of the divider with the electrical loading of the circuitry for characterising the device. Having a divide stage drive a 200fF pad capacitance (plus more for any ESD protection) as well as another divider in the actual divider chain can halt the running of the circuit. In order to isolate the actual node of interest, a SCL-to-CMOS level translator, together with a 3-stage tapered CMOS buffer, connects the output of the dual modulus divider circuit to the bond pad. The translator runs on a 50μA current source within a 6.8V supply and drives a negative feedback inverter with low input capacitance. The term 'tapered' is used to refer to the aspect ratio and current drive of each inverter [2] [17]; the aspect ratio increases by a factor of 6 per stage. Whilst recent literature suggests

that the tapering factor should be 3-4 for an optimum power-delay figure-of-merit [2], the buffers in this SOI design have a factor of 6 between CMOS inverters. The reason behind this conservative figure, is uncertainty in the pad capacitance and load connected to the output. Although the power-delay product is not optimum, the delay is outside of any critical feedback loop, and the actual divider does not suffer. Simulations have shown this stage to be capable of driving 10pF at 200MHz rail-to-rail. The power supply of this buffer is 2.5V, supplied from an off-chip source.

5.1.11 Bias generation

In order to the set the static currents through all the running dividers and the divider stack, as well as the bias voltage generators, a current bias (external to the chip) is applied. Using a variable resistor (see Fig. 5.12) between the diode connected PMOS and NMOS transistors of their current mirrors, the source current through them is controlled by varying their gate-source terminal voltages. The nominal current of this branch is 1mA and a 1kΩ potentiometer will suffice (700Ω should deliver the correct bias current).

5.1.12 Complete circuit

Figure 5.15 brings together the full circuit diagram (albeit with blocks), showing the arrangement and interconnections between the blocks. In this diagram, the feedback loop controlling the phase selector is shown explicitly . The only missing components are the bias and 50Ω networks. As mentioned previously, the current in the second off-stack divider can be sunk into the first divide-by-2 stage, shaving off 1mA from the current budget. A complete set of schematics can be found in Appendix B.

5.2 Layout

The task of laying out the cells and compiling the final chip has been performed using Cadence's Virtuoso Layout software. A Design Rule Check (DRC) and extraction file has been written to allow physical verification of the design by submission. The processing is by Honeywell Corp, USA, using their four layer metal 0.35μm PDSOI CMOS design rules.

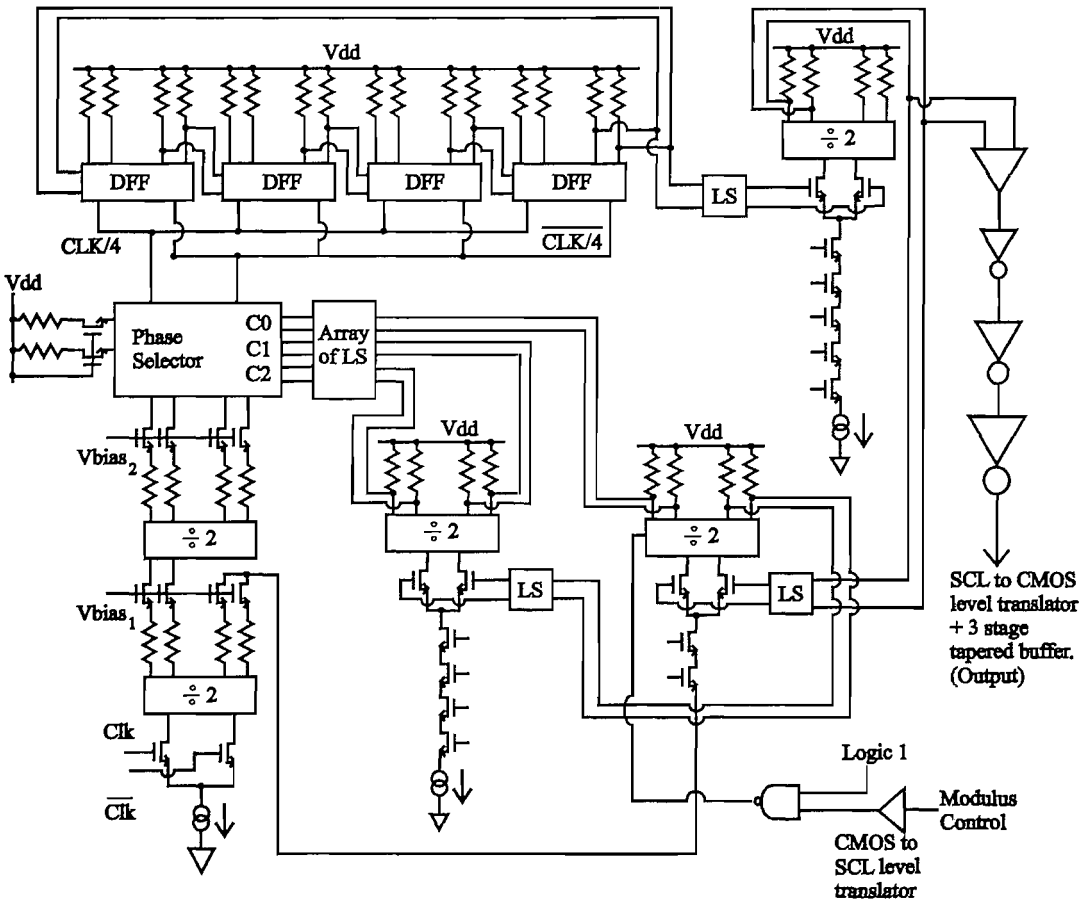

FIGURE 5.15: Schematic of the full dual modulus divider.

5.2.1 Transistors

The basic active unit of this circuit is the MOS transistor. The arrangement used is a 'H-gate' style transistor which allows the channel to enhance more uniformly due to the applied gate voltage than if the arrangement were a T-gate configuration. An important design rule to highlight is the maximum limit of the channel width associated with a particular gate length. For example, the maximum channel width of the NMOSFET for a gate length between 0.35μm and 0.7μm is 7μm. The reason for this is that the channel-body tie resistance becomes unacceptable outside this bound due to the constriction of the conducting path under the channel. Large aspect ratio transistors must be (and therefore are) split into an integer number of smaller transistors with a 'finger' type characteristic. Body-tied devices rely on ohmic contacts placed close to the channel ends to reduce this resistance. Design rules exist for different lengths and widths.

PMOS devices are drawn in an N-well polygon whose local body is connected to a node brought out to the surface of the silicon film. A p-type source/drain implant layer exists and must be spread over the cut in the field oxide. The poly gate masks the channel un-

derneath, keeping it n-type, and body tie contacts must also exist outside this definition. It is also noted that PMOS transistors have different rules concerning the channel-body tie resistance (and tend to allow wider channels).

Looking at the body contacts, various types have been incorporated in the design. For the majority of transistors in the layout, stand alone body contacts exist, and these are placed directly in front of the gate ends. These are either shorted to drain/source, or remain unconnected for floating body operation. More sophisticated body tie connections have been employed in current mirrors, with the help of 'butting' implant regions at the ends of the channel. Although symmetry is reduced, there is a certain amount of silicon space saved. Contacts are also embedded into the ends of the source implant regions for (body-tie)-to-source connections. A number of these cut into one of the contacts at the source node, thereby giving it a more compact connection. However, the connection to the back of the channel from these nodes is not as good as for the previous two, and their merits have been weighed against their flaws. High current carrying transistors and those operating at high frequencies have used the former two configurations.

The transistors also have appropriate dimension metal connections for current carrying lines. This is so that electromigration can be mitigated as well as reducing *IR* drops. IR drops can become severe if long narrow power lines connect to certain blocks, although with such a low sheet resistance (<100mΩ/□) this is not so much of a problem. Nonetheless, DC current lines have been widened to ensure this effect does not make itself a problem. Likewise for the non-silicided poly resistors, the widths of the polysilicon tracks are wide enough to ignore the effects of electromigration, even though the distributed capacitance would increase (this incidentally cannot be extracted using the current extraction rule setup).

5.2.2 Divide stack

Figure 5.16 shows the layout of the first divide-by-2 stage, which operates at the highest speed. The inputs to the clock come from either side of the current source to keep things symmetrical. The decision to route these clock lines on metal 3 is made to reduce the capacitance to the back substrate. By having a buried oxide in the wafer, metal lines across the field tend to couple less than in bulk processed designs. For reasons of matching, a common centroid layout has been adopted wherever possible. The transistors in the clock section are arranged so that the clock lines feeding them enter through the middle of this divider stage. The divider is turned on its side as only one pair of outputs is required to feed subsequent stages further up the stack. The remaining output acts as a sink for an off-stack divider. The 300Ω poly resistor loads (non-silicided poly-silicon) are sized slightly wider than the design rules, as the

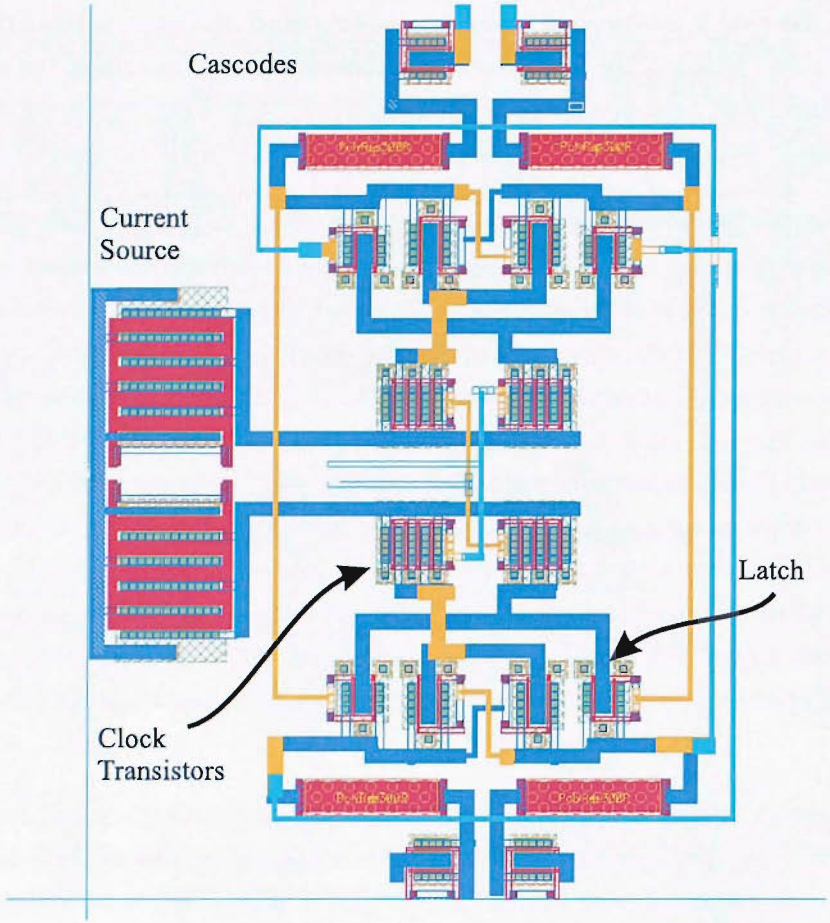

FIGURE 5.16: Layout of the first divide by 2.

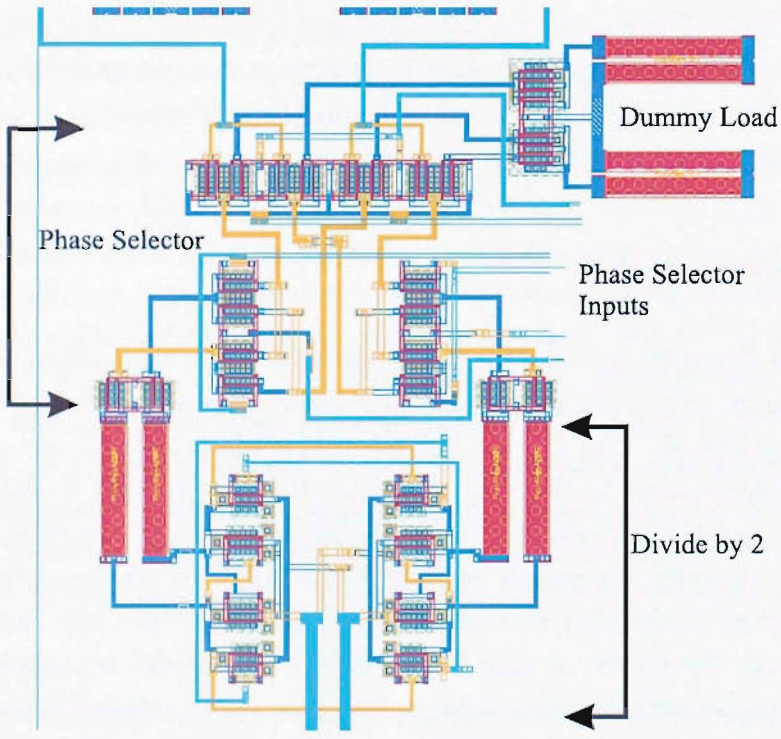

FIGURE 5.17: Layout of the second divide by 2, with the phase selector.

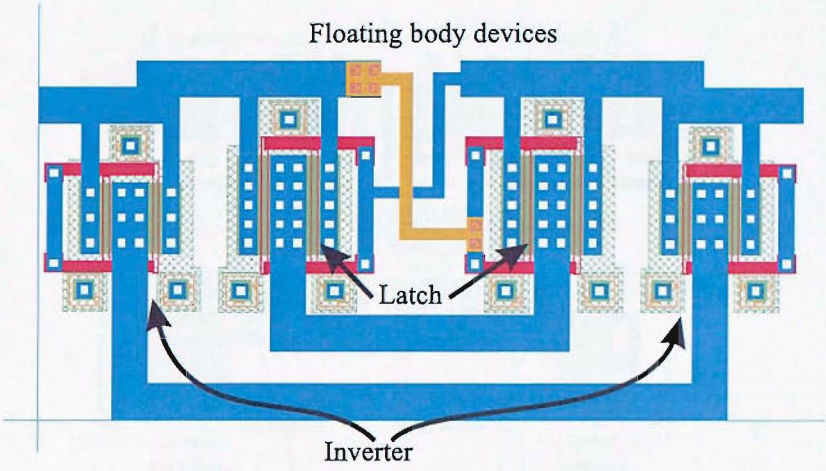

FIGURE 5.18: Layout of the latches within the first two divide stages.

foundry has discovered that the fabricated resistors tend to be larger in electrical value. Incidentally, the resistor layouts were extracted individually before being placed and routed.

The next stage in the divider stack is shown in Fig. 5.17. This diagram shows both the second divide-by-2 and the phase selector. The input signal current is fed through the middle and the divided outputs emerge from the sides. Cascodes tie the ends to low impedance nodes before the signal reaches the phase selector. Each of the latches within the first two divide stages is composed of floating body devices; an example is given in Fig. 5.18.

One clear observation is that the coupling between the outputs of the master latch and the inputs of the slave latch are different to those between the output of the slave and the input of the master latch (when configured as a divide-by-2 stage). This is unavoidable owing to the design, but this discrepancy is minimised by promoting the longer signal lines to metal 3 as soon as possible, for lower capacitance to the back substrate.

The phase-selector has the first stage clock inputs ('I' and 'Q' from the previous chapter) entering from either side, and the results are channelled through the middle and up in to the second tier of the selector (Fig. 5.19). Here, the decision of whether to pass 'I' or 'Q' determines the phases present on the clock input of the synchronous divider. This unit is symmetrical in current flow, but not in control line input. These inputs couple from the right hand side, although they remain in metal 3 for as long as possible. All devices are body-tied.

The last part of the divide stack is the synchronous divide-by-8. Local symmetry is approxmately the same as that achieved in the afore-mentioned divide stages. The clock lines in Fig. 5.20 resemble a 'forked' topology, where the phase-selected input pair ar-

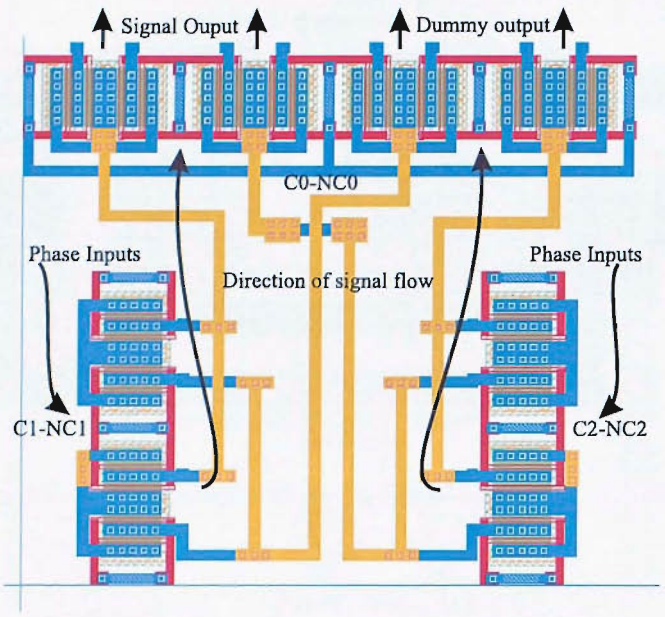

FIGURE 5.19: A close up of the phase selector.

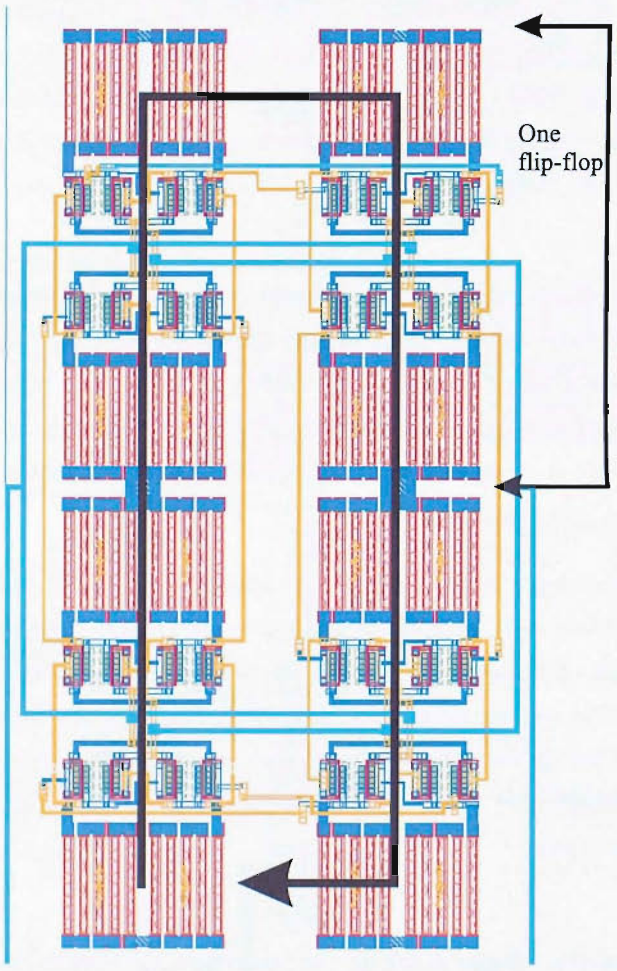

FIGURE 5.20: Layout of the synchronous divide by 8 unit.

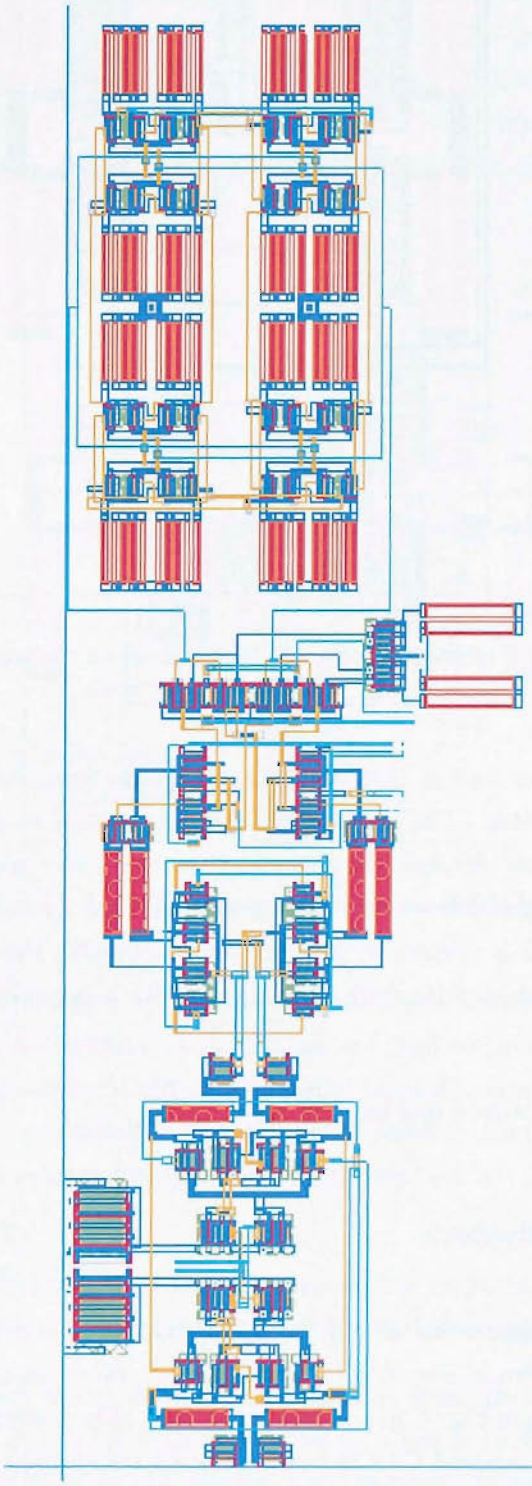

FIGURE 5.21: The divider stack, consisting of (bottom-up): an asynchronous divide-by4, phase selector and the synchronous divide-by-8.

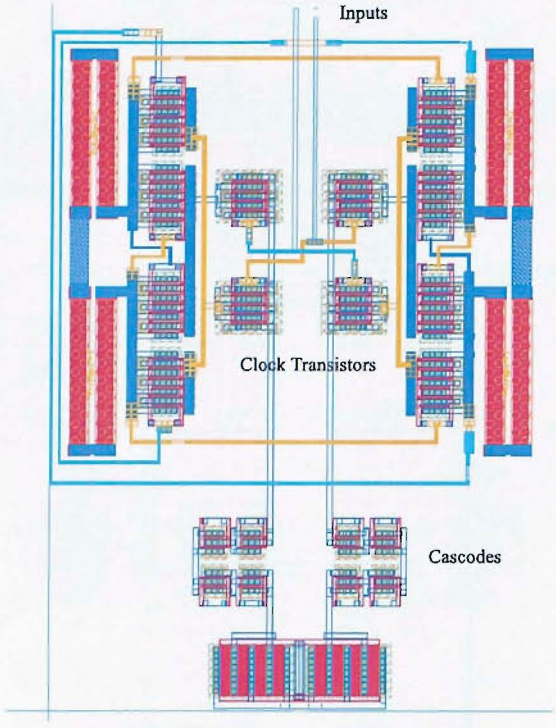

FIGURE 5.22: Layout of the first off-stack (not present on the main divider stack seen in Fig. 5.21) divide-by-2 block.

rives through the sides and is then fed into the clock transistors with approximately the same length of metal. The power supply line is routed to all four corner resistors and the centre resistors through the top and then down the middle of the figure. The propagation of the signal follows the black arrow, and the fortunate aspect here is that when positive feedback is eventually needed (see schematic), the distance is much closer than if the synchronous divider were constructed with a lateral signal flow.

The whole stack is captured in Fig. 5.21.

5.2.3 Off-stack dividers

The design then progresses to an off-stack divide by 2 block that runs from a 1mA current source, as shown in Fig. 5.22. (By 'off-stack', we mean not placed on top of the stack of circuits shown in Fig. 5.21). The output from this divider is the intended output before it couples to a tapered CMOS buffer (via a level shifter). The load resistors are 1.2kΩ in value and consist of two series-connected 600Ω non-silicided poly resistors. The inputs are channeled through the middle of the layout, with the outputs emerging from one side. In the floorplan, this unit is turned upside down so that signals can flow round in a clockwise direction.

Figure 5.23 shows the second divider stage (also not on the main divider stack) and immediately one should notice the lack of a current source. As described in the schematic section, current is used efficiently in order to minimise the power consumption, and this section sits on top of the slave latch of the first divide-by-2 unit in the divider stack. Three cascodes are placed between the long-tailed pairs of this divider and the resistor terminals of the first divider found at the bottom of the divider stack. These cascodes are biased using voltages derived from a MOS diode bias network (see later).

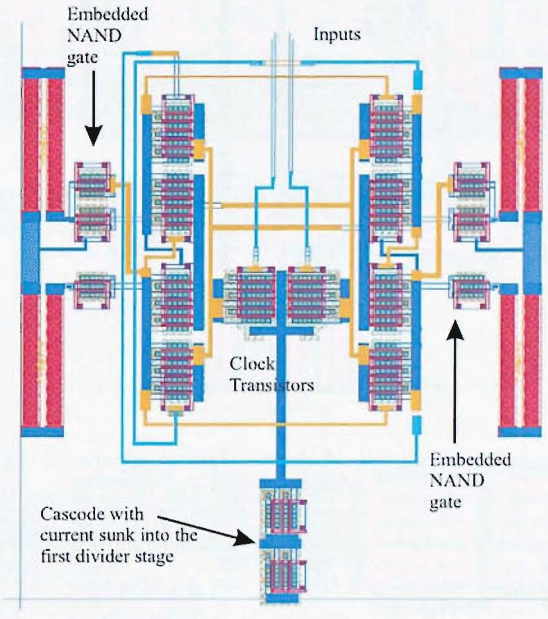

FIGURE 5.23: Layout of the second off-stack divide by 2 block.

Another striking feature of this divider is the NAND gate structures embedded between the latch transistors and resistive loads. These are used to provide the function of the modulus control using a 600mV differential signal biased around Vdd-300mV. The poly resistors are drawn as two series- connected 600Ω resistors and this allows them to be folded back so that the metal 1 connection can be kept short.

The final sequential divide-by-2 stage is shown in Fig. 5.24 and this happens to be a crucial block in the operation of the dual-modulus divider at high frequencies.

Simulations have shown that this divider requires at least 2mA in order to toggle the phase selector inputs in the right time window to allow it to function normally. This should explain why the resistors are smaller than the previous section and current source transistors have a greater current handling capability.

The design also relies on the use of source followers, like the one presented in Fig. 5.25(a), to couple the signals without their output common mode levels. The high power supply

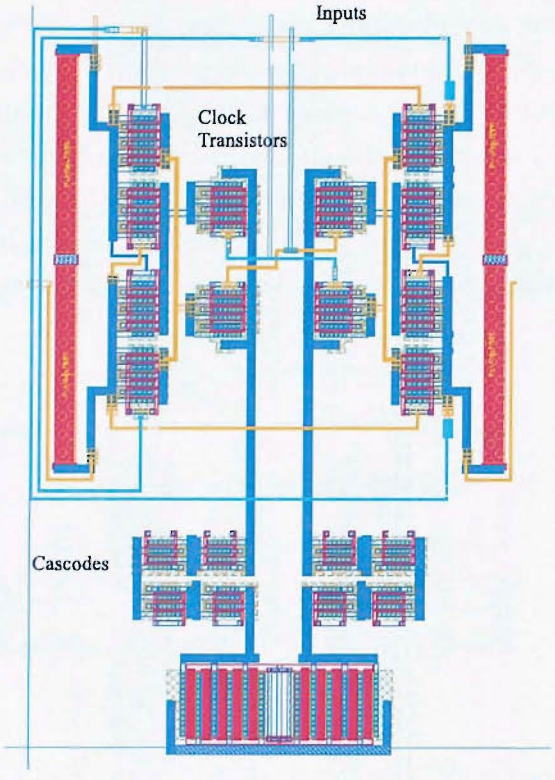

FIGURE 5.24: Layout of the final off-stack divide by 2 block.

associated with this divider has forced the use of cascodes biased from a network of series connected MOS diodes. These diodes run all the way to the top power supply rail and this is consistent with all the dividers, providing less sensitivity to power supply variation (up to a limit). The whole structure moves accordingly and the variations appear at the drains of the NMOS current sources that sit well into saturation. Bias blocks 1 and 2 are shown in Figs. 5.25(b) and 5.25(c).

5.2.4 Input and Ouput cell layouts

The input to the circuit requires a capacitively coupled differential input signal; this is terminated on chip with 50Ω non-silicided poly resistors connected to a common RF ground. The resistors are actually drawn as four lots of 200Ω resistors connected in parallel (see Fig. 5.26(b)). It is hoped that, a closer match to the intended 50Ω terminations will be achieved by splitting the resistors this way, thus reducing reflections in the balanced feed, and minimising distortion of the input signal (sinusoidal for testing.) The RF ground is in fact a bias voltage derived using NMOS transistors connected as diodes fed from a PMOS cascode current source. A high resistance of 20kΩ connects the bias to the common resistor terminals, thus reducing any additional current flow out of the feed. This bias voltage is brought out to a wirebond pad, enabling that point to be adjusted externally by 'over-driving' the bias from a low resistance voltage source (Fig. 5.26(a)). Looking at this block, it is apparent that any power supply variations would

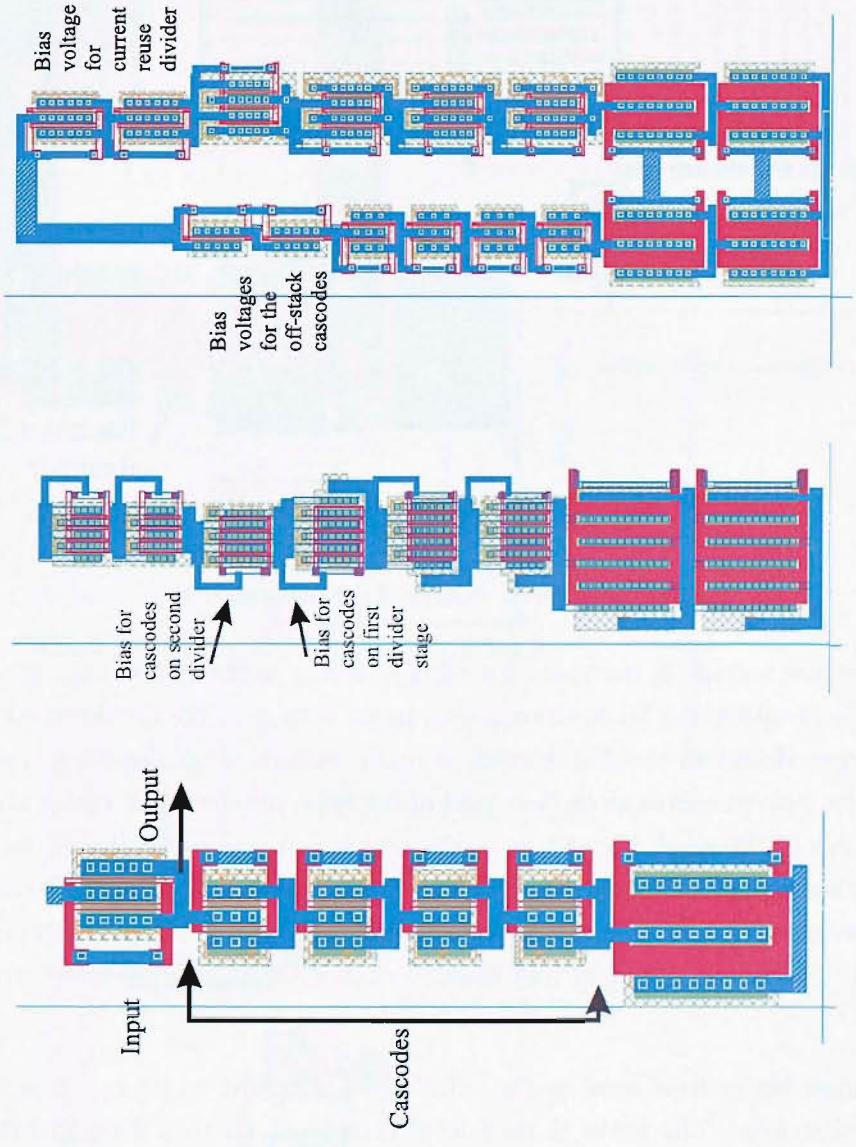

(a) Source follower (b) Bias block 1 (c) Bias block 2

FIGURE 5.25: Screen captures of miscellaneous layouts.

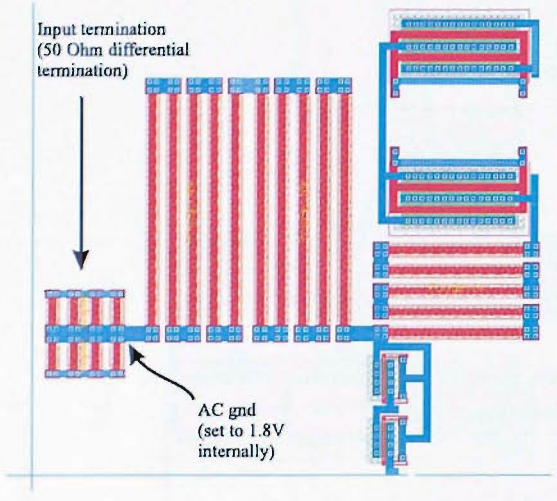

(a) Balanced input section

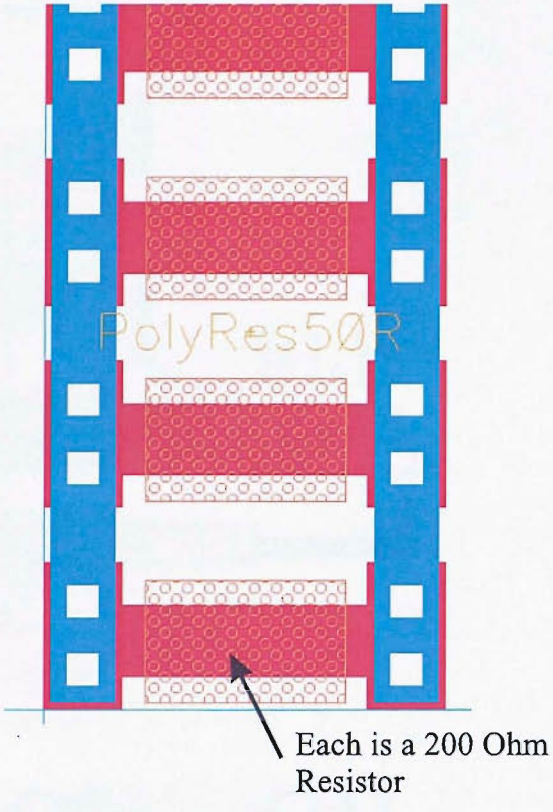

(b) 50Ω termination resistor

FIGURE 5.26: Screen capture of the input section.

not cause the bias voltage to fluctuate accordingly, owing to the PMOS current sources. With the help of subsection 5.1.6, the common input voltage of the divider stack is kept relative to ground and so the divide stack is really cushioned on the clock transistors rather than its current source alone (any part of the main divider stack sitting above the drain terminals of the clock transistors in that stack can move with variations on the upper power supply rail). Simulations that varied the power supply showed no complete failure in operation, but the design may be changed so that it is relative to Vdd rather than ground. This allows the input bias stage to move with any power supply variations.

The input modulus control seen in Fig. 5.27(a) is designed to be run from a 2.5V CMOS signal source. The input is then level translated up to a 600mV differential signal balanced around V_{dd}-300mV to match the input bias requirement of the embedded NAND gate, where it is DC coupled to the modulus control inputs of the second off-stack divider. The input to the level translator is preceded by a two stage CMOS buffer, that generates the modulus control and its complement. The delay in signals is irrelevant because this input is not time critical. One could even say that the input buffer is irrelevant, as one input of the long-tailed pair can be fixed to a bias voltage with the other being driven hard. To lower the dynamic power consumption, this would probably

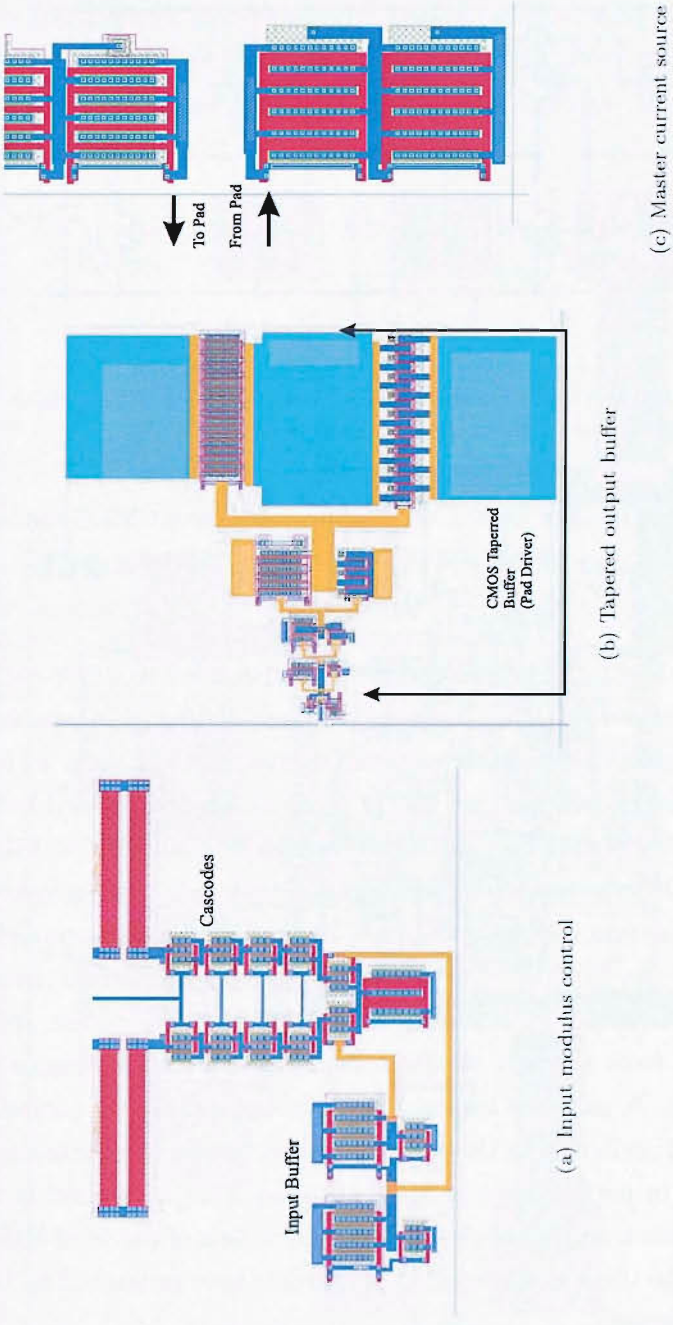

(a) Input modulus control (b) Tapered output buffer (c) Master current source

FIGURE 5.27: Screen captures of various sections of the layout.

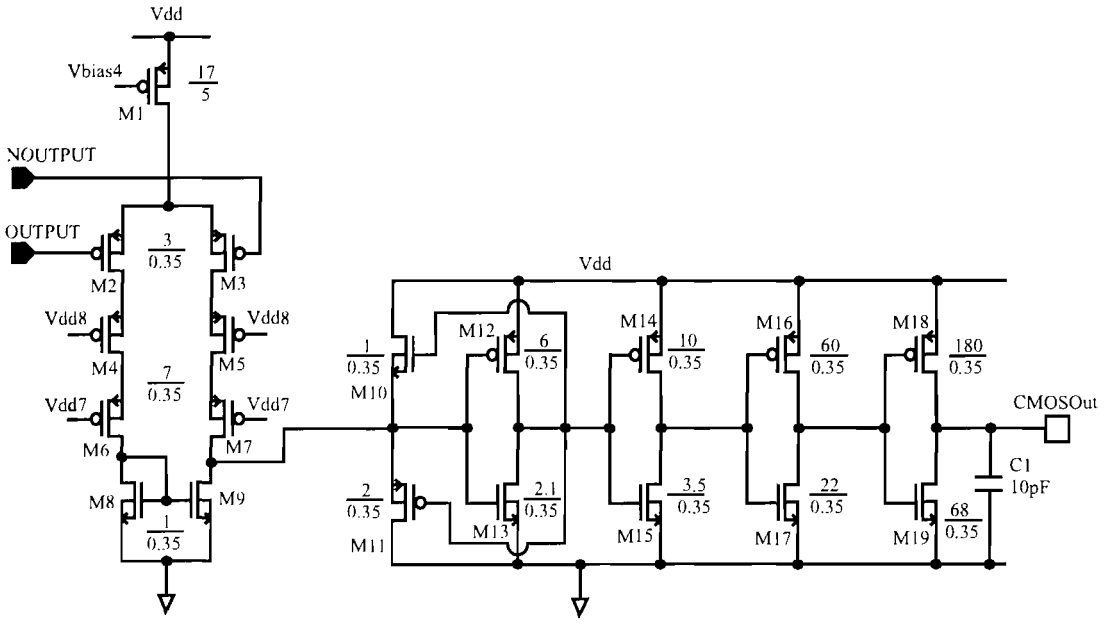

FIGURE 5.28: Schematic of the level translator followed by a tapered buffer to drive the bondpad and additional load capacitance.

be removed if incorporated in a sythesiser design, so that the low power CMOS modulus control signal (from a modulator) can be coupled directly to the translator.

The test chip has been designed with the intention of being wirebonded to a chip carrier rather than being probed. The output from the first off-stack divider is incapable of driving the relatively large capacitive loads seen on the output pin of a chip at high frequencies. Therefore, the tapered buffer shown in Fig. 5.27(b) has been added before a wirebond pad, in order to charge those capacitances at high frequencies. A 3-stage buffer is used, comprising CMOS inverters which are cascaded with each successive stage three times larger in active area than the preceding one. Each stage has an appropriate width metal line to ensure that the dynamic current does not exceed the current density capability of the metal inconnects. A level translator is placed before the buffer, transforming the differential output from the first off-stack divider into a CMOS compatible single-ended signal (Fig. 5.28). A negative feedback input stage couples its output to the input of the tapered buffer, minimising the swing required by the level translator, and giving an apparent boost in performance. PMOS cascodes have been used in the translator to lower the voltage seen across the drain-source terminals of the final NMOS active load. The bias voltages to those cascodes are taken from those generated by bias block 2, making the biasing simpler.

The last important unit is the layout of the master current source pictured in Fig. 5.27(c). Cascode current sources are connected between V_{dd} and ground and are sized according to the current density and input current. The two floating nodes in the current

source are brought out to a pair of pads so as to allow a variable resistor to set the master current externally.

The wirebond pads used in the chip have a passivation cut that is 123μm × 123μm and have a nominal pitch (distance between centres) of 200μm to ease the bonding. Only the last layer of metal (metal 4) exists, with nothing below it. Protection against electrostatic discharge (ESD) is solved using Honeywell proprietry diodes close to the bond pads. They consist of p-type regions inside n-type 'donuts', interconnected with metal

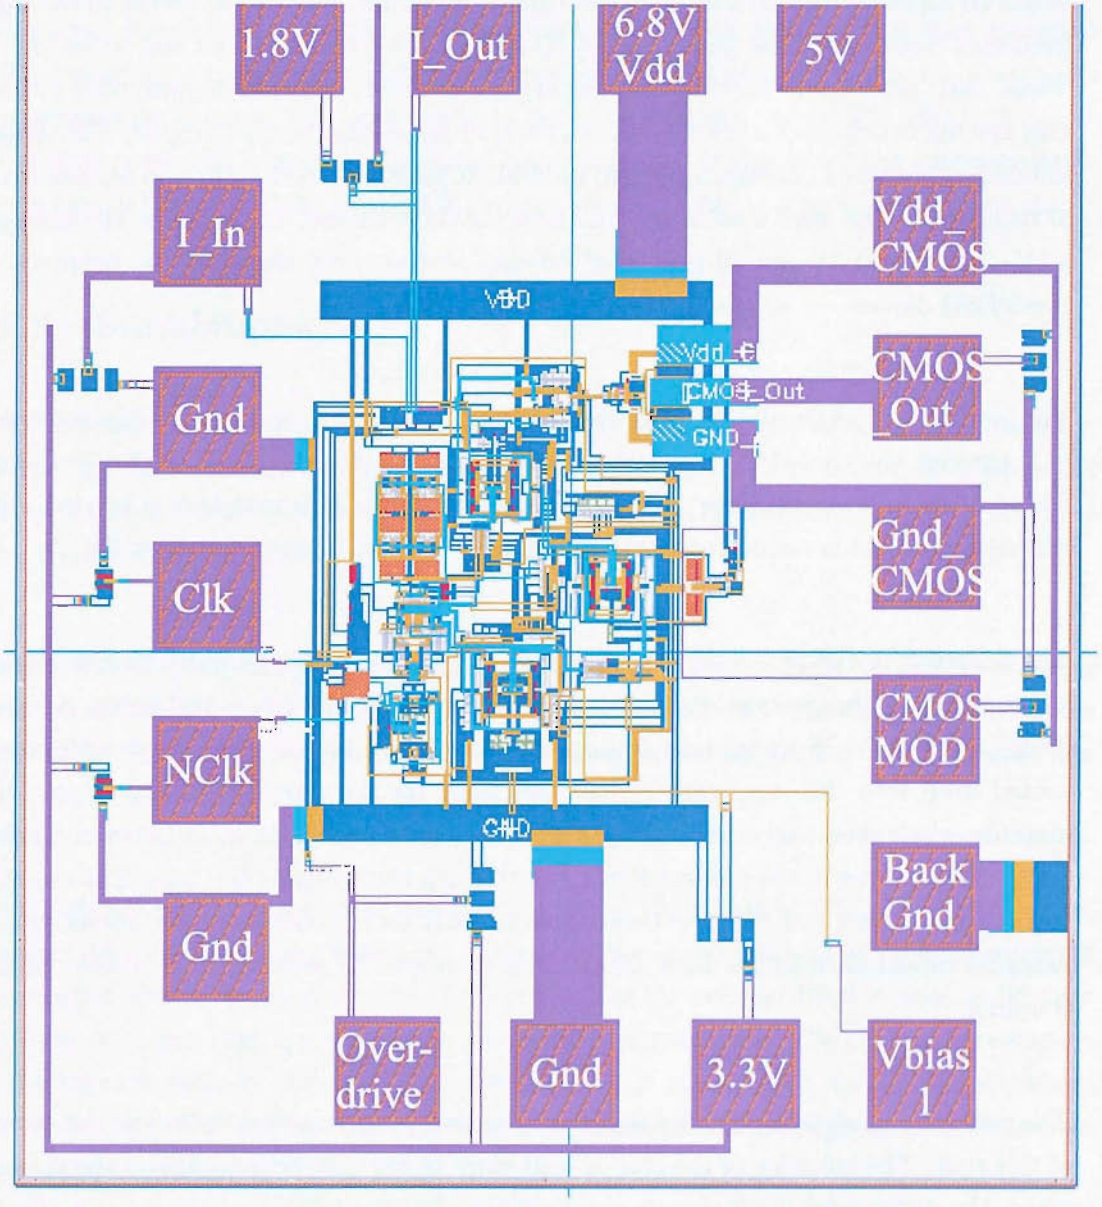

FIGURE 5.29: Layout of the final test chip.

lines to heavily doped ohmic contacts. Each protection circuit consists of 32 diodes (2 lots of 18) between two power rails, and this ensures a better clamping characteristic. The drawback is the reverse bias depletion capacitance which, coupled with any series resistance, can severely reduce the performance of the connected circuit block. Pads without time-critical inputs have had these diodes placed between 3.3V and ground. The signal connects to the midpoint, except for the current source which has one pad connected to the midpoint of the diodes between 5.1V and 1.7V. The CMOS output has been simulated with models of these diodes (not extracted) and the results have shown neglible attenuation with the divider running on a 5.5GHz input signal. The diodes on the CMOS section of the test chip are placed between 2.5V and ground. However, the balanced inputs to the divider have had their own diodes constructed owing to the high frequency nature of these pads. $\frac{50\mu m}{0.35\mu m}$ NMOS diodes have been laid out to provide a 'weak' ESD protection. As these designs are not commercial grade designs, only protection against bench discharges is required, without impeding the input signal. The diodes are series connected between 3.3V and ground. It must be stressed that no human body or machine models were used in the simulation of these diodes. Incidentally, the voltages 5.1V, 3.3V and 1.7V are all generated off-chip and are only there for the purposes of these ESD diodes.

The final feature of the design is the contact ring. The handle wafer has a contact made to it through the buried oxide, allowing electrical connection on the surface of the chip. This feature is there to allow on-chip modulation of the substrate bias if needed. For this design, a pad is connected to the ring and is shorted to ground externally.

The floorplan of the proposed test chip is shown in Fig. 5.29, detailing the placement and routing of the macros discussed. Unlike standard cell place and route designs, complementary MOS transistors are actually distributed around the layout with power routed deep into the core, rather than just lying on the top and bottom edges. To minimise clock skew and overall delay, signal lines have been kept short between blocks, whilst control lines to the phase selector run thorugh the centre of the core. The cascode voltages and power routing have been routed mostly on the outside of the blocks, where parasitic capacitance to the back substrate is employed to provide some additional decoupling.

The pad ring consists of 18 wirebond pads, but with the core not sitting in the centre of this ring. The left edge of the chip is kept close to the core because this is the side on which the differential clock signals enter, minimising the lengths of these lines. These clock lines are also shielded with ground pads on each side of the pair. The foundry has also offered probe pads for the design, but experience has shown those pads to be too small (for the probing machine at Southampton), with the bondwire metal contact

engulfing the whole passivation cut. Minimum separation between pads has also been avoided to reduce the chance of shorted bondwires.

Prior to taping out the design, a fill metal pattern is added to the design on the first three metal layers; this is to give a better uniformity after a chemical mechanical polish of a dielectric layer. The pattern is drawn in blank areas of the design which raises the interlayer dielectric (ILD) so that the final surface is more planar than a surface without such a feature. The option exists to let the designer add patterns to sensitive areas so as to deter the automatic 'filler' from interfering with the surrounding space. Fill metal patterns in fact enhance coupling capacitances and therefore crosstalk between rapidly changing signal lines. The fill metal pattern can be grounded, but this then increases the capacitance to ground which is detrimental to the speed of the circuit. Instead, a conscientious decision has been made to leave the pattern electrically 'floating', as it is hoped noise will couple to the signal lines in phase. If this is so, the common-mode rejection inherent in differential SCL designs should tackle this problem of crosstalk.

5.3 Simulations

All simulations have been carried out using the SIMetrix circuit simulator, together with the 'Southampton Thermal Analogue' (STAG) model. The model parameters are a set extracted from 0.35μm PDSOI (NMOS and PMOS) measurement data provided by Honeywell Corp, USA.

Figures 5.30, 5.31 and 5.32 show the outputs at all divide stages as well as showing the phase selector control inputs. The traces labelled '`:f-:nf`' correspond to the divide by 64/65 output. These simulations are all taken before layout and hence exclude any details of layout parasitics.

Though not explicitly shown in this thesis, jitter in the region of 17% of one clock cycle, either side of the 65 ratio, is present. This error becomes apparent when measurements are performed on one output cycle. Taking a look at the average division ratio, a different picture emerges that questions this divider's performance at the declared frequency. The average ratio is obtained with the use of $2{\cdot}P$ consecutive output cycles (where $P \in \mathbb{Z} \geq 1$). If P=1, the following equation computes the average division ratio over two output cycles:

$$division\_ratio = \frac{f_{in}{\cdot}(T_{out(2{\cdot}P-1)} + T_{out(2{\cdot}P)})}{2} \tag{5.6}$$

When computed this way, the division ratio comes out as 65. The problem is thought

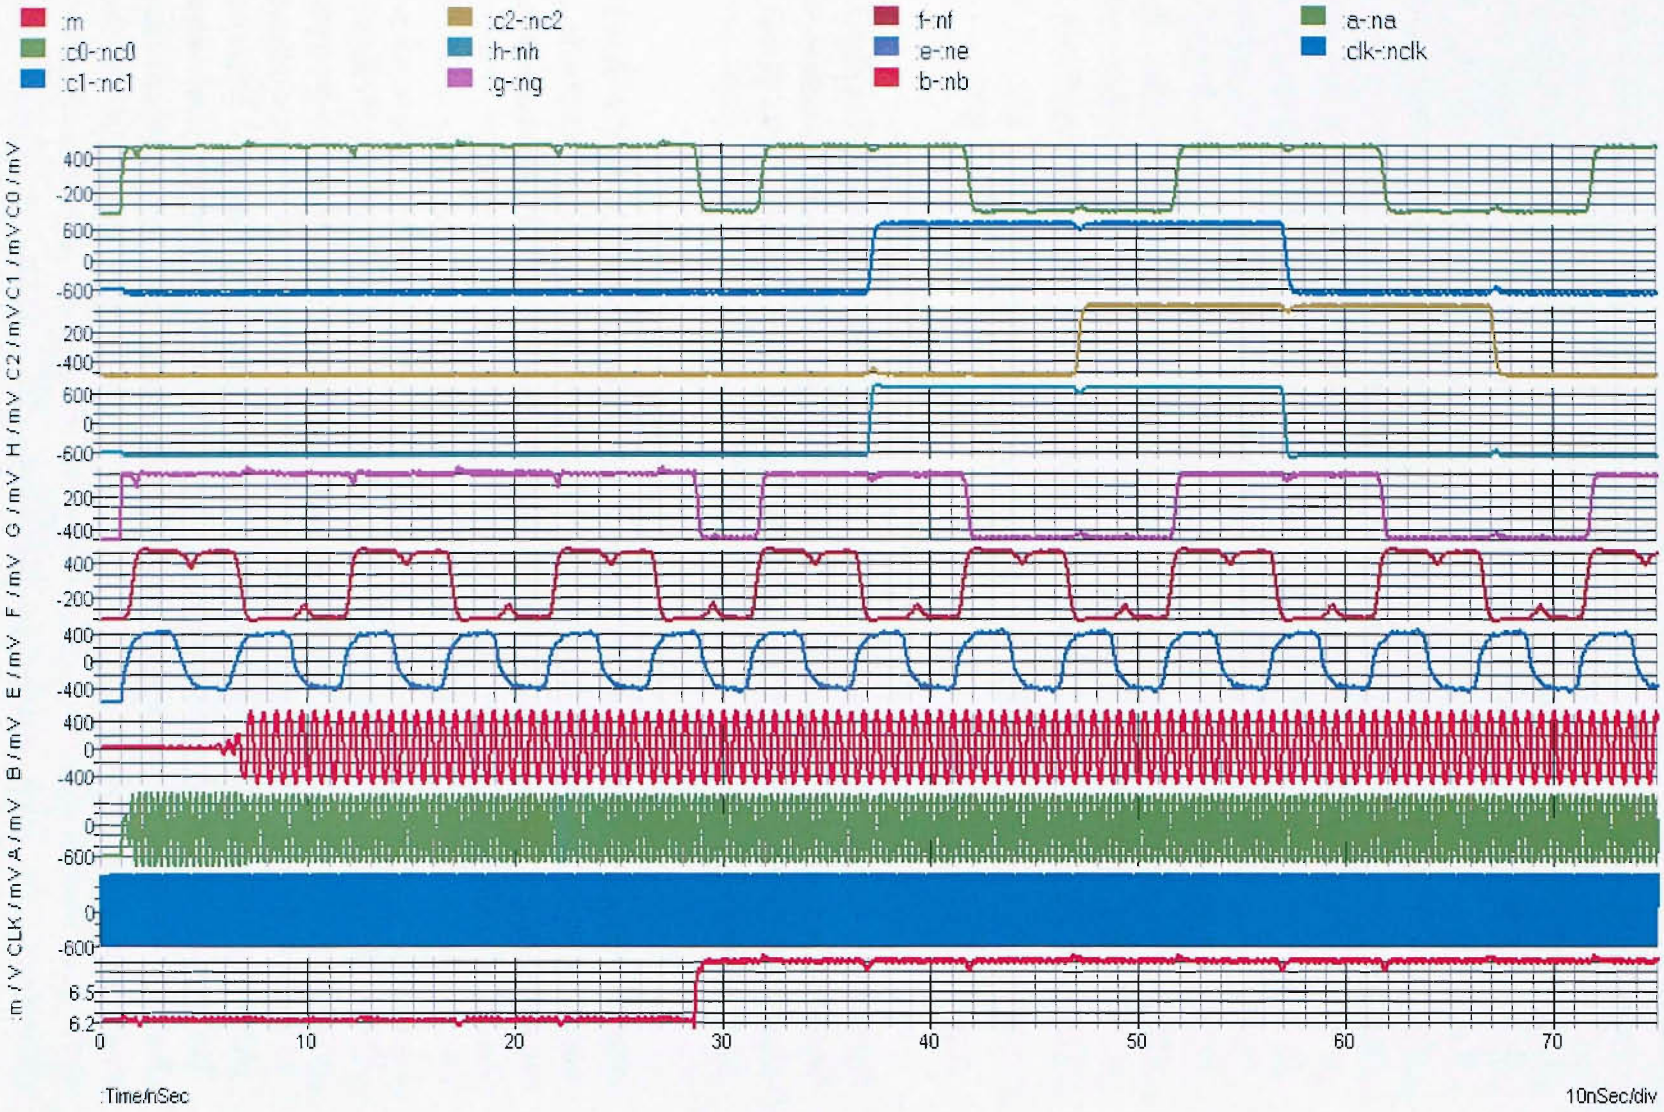

FIGURE 5.30: Screen capture of a 6.5GHz simulation.

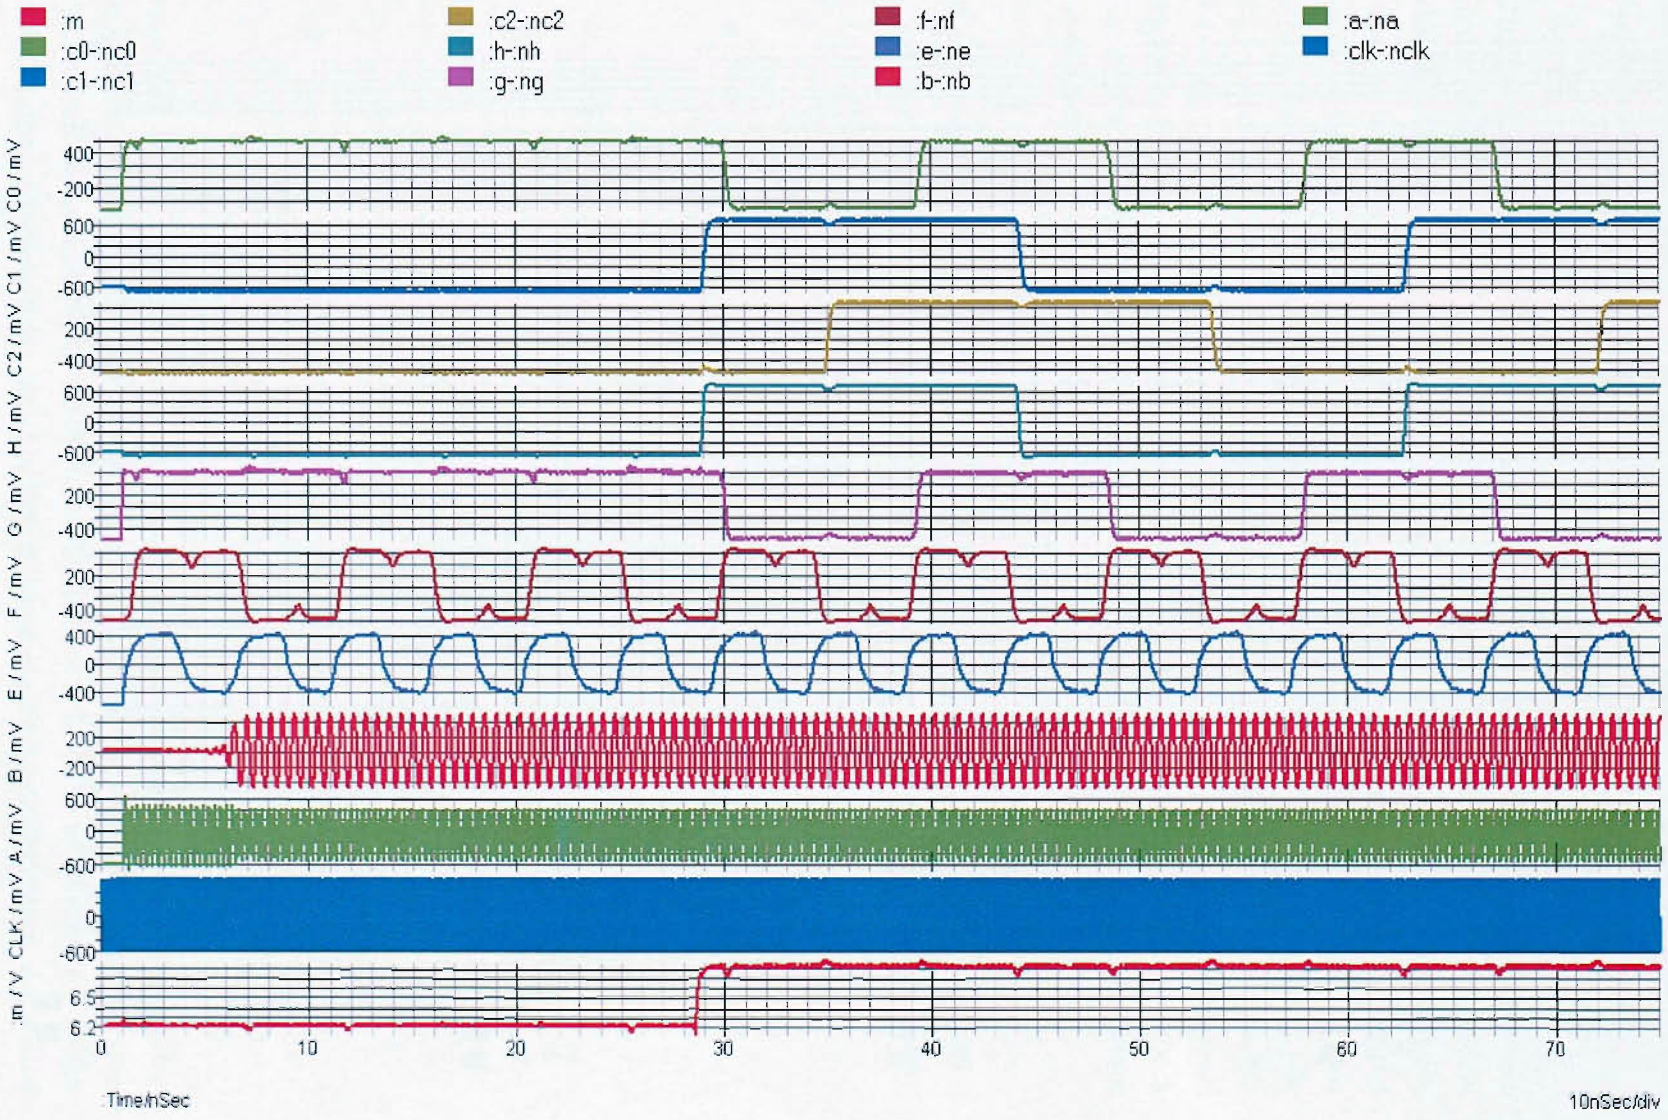

FIGURE 5.31: Screen capture of a 7GHz simulation.

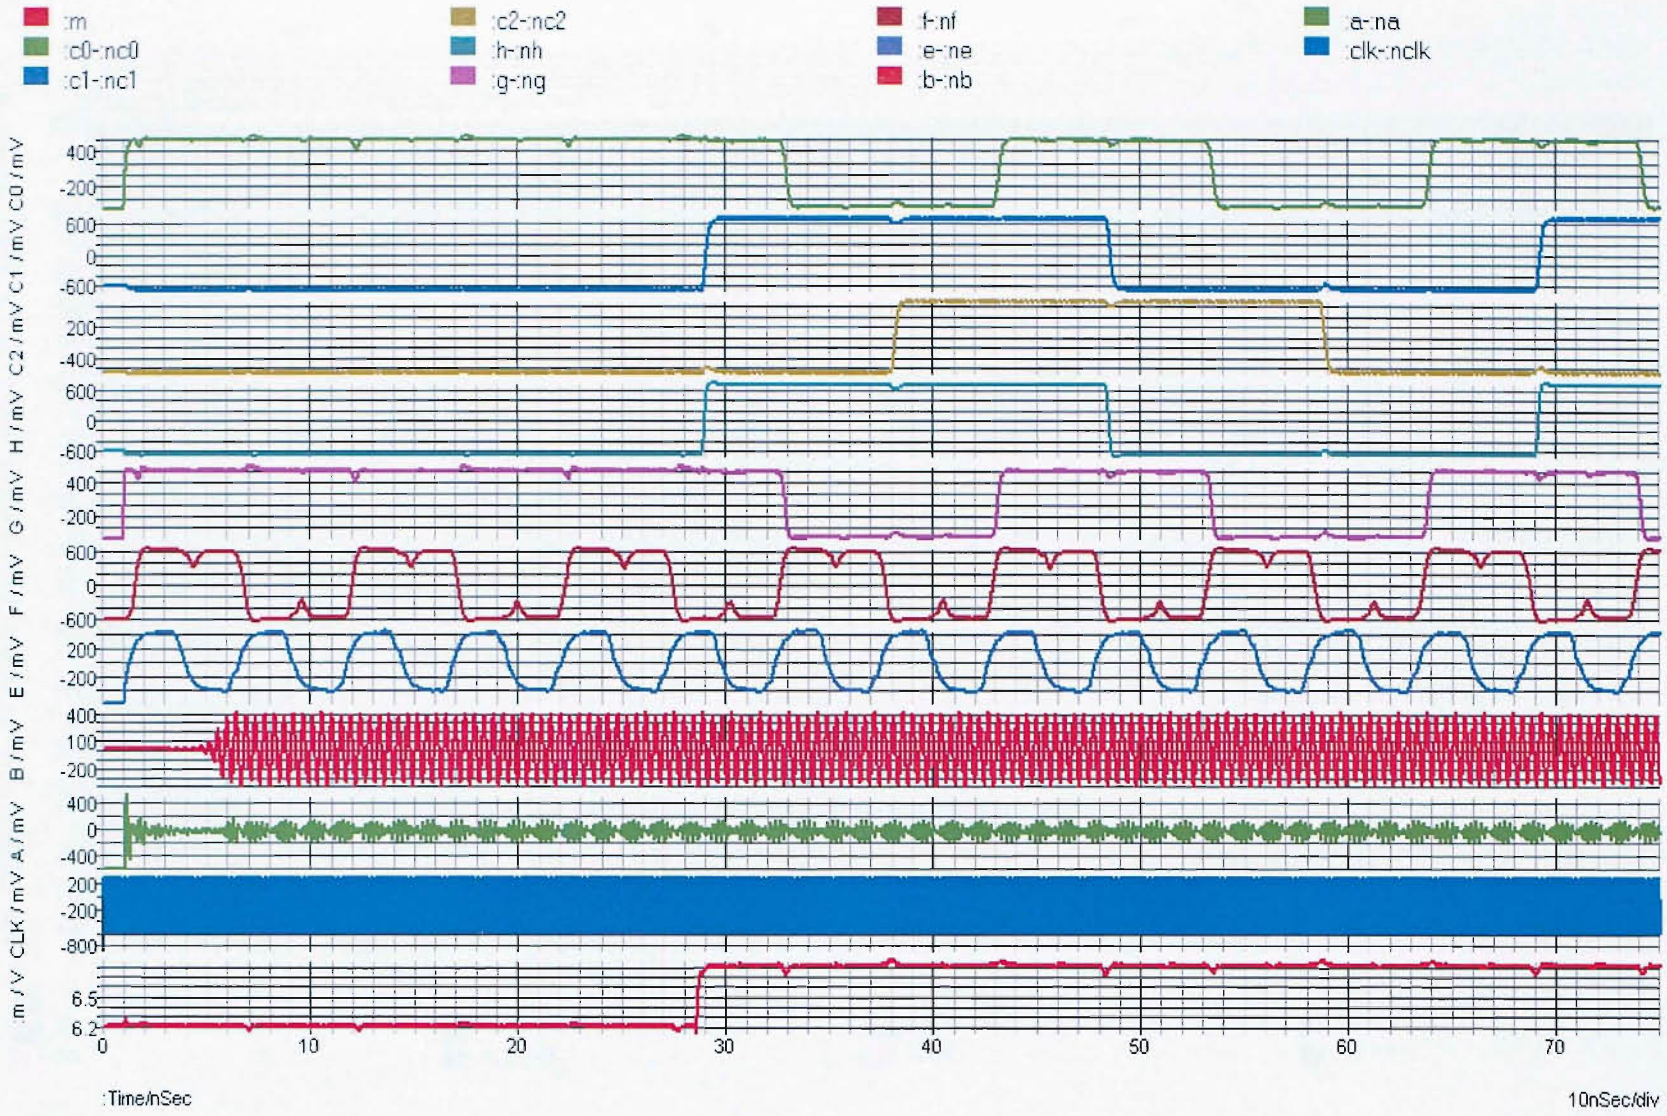

FIGURE 5.32: Screen capture of an 8GHz simulation.

to lie with the dummy load sitting on the spare phase selector outputs. The dimensions of the transistors are chosen to emulate those in the synchronous divider but, unfortunately, the simulations seem to point to a mismatch on the outputs of the phase selector.

Figures 5.33(a) and 5.33(b) are included to show how a divide-by-2 stage fails when operated at speeds beyond its capability. As one looks through the figures, it is evident the first divide-by-2 stage cannot toggle fast enough, with the 8GHz simulation showing unsynchronised oscillation (as a result of incomplete toggling). With the output swing of the first divide-by-2 stage being insufficient to drive the clock input of the next stage, the oscillation seen at the output of the second divide-by-2 stage is likely to be its self-oscillating frequency, which incidentally looks noisy. Any improvements with the intention to upgrade its speed require a thorough analysis of the output and internal nodes associated with this divide-by-2 stage. Also, one needs to examine the feedback loop when increasing the speed to ensure that the control signals reach the phase selector at the right instances.

Lastly, this design has been simulated to deduce its immunity to process and electrical variations; Table 5.1 shows the results from these tests. Although an observant reader

Test	Tolerance Bounds		Result
Temperature			
	-10°C		Passed at 5GHz
	+70°C		Passed at 5GHz
Power Supply			
	6.8V and CMOS 2.5V +10%		Passed at 5GHz
	6.8V and CMOS 2.5V -10%		Passed at 5GHz
Threshold Voltage			
	V_{TP} +75mV = -0.71V		Passed at 5GHz
	V_{TP} -75mV = -0.86V		Passed at 5GHz
	V_{TN} +75mV = 0.59V		Passed at 5GHz
	V_{TN} -75mV = 0.45V		Passed at 5GHz
Current Variation			
	1mA +10%		Passed at 5GHz
	1mA -10%		Passed at 5GHz

TABLE 5.1: Table showing the results from the tests on electrical and process variation.

may notice the discrepancy between the quoted speed and intended speed of the divider, it must be stated that these tests were carried out to check the robustness of the divider. By varying the parameters listed in the table, an idea is given as to whether the design will fail due to its sensitivity to variations in those parameters. Having an operational divider that doesn't require such stringent operating conditions reduces the demand on the test setup.

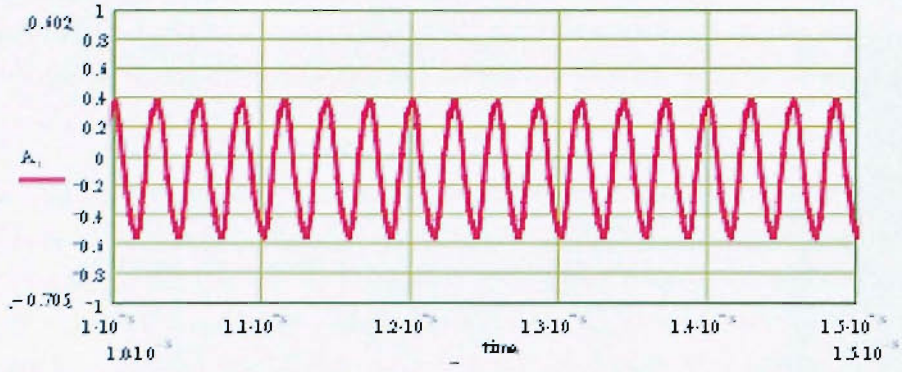

(a) clocked with a 7GHz clock

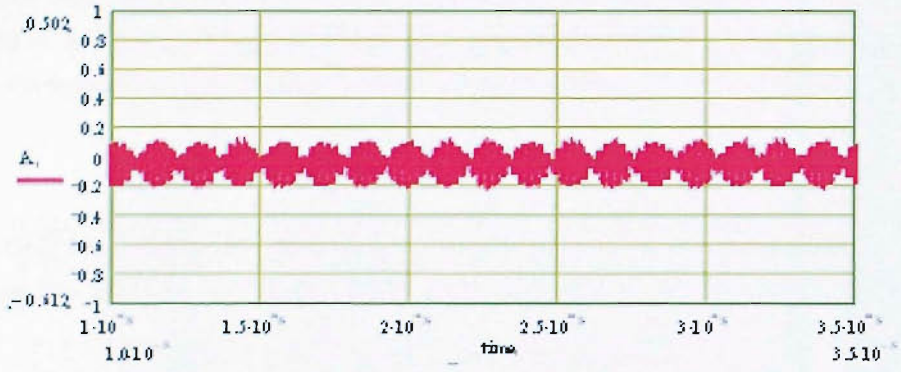

(b) clocked with a 8GHz input clock

FIGURE 5.33: Simulation results of the 64/65 dual modulus divider, showing the output of the first divide by 2 stage in the divide stack with different input frequencies.

In the dual-modulus divider design, the test chip layout has been accomplished with Cadence's *Virtuoso* tool (including design rules) whereas simulation jobs have been resolved using the PC based SPICE simulator, *SIMetrix* [14]. After the layout has passed all design rules, the extraction tool along with an extraction file, works on the design to produce a text description (netlist) of the design, with or without parasitics. This netlist must be rewritten so as to be compatible with the PC based simulator before it can be referenced by a top level schematic that incorporates signal sources and decoupling capacitors. Unfortunately, a *layout-versus-schematic* (LVS) check is not performed explicitly, and is instead accomplished using *SIMetrix's* [14] electrical rule checker as well as performing a functional simulation (on extracted netlist without parasitic information) that should be identical to the original schematic. Once this has been verified, the layout must be extracted again to include parasitic information not included in the original schematic. A simulation after this stage can illustrate the effects of parasitic elements on performance. It must be stated, however, that the extraction file containing

device identification rules, has become primitive compared to the version used by the foundry. The author has had to remove the calculation of fringing capacitance associated with interconnect lines due to anomalous capacitances. Regardless of this fact, it is thought that some parasitic calculated capacitances may be duplicates of those present in the MOS models, giving a netlist that is undermining the true performance of the fabricated device. This file would need further work and assistance with the help of foundry layout examples.

The traces shown in Fig. 5.34, focus on the output of the dual-modulus divider when dividing by 64 and then 65. Figure 5.35 captures the whole simulation together with the correct timed control inputs to the phase selector.

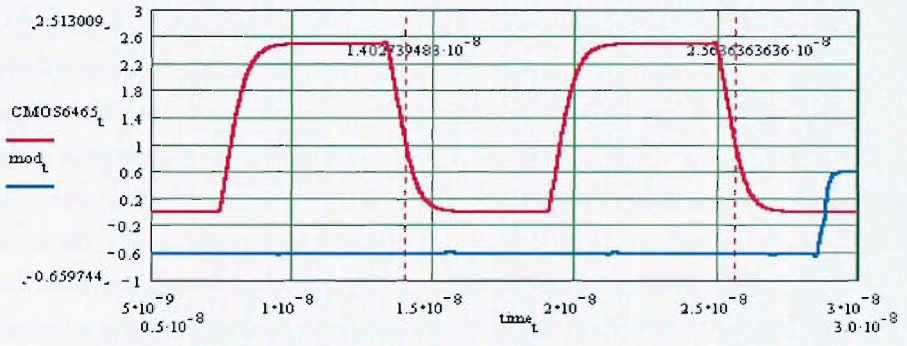

(a) driven with a 5.5GHz clock, when dividing by 64

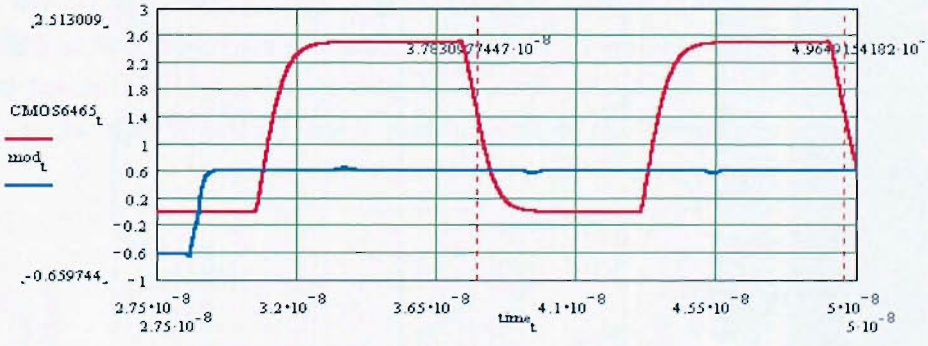

(b) driven with a 5.5GHz clock, when dividing by 65

FIGURE 5.34: Simulation results of the 64/65 dual modulus divider, showing the output of the whole divider. These simulation results are generated from the layout extracted netlist with parasitics.

5.3.1 Self-heating

All simulation results, have been generated from schematics that use MOS models with their *self-heating* code activated. This is an extremely useful component of the model, allowing a designer to characterise their circuit and model the design just that little closer to reality. Although not included as a figure or tabulated, the dual-modulus output was seen to improve when this feature was deactivated, resulting in a slightly better accuracy of the divide by 65 operation. All transistors had *self-heating* turned on except for the floating body devices present in the first two divide stages located at the bottom of the divide stack. This is because simulations failed to converge when calculating a DC operating point. Although the STAG model had been simulated with the self-heating component, the model still lacked robustness (at that time) when simulated in devices with floating bodies, especially in a circuit with a large number of circuit nodes. One argument, insisting on its inclusion, is that the transistors with floating-bodies happen to switch the largest amounts of current. The other side to this argument is that there should be a symmetrical waveform in the switching transistor's drain current. The frequency of operation is much greater than that associated with the thermal time constants of the MOSFETs. Had this been a high frequency divider based on a CMOS topology, then current would only be conducted during the transitions (ignoring leakage) and the local temperature of the devices would rise to an equilibrium, as there are more transitions per unit time when the divider is driven with a very high frequency input. One would expect the average temperature to be low in the case of a low frequency input because there would be adequate time during no current conductions, for the heat to dissipate. However, in the case of our static SCL divider, the switching transistors can be 'off' for equally long as them being 'on'. The transistors running at a lower frequency higher up in the divider stack, are in fact conducting much less current than the floating body devices near the input and hence should not be susceptible to self-heating effects. By this reasoning, it is implied that the lower frequency dividers running off the main divider stack are vulnerable to self-heating effects, as is the bias network, though this phenomena has not been characterised. Even with symmetrical cycles, there must exist a mechanism by which to remove any local heat generated because the buried oxide layer and field oxide regions will act as thermal insulators. Further details, along with an explanation of thermal resistances and capacitances can be found in [8] [3].

5.4 Measurements

The design, in its own pad ring including input bias and output buffer, was placed inside 3.2mm x 3mm area of free silicon. A closeup of the fabricated design is shown in Fig. 5.36. It was realised after the receipt of the dice that the design was placed incorrectly in the space allocated, for which the author accepts responsibility. The problem is that

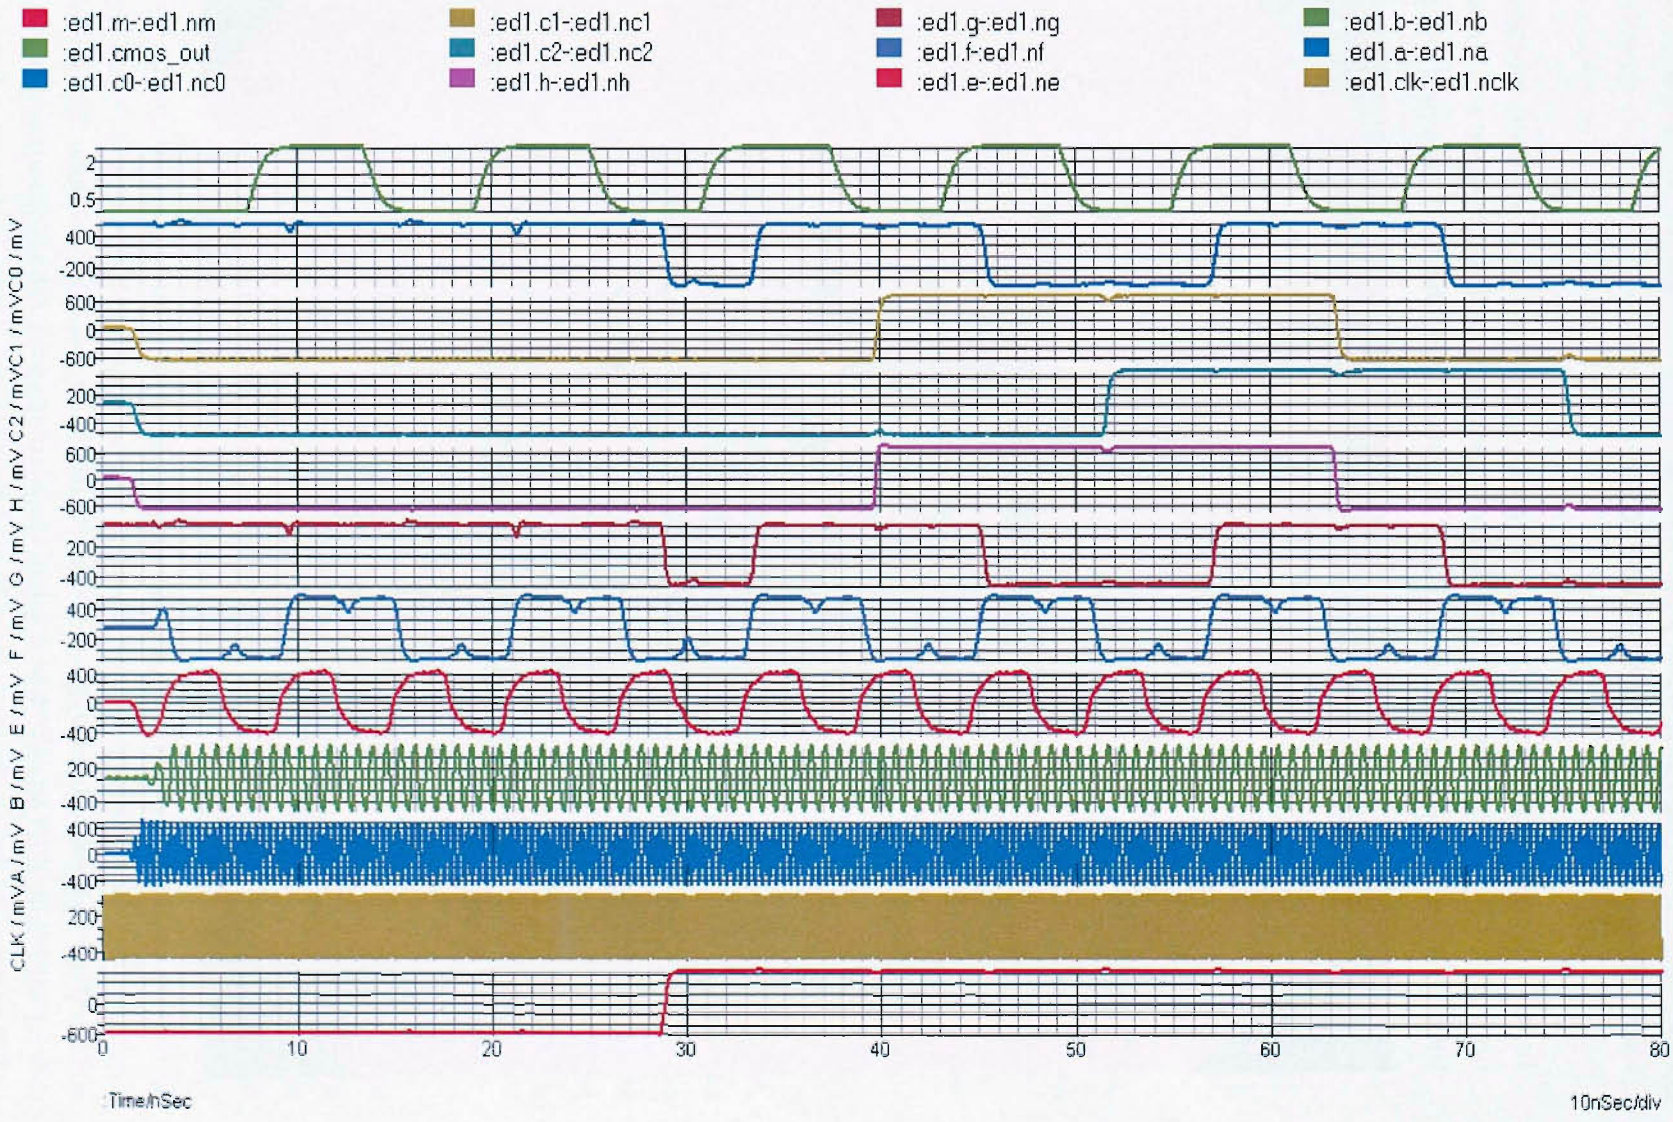

FIGURE 5.35: Screen capture of a post-layout 5.5GHz simulation with parasitics.

the high frequency inputs to the chip should have been placed as close to the saw plane as possible. This results in a lower bondwire inductance, leading to less input power being wasted in the form of reflected power.

In order to test the IC, each individual die must be mounted onto a separate substrate, upon which lie several discrete components and the connectors for powering and biasing the IC, as well as applying and retrieving signals. With the high frequency inputs

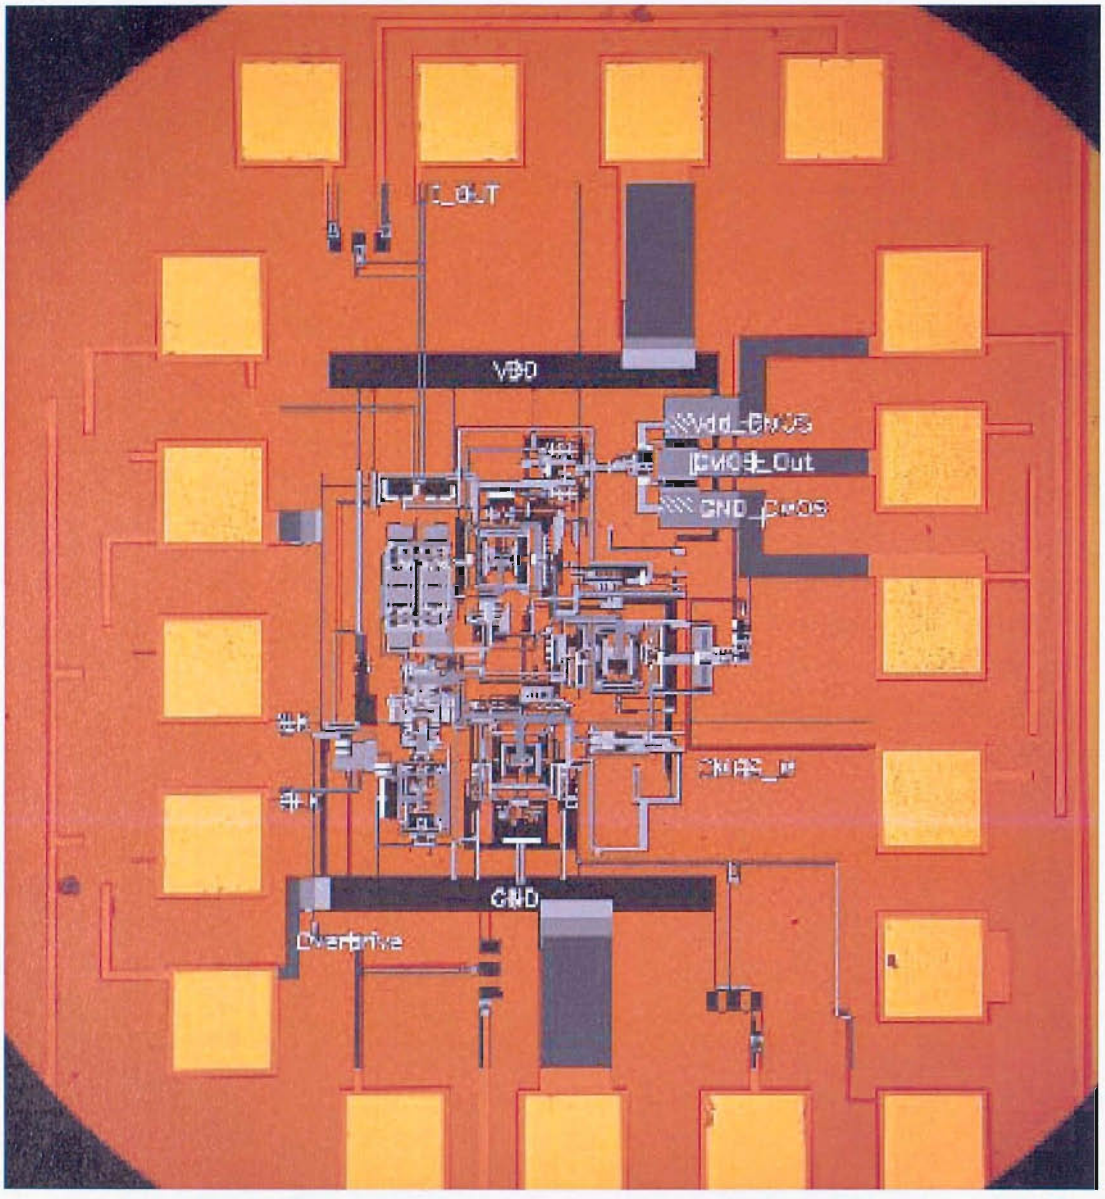

FIGURE 5.36: Photo of divider die.

supplied from a high frequency source (with the use of a hybrid coupler), matching is also required in order to deliver as much of the source power to the 50Ω terminations on-chip as possible; hence the motivation behind transmission lines (on the test substrate) and their design. The design of the board, together with the justification behind some of

the decisions made, as included in Appendix A. Once the die has been mounted into its correct position, gold wirebonds connect the bond pads on the chip to the gold landing sites on the alumina tile. DC supplies and testpoints are connected using on board SMC connectors.

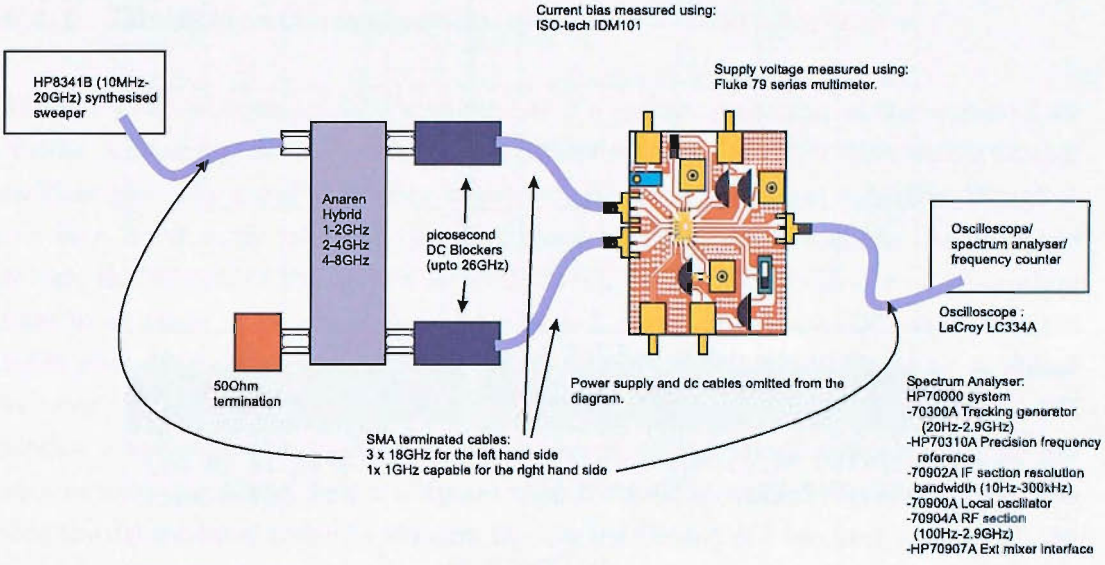

FIGURE 5.37: Test arrangement for the frequency measurement of the SOI dual modulus divider.

Initially, the DC conditions were measured where possible. Measurements showed that, in all five boards constructed, with a 6.8V supply and a 1mA nominal current bias, the on-chip input common-mode bias was 1.8V, the bias voltage on the gate terminal of the current source transistors close to the circuit ground was 1.3V, and the DC voltage at the output of the divider was 1.25V (half way between the 2.5V and 0V rails). These values matched those extracted from simulations. Having the correct DC parameters is crucial feedback for the STAG MOS model in terms of its large signal modelling.

In order to verify the operation of the dual modulus divider, as well as obtaining a value for the top input frequency, the setup shown in Fig. 5.37 was constructed. Here, a single-ended signal arrives at the input port (marked 'Δ') to the hybrid coupler. This four port device has an unused port terminated with a 50Ω broadband load (port marked 'Σ') with the differential signals emerging from the pair of '3dB' ports. These differential signals have a nominal π rad phase relationship between themselves as well as each having a power 3dB less than was delivered to the input port (assuming correct matching) whilst driven with frequencies in the pass-band of the coupler. High quality DC blocking capacitors are fixed to the outputs of the coupler, ensuring that no DC signal arrives at the inputs to the divider IC, which might otherwise upset the on-chip bias. The loss in the test setup from the signal source to the SMA connectors on the

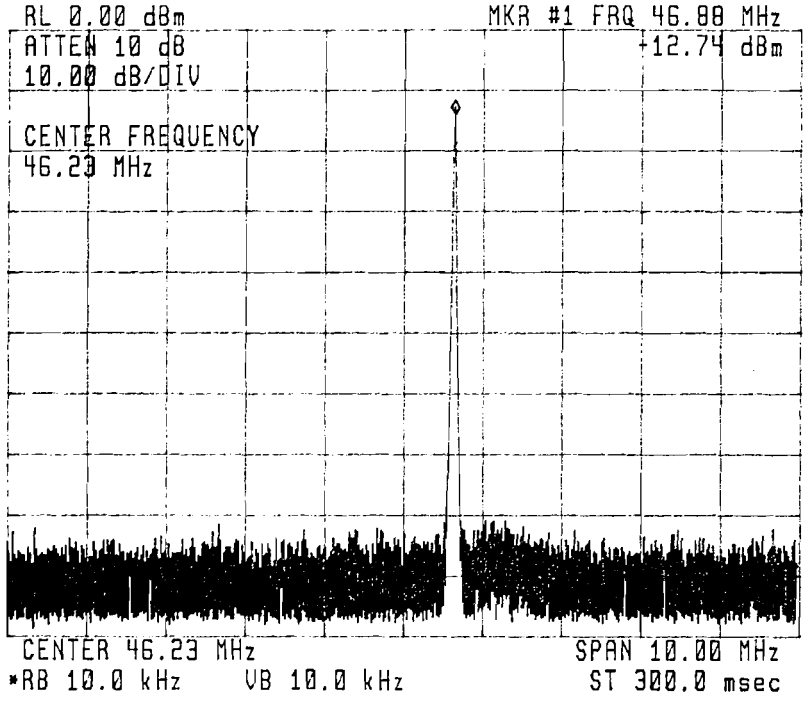

(a) Divide by 64.

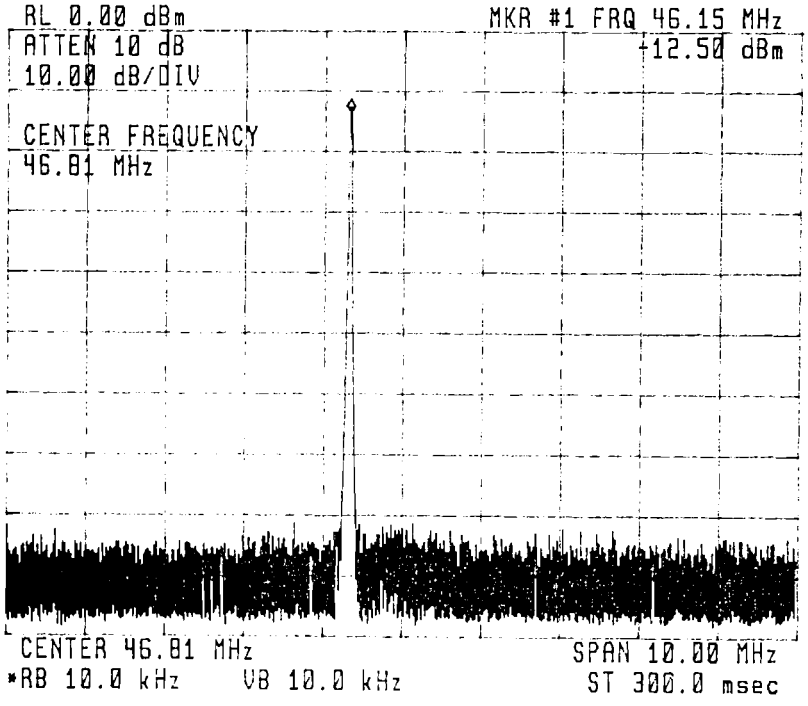

(b) Divide by 65.

FIGURE 5.38: Plots showing the output after dividing a 3.0GHz input (NOTE: a 20dB attenuator is placed before arriving at the RF input of the spectrum analyser: V_{dd}=6.82V, I_{bias}=1.15mA, input trigger power=2.25dBm.)

test board has been measured to be 4.5dB at 2GHz (1 x SMA RF cable, 1 x hybrid coupler, 1 x DC blocker and 1 x SMA RF cable).

5.4.1 Division ratio evaluation

The operation has been verified with the use of a spectrum analyser at the output of the divider. Obtaining the output frequency in this way is a lot easier than with a storage oscilloscope. The input frequency is simply read from the signal generator though it has been fed directly into a spectrum analyser to ensure its calibration. As described before, the output of the divider is a rail-to-rail 2.5V CMOS output and precautions have to be taken to ensure that no large signal, more importantly DC, arrives on the input port of the spectrum analyser. Hence, a broadband 20dB attenuator is placed between the output of the divider and the input to the spectrum analyser, and no DC blocker is required. In the following spectrum plots, only the fundamental signal has been captured, as the output is a square wave centered around a DC voltage. Note also that the die mounted onto the alumina tile marked 'Board #2' has been used for all the frequency characterisation.

Figure 5.38 shows the plots for both division ratios, 64 and 65, when driven with a 3GHz input. A marker has been placed on the apex of the tone, with which the input frequency can be verified once multiplied by the correct division ratio. The test conditions have been included in the figure caption and this is quite rigid across all three boards tested. With the current bias set approximately to 1.1mA, the current through the divider circuit including the translators and cascode biasing is calculated to be close to 11mA (excluding the current through the current bias and the CMOS switching circuits). A mistake in the design of the testboard accounts for the omission of measurements for the die's current consumption. The VDD track on the board powers the regulators as well as the test die. In order to measure the current, the track leading to the supply pad of the divider would need to be broken and routed through a digital ammeter. Unlike traditional copper printed circuit boards, metallisation on an alumina tile is 'fused' into the ceramic, hence it cannot (theoretically) be removed using an abrasive method.

Next, Fig. 5.39 captures the outputs of the divider, both divide-by-64 and divide-by-65, when driven with a 3.2GHz input signal. Again, the test conditions have been stated explicitly in the figure caption. The power supply and current bias had to be increased in order for the divider to toggle with the high input frequency. As the circuit is so large, it is hard to extract the root cause behind its inoperability at frequencies higher than 3.1GHz under the same bias conditions used for the simulations. On this occasion, the circuit is expected to draw 15mA from the supply rails.

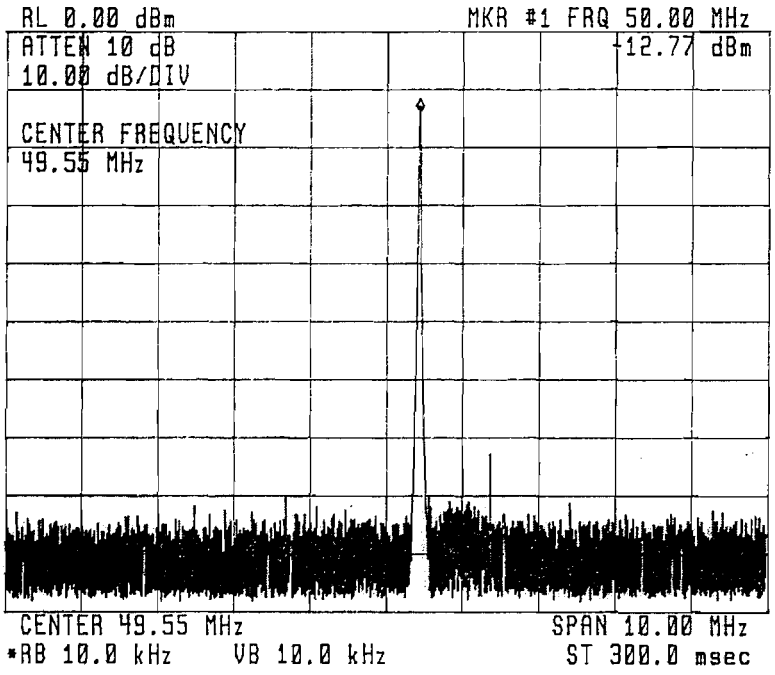

(a) Divide by 64.

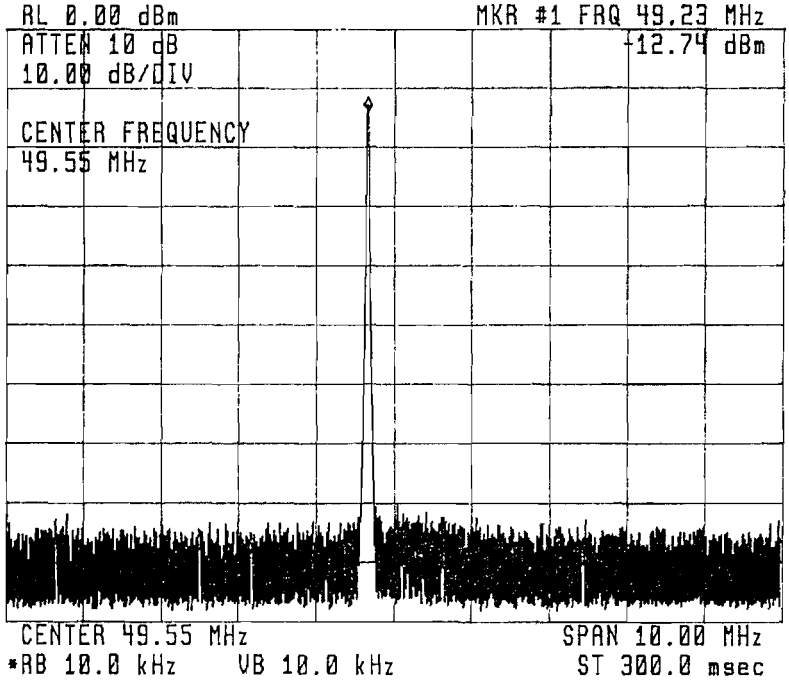

(b) Divide by 65.

FIGURE 5.39: Plots showing the output after dividing a 3.2GHz input (NOTE: a 20dB attenuator is placed before arriving at the RF input of the spectrum analyser; V_{dd}=7.09V, I_{bias}=1.49mA, input trigger power=2.6dBm.)

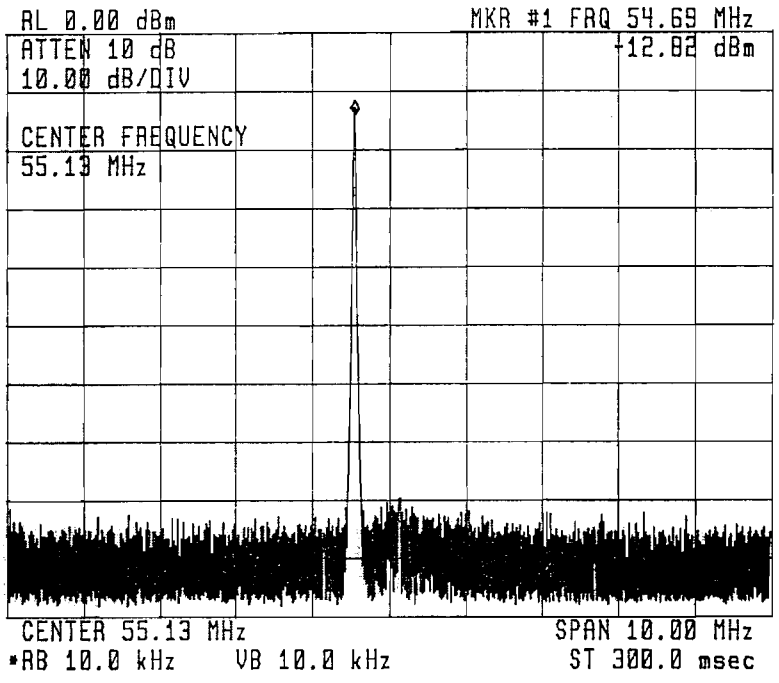

(a) Divide by 64.

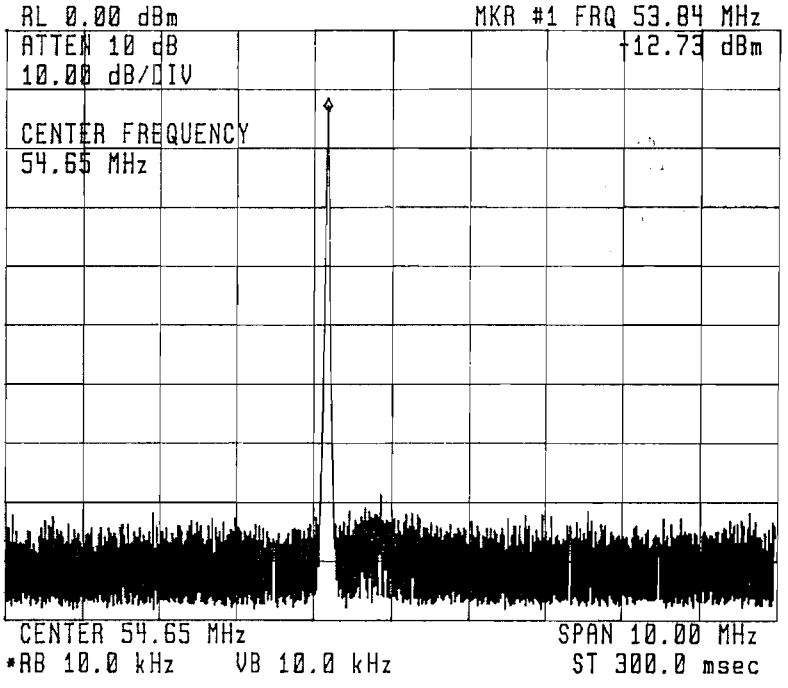

(b) Divide by 65.

FIGURE 5.40: Plots showing the output after dividing a 3.5GHz input (NOTE: a 20dB attenuator is placed before arriving at the RF input of the spectrum analyser; V_{dd}=8.13V, I_{bias}=2.13mA, input trigger power=-4.25dBm.)

The last pair of frequency plots is for a 3.5GHz input signal. The reader should take note of the approximate 20% increase in power supply voltage and the 100% increase in current bias stated in the figure caption. With these values, the circuit draws 21.4mA, assuming no strange effects with the current mirrors. This measurement in particular is testimony to the possibility of robust design in sub-micron SOI CMOS technology.

The time domain plots in Fig. 5.41 have been captured using a digital storage scope (1GS/s bandwidth). Like the spectrum analyser plots, only the output after division-by-64 or division-by-65 is on view. The intention was to capture the input and output waveforms on the same screen (or have the data points streamed to a file). Unfortunately, this was hindered by the lack of a high bandwidth scope. The images in this figure show that the CMOS tapered buffer copes well with the sub-100MHz divider outputs swinging between 0V and approximately 2.5V.

5.4.2 Input sensitivity

The input sensitivity plots have also been included in this chapter for different currents and supply voltages. Figure 5.42(a) captures the sensitivity for the 3 working boards. There is clearly a consistency upto 2.5GHz, after which boards #2 and #4 have differing sensitivities. The encouraging news is the trend and the frequency at which the higher sensitivity occurs matched well. Unfortunately, the results for board #5 end before reaching 3GHz, on the nominal supply voltage and current bias values. It is uncertain why this is the case, other than the re-attachment of bondwires which were damaged whilst constructing the test boards and different SMA connectors at the input to the board.

Figure 5.42(b) examines the behaviour of the divider IC (in particular, that mounted on board #4) with varying supply voltages and current biases. With the current bias held to the value used in the simulation, the variation in voltage was not reflected in any sensitivity of the divider between 1.6GHz and 3GHz. At lower frequencies, the divider running under nominal conditions tends to have a slightly higher need for a larger input voltage, whilst the higher supply benefits the divider working beyond 3GHz. When the current is increased at higher supply voltages, the curves move further to the right, indicating a higher maximum frequency with a lower input power, at the expense of greater power consumption and lower sensitivity at lower frequencies. Although a higher operating frequency is possible, more input power is needed at 3.0GHz for the divider to function correctly compared to when the power consumption was lower. From

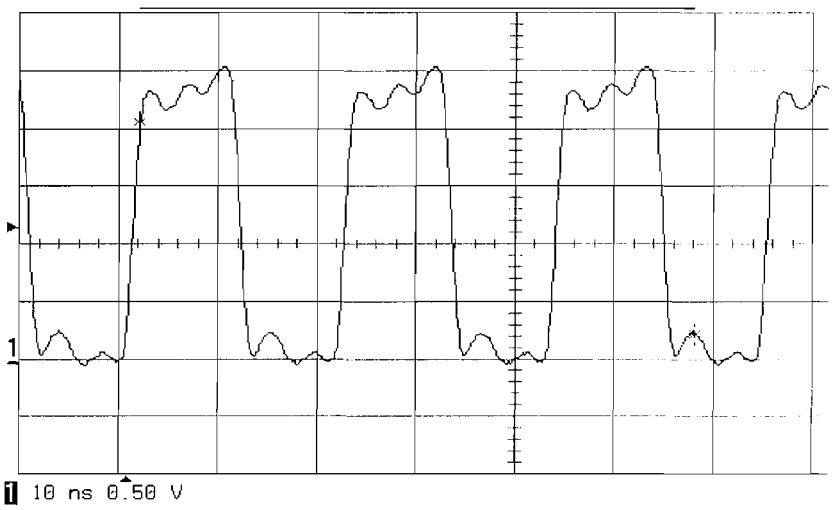

(a) Divide by 64.

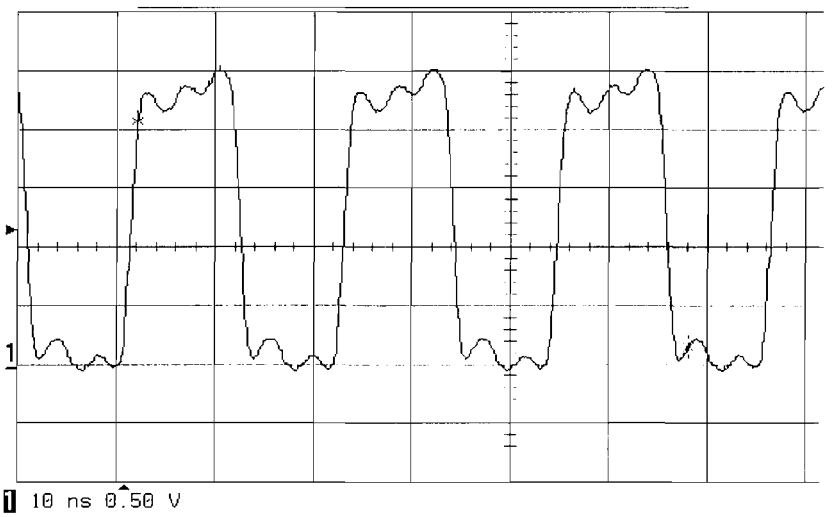

(b) Divide by 65.

FIGURE 5.41: Plots showing the output after dividing a 3.0GHz input: V_{dd}=6.82V. I_{bias}=1.15mA. input trigger power=2.25dBm.

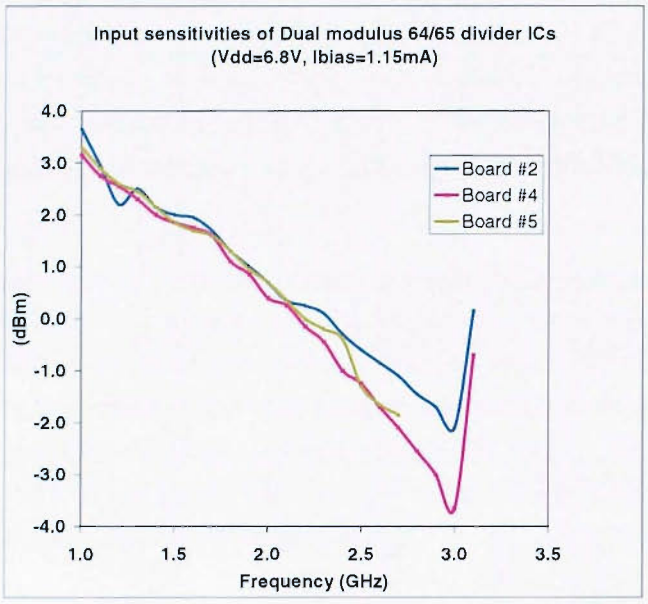

(a) Input sensitivity for different test boards.

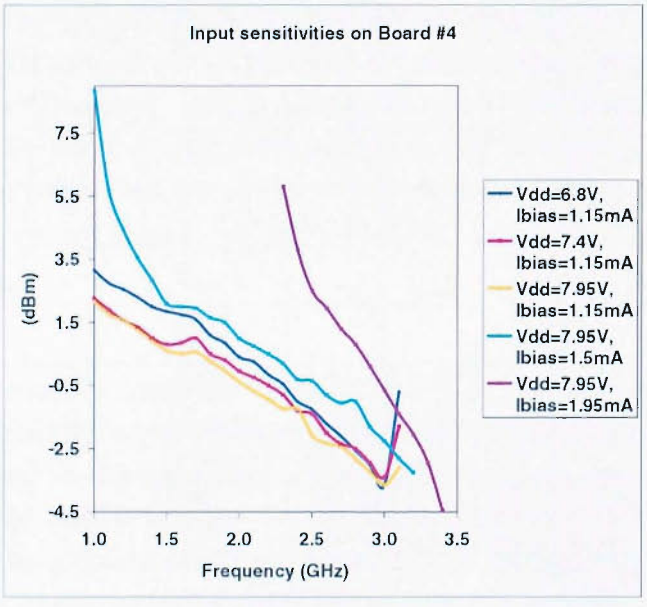

(b) Input sensitivity for 'Board #4' with different bias conditions.

FIGURE 5.42: Input sensitivity of the SOI divider design.

this measurement, it is hard to ascertain where the minimum is, as the divider loses synchronicity with the input after 3.4GHz, for reasons unknown. The observant reader will recall that Fig. 5.40 shows one of the samples to function at 3.5GHz with increased supply voltage and current consumption, but this was achieved on a different die and board. That particular die and board was found to require a higher minimum input voltage in order to function correctly as shown in Fig. 5.42(a).

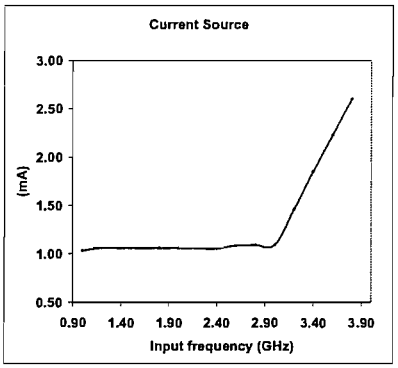

(a) Current bias versus input frequency.

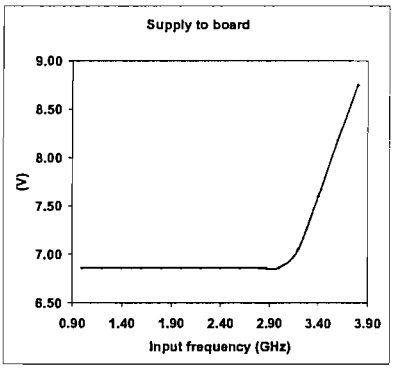

(b) Supply voltage versus input frequency.

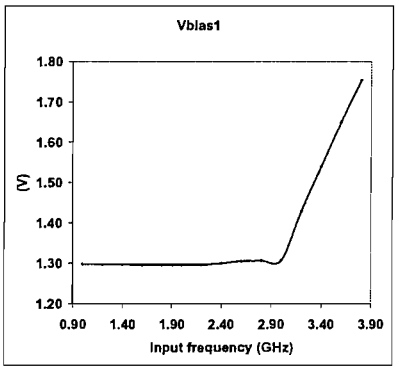

(c) Bias voltage for the current source transistors versus input frequency.

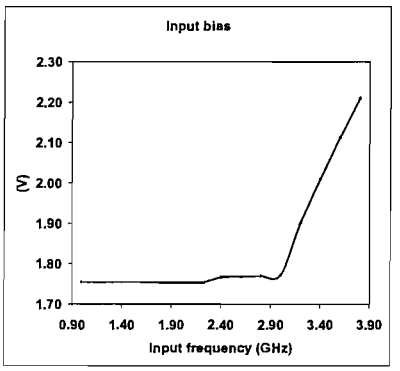

(d) Input common-mode voltage versus input frequency. (Note, this is an observation of how this parameter varies as a result of varying the power supply.)

FIGURE 5.43: Graphs showing the trend of various parameters (required for correct operation) with increasing frequency.

The graphs in Fig. 5.44 show the power at which the divider fails, and where it triggers. In order to push the divider to function at a higher frequency than observed, the DC conditions need to be changed to values outside of those used during the simulation.

Figures 5.43(a), 5.43(b), 5.43(c) and 5.43(d) illustrate how the bias current, supply voltage, voltage on current sink transistors and input common-mode voltage, respectively, need to be adjusted in order to obtain such operating frequencies.

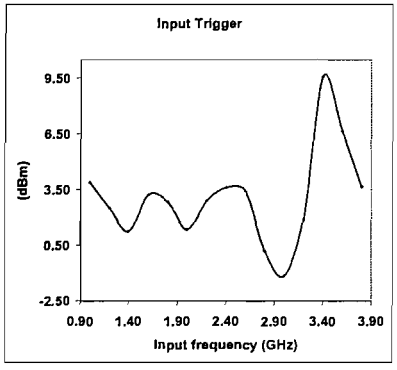

(a) Input trigger power (required to start the divider for correct operation, whilst ramping up this value) versus input frequency.

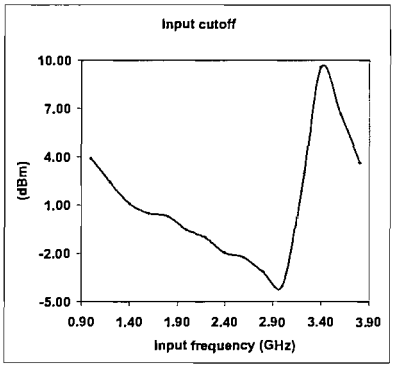

(b) Power at which the divider fails (power is ramped down showing minimum power to sustain operation)versus input frequency.

FIGURE 5.44: Graphs showing the input 'startup' and 'stopping' conditions with increasing frequency.

For completeness, the input impedance for the test board has been characterised too. As Fig. 5.45 shows, the network analyser is characterised up to the end of the SMA input cables. Hence, the measured port impedance takes into account: the SMA connector, the interface between the needle of the connector and the silver palladium transmission line, the intended 50Ω transmission line, the interface between the gold and the silver palladium tracks the bondwire and any contact resistance at either end, the pad and ESD capacitance, and finally the 50Ω termination. After performing the calibration (using open, short and broadband 50Ω load standards), the impedance characteristic of each input port was captured from the S_{11} and S_{22} data.

Ideally, a 50Ω resistance should be seen with no reactive component. Hence, the power stated on the signal generator should be the power delivered at the 50Ω termination. Looking at these plots, mismatch from the SMA connectors onwards is clearly visible. The reactance swings between capacitive and inductive behaviour within the frequency range 1GHz to 5GHz. The resistive component of the impedance falls short of 50Ω for frequencies lower than 3GHz. With no simulation of the matching networks together with the pad capacitances and bondwire inductances (only calculations of some line parameters using equations and public domain software mentioned in Appendix A),

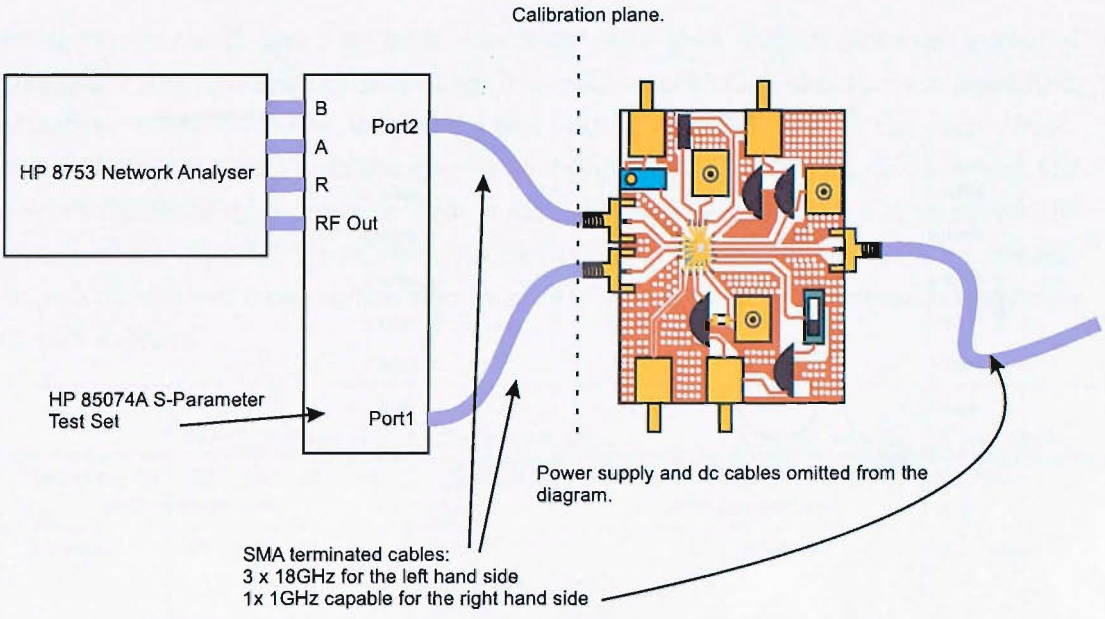

FIGURE 5.45: Test arrangement for the input impedance measurements of the SOI dual modulus divider and testboard.

poor matching between the 50Ω feed from the hybrid coupler and the input ports of the microstrips would go someway to explaining the behaviour observed in the reactance plots. Without 'Ground-Signal-Signal-Ground' probes, it is hard to quantify how much of the mismatch is due to the input parasitics on-chip and how much is due to the design and manufacturing of the testboard.

Subcircuit	current consumption
Main divider Stack	2mA
Off-stack divide-by-2 stages	3mA
SCL-CMOS level-translator	50μA
CMOS-SCL modulus control	250μA
Bias networks for cascodes	1.5mA
Master current source	1mA
Source-followers	3.4mA
Common-mode input voltage generator	500μA

TABLE 5.2: A breakdown of the static current consumption in the key blocks (V_{DD}=6.8V).

Lastly, the design has been broken down into various section with respect to current consumption and has been tabulated in Table 5.2.

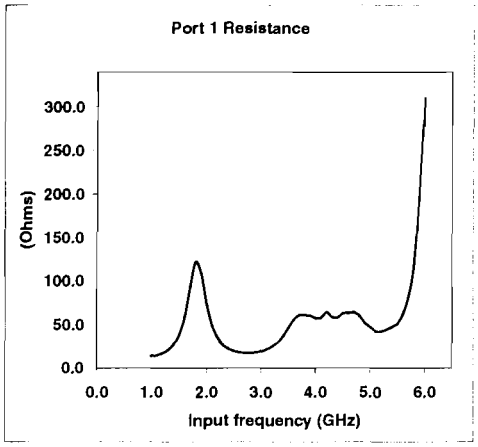

(a) Port 1 resistance.

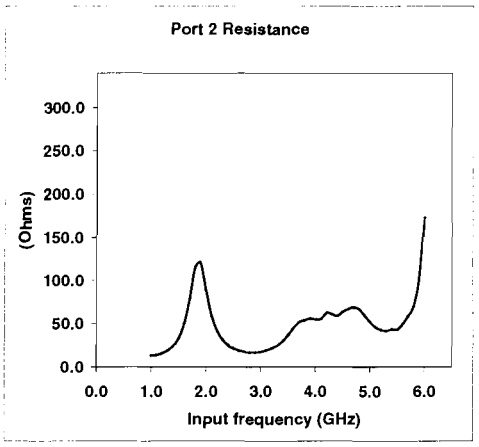

(b) Port 2 resistance.

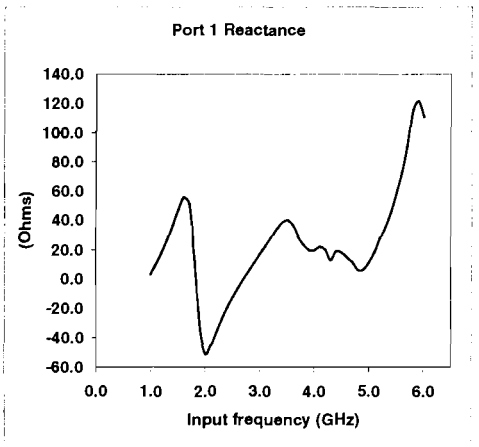

(c) Port 1 reactance.

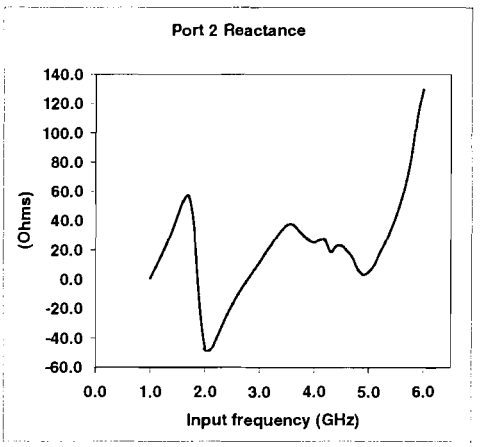

(d) Port 2 reactance.

FIGURE 5.46: Plots showing input impedance of the divider test board; V_{dd}=6.82V, I_{bias}=1.15mA.

5.5 Discussion

The previous section has proved a number of concepts with experimental data in a high frequency test setup. A new circuit topology of the phase selector, and a glitch-free controller have been successfully demonstrated. In addition, a process technology has been exploited to reduce power consumption for a large section of the divider design. Although the power supply of the analogue sections is at 6.8V, the aforementioned design has proven there to be much redundancy in an SCL implementation. With it

being current-mode logic, we have done away with level shifters and even a level of transistors amongst each divider stage (current source). One should draw immediate attention to the low power usage (less than 14mW at 3GHz) within the main divider stack. When compared with the current consumption of the off-stack divide stages, the level of functionality running on 2mA is quite remarkable. Table 5.2 also reinforces the argument for stacking, when observing the current consumption of the level shifters, though their power consumption may be pared back by focusing attention on the design of each follower.

	Krishnapura [15]	Craninckx [9]	Lam& Wu [6]	This work
Modulus	220-224	128/129	2	64/65
V_{dd}	2.2V	3V	3V	6.8V
Current	26.8mA	8mA (1^{st} div: 2.5mA 2^{nd} div: 1.5mA phase select: 2.5mA last div-by-32: 1.5mA)	0.8mA	10.7mA
Max freq.	5.5GHz	1.75GHz	5.7GHz	3GHz
Technology	0.25μm Bulk CMOS	0.7μm Bulk CMOS	0.4μm SOS CMOS	0.35μm PDSOI CMOS

TABLE 5.3: Comparison with other dividers.

5.5.1 Comparison

Clearly the design doesn't perform as well as some of the cited publications, with Table 5.3 providing a sample of other dividers. However, owing to the rather basic design setup for SOI CMOS here at Southampton, one of the main aims was to get working silicon as well as a divider operating in the gigahertz region. The STAG MOS model has never been characterised for high frequency operation within a circuit. Although the DC aspects of the model seem to be correct, the > 40% reduction in top speed could be attributed to the lack of good high frequency modelling.

The table only shows one publication [9] where the authors have divulged a breakdown of the current consumed within some of the major blocks. Though their process was dimensionally larger, the reduction in current between the first and second divide stages falls short of a factor of 2. One reason explaining this is their choice of a large output swing and the lack of a current source. The phase selector circuit also consumes a notable current. Their decision to use CMOS logic for the latter divide stages pays dividends with a reduction in power when one takes into account the fact that they have

a division modulus 2 times higher than that reported for this chapter's design.

The *silicon-on-sapphire* (SOS) design [6] has been measured to run with a higher input frequency. However, their design is merely a divide-by-2 circuit with no long feedback loop, thus allowing them to drop the current consumption whilst keeping the power supply the same.

5.5.2 Analysis of issues for improvement

The first issue to arise is the low number of body contacts per device. During layout, the importance of tying the periphery silicon around the active area was not appreciated nor was the concept of a complex RC network within the silicon film. With hindsight, the author acknowledges the need for using body-ties more extensively for the high frequency portions of the design (all those divide stages operating up to say 200MHz). Whilst the footprint for each device would increase, the ability to remove and supply mobile carriers via this node is one that should not be compromised as would the reduction in parasitic resistance to the 'electrical' contact of this body node [4].

Other sources contributing to the lower observed top frequency are: mismatch and large attenuation in the transmission lines on the alumina test substrate, no negation of the reactive component of the input impedance, lack of decoupling on chip, and lack of robust high frequency ESD pads for the differential inputs. Prior to placing the design in its own pad ring, a conscious decision to wirebond the device was made and the author still maintains that, although probing the design with an RF probe station may have yielded better results, the cost of GSSG probes was and still is prohibitive (at the time of writing this text.)

A major flaw in the layout, especially being a high performance design, is the lack of 'current direction' matching in the transistors themselves. A redesign should consider having the current through any active area travel along the same vertical/horizontal line. Although it is unclear how much degradation there is due to the lack of systematic matching, better layout can only serve to rule out its contribution to deviation from the simulations.

If a second tape out was planned, then the current in some of the low speed sections could be reduced considerably, as the current budget is quite conservative. With the main stack operating on 2mA, it would not seem unreasonable having the bias network,

CMOS-to-SCL level translator, SCL-to-CMOS level translator plus two more divide-by-2 stages consuming more than 2mA, giving a total current consumption of 4mA.

Testing at high frequencies with low power levels compared with high power high frequency design is relatively difficult. Factors such as skin effect, dielectric loss, poor-grounding, and decoupling power supplies all need to be given careful consideration. Improvements are certainly needed in order for testing to be more effective. One such improvement is to have the constructed alumina board (minus the SMA connectors) fixed onto a mechanical substrate such as an aluminium block and held down by screws. The SMA connectors would be fixed to removeable metal pieces, allowing them to slide on and off the block. Contact from the centre of the connector to the high frequency transmission lines would be achieved by means of contact pressure. At gigahertz frequencies, energy propagates to a greater extent by means of radiation and the E- and H-fields would couple to the transmission line. Soldering the needle to the tile doesn't achieve much in the way of energy coupling as skin effect dominates at these frequencies with tin/lead sharing poor conductance at such frequencies. The test board would need a slight redesign in that some of the plated through vias on the DC connectors would have to be removed. Also, holes for the mounting screws would have to be placed, though this requires thought as they can reduce the effectiveness of the vacuum when fabricating the desired plated through vias. Provision for measuring the current through the die would also be an amendment to the current test board.

Another worry with regard to testing is the relationship between bias current to the chip, power supply, and the input common-mode voltage. As this voltage is set by the current and/or power supply, any change in their values during testing also changes its value. However, during the layout phase, a conscious decision to override this voltage influenced a centre-tap between the AC ground at the 50Ω termination end, and the DC voltage set by the on-chip bias network consisting of MOS diodes, and this can be set using the 'overdrive' SMC connector on the alumina testboard. However, during measurements, no change in its maximum operating frequency was observed. Figure 5.43(d) does appear to contradict that statement, but it must be stated that the common-mode input voltage rises with frequency (after a 'knee' on the frequency axis) as a result of the power supply being increased.

With there being a SCL-to-CMOS buffer between the output bond pad and the output of the core divider, any phase-noise measurements would mask the true performance of the dual-modulus divider. A redesign of the test substrate should allow for a second die to be placed (together with a high frequency $\frac{\pi}{2}$ rad phase shifter), providing a means for characterising this performance metric [19]. Such measurement could show the effect on the spectral purity of the divided output (if any), of the current bias resistor, and any

noise on the power supply.

Although not apparent in this design, there exists a design flaw if this circuit were to be deployed within a programmable divider architecture. The problem centres around the embedded NAND gate in the divide-by-2 unit driven by the output of the dual-modulus divider. The problem is that when this gate stops being 'transparent' and instead blocks further outputs driving subsequent divide stages, its output forces the outputs of the divide-by-2 stage to assume a particular logic state, regardless of the state of the actual flip-flop upon which it sits. This means that, when the dual-modulus divider is set to divide-by-65, the chain functions as normal with the state machine running through the phase selector sequences as normal. However, as soon as the modulus control line is toggled requesting divide-by-64, the second off-stack divider has its outputs overrun by the NAND gate. This asynchronous switching may cause the subsequent divide-by-2 stage in the FSM to clock and toggle its outputs, forcing the next state to be activated prematurely. After this erroneous state, the dual-modulus divider then resumes normal operation, giving a divide-by-64 output. Unfortunately, in a programmable divider based on a dual-modulus divider, the dual-modulus divider drives other dividers and such spikes can cause them to toggle too, resulting in an erroneous division ratio in the overall programmable divider. To correct this, the NAND gate has to be removed and an actual D-latch must be placed before the second off-stack divider. The solution to this problem will be provided in chapter 7.

Bibliography

[1] A.Abidi. CMOS-only circuits for a 900MHz wireless transceiver (invited). In *IEEE Bipolar Circuits and Technology Meeting*, pages 35–42, Minneapolis, USA, Sept 1996.

[2] B.Cherkauer and E.B.Friedman. A unified design methodology for CMOS tapered buffers. *IEEE Trans. on Very Large-Scale Integration (VLSI)*, 3(1):99–111, March 1995.

[3] B.Tenbroek et al. Self-heating effects in soi mosfets and their measurement by small-signal conductance techniques. *IEEE Transactions on Electron Devices*, 43 (12):2240–2248, December 1996.

[4] C.F.Edwards et al. The effect of body contact series resistance on SOI CMOS amplifier stages. *IEEE Trans. on Electron Devices*, 44(12):2290–2294, December 1997.

[5] J.P. Colinge. *Silicon-On-Insulator Technology : materials to VLSI.* Kluwer Academic Publishers, 1991.

[6] F.Lam and G.Wu. 5+ GHz CMOS prescaler. In *2001 IEEE International SOI Conference proceedings*, pages 65–66, Durango, CO, October 2001.

[7] G.Shahidi. SOI technology for the GHz era. *IBM Journal of Research and Development*, 46(2/3), 2002.

[8] J.Benson. STAG v2.7 model user guide. http://www.micro.ecs.soton.ac.uk/stag, 2002.

[9] J.Craninckx and M.Steyaert. A 1.75GHz/3V dual modulus divide-by-128/129 prescaler in 0.7μm CMOS. *IEEE Journal of Solid-State Circuits*, 31(7):890–897, Jul 1996.

[10] K.Bernstein and N.Rohrer. *SOI Circuit Design Concepts.* Kluwer Academic Publishers, 2000.

[11] K.Mistry et al. Scaling revisted : 100nm PD-SOI transistors and implications for 50nm devices. *IEEE VLSI Technology Symposium*, pages 204–205, June 2000.

[12] K.Mistry, W.Redman-White, J.Benson, and N.V.D'Halleweyn. Low power multi-gigahertz divider architectures in SOI CMOS employing series current reuse. In *European Conference on Circuit Theory and Design proceedings*, pages III.161–III.164, Finland, August 2001.

[13] M.M.Pelella and J.G.Fossum. On the performance advantage of PD/SOI CMOS with floating bodies. *IEEE Transactions on Electron Devices*, 49(1):96–104, January 2002.

[14] *SIMetrix release 3.0.* Newbury Technology, 1999.

[15] N.Krishnapura and P.R.Kinget. A 5.3GHz programmble divider for HiPerLAN in 0.25μm CMOS. *IEEE Journal of Solid-State Circuits*, 35(7):1019–1024, Jul 2000.

[16] P.Lu et al. Floating-body effects in partially depleted SOI CMOS circuits. *IEEE Journal of Solid-State Circuits*, 32(8), Aug 1997.

[17] R.C.Jaeger and L.W.Linholm. Comments on 'an optimized output stage for mos integrated circuits' [and reply]. *IEEE Journal of Solid-State Circuits*, 10(3):185–186, June 1975.

[18] S.Krishnan and J.G.Fossum. Grasping SOI floating-body effects. *IEEE Circuits and Devices magazine*, 14(4):32–37, July 1998.

[19] W.F.Egan. Modelling phase noise in frequency dividers. *IEEE Trans on Ultrasonics, Ferroelectrics and Frequency Control*, 37(4):307–315, July 1990.

Chapter 6

Modulo 2 frequency divider in 0.18μm bulk CMOS

In this chapter, we present a divide-by-2 circuit capable of running with an input frequency higher than 12GHz. The circuit has been designed and fabricated in a Philips 0.18μm bulk CMOS process, regarded as a process predominantly for digital designs. This work is aimed at an IEEE802.11a wireless LAN application whose channel frequencies sit in the RF band between 5GHz and 6GHz. For a zero-IF receiver, there is a requirement for quadrature local oscillator inputs to a pair of mixers, and so the design is a divide-by-2 circuit with quadrature outputs driven with a 10GHz VCO. The design represents state-of-the-art performance at the time of the design work.

The requirements in this radio design have been quadrature outputs as well as low loading on the output for the VCO. It is assumed that a differential output VCO with a reasonably balanced duty cycle will be used to drive the circuit, resulting in good master-slave action. The divider will also lead to subsequent stages consuming less power.

6.1 Circuit design

The overall circuit architecture for the final divide-by-2 quadrature generator is given in Fig. 6.1. A breakdown of the circuit, as well as the choices made in its design, is presented below.

6.1.1 Divide-by-2

The basic cell is still the SCL topology used in the previous chapter. The D-latch accepts a differential input directly coupled from a 10GHz VCO with no bias adjustment. The output specification of this cell is dictated by the LNA-mixer-low frequency amplifier chain, as well as the need to drive another programmable divider, all at 5GHz.

The schematic shown in Fig. 6.2 is the basic cell in this divider. In this bulk design, the current source has been omitted in this schematic and is added in an upper level of the design hierarchy. An explanation into its operation is avoided here, as its operation has been discussed at length in a previous chapter.

With the bulk technology, there is a subtle but important design change in the appearance of the schematic. With the local isolation of the PDSOI transistors unavailable, all body connections on every NMOS transistor are now connected to the same potential, which in our case is the circuit ground. The body effect must now be taken into account during hand calculations and simulation, especially in our chosen circuit topology.

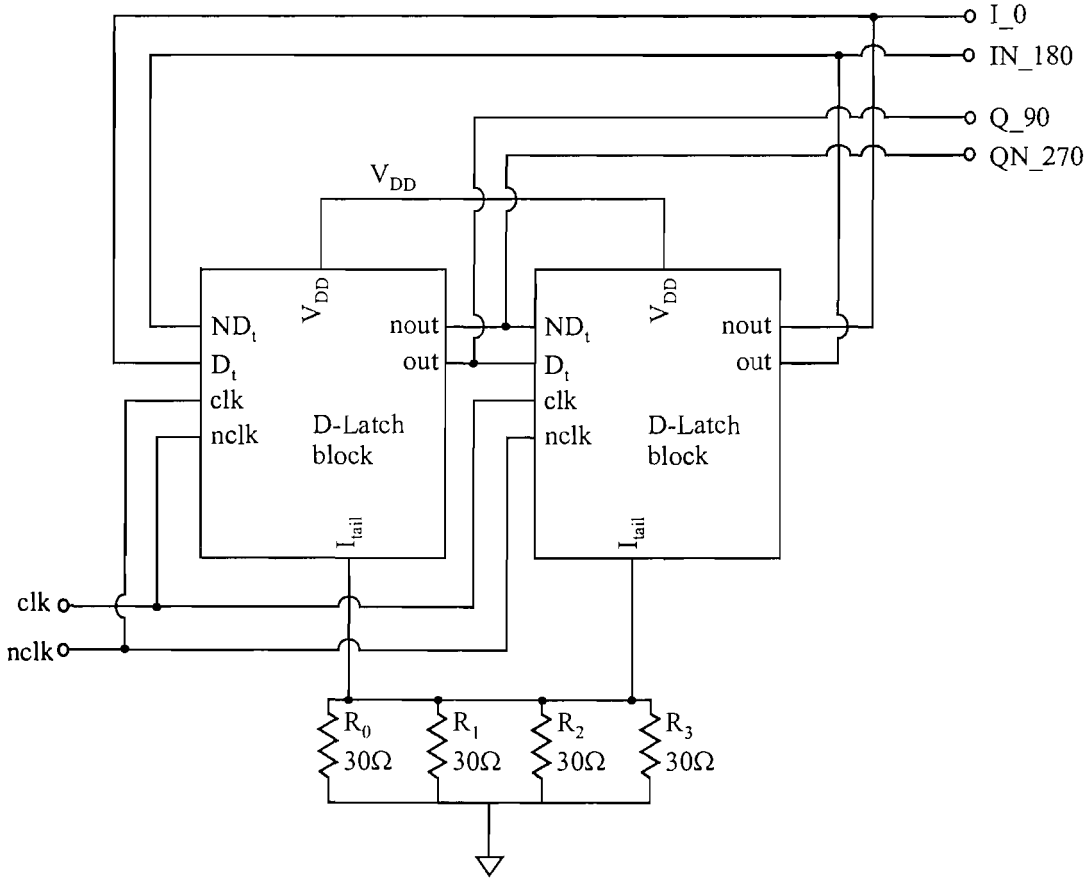

FIGURE 6.1: Schematic of the divide-by-2 circuit (higher level of abstraction.)

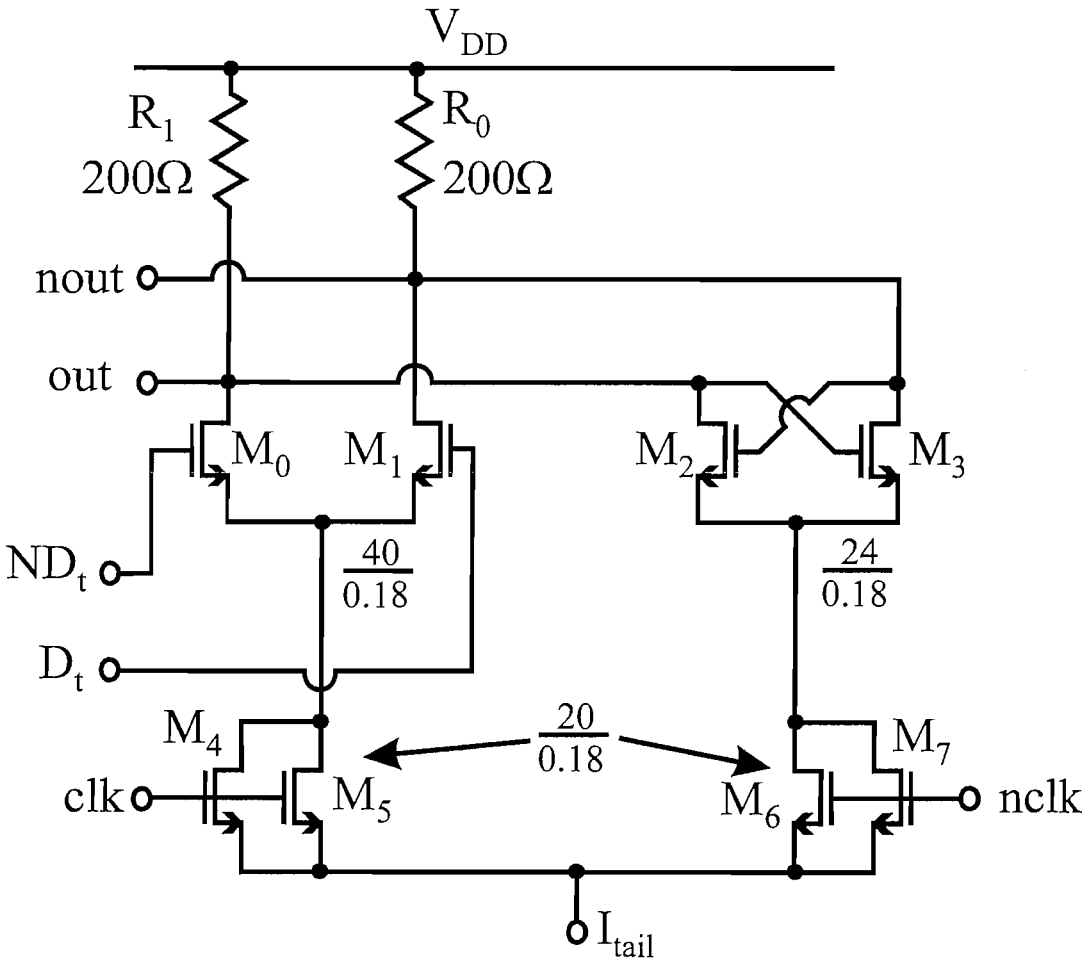

FIGURE 6.2: Schematic of the D-Latch macro circuit.

The divider as a whole is driven by a VCO with a signal swing of 700mVpp around a very stable 900mV common-mode voltage (with no load), between the frequencies of 10GHz and 10.9GHz. At the output of the whole divider, iterations of noise simulations of the whole receiver chain have yielded the requirement of a 700mVpp signal (or higher) around a 1.2V common-mode voltage, justified by the mixer specifications [5]. A number of design and simulation iterations have been performed in order to achieve the desired result. Each output of the VCO is loaded with two 40μm wide transistors and together with the output resistance of the VCO, generates a pole lower than the 10GHz.

The whole divider circuit is illustrated in Fig. 6.1 and the first observation is the absence of an active current source. Again, simulations showed the unit to achieve the desired operation using a very low ohmic resistor as a current sink. The DC voltage at the source terminal of the clock differential pair has dictated the choice of resistor. Being only 7.5Ω, it is feasible to do away with the resistor and have the clock transistors sitting directly on the ground terminal. For this circuit to work, a large signal swing is

called for, which is not really a problem for the VCO as it has been optimised for low phase noise operation. The choice behind the passive current source is one of contention, being so low in resistance and having a large gate overdrive associated with the clock transistors. With the input voltage 'spilling over' on to the current source, there can be considerable modulation of the supposed DC current, yielding a 10GHz harmonic in the output.

A second observation, and cause for concern, will be the lack of symmetry between the transistors in the latch pair and the sense pair. With such differences, a different current will exist when switching between the pairs. If one recalls, each SCL latch in a master-slave flip-flop toggles (in an opposite fashion) between the latch and the sense transistor pairs. Therefore, when the master latch pair is active, so to is the slave sense pair. The output of each latch sees the capacitance of the drain nodes of both pairs as well as the gate capacitance of the sense pair in the master/slave latch. At high frequencies, the output of the divide-by-2 circuit must toggle quickly by way of a transconductor with a high current drive, necessitating high aspect ratio transistors in the sense pair (can also increase the bias current, but it must be noted that the metal lines carrying the current is usually widened to mitigate the effects of electromigration). The latch pair, however, already has its inputs and outputs ready and, when 'energised,' the positive feedback of the pair drives the inputs apart. This implies that lower aspect ratio devices can be employed, resulting in lower capacitance on the output nodes of the divide-by-2 cell. From a divider point of view, this is not an issue provided that division is sustained. However, being a quadrature generator, quadrature mismatch is an issue with this changing shift in output voltage during a stable voltage. The divider was found to oscillate with a higher input frequency and extended swing when the capacitance on the drain terminals of the sense pair was minimised. The dimensions of the sense pair transistors remain as they are because of their current drive during transition. Although gate folding was used (where a 'wide' polysilicon gate is divided into an even number of parallel connected gates in order to lower gate resistance and halve the capacitance on the drain terminal), it was still not enough to have the latch pair transistors the same size as the sense pair. If run with a lower input frequency, switching between the sense and the latch pair will result in a distinct 'step' during the latch phase of the D-type flip-flop. Another argument for folding the gate is the reduction in thermal noise attributed to the gate resistance, with considerable effort spent refining MOS models to account for this behaviour.

Remaining on the subject of folding (though not strictly 'folding'), the clock transistor pair in Fig. 6.2 have had their gate widths divided explicitly amongst a pair of equal transistors, in order to facilitate LVS during physical verification. These transistors, as will be seen in the layout section, are split up for symmetry purposes.

Simple poly resistors were used for loads in the D-latch circuits. Obviously, when compared with active loads such as PMOS transistors, they can possess a great deal of silicon area. However, over a large signal swing, their linearity is better than that of a triode-operating PMOS. Another problem is the poor transconductance of the PMOS, thus yielding a larger device and more capacitance on the output node of the divider. One problem with the resistors in terms of noise simulation is the lack of noise models for this particular device. Superior noise models exist for MOS transistors for integration (taking into account the polysilicon metal and the gate resistance associated with it) but sadly, in the case of polysilicon resistors, only the classical small-signal noise model is used, without taking into account the parasitic and intrinsic effects (such as parasitic capacitance from the gate metal to RF ground) associated with the polysilicon itself. An alternative when using PMOS transistors as loads, is to resonate out the capacitance on the output node using an inductor tuned to the centre frequency of interest. The obvious penalty here is larger area and the need for a good scalable inductor model.

The output DC level in this circuit is crucial, with the aim to directly drive the mixers without any external level shifting. The large DC excursion from the power supply rail is the first challenge with this circuit. The output specifications required by the mixer stage require that the quadrature generator drop the DC level to 0.6V below the power supply rail. When operating the divider at a frequency below the dominant pole of the system, the result would be a swing that is twice as large as the drop in output common-mode voltage from the upper rail. With a 1.8V supply, this swing would be 1.2V and hence would thus force many of the transistors into their triode regions, thereby reducing speed. With the presented design, the current is never really switched through one transistor of a differential pair. Instead, there is a reduction in current in one half of the pair as current increases in the other. With the requirement for a 700mVpp swing, there is still 850mV left across the remaining transistors, when the uppermost transistor is conducting. Trying to push bulk CMOS to such limits is partly the reason behind a purely resistive current source as discussed earlier. The value of this output common-mode is set using half the peak current (which is the bias current through each latch) multiplied by the chosen load resistor.

6.1.2 Test buffer

For measurement purposes, buffers are required to provide a means to drive loads present on the output of the divider. The input buffer presents a 50Ω load to the measurement equipment and the benefits of this are the calibration standard and power matching. Adhering to a 50Ω standard from the high frequency signal source to the load allows the

determination of the voltage arriving at the inputs to the die. Without such a match, one is left with the dilemma of not knowing how much signal power is reflected back to the source.

The input buffer shown in Fig. 6.3(a) is a very simple differential driver and has 200Ω load converting the differential current into a voltage. This circuit has been designed for use in the stand-alone divider circuit with no intention for it to be driven by any circuit on the same die. Hence the NMOS transistors had quite an open specification in terms of the input capacitance as it was driven by an external source. However, this can only be abused to a certain point as the drain capacitance of the transistors can affect their AC performance. With the 200Ω resistor, the need for a 900mV common-mode voltage arriving at the input terminals of the raw divider calls for a 9mA static current through this buffer cell, although simulations forced the current to be increased to 13mA in order to drive the output. It is acknowledged that the common-mode voltage at the output of this buffer will drop to between 500mV and 600mV, implying a drop of less than 150mV across the latch circuit's current source (gate voltage minus NMOS threshold voltage.) Though one may argue that this is 'wasteful' designing, showing a functional divider with a low input common-mode voltage should demonstrate its ability to work with the direct input of the VCO clock. In order to drive the 50Ω output load, the output

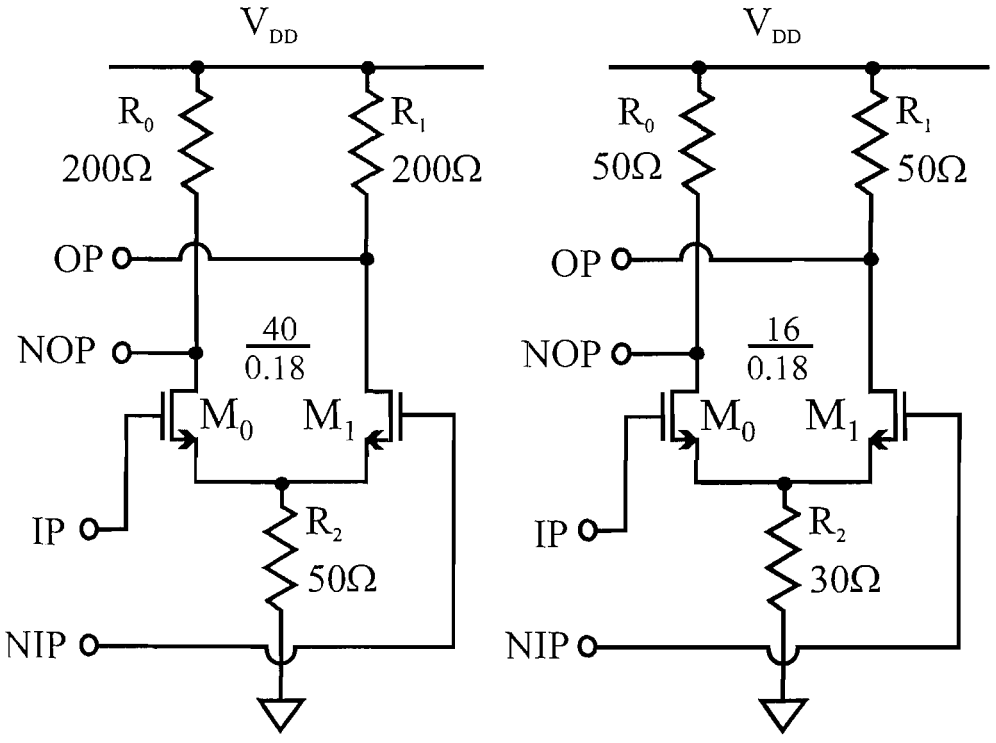

(a) 10GHz buffer. (b) 5GHz buffer.

FIGURE 6.3: Schematics of buffers.

buffers had to be driven by a higher output common-mode voltage By having a lower common-mode input voltage at the divider clock inputs, the input bias to the buffers could be raised, albeit at the expense of a lower divider output swing. The purpose of the buffers to couple the divider inputs and outputs to the pads and outside source and load impedances. The primary aim of this tapeout (with regards to the diagnosing a fault in a chain of circuit blocks) is to show that the divider is functional with a 10GHz input, with the VCO-divide-by-2 combination being a more rigorous and meaningful test.

Remaining on the topic of the input buffer, simulations showed that the input impedance would approach 50Ω, with very little reactance by shunting a 50Ω resistor between the upper power rail and the input terminals of the buffer. However, this turns out to be a rather poor solution and will be commented upon in the 'Discussion' section of this chapter.

In the case of the output buffer, it is required because the need to isolate the outputs of the divider from the contact pads. With the output still being in the GHz range, it is important to get as much of the output power to the spectrum analyser inputs as possible. Again, a 50Ω standard is used for the output impedance to match with the external cables, hybrid couplers, and so on. Without this stage, the high speed divider would need to drive a very low impedance compared with its own load resistors and this would cause a disastrous imbalance in its operation. The schematic in Fig. 6.3(b) shows the identical arrangement as in the 10GHz case, but with the need to lessen the load capacitance on the output of the divider. As there is no current budget at this stage of the development of this IC, the output swing can be left to the designer's choice. Without thorough characterisation of the buffers themselves, it is hard to obtain values for the output voltage (at the output of the divide-by-2), especially without the matching load capacitances as have been planned for this divider. After simulation, the static current through each buffer is set at 14mA, requiring 15μm width for each transistor in the differential pair in order to switch such a level of current.

Owing to a tapeout deadline, the high frequency peripheral cells had to be designed heuristically. With more time, a more linear amplifier at both input and output should be made available to allow better extraction of the performance of the high speed divider circuit.

6.1.3 VCO and divide-by-2 combination

While not by any means a commercial product, the 10GHz divider core has to be simulated with the VCO, in order to check the speed and tuning range of the combination.

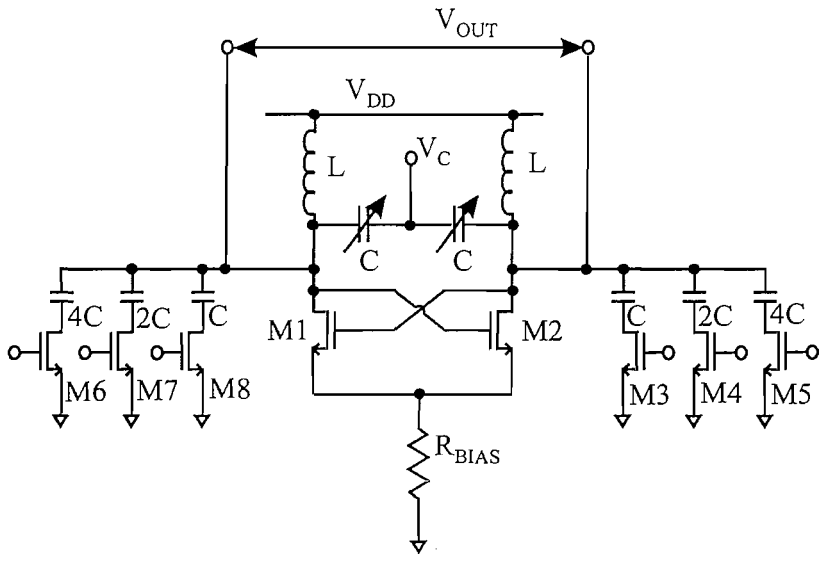

FIGURE 6.4: Schematic of the VCO circuit used in the VCO-Divide-by-2 combination.

The schematic of the VCO circuit is given in Fig. 6.4. Although the VCO has been designed with an estimate of the capacitance present on its outputs, the unit has to be coupled to the actual load in a simulation environment as the capacitances on certain nodes do not behave in a linear manner, with the large-signal output of the VCO driving the divider clock input transistors through more than one mode of operation. Due to the large gate width transistors on each output of the VCO, hand calculations based on model parameters estimate the nominal capacitance to be around 60fF on each load. With one design iteration involving the reduction of the fixed capacitance in the tank, the design was simulated to run at the appropriate centre frequency with an adequate tuning range.

At this opportunity, the fixed capacitances connected to the output node of the VCO deserve an explanation. Owing to the choice of varactor, the analogue tuning range of the VCO was simulated to be insufficient to cover the required bandwidth at the intended centre frequency, without violating the phase noise requirements of the application. Hence, a decision was made (by Philips engineers Nenad Pavlovic and Domine Leenaerts) to switch fixed capacitances onto- and off the load terminal of the VCO output, thus changing the resonant frequency of the tank in discrete steps. As can be seen, a MOS device performs the role of the switch. The values of the capacitances increase geometrically, so that a digital signal can directly programme the load capacitance (from 0C to 7C inclusive, where C is the 'unit' capacitance) without the need for some kind of thermometer decoding.

Figure 6.5(a) shows the arrangement of the cells constituting the second configuration of this chapter's core divider circuit, as taped out. For measurement purposes, 5GHz

buffers have been included on both outputs of the divider in order to drive the off-chip 50Ω load. As before, only one output pair will be measured, but the divider still needs to see a balanced load on its outputs.

As part of the integration, a 10GHz buffer was inserted between the VCO and the quadrature generator and is shown in Fig. 6.5(b). This buffer is exactly the same as that used at the input to the divide-by-2 cell, with the omission of the 50Ω shunts between their inputs and the upper power rail. The problem, however, is related to the critical design of the VCO and its centre frequency's dependence on the capacitance present (fixed or variable). A LC-tank based VCO running at high frequencies needs to have a low value for the capacitance and thus its operation can be very sensitive to parasitic capacitance especially when is comparable in magnitude to the desired capacitance. After simulating, it was shown that the tuning band had shifted above the IEEE802.11a frequency band (lower half). Though good news from the perspective of a publication, missing the specification is a problem for the project.

6.1.4 Divide-by-2 used in a 10GHz OC-192 compliant PLL

As well as the wireless LAN project, the divide-by-2 has also found its way into a PLL (also referred to as a 'clock multiplier unit' (CMU) in the context of this SONET-compliant work) designed for optical networking applications. A block diagram of the transmitter-receiver is given in Fig. 6.6, with the low-jitter CMU, clocking the retimer block in the transmitter. Apart from the phase-frequency detector, most of the sub-blocks within the CMU can be found in the receiver section, where the PLL is used as a clock and data recovery module. This work has been co-ordinated by Prof. Bram Nauta at University of Twente, together with the foundry and design help Dr. Domine Leenaerts, Nenad Pavlovic, and Cicero Vaucher at Philips Research NV, Eindhoven. The standard driving this design is an optical networking application with a need for an extremely low-jitter clock [4].

The highest speed sections of this chip (VCO and divide-by-2) were cells used from the previous section of this chapter, and hence the text below will concentrate on the remainder of the loop.

The merits of this design, from an architectural point of view, lie in the implementation of the phase-frequency detector. By using the quadrature signals already present from the divide-by-4 circuit (a pair of divide-by-2 circuits immediately after the VCO shown in Fig. 6.7), it is shown that this implementation is faster than the traditional tri-state

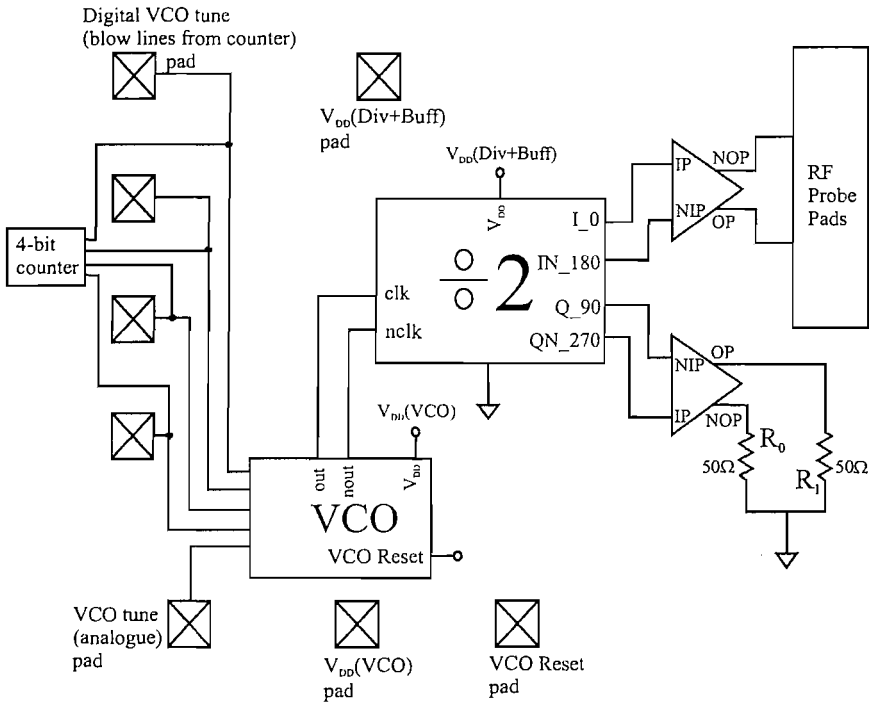

(a) VCO driving the divider loaded with the 50Ω drivers on each output pair.

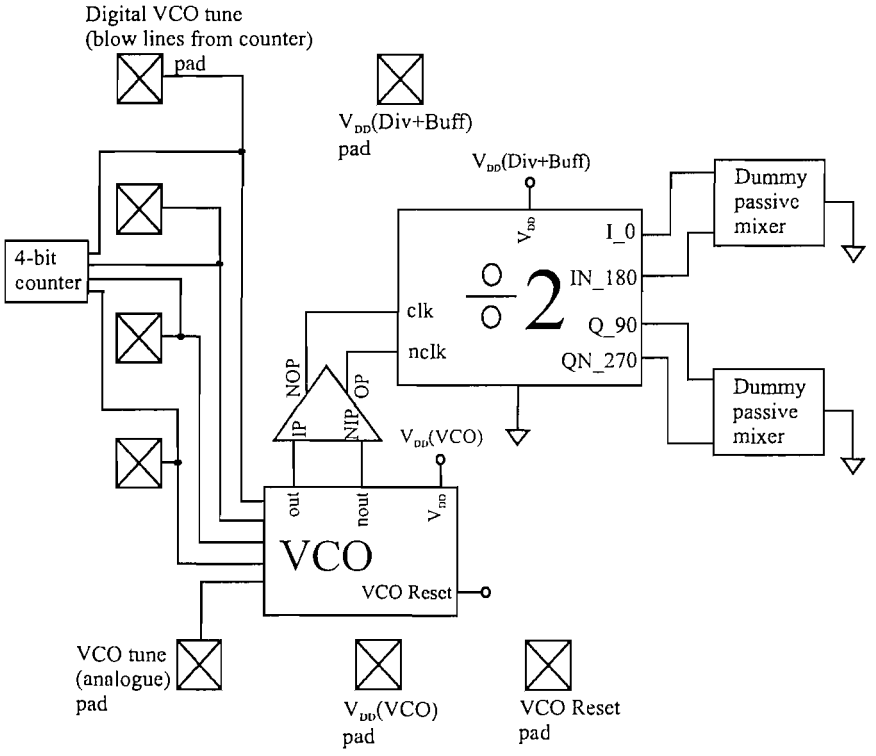

(b) VCO driving the divider with the help of a 10GHz buffer and being loaded with the dummy mixer.

FIGURE 6.5: Schematics of the various arrangements of the VCO-Divide-by-2.

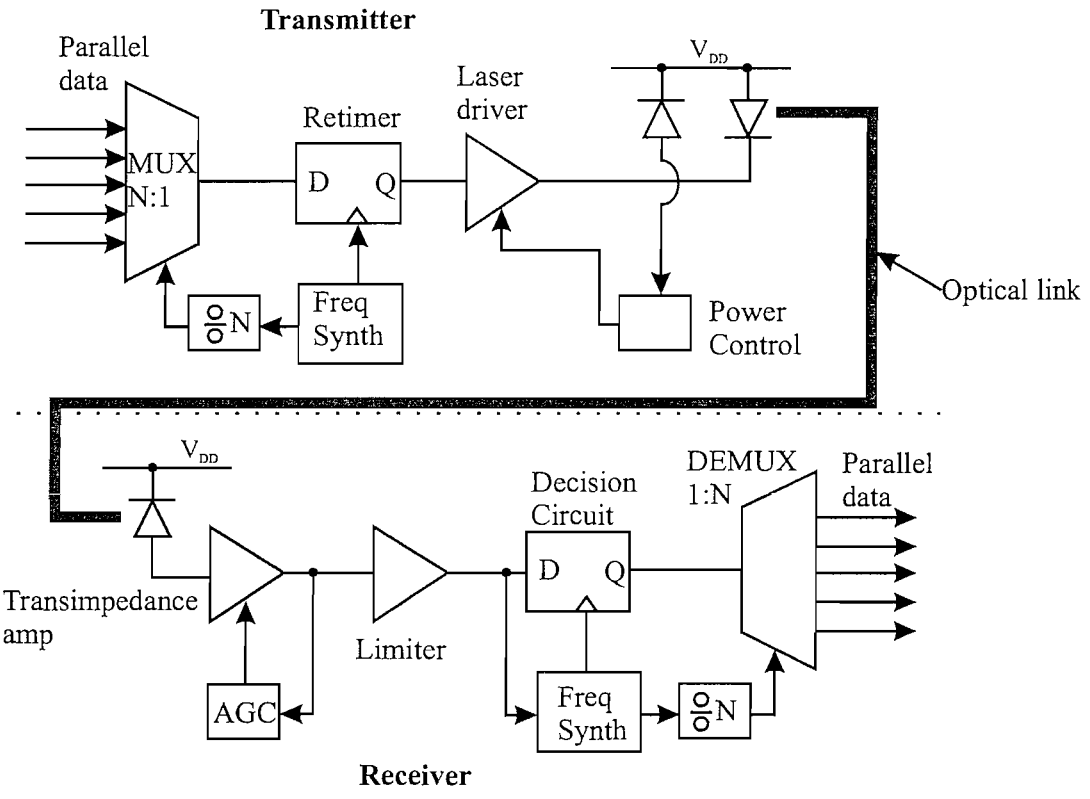

FIGURE 6.6: Transmitter-receiver block diagram, used in SONET OC-192 applications.

PFD as it doesn't use a feedback loop to reset any flip-flops asynchronously. The gain of the PD charge-pump combination does not drop significantly around zero radians of phase difference and therefore it doesn't suffer from a dead-zone problem. The duty cycle of both the Up and Down signals with respect to the charge-pump is 25% when the PLL is locked. The low spurious output is a result of the charge-pump loop filter design, with the PLL locking to a phase error of 0 rads and the charge-pump signals now cancel one another. The phase detector also has a stable gain around the lock point when a mismatch in the quadrature signals driving the detector is taken into account, even at the expense of a skewed lock range. With a correct quadrature coherency between the divided VCO phase detector inputs, the pull-in range is between $-\frac{\pi}{2}$rad and $\frac{\pi}{2}$rad.

The frequency detector also contributes to the low-noise performance of the circuit by signalling the charge-pump not to deliver current to the loop filter capacitor when in frequency lock. Before frequency lock, the phase difference between the reference input and the 'DivI' signal (refer to Figs. 6.7 and 6.8) varies almost linearly with time, implying that the average phase detector charge-pump current is the average of the PD/CP transfer curve. Ideally, this should be 0A, but is not so in practice, and integration of this current by the loop filter will cause the VCO output frequency to drift away. The frequency detector, together with another charge-pump, delivers a mean output current that is at least as high as the drift current mentioned previously

(this 'mean' charge-pump current is 50% if it maximum current owing to the 50% duty cycle of the frequency detector output). The MATLAB/Simulink simulation in Fig. 6.9 shows this benefit from the addition of the frequency detector.

The charge-pump circuit in Fig. 6.10 has been designed to respond to the 2.5GHz Up-Down signals generated by the preceding detector. This immediately calls upon the use of NMOS transistors at the input of the pump. The resulting currents from the input stage are processed by current mirrors, and though the profile of the output signal is different from the input waveform, the total charge transferred to the output is accurate. In the case of the frequency detector charge-pump, two pull-up current sources (shown dashed in the figure) are there to switch off the output current when the PLL is in lock, tracking the input.

The maximum output current delivered by the phase detector charge-pump was found to be 100μA, close to the maximum switched by the phase-detector. The maximum output current of the frequency detector charge-pump was given as 40μA and is enough to guarantee correct locking. These values are obtained after simulation.

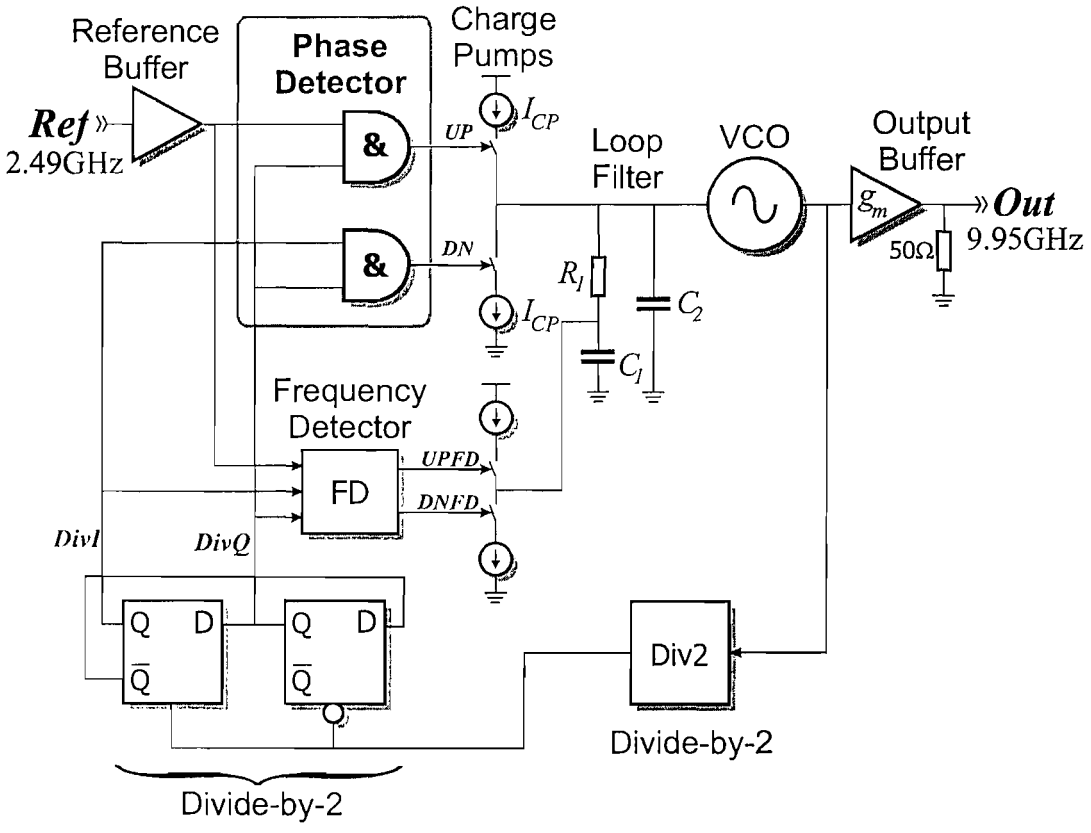

FIGURE 6.7: Architecture of the clock multiplier unit.

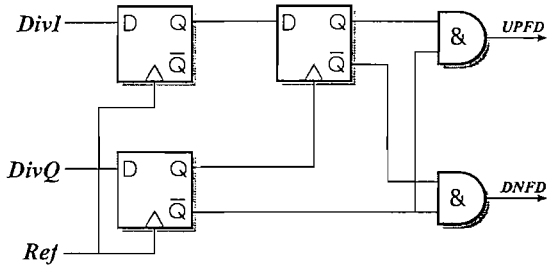

FIGURE 6.8: Implementation of the frequency detector functional block.

The decision behind the high reference frequency was motivated by the need for a low frequency division ratio within the loop. This in turn was influenced by the following equation :

$$\mathcal{L}_{CF} = \frac{1}{2} \cdot S_i \frac{4\pi^2 N^2}{I_{CF}^2} \tag{6.1}$$

This expression quantifies the phase-noise that is present on the output of the charge-pump, where S_i is the power sprectral density of the CP noise, N is the frequency divider ratio and I_{CF} is the current of the CP current sources. An assumption is made that the close-in phase noise of the clock multiplier unit would be dominated by this noise, which aids the calculation of the optimal loop bandwidth. When implementing the loop filter, the capacitors were dimensioned in order to keep the zero and pole sufficiently away from the loop bandwidth frequency. This prevents jitter peaking, as well as giving flexibility in the loop bandwidth by tuning the charge-pump current without risking an

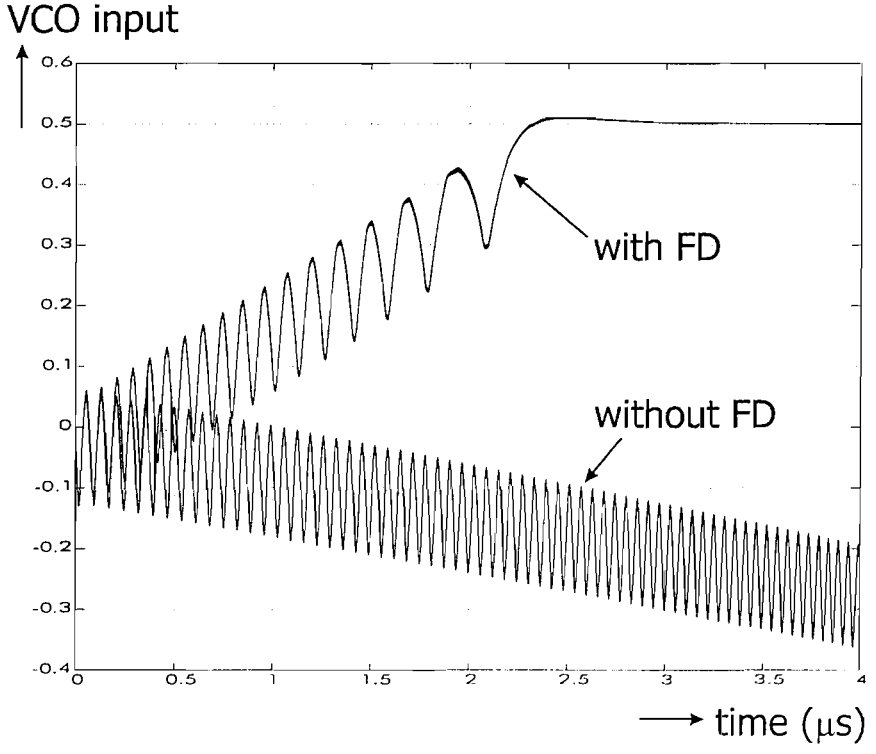

FIGURE 6.9: Simulation justifying the need for the frequency detector.

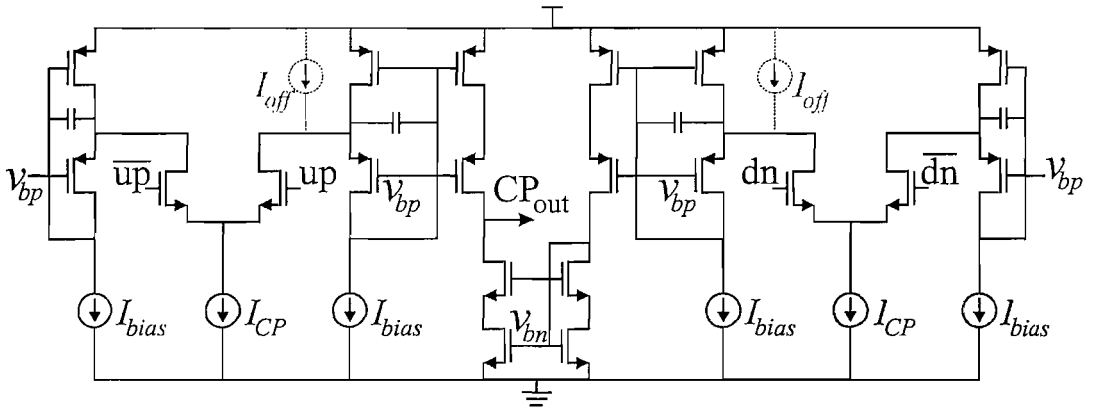

FIGURE 6.10: Charge-pump circuitry.

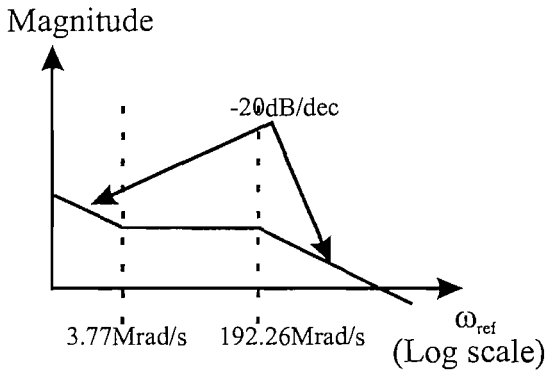

FIGURE 6.11: Bode diagram of the second order loop filter, with a pole at 0Hz.

unstable loop. With reference to the second order loop filter in Fig. 6.7 (Type-II third order PLL), R_1 is 5.3kΩ, C_1 is 50pF and C_2 is 1pF. These values result in a loop filter transfer function zero at 600.58kHz, and two poles at 0Hz and 30.6MHz. Its magnitude/frequency diagram can be seen in Fig. 6.11.

6.2 Layout

With these remaining chapters reporting on the work conducted at Philips Research Laboratories, NV, the process shifts from a SOI technology to a plain digital bulk CMOS technology, and there are some differences in the basic layout of the active circuit elements. Unlike PDSOI transistors, the bulk versions do not have the same isolation. Therefore, the bulk silicon region under the enhanced channel is electrically connected to the same potential. Depending on the application, the number of body taps can be as low as one. Unfortunately, with the high resistivity substrates used for these ICs, the complex distributed resistance capacitance network of the body gives rise to a varied potential from the body tap to the region under the channel. An example of the network to a first-order approximation is presented in Fig. 6.12.

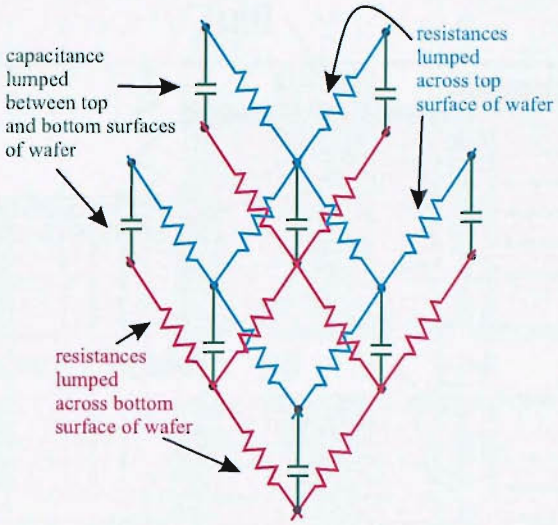

FIGURE 6.12: An electrical equivalent circuit of the substrate network.

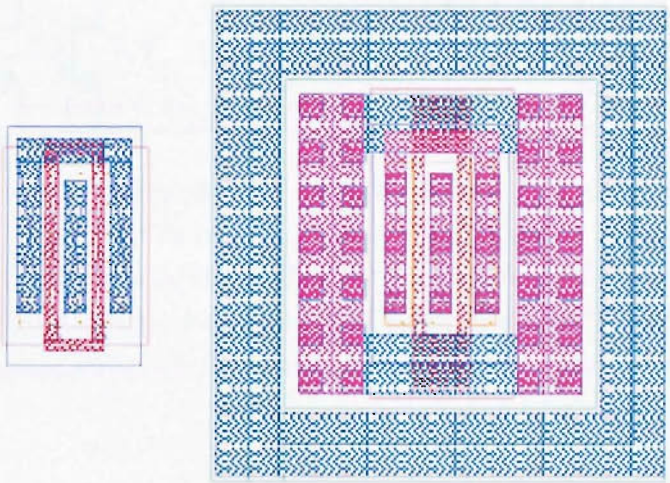

FIGURE 6.13: Footprints of a standard 5μm wide transistor against an RF version.

Figure 6.13 is a screen capture showing the difference between a standard digital NMOS transistor, and one designed for RF operation. The required area is significantly greater due to the need to clamp the substrate in the immediate vicinity to a common potential with the help of a considerable number of body taps all shorted with a low resistance metal connection. With the help of Fig. 6.12, it isn't too difficult to suspect a semi-floating body node associated with each transistor, leading to the possibility of high frequency feedback as well as a modulated body node. The models used at Philips are augmented by extracted data based on such a layout and deviating away from such a layout is at the designer's peril. As well as lowering the body resistance, the gate

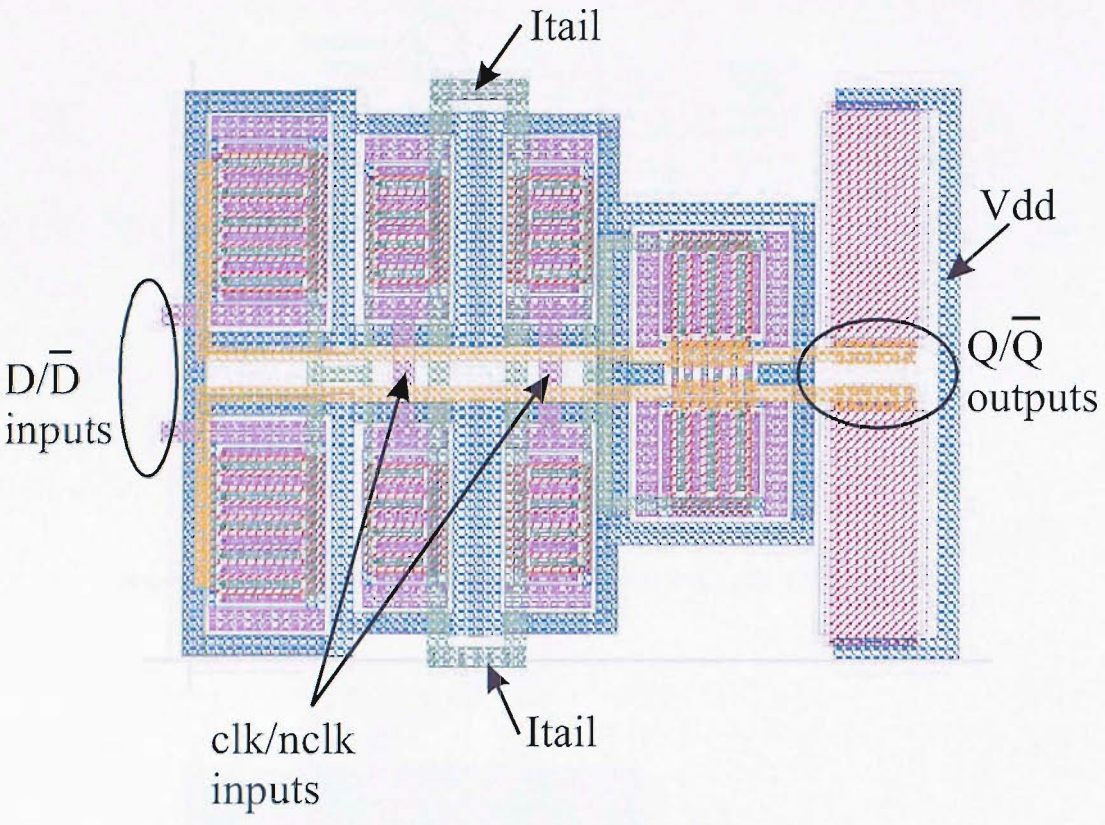

FIGURE 6.14: Layout of the taped out D-Latch cell.

resistance is also reduced, changing from the polysilicon gate to a metal material at the earliest opportunity. At first, one may be concerned at the gate-substrate capacitance, but this is tolerable when compared to the peformance gained with a lower gate resistance on such a terminal (both in terms of bandwidth and noise). Also, all layout has been done with the assumption that the on-chip temperature is 70^{o}C (required for the choice of SPICE parameter deck).

Figure 6.14 captures the final layout for the D-latch cell. Being a quadrature generator, symmetry is the key here to minimise phase mismatch. The latch is symmtrical about the centre line running the length of this cell. As mentioned in the previous section, each clock transistor is split into pairs so that no metal interconnect line crosses over an active area. The cross-coupled latch to the right of the plot, behind the polysilicon load resistors, has a balanced capacitance on each transistor's terminals using a clever interdigitated cross-coupling method with an even number of ties. The sense pair lies on the left of the plot with the load resistors at the opposite end. There is more than one possibility for the floorplan of this cell, such as the load resistors lying in the centre of the diagram, keeping the sense and latch pairs equidistant. The drawback is that the separation of the common source point of the clock transistors increases further. The differential outputs run along the centre line of the D-Latch. With this chosen layout, lumped parasitic capacitances between electrical nets are extracted and manually in-

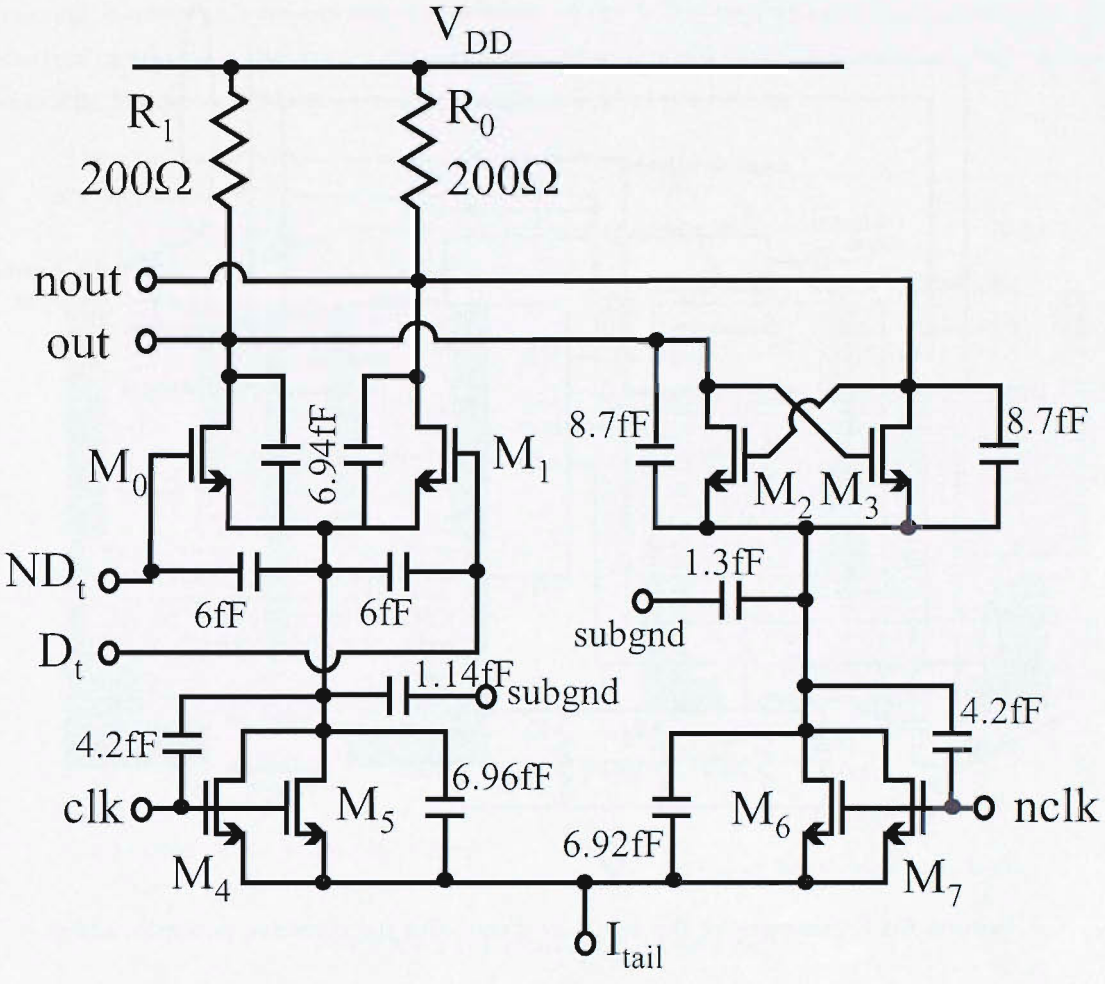

FIGURE 6.15: Schematic of the D-Latch core with the extracted parasitics added.

serted into the original schematic, with the final result given in Fig. 6.15.

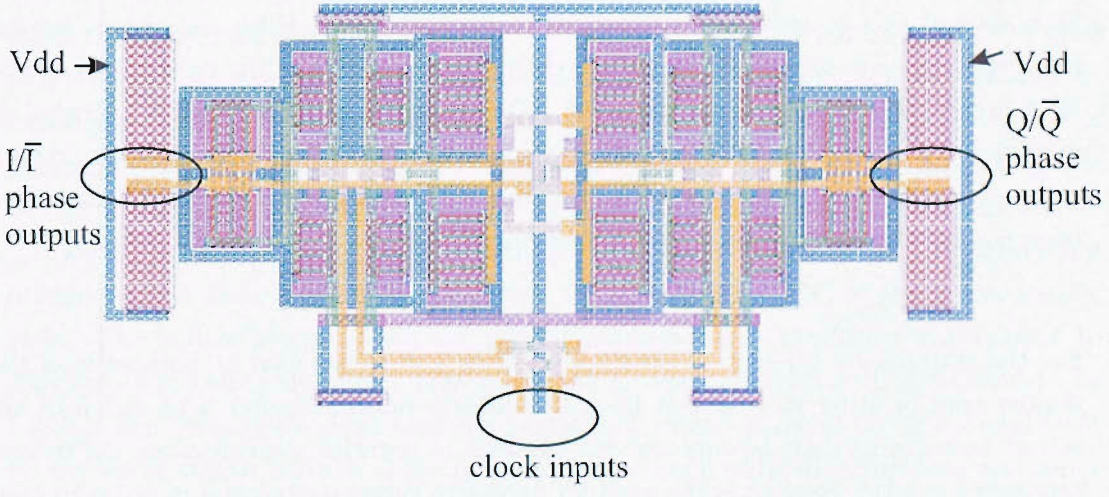

FIGURE 6.16: Layout of the taped out divide-by-2 core.

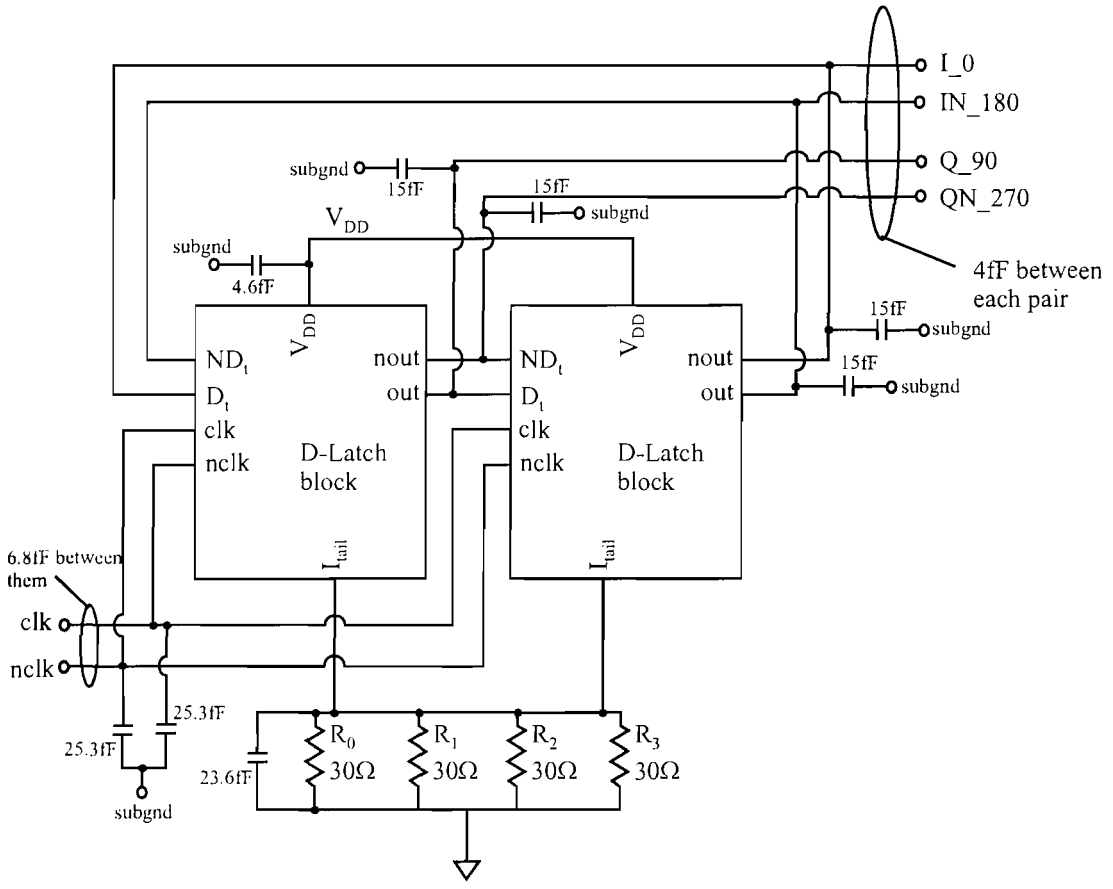

FIGURE 6.17: Schematic of the divide-by-2 core with the extracted parasitics added.

Moving up one level in the design hierarchy, the divide-by-2 layout cell is assembled as in Fig. 6.16. Taking the basic D-Latch cell, there is more than one possible configuration for this quadrature generator. One could place two identical latches side-by-side facing in the same direction. Although the capacitance to substrate on the high frequency clock lines is lower than the taped out cell, the penalty would be a considerable increase in capacitance between the output terminals of this master-slave flip-flop. Simulations have showed this to reduce the swing on the output node. With the chosen layout, this capacitance is minimised to an acceptable level, such that the noise figure of the LNA-mixer-LF amplifier is not degraded. The cross-couple in the centre is crucial in minimising the load capacitance on these nodes and keeping the latches as close to each other as possible. As before, a schematic with extracted lumped capacitors that have been manually placed between nets is shown in Fig. 6.17.

For the stand-alone tapeout of the divider, buffers were included at both ends of the divider core in order to isolate it from the outside world. Figures 6.18 and 6.19 are screen captures of their layouts and the network of parasitic capacitances. As before, symmetry is a key issue as is the need for extensive substrate tapping in order to bind the local substrate to the circuit ground. The low ohmic resistors (in the case of the

current sources and the output driver loads of the 5GHz buffer) carry a large amount of current, and hence the aspect ratio of such devices make them look unusual, but this is down to the process sheet resistance of the silicided polysilicon.

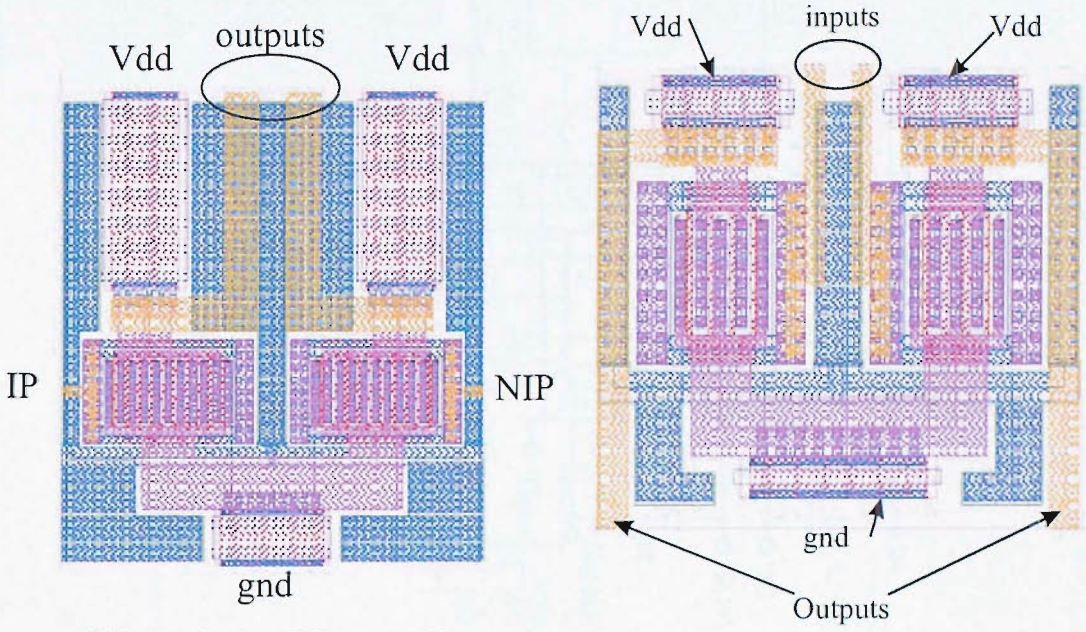

(a) Layout for the 10GHz input buffer.

(b) Layout for the 5GHz input buffer.

FIGURE 6.18: Layouts of the buffer cells.

Note that the 50Ω polysilicon resistors shunted between each of the inputs and V_{dd} in the 10GHz capable buffer, are placed in an upper level of the layout hierarchy.

The first instance of the cell is taped out in the standalone divider as in Fig. 6.20(a), with the parasitic capacitances placed in the schematic of Fig. 6.20(b). The layout has taken the shape indicated because of the compactness of the divider core, and the need to get the RF signals into- and out of the the divider circuit (plus buffers), without the signal power having to travel a large distance. The four centre pads on each side of the die are the ground-signal-signal-ground probe landing sites, with the right hand side being the differential input. The two bottom pads are the unused outputs of the divider and should have DC probes placed on them to match the capacitance and inductance present on the electrically used RF pads. The upper two pads are the power supply pads. The yellow blocks shown are metal-insulator-metal capacitors and amount to 20pF on-chip, before the power lines reach the bondwire inductances (estimated to be 1.2nH). The considerable blue area is a perforated ground plane in the first metal layer for the static return current. This approach has a much reduced inductance and series resistance and lessens the effect of an unwanted IR drop in the already low power supply.

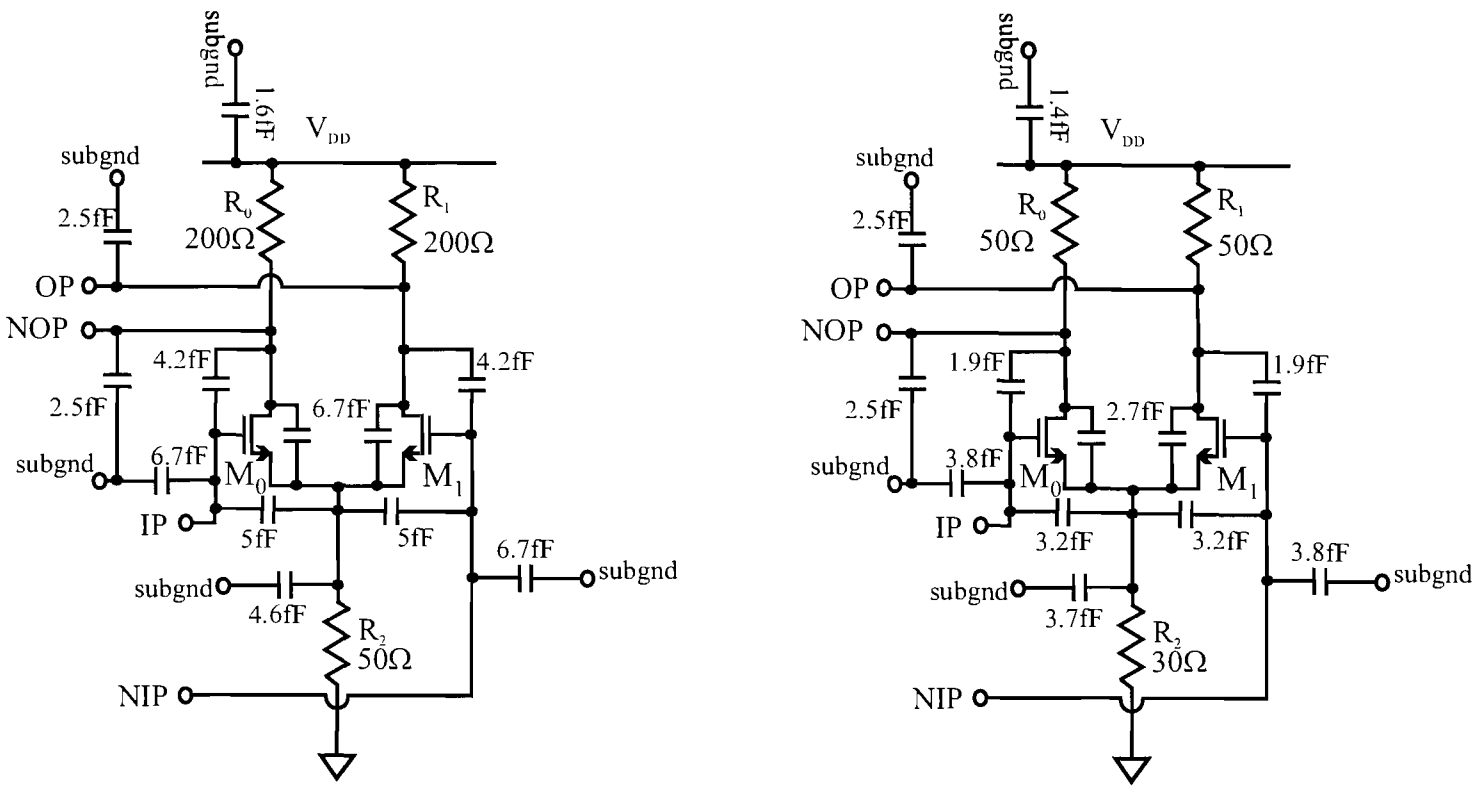

(a) 10GHz buffer.

(b) 5GHz buffer

FIGURE 6.19: Schematics of buffers with extracted parasitic capacitances added.

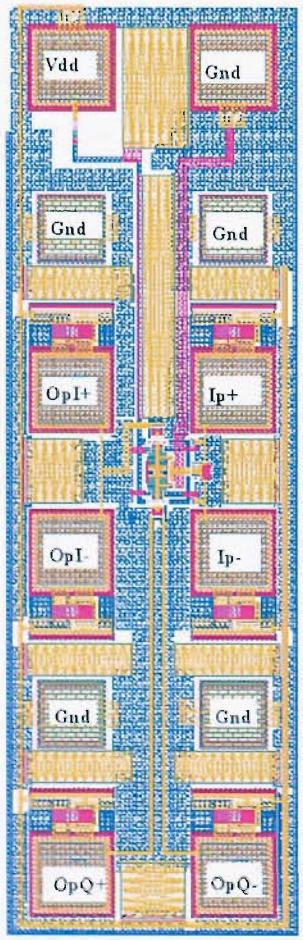

(a) Layout.

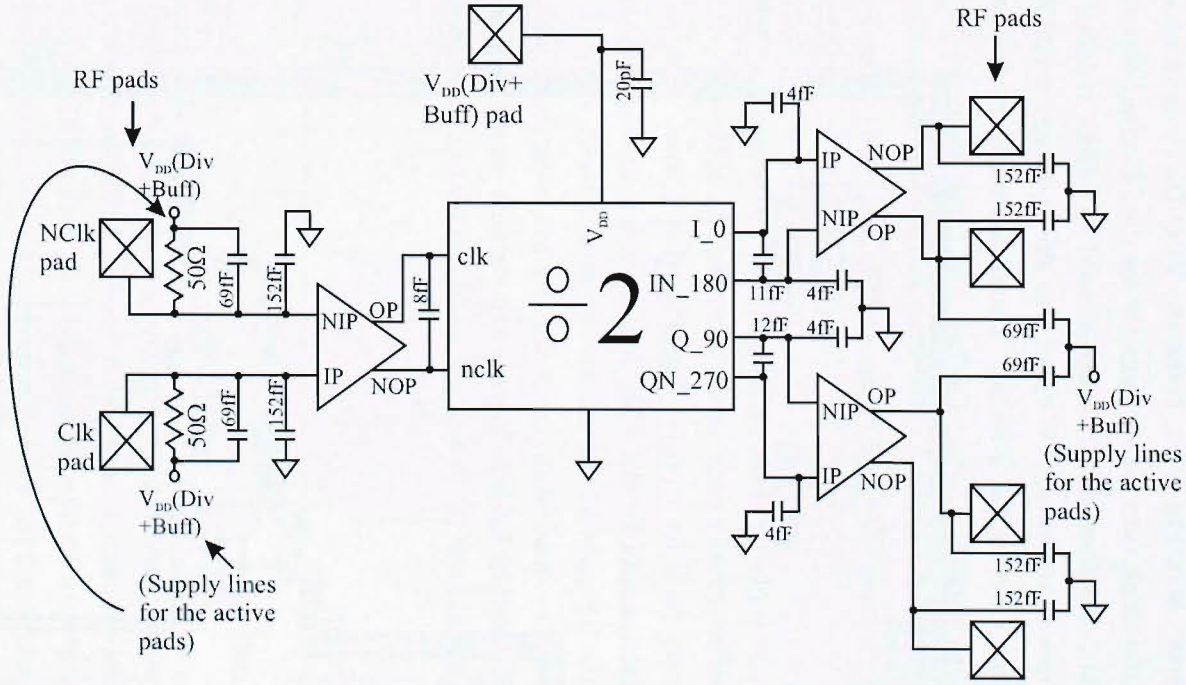

(b) Circuit with buffers, pads and extracted parasitics.

FIGURE 6.20: Layout of the standalone divide-by-2.

The final layout shown is in Fig. 6.21, being the VCO-Divide-by-2 combination. The circular element is the single turn planar inductor for the LC-tank, with the red polygon being the patterned polysilicon shield, which is used to reduce induced eddy currents and thus coupling into the substrate. Without this shield, energy would by lost into the bulk substrate, hence lowering its Q-factor. Through a number of inductor tape-outs (performed by Philips engineers Nenad Pavlovic and Luuk Tiemeijer [2]), it was found that the quality factor of the inductor could be increased with this patterning, which in turn results in better phase noise performance of the proposed synthesiser.

The divide-by-2 circuit is placed to the right of the VCO structure, just before the vertical RF ground-signal-signal-ground pads on the bottom right of the plot. 5GHz buffers are coupled to its outputs in order to drive the 50Ω off-chip load. The pads around the design comprise of the inputs for the 4-bit digital tuning, analogue tuning, power, and ground, as well as the RF pads for the divider outputs. As with the standalone divider, the two outputs from the quadrature generator are left redundant, with their outputs terminated into 50Ω on-chip loads. Having such outputs would be costly in terms of area, as the design would have to expand in two directions to accomodate the RF pads.

6.3 Simulations

6.3.1 Divider cell

With the divider circuit requiring an iterative procedure for its design, the need for robust, adequate models cannot be stressed enough. The simulations of the 0.18μm bulk CMOS divider use the Philips MOS model 9 [3] with an RF extension. The basic model without the additional RF extension models the $\frac{1}{f}$ behaviour of MOS transistors for circuit designs that require the noise to be quantified and analysed. With the RF extensions, four levels of complexity are available to the designer. Depending on the complexity and size of the design in hand, as well as design time available, the level of accuracy is determined by the choice of extension.

The extensions for the model are as follows:

- RFA : includes the gate resistance of the polysilicon material,
- RFB : as RFA, also including extrinsic source/drain resistances,
- RFC : as RFB, also including the bulk resistance network,
- RFD : as RFC, also including modelling of non-quasi static effects (very computer intensive version.)

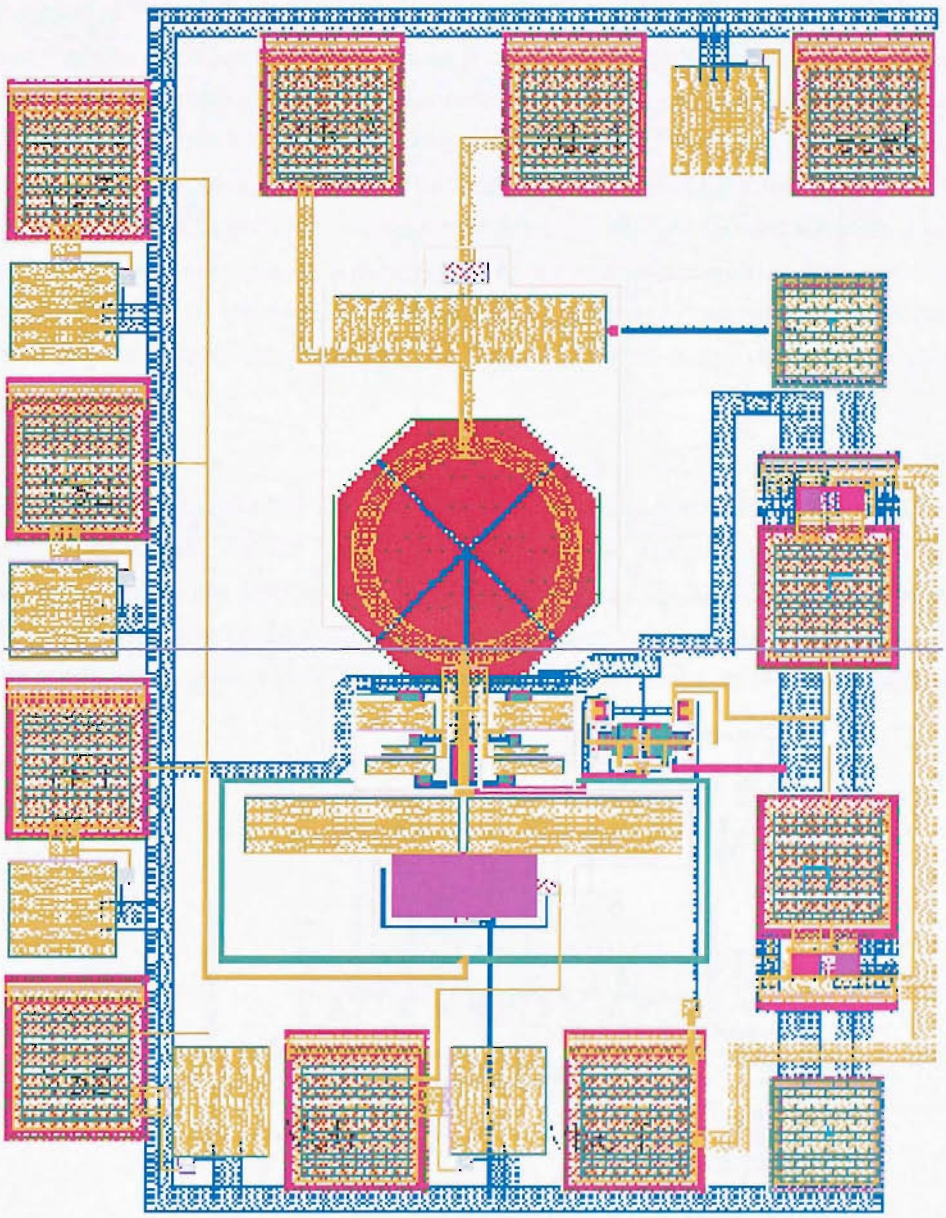

FIGURE 6.21: Layout of the taped out VCO-Divide-by-2 combination. The pads names can be found in figure 6.32(a).

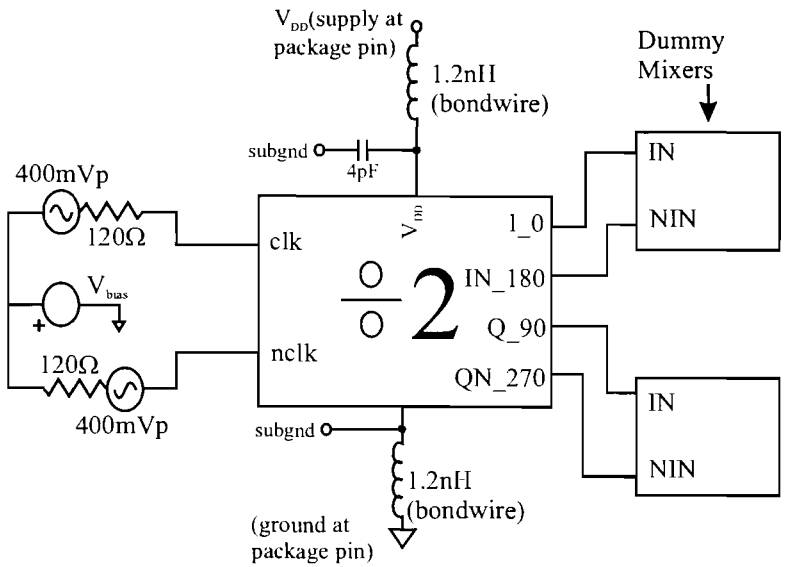

(a) Testbench for the simulation.

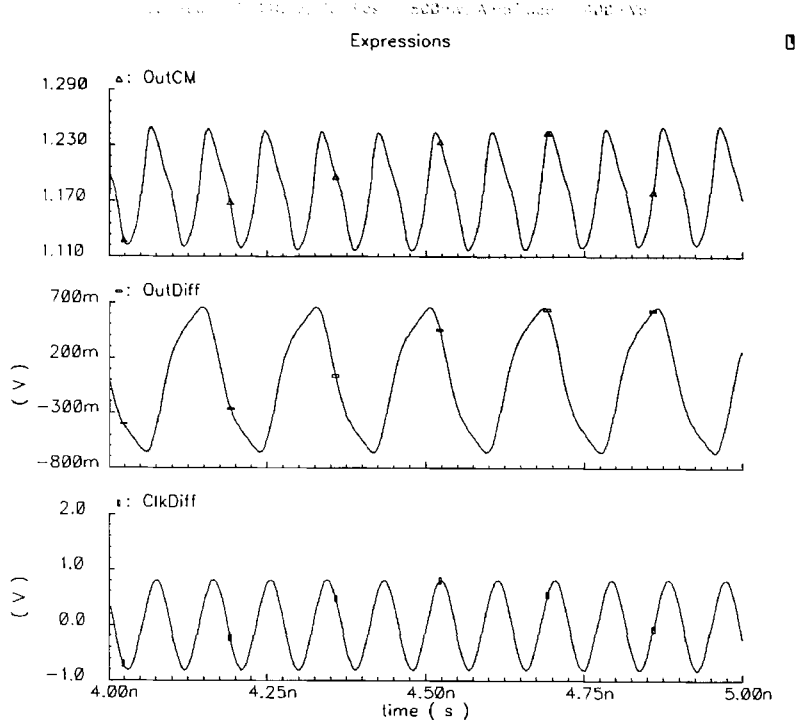

(b) Plots from the testbench simulation (from top to bottom: output common-mode voltage; differential output voltage, and input differential clock.)

FIGURE 6.22: Simulation of the divide core.

For our design, the Philips MOS model 9 with a RFC model extension is used in the time domain simulations. The length of the simulation has prohibited the choice of extension RFD.

The diagrams in Fig. 6.22 show the simulation of our divide-by-2 core. The stimulus for the divider along with loading is shown as if on the chip with realistic test conditions. The input clock is an 800mVpp signal running at 11.14GHz and the output signal shows a 700mVpp signal around a 1.2V common-mode output. This input frequency results in a quadrature LO that is above and outside the lower 5GHz IEEE802.11a band. With hindsight, this is a bad approach, owing to pushing the divider to operate with a higher input frequency, and resulting in a design that is not very sensitive (i.e. low input signal power) in the middle of the band of interest. One can also end up with a higher output swing and possibly hard saturation at the 'low' logic level if the output bias value is very low.

Figure 6.23 has been included for completeness, showing the passive mixer [5] when biased in a DC state. No RF signal is present on the drain terminals of the mixer transistors, and these are instead biased at approximately the centre point between the power supplies. The size of the transistors in the dummy passive mixer of figure 6.23 are oversized by 10μm in gate width, to account for the input stage of the programmable

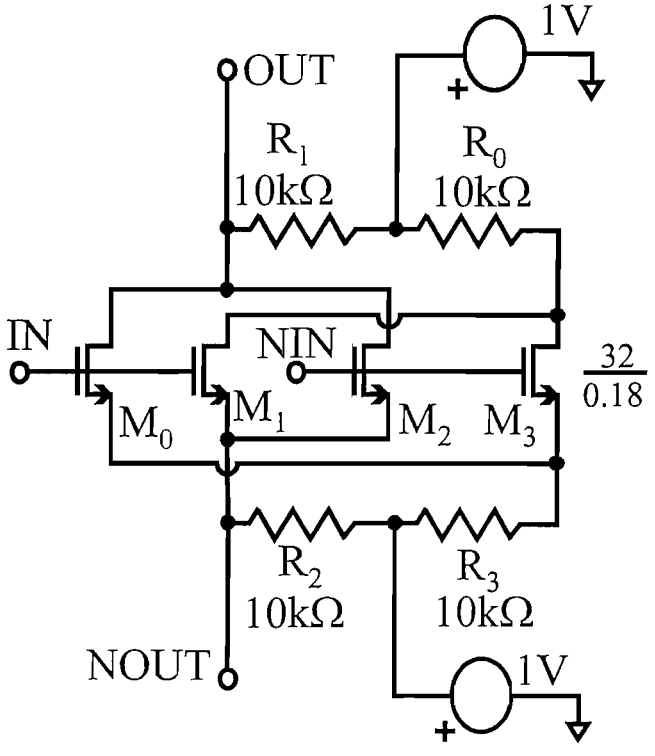

FIGURE 6.23: Details of the mixer circuit.

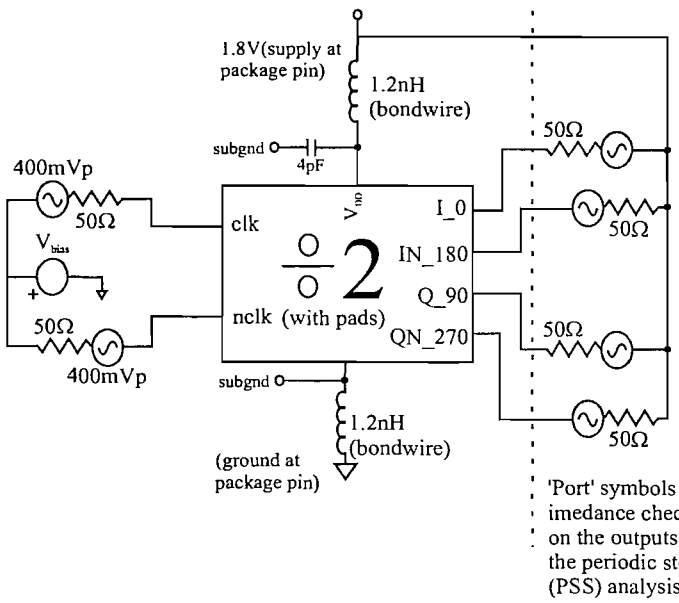

(a) Testbench for the simulation.

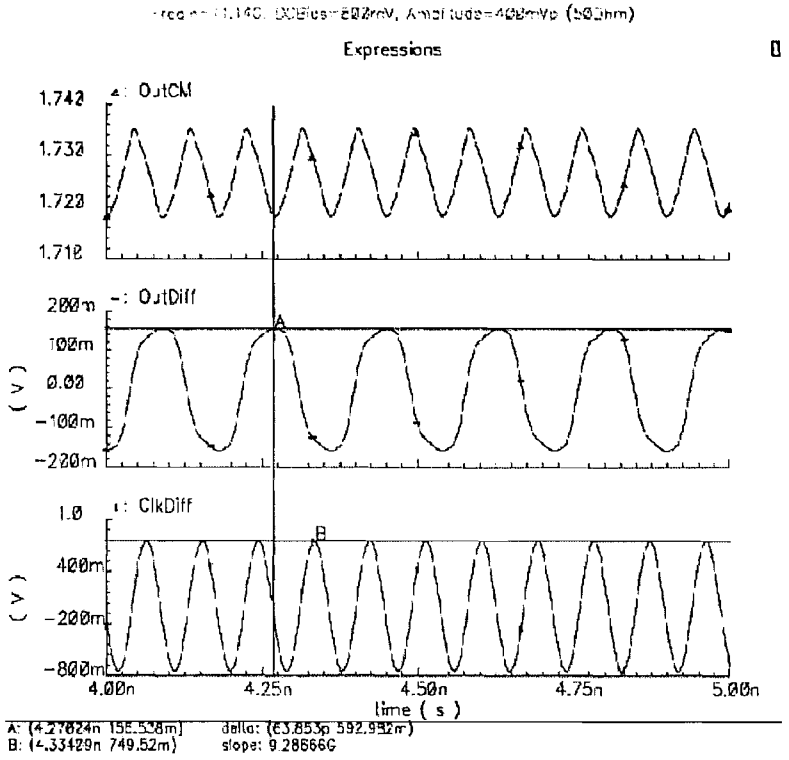

(b) Plots from the testbench simulation (from top to bottom: output common-mode voltage; differential output voltage, and input differential clock.)

FIGURE 6.24: Simulation of the taped out divide-by-2 with pads.

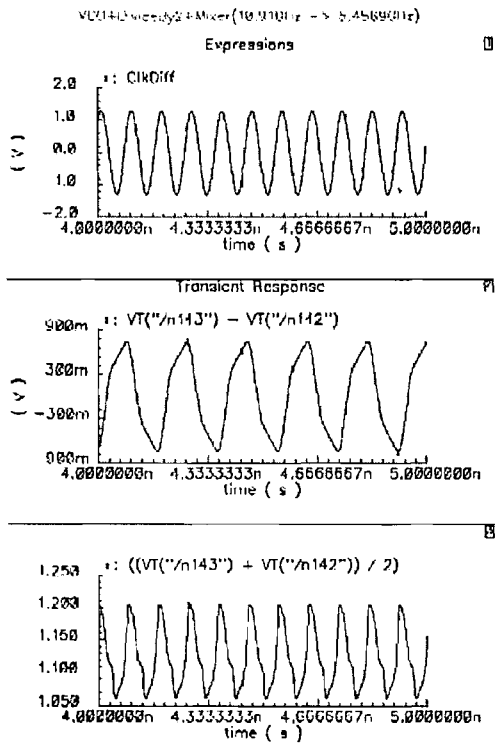

(a) Plot without a buffer between the VCO and divide-by-2.

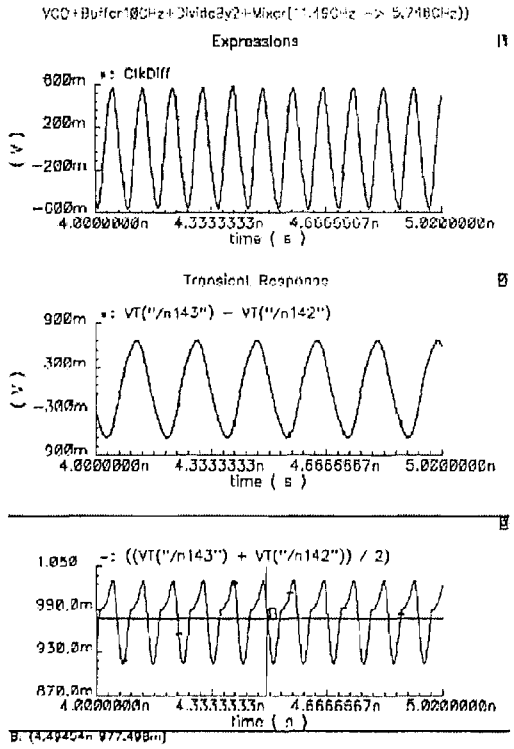

(b) Plot with a buffer between the VCO and divide-by-2.

FIGURE 6.25: Simulation of the VCO-divide-by-2 with dummy mixer loads showing the effects of the high speed buffer Plots from the testbench simulation (from top to bottom: input differential clock; differential output voltage, and output common-mode voltage.).

divider (which at the time still remained to be designed). Hence the need to drive the divider at 11.14GHz, gives that extra headroom for when the programmable divider loads the quadrature generator.

Figure 6.24 shows a simulation of the standalone divide-by-2 design. The testbench shows the divide-by-2 cell wired up with bondwire inductances and 50Ω driving- and output impedances. The simulation output shows a 160mVpp signal output around a 1.725V common-mode voltage. This is clearly outside the intended specifications of the receiver chain, but the motivation behind this tapeout was not to check the output amplitude which would then be directly applied to the mixer inputs. The testbench schematic is a little misleading by not including the implementation of the balun and the bias tees, thus neglecting the attenuation along the signal path. However, being an off-chip signal source, the power can be increased at will.

Lastly, the design of such high performance circuits relies heavily on accurate device models as well as circuit simulations that include parasitic effects such as the impact of the bondwire inductance on the initial startup of the circuit. As on-chip decoupling capacitors are included, a series LC tank is created with the bondwire. Classical network theory shows that the voltage across either the inductor or capacitor at resonance is equal to the product of the RLC tank Q factor and the voltage across the resistance. By lumping any series resistance associated with the bondwire and bypass capacitor, the amplitude of the component equal in frequency to that set by the LC combination (during startup, assuming a step change on the power supply) appearing across this resistance is magnified by the Q and appears across the circuit supply. Simulations show that a large spike exceeding the supply voltage is present on the supply lines during startup, even though it is only for a short time duration. As further work, it is suggested that one looks into whether such effects stress any part of circuit and whether it has implications on the mean time before failure (MTBF).

6.3.2 VCO and divide-by-2 combination

The last set of figures in this section refer to the VCO-Divide-By-2 combination. The plots in Fig. 6.25(a) show the output of the combination when the VCO couples directly to the divider inputs. The output characteristics approach the desired values, when the VCO is tuned to oscillate at its highest frequency, in this case, 5.5GHz (i.e. with the lowest capacitance in its tank.) With a buffer as illustrated in Fig. 6.25(b), the performance is affected in more than one way. First, the highest frequency is beyond 5.7GHz, outside the band of interest, and this is attributed to the reduced loading on the VCO outputs, raising the resonant frequency of the tank. This is expected, as the

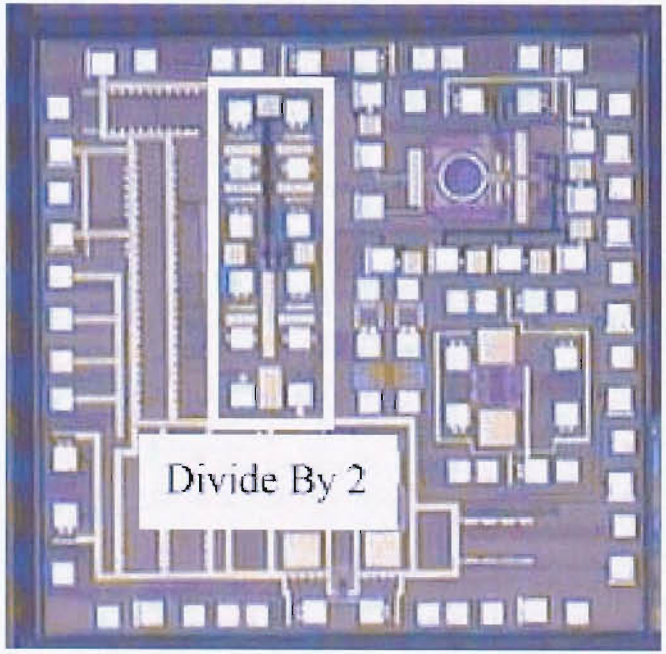

(a) Die with collection of circuits.

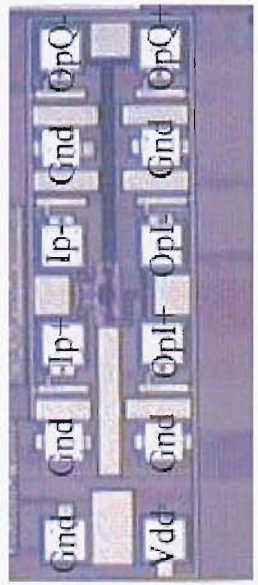

(b) 10GHz stand-alone divide-by-2.

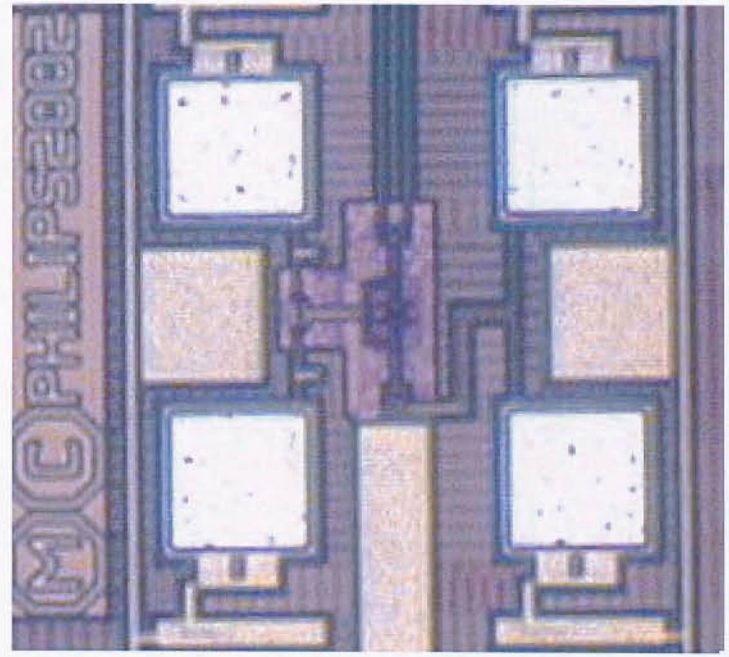

(c) Close-up of the core.

FIGURE 6.26: Photos of the stand-alone divide-by-2 cell.

VCO centre frequency is highly sensitive to the femto-farad load capacitances at this frequency. The buffer also generates a smaller signal swing which switches less current through the divider circuit. With the higher output common-mode voltage of the buffer amplifier, a higher DC current exists for the divider (there is a greater DC voltage drop across the resistive current source after subtracting the gate-source voltage of the clock transistors), yielding a lower output common-mode voltage from the quadrature generator. Owing to the *out-of-specification* top frequency, this combination was not taped out.

6.4 Measurements

This section brings together the actual results obtained from probing the various fabricated versions of the divider circuit described in the first section of this chapter. It is written such that an 'evolution' of results is presented, starting with the basic divider, and ending with a look at a 10GHz clock multiplier for a different application.

6.4.1 Standalone divide-by-2

The first set of measurements concentrates on the 10GHz divider tapeout shown in Fig. 6.26. This chip has a collection of other circuits designed for characterisation by other members of the wireless team. The test arrangement for this circuit is illustrated in Fig. 6.27. The differential signal is generated off-chip with the help of an octave hybrid coupler. Each output on this balun is 3dB lower in signal power compared with the input, as expected, because the differential output is generated from a single input feed. A 50Ω standard is adhered to, from the generator to the probe tips, and this not only minimises unwanted reflections running back and forth between the generator and load, but allows the designer to estimate (to some accuracy) the voltage arriving at the input terminals of the divider. Unfortunately, ensuring a broadband resistive match at high frequencies is a challenge and thus it is contentious what signal voltage is impressed on the input of the core divider circuit.

DC blockers are used to decouple any bias from the signal generator, as well as preventing any DC voltage arriving at the spectrum analyser inputs. The high quality bias tees in the diagram set the bias using a VDD-referred bias voltage. Having a ground referred bias source would result in current sunk into the signal source and thus an incorrect bias voltage. With the differential outputs, only one is connected to the analyser, with the other terminated in a high frequency 50Ω resistive load. The RF probes used are 40GHz ground-signal-signal-ground air coplanar probes, with a probe pitch of 200μm

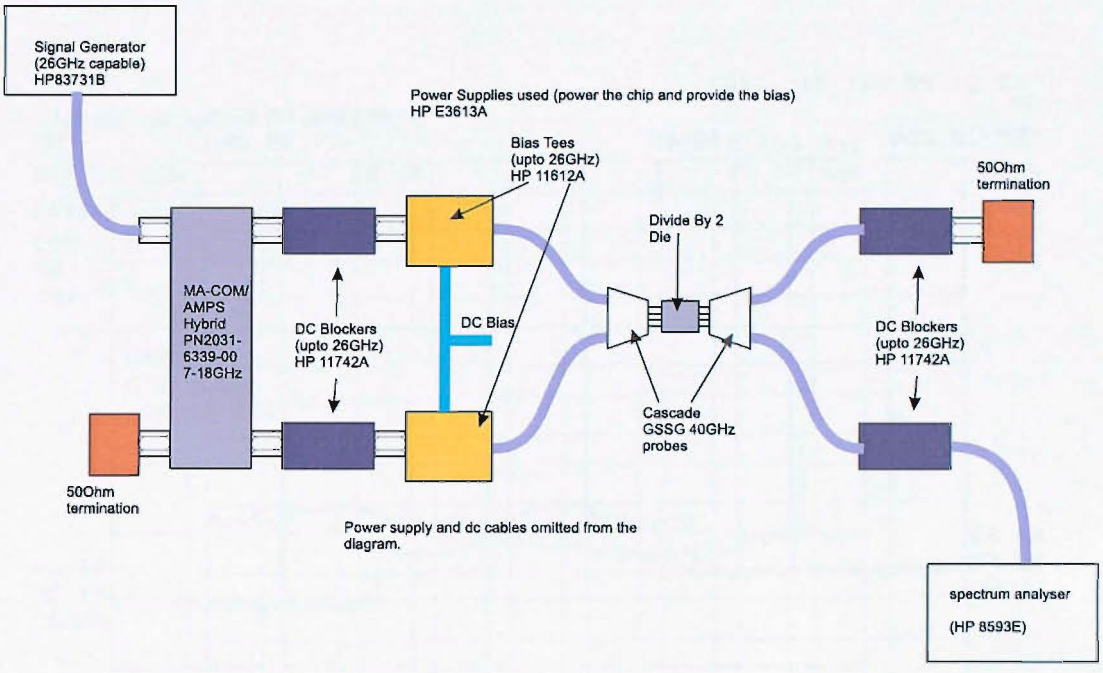

FIGURE 6.27: Arrangement for testing the bare divider IC.

and require a passivation window of 180μm x 180μm within each pad. DC probes are used to supply the low frequency inputs to the die.

The attenuation between the signal generator and the output of the bias tees has been characterised to be approximately 9dB in the frequency range 8.5GHz to 12.5GHz and this is important for the input sensitivity measurements.

The diagrams in Figs. 6.28, 6.29 and 6.30, capture the output spectrum of the running divider, with close-ups of the output spurs. The outputs are shown for 10GHz, 11GHz and 12GHz inputs. The estimated input signal power (assuming power match) can be deduced from the input sensitivity plots of Fig. 6.31. The output signal power distorts the true output of the divider with no simple linear mapping between the input and output at such frequencies, and thus only the accuracy in the frequency of the output fundamental tone is of importance to us. As one reduces the span of the analyser whilst increasing the resolution bandwidth, the accuracy of the divider becomes more apparent as the 'skirts' begin to fall, showing its stable operation and low phase noise property. One noticable and expected result is the reduction in output amplitude as the frequency is increased. This effect can also be exascerbated by any frequency dependent impedance mismatch in the output stage.

The observant reader will notice a distinct step in the noise floor of the plots in Figs. 6.28(a), 6.29(a) and 6.30(a). It is the author's opinion that this was caused by the

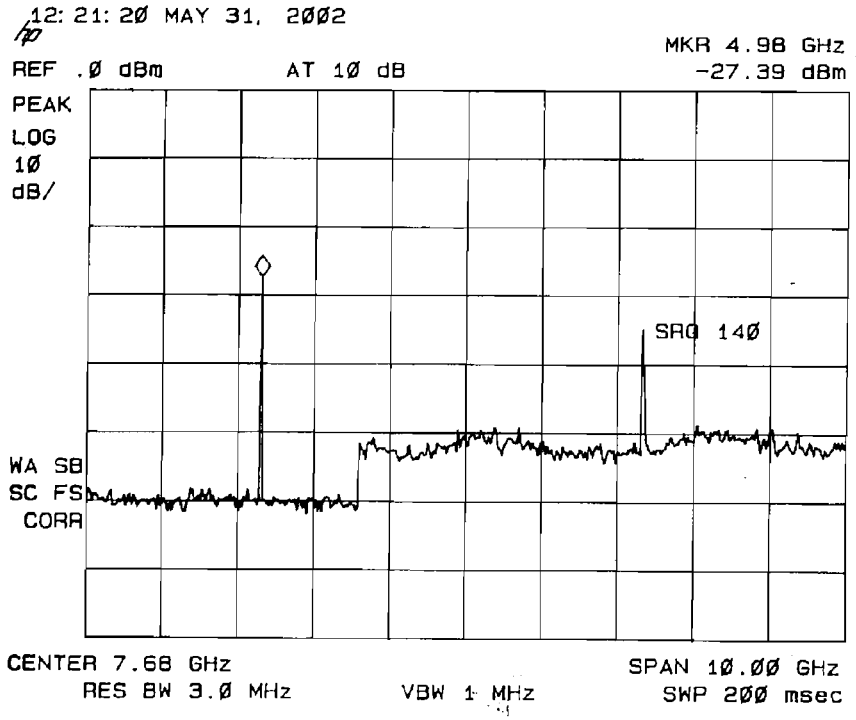

(a) Output with a 10GHz input.

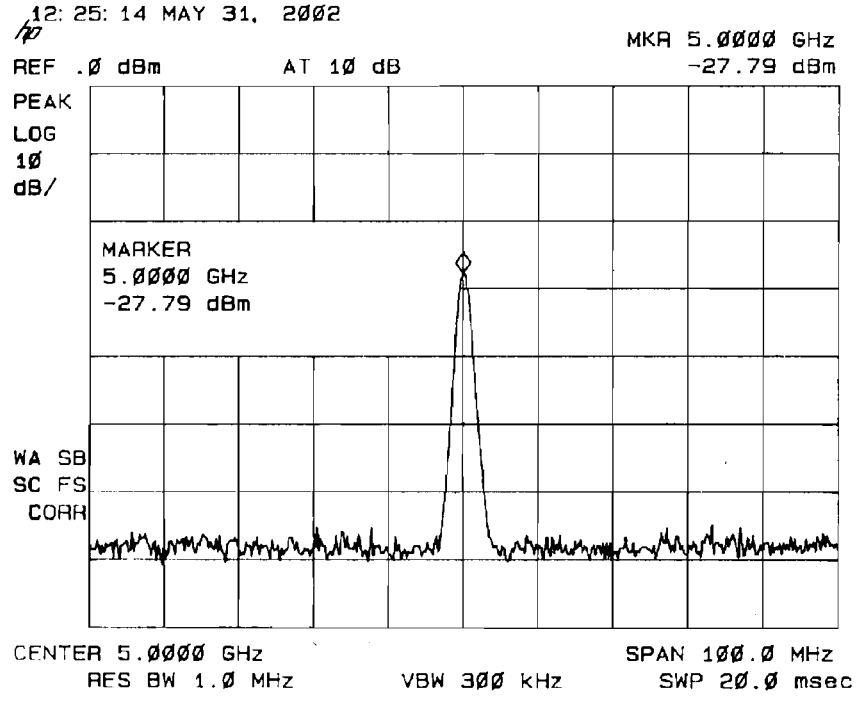

(b) Close-up of the output spike.

FIGURE 6.28: Output spectrum of the stand-alone divide-by-2 (frequency of diamond marker at the peak is given in the sub-box headed 'MARKER'.)

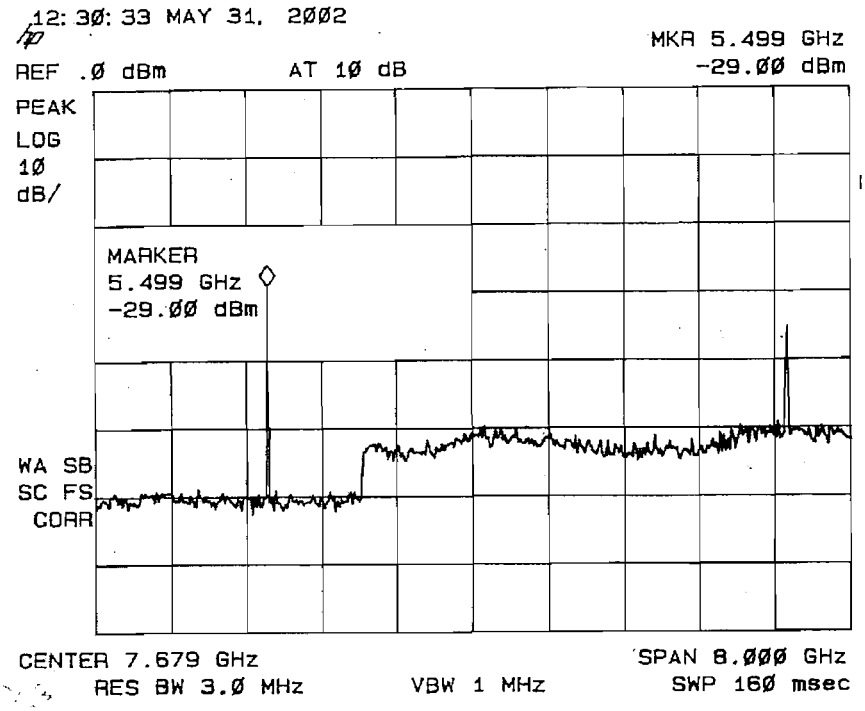

(a) Output with an 11GHz input.

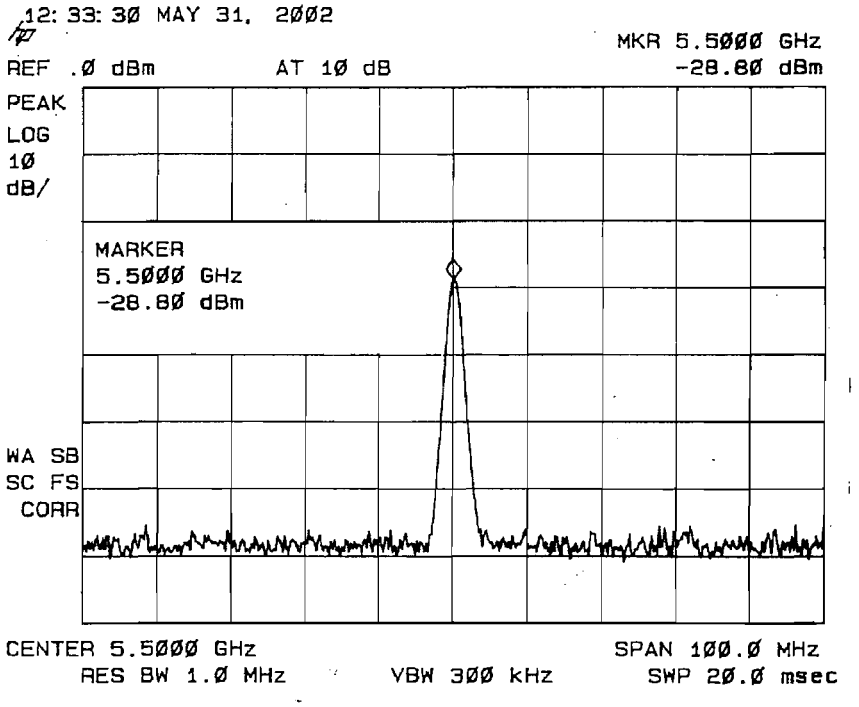

(b) Close-up of the output spike.

FIGURE 6.29: Output spectrum of the stand-alone divide-by-2 (frequency of diamond marker at the peak is given in the sub-box headed 'MARKER'.)

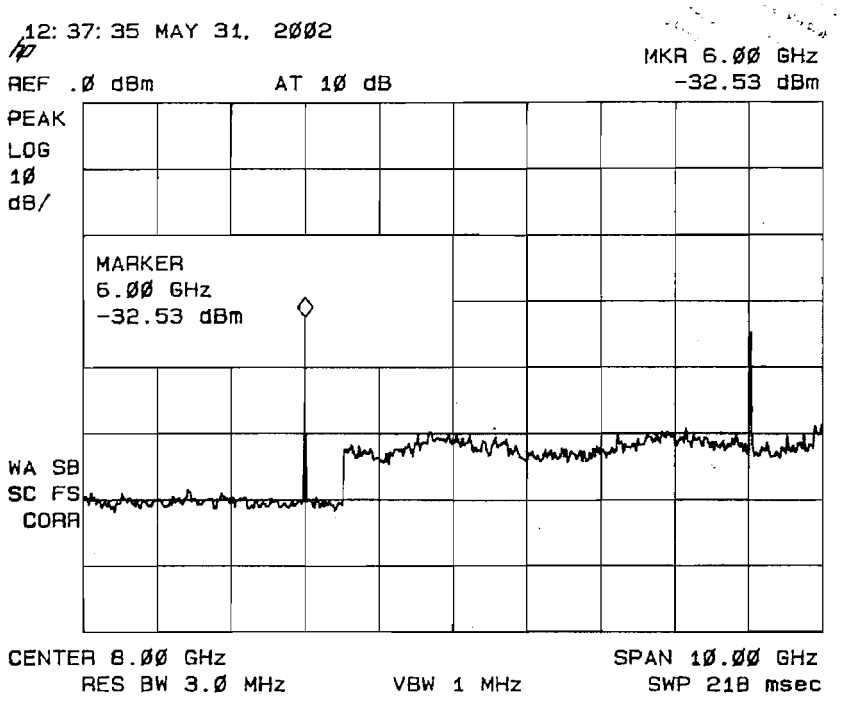

(a) Output with a 12GHz input.

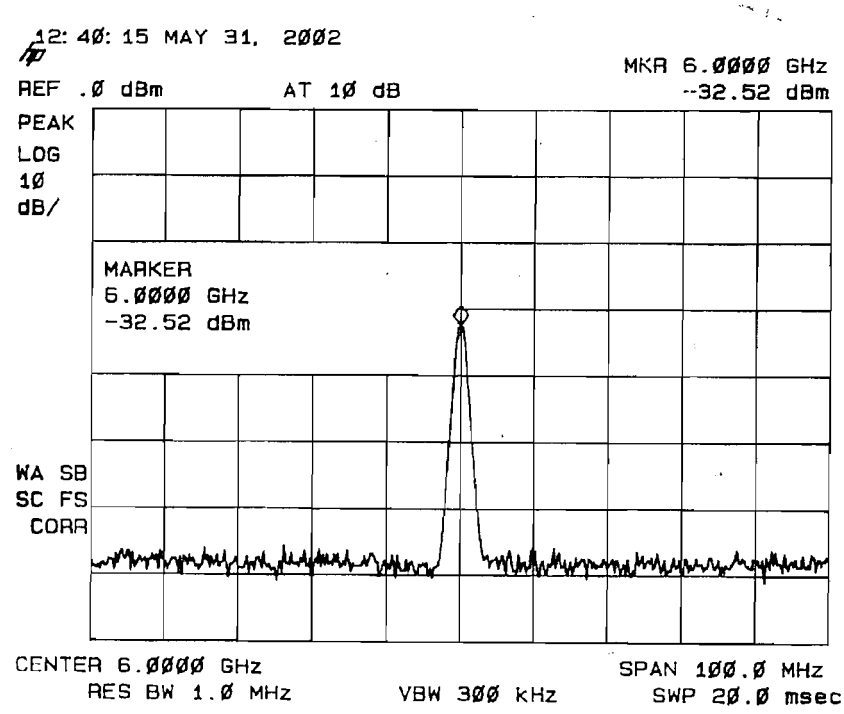

(b) Close-up of the output spike.

FIGURE 6.30: Output spectrum of the stand-alone divide-by-2 (frequency of diamond marker at the peak is given in the sub-box headed 'MARKER'.)

10GHz frequency sweep of the spectrum analyser's input port. A local oscillator within the analyser is connected to one input terminal of a mixer, with the signal arriving on the input of the analyser applied to the second input of the internal mixer. The local oscillator is swept, downconverting individual bands of frequencies (determined by the resolution and video bandwidth settings) where its power can be measured. However, instrumentation designers may decide to switch between local oscillator sources in order to sweep such a large frequency range up to a high frequency. It is possible that a compromise has to be made when selecting a higher frequency source, that unfortunately raises the noise floor of the measurement, unlike in the case where the low frequency portion of the bandwidth is swept.

The sensitivity plots in Fig. 6.31 show the minimum signal required to sustain correct division. A higher signal power is required to kick-start the divider away from its free-running frequency. If running at a frequency other than this value with a high enough signal power, the divider output will track the input within a wide range of frequencies. The 'free-running frequency' is one where the divide-by-2 circuit operates on the smallest input power. With the complex network of parasitics, it is common to see an output frequency tone between 5 and 5.5GHz without any direct application of a stimulus. By having a higher input bias voltage, the sensitivity of the divider increases with the free-running frequency dropping towards the 10GHz mark. What is interesting is that with the input common-mode voltage of 0.8V, the 'notch' in the characteristics seems to be at the high end of the IEEE802.11a lower band (after division), which demonstrates the accuracy of the transistor models. In both, the simulation of the sensitivity curve is superimposed on this image for a divider biased with a 0.9V input common-mode signal. Sadly, owing to a large simulation time for this curve, only one trace has been obtained. It is not yet clear why a spike in the simulation occurs at 2.5GHz, although this has been witnessed in other dividers [1], also at the low frequency end of the plot.

Finally, the current consumption of the standalone divider was measured to be 59mA. Without having separate nodes by which to power the input buffer, divider core, and output buffers, simulations had to be used to estimate the power consumption amongst the individual cells. The divider core is estimated to be using 13mA whilst running beyond 10GHz, which is an admirable figure despite it being the first attempt. The input buffer is estimated to be consuming a static 13mA current, justified by the need to take a 10GHz signal from a 50Ω signal source and drive the divider inputs at the same frequency with adequate signal swing. The 5GHz buffers are also power hungry, consuming an estimated 14mA each. Though large, the reader must appreciate the challenge of driving a 50Ω termination, as well as parasitic capacitances associated with the bond pad.

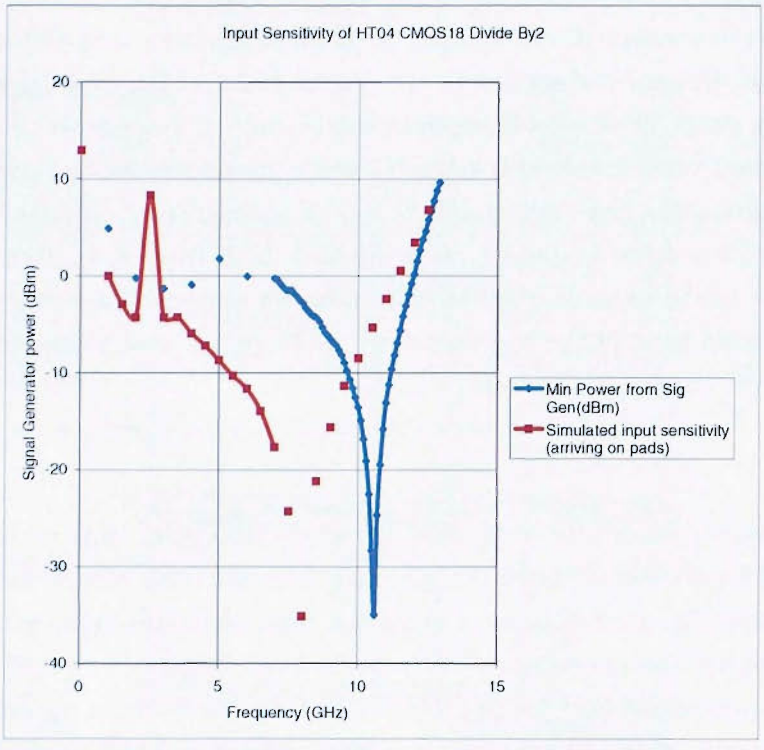

(a) Sensitivity of the Divide-by-2 circuit (input common-mode voltage is 1.0V)

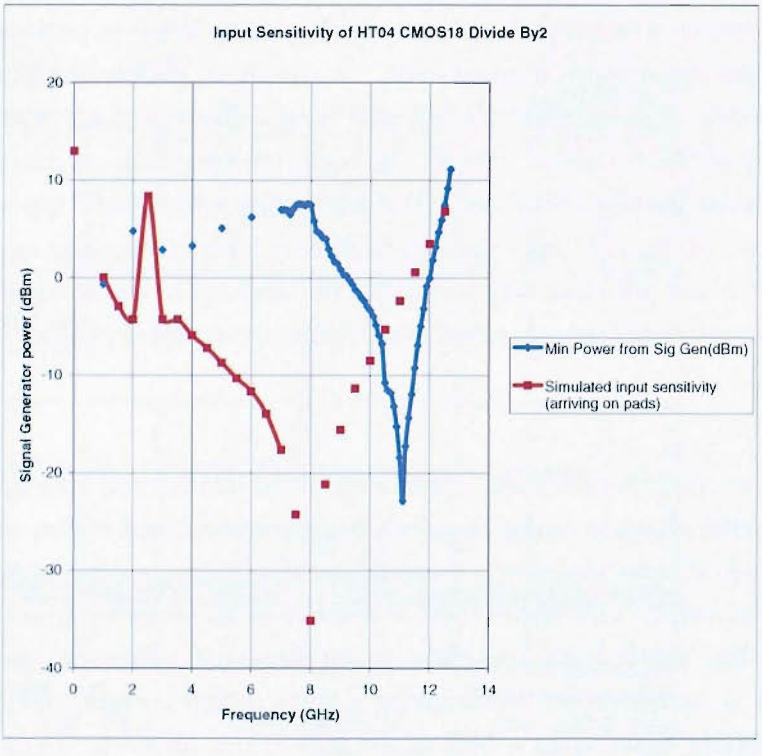

(b) Another sensitivity of the Divide-by-2 circuit (input common-mode voltage is 0.8V)

FIGURE 6.31: Sensitivity of the stand alone divide-by-2 cell.

6.4.2 VCO divide-by-2 combination

The second instance of the divide-by-2 cell can be found in the VCO divide-by-2 combination. Two chip photos are shown in Figs. 6.32(a) and 6.32(b). The phase noise at the output of the VCO, without any divide-by-2 circuit, is given in Fig. 6.33, and the corner frequency between the $\frac{1}{f^3}$ and $\frac{1}{f^2}$ is at approimately 20kHz. The quality factor of the tank dominates the noise after that corner frequency. The deviation seen after 1MHz is attributed to the RF probes and the contact made to the metal pads. The output characteristics of this chip have been captured in the plot shown in Fig. 6.34. This graph shows both the analogue and digital tuning characteristics of the circuit, with a 4-bit binary input (CMOS logic levels) setting which of the curves the analogue tuning voltage is applicable to. From the measurements, it is clear that although the tuning range (i.e. lowest frequency possible [with all capacitances at their maximum] to the highest frequency [where all the digital capacitances are disconnected and the varicap imparts a minimum capacitive load]) is enough to cover the 200MHz band of the IEEE802.11a lower band, the centre point of this characteristic is sadly too low due to insufficient estimtation of VCO load capacitance. Another unfortunate aspect of the circuit is there is no clear overlap between the digital ranges. Its ramifications are that the eventual frequency synthesiser, with combined analogue-digital tuning, would jitter heavily, as it tries to lock to a frequency not covered by one of the digital ranges; for example, 4.8GHz is not programmable. These results show the need for a redesign, where the

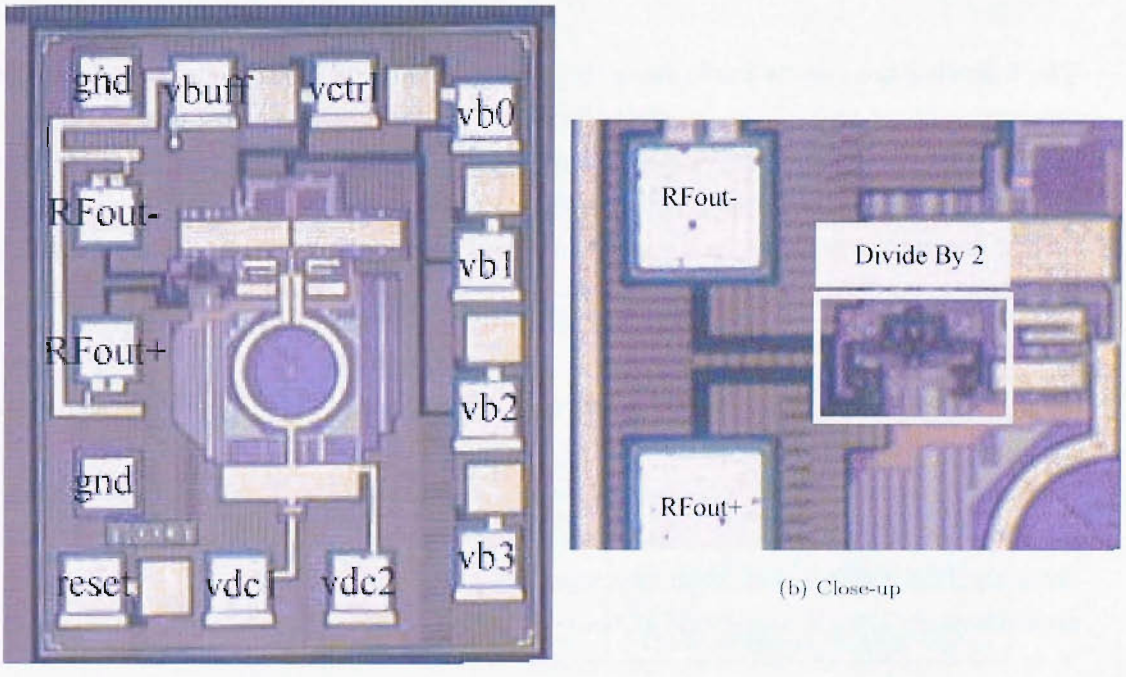

(a) Pad arrangement

(b) Close-up

FIGURE 6.32: Chip photo of the VCO-Divide-by-2 combination.

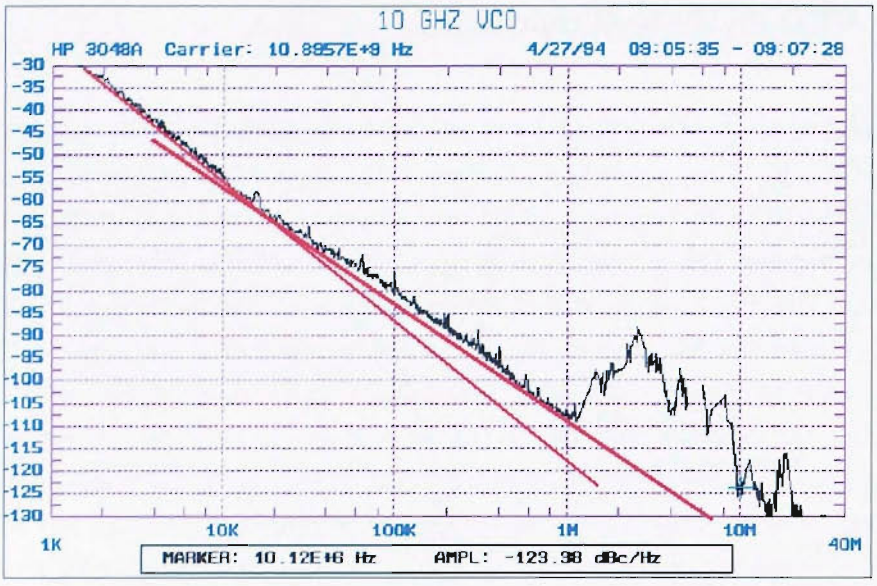

FIGURE 6.33: Measured phase noise of the VCO output (no divide-by-2) using the HP3048A phase noise system together with a high frequency probe station.

excellent phase noise performance of this circuit is traded for a larger analogue tuning range. A better quality varicap is thought to solve this problem (in terms of Q-factor, range over which the capacitance varies monotonically and a good RF model for use in the 'Cadence' circuit simulator.)

6.4.3 10GHz PLL

The following are results taken from the optical network PLL introduced earlier. The measurements of this IC are presented below. This PLL IC was fabricated in the same process technology as the divide-by-2 circuit mentioned earlier. A photo of the die is given in Fig. 6.35. As before, owing to the compactness of the divide-by-2 design, the pair of dividers is dwarfed by the rest of the circuit, especially circuit elements such as the inductor and loop filter capacitor.

The VCO is the same as that used in the combination circuit presented in Section 6.1.3. A 0.6nH planar single-turn coil is used in the tank; this has a Q-factor of 17 at 10GHz. Figure 6.36 shows the clock multiplier unit's input and output traces. The multiplier's input is driven by a Marconi 2042 signal generator. Looking carefully, the output is certainly four times faster than the input to the PLL, generating a waveform with approximately 130mV amplitude at the output.

The plot in Fig. 6.37 was generated using a HP3048A phase noise measuring system set up in a PLL configuration and is used to characterise the jitter. Between the integration

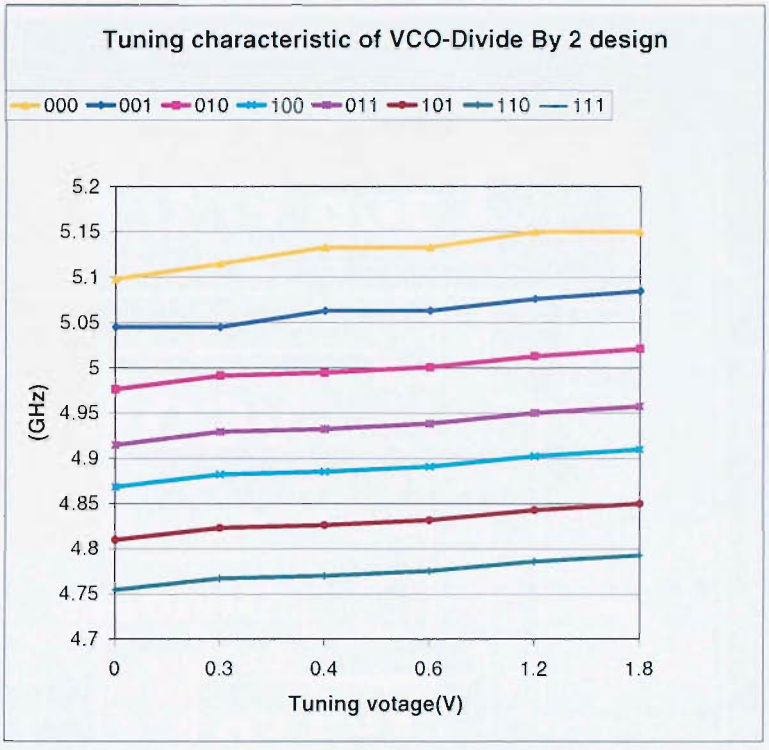

FIGURE 6.34: Measured tuning curves for the VCO-divide-by-2 circuit, with the horizontal axis showing the analogue voltage on the input to the VCO and the ordinate axis showing the frequency at the output of the divide-by-2. The different curves are for different digital VCO inputs.

limits of 50kHz and 80MHz, the rms-jitter is found to be 0.22ps, which is roughly one fifth of the OC-192 specification.

This IC runs from a 1.8V power supply, drawing 55mA. The pertinent characteristics have been summarised in Table 6.1. It should be noted that any reference to 'SONET filter' implies integration over the bandwidth 50kHz-80MHz from the 10GHz carrier, and it is the energy within this bandwidth that is used to calculate the jitter.

Output frequency	9.953GHz
Reference frequency	2.488GHz
Technology	0.18μm bulk CMOS
Transmitted clock jitter (RMS)	0.22ps with SONET filter
Transmitted clock jitter (peak-to-peak)	2.2ps with SONET filter
Supply voltage	1.8V
Chip size	0.83x0.86mm^2 (active area)
Power consumption	99mW 81mW without output buffer

TABLE 6.1: Performance of the PLL.

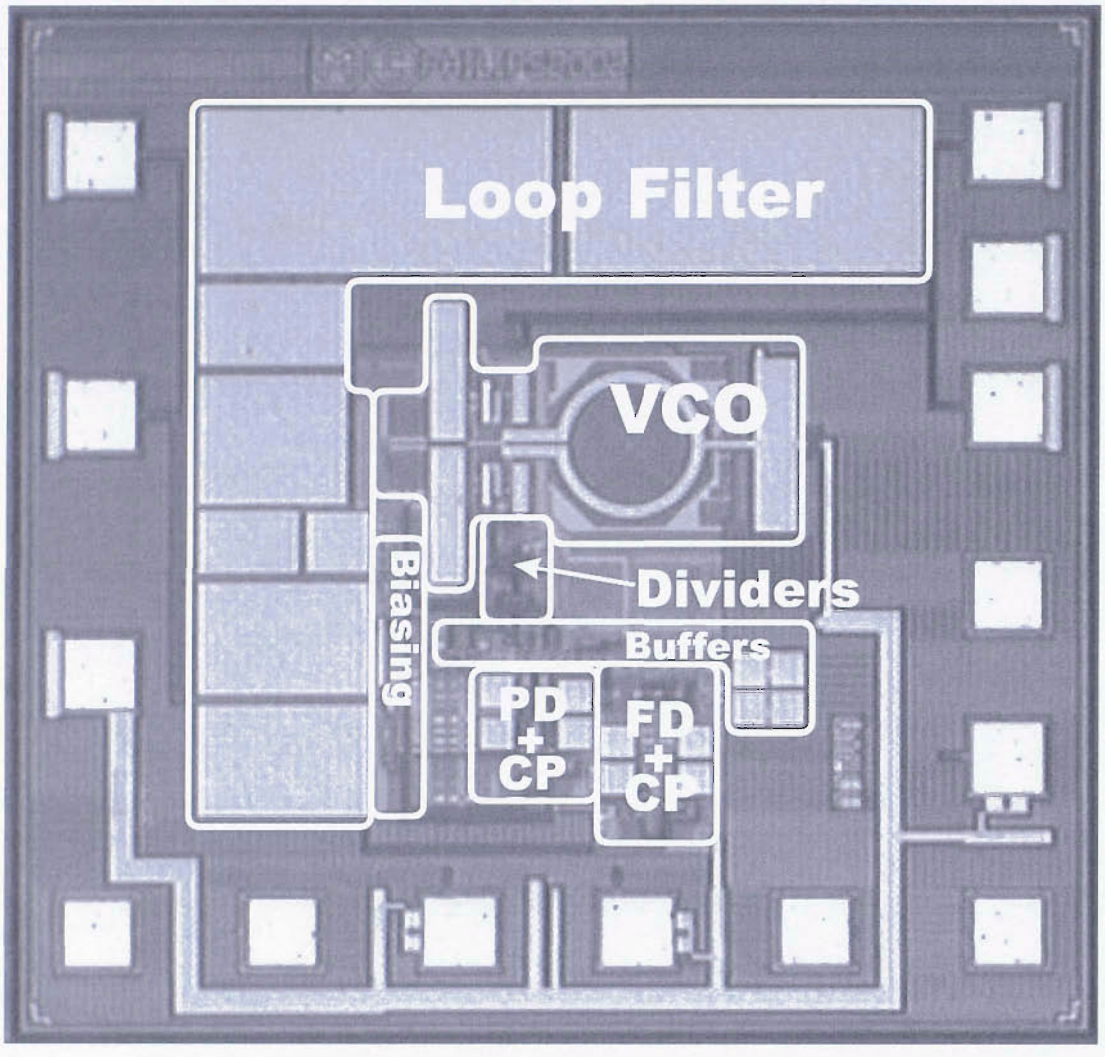

FIGURE 6.35: Photo of 10GHz PLL die.

6.5 Discussion

The preceding measurements clearly show the divider operating up to 12.5GHz, beyond the required frequency which has been the aim of the circuit, and represented state-of-the-art at the time of measurement. The probed measurements have been successful to the point where no special startup arrangement has been required for this active quadrature generator. Between 5GHz and 5.5GHz (output frequency), the output power does drop by 1dBm, corresponding to a peak-to-peak drop of 2mV (at those power levels). Although not shown explicitly, the divider remains stable when excited with an input, be it from an on-chip VCO or external signal generator.

On the topic of power consumption, there is a need to optimise this value for lower power drain. At present, it is estimated that 13mA flows through the actual divider cell with the remainder of the measured current passing through the buffers. With a 1.8V

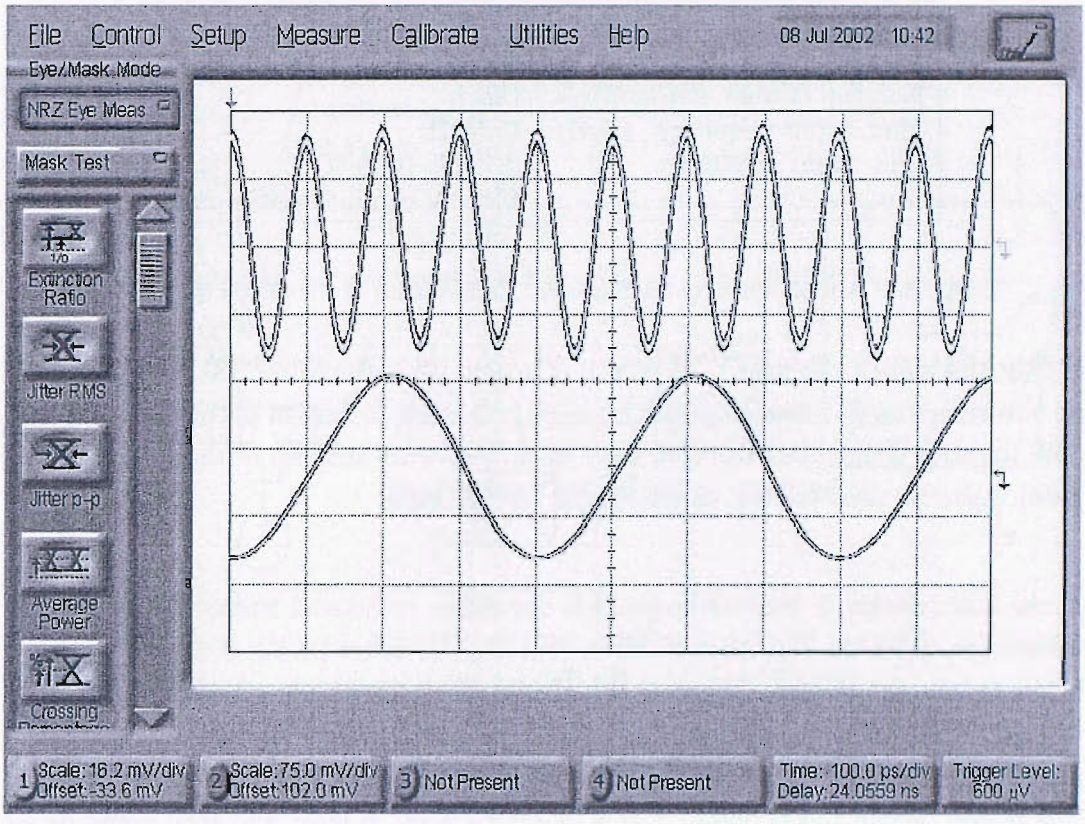

FIGURE 6.36: Oscilloscope capture of the output (top) and input(bottom) signals.

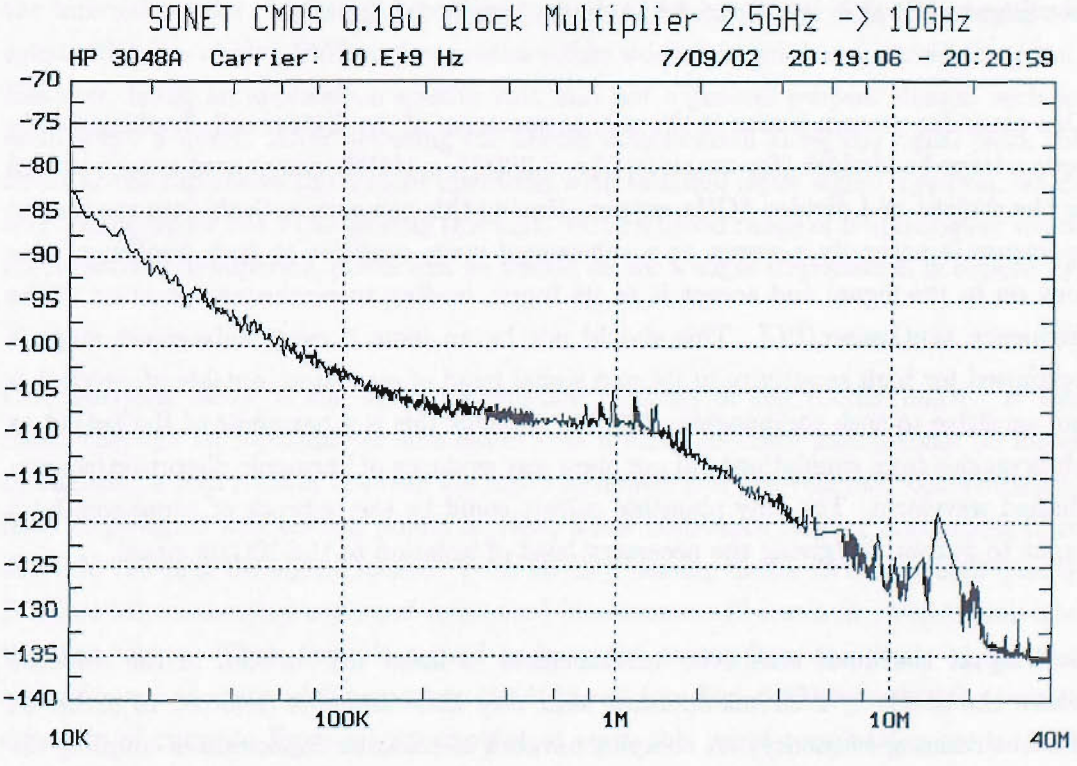

FIGURE 6.37: Phase Noise response of the 10GHz PLL in lock.

Performance	Value
Nominal power supply	1.8V
Nominal current consumption	13mA
Max. input frequency	12.5GHz
Min. input sensitivity	-35dBm @ 10.5GHz with 1V common-mode input

TABLE 6.2: A table showing the measured performance of the divide-by-2 circuit.

supply, this constitutes more than 20mW of power. No optimisation has been performed on this structure (ie speed vs power trade-off); in order to find an optimum point, with only iterative simulations together with basic first-order analysis of the *load resistor - drain capacitor* combination, as the formal design stages.

In the basic design, a noteworthy point is the choice of channel lengths for these clock transistors. With the resistance of the current source resistor (in the divide-by-2 circuit) being so low, the channel length of the divider clock transistors should have been set to a value equivalent to the $energy - restorers$ (cross-coupled transistors acting as a negative transconductance driver) in the LC-VCO, also sitting on the ground terminal. The reason behind this is purely from a matching point of view, and thus would be set to 0.22μm. An immediate repercussion is approximate 25% increased gate capacitance which will no doubt cause a severe problem from the VCO signal amplitude. This was not raised until after the design had been submitted for tape-out.

One cause for concern is seen in the output spectrum of the divider cell. Looking closely over a large bandwidth (for example, Fig. 6.28(a)), a 10GHz component is seen as well as the desired and divided 5GHz output. Having this component 'leak' into the output spectrum is naturally a worry, as a subsequent stage sensitive to such frequency, may lock on to the signal and accept it as its input, leading to erroneous operation of the frequency synthesiser/PLL. This should not be an issue if every subsequent stage is optimised for high sensitivity in its own signal band of operation, outside of which it is not sensitive to such components. It is not thought this is a harmonic of the output as observations from simulations did not show any evidence of harmonic distortion/square-shaped waveform. The only plausible culprit could be the network of capacitors from input to output not giving the necessary level of isolation to the 10GHz signal.

Looking at the input sensitivity measurement, a large 'dip' is seen in the response where the divide-by-2 circuit operates with very little stimulus (referred to earlier as the free-running frequency). A complex network of parasitic capacitances coupling the output nodes to the input nodes with a positive feedback characteristic would explain the oscillatory behaviour mentioned in the 'Measurements' section. It has been assumed

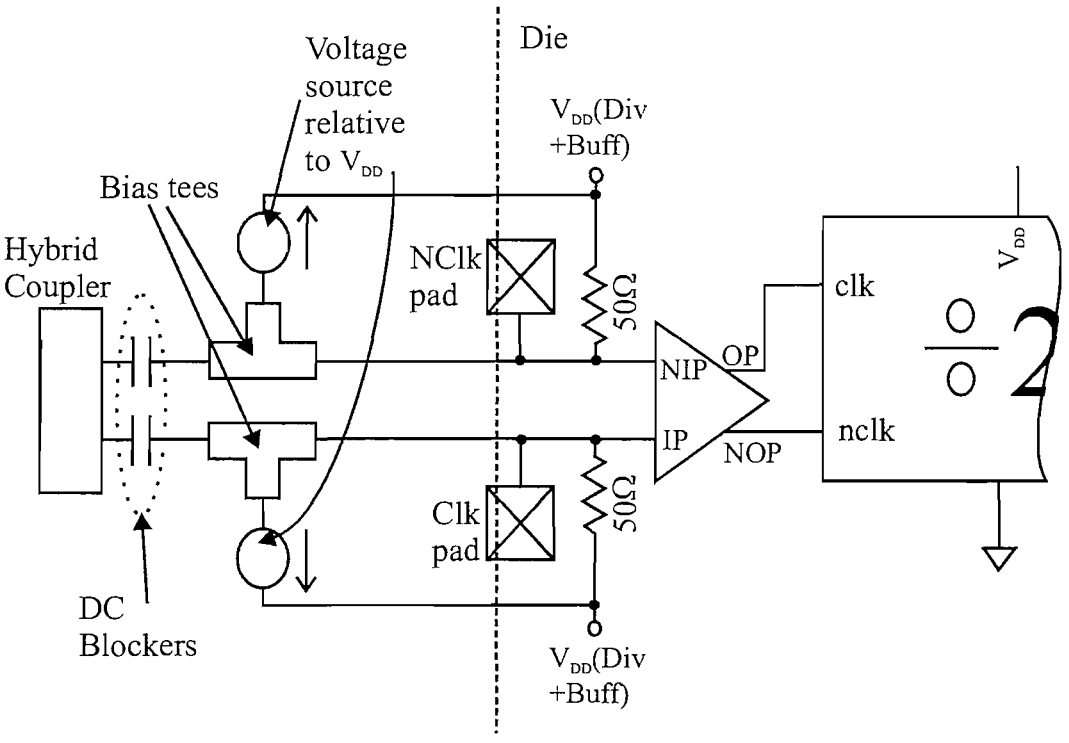

FIGURE 6.38: Test arrangement showing the bias tee and the issue of referring the input bias voltage to V_{DD}.

that adequate matching allows one to determine the voltage appearing on the inputs of the internal divider. Sadly, no broadband match exists on chip and the S_{11} magnitude eventually rises above -10dB at frequencies either side of the minimum sensitivity point. However, being an application specific cell, and not a general purpose design, such an issue is not a worry. After deducing the known attentuation along the signal path, the result at the dip shows the divider operating with minimal input signal injection, which is encouraging for the VCO driving this unit. With a broad range of frequencies in which the sensitivity is superior, power can be traded off for a slight degradation in sensitivity.

One worrying factor is the input termination resistors of the 10GHz buffer. It was found through an S-parameter simulation that a more accurate match could be found by shunting a 50Ω resistor between the input terminals and the Vdd rail. Unfortunately, from large-signal and testing points of view, a low resistance current path exists from the Vdd rail into the signal source. With off-chip biasing in the form of a high quality DC bias tee, connecting a ground referenced bias source will result in current sunk into the bias source. Thus, there is a need to hang the source from the upper supply rail and reference the bias with respect to the Vdd rail (see Fig. 6.38), ensuring the correct direction of current. From a testing point of view, this input termination needs to be carefully implemented and ground-referenced.

Looking at the phase noise measurement of the VCO, the results meet the specifications laid out for the IEEE802.11a project. At 1MHz offset from the carrier, the noise power is 105dB below the carrier. After the divider, it is 111dB below the carrier, though it was not possible to capture the phase noise plot at the time of the measurement. MOS transistors have poor flicker noise performance compared with bipolar transistors, and hence the cross-coupled transistors in the VCO would be responsible for the degradation close-in to the carrier, masking the quality of the tank circuit. However, not being a standalone product, it is debatable what the measurements really mean. From a system point of view, it is important to know what the value the phase noise is on the local oscillator output. Unfortunately, adding a buffer amplifier on the end of the VCO or divider masks the useful value of the phase noise. Obviously, one can simulate the measured phase noise before probing the output of the device under test before any buffer, in order to obtain some idea.

Measuring the output amplitude of the divider is an issue as to what it really refers to.

Performance	Value
Nominal power supply	1.8V
Nominal current consumption	29mA
Analogue tuning range	50MHz
Digital tuning	7 bands (spanning 4.75 - 5.15GHz)
Phase noise before divider	-105dBc @ 1MHz offset
Phase noise after divider	-111dBc @ 1MHz offset

TABLE 6.3: A table showing the measured performance of the VCO-divide-by-2 circuit.

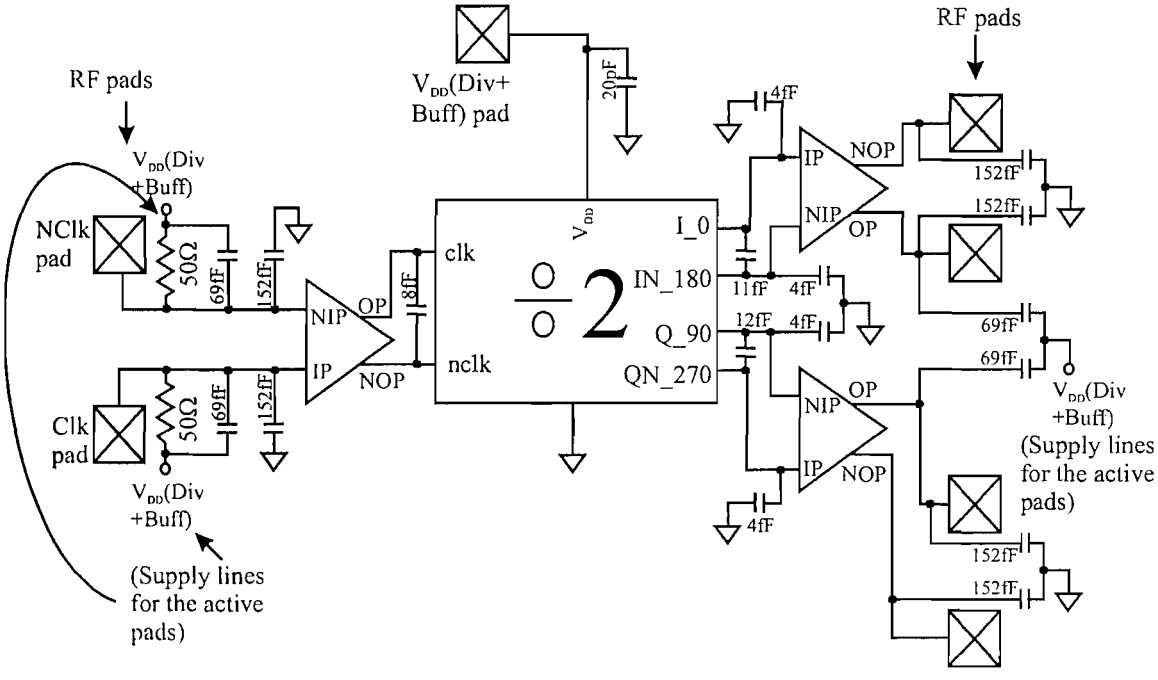

FIGURE 6.39: Circuit with buffers and pads.

If one recalls, the output of the divide-by-2 stage has an output buffer (Fig. 6.39) which is capable of driving a low impedance and not significantly affecting the operation of the core divider cell. Unfortunately, owing to time-constraints, this output driver was never designed for linear operation and hence extracting useful information about the output drive is not a trivial task without characterising the output buffer.

The same is true of the input buffer, where not only is there a problem of power mismatch owing to poor broadband load, but also the fact that this amplifier masks the voltage impressed on the input connections of the core divider. Hence, the results from the input sensitivity measurements for this divider in this project should be interpretted bearing the above in mind.

As the sole purpose of this divider stage was to act as a quadrature generator, a key performance metric is the quadrature relationship between the outputs. Any deviation in such a coherent relationship is translated into skewing of the constellation pattern associated with the baseband modulation scheme. More importantly, the image-rejection ratio can deteriorate. This, in turn, has repercussions on the retrieved symbols and can be seen as a degradation of the bit error rate.

Unfortunately, at the time of layout, it was decided (by Philips engineers Dr. Domine Leenaerts and Nenad Pavlovic) not to characterise such an image rejection metric, as the necessary equipment was unavailable. The motivation behind the standalone divider was simply to check its functionality and then its speed capability, being such a high speed design in a predominantly digital process technology. It is hoped that when the receive chain, together with the divider acting as a LO, can be taped out (twice - for a phase noise measurement) and characterised, such a measurement can be performed.

The comments above may harshly criticise the design discussed in this chapter, and this would be unfair without equally highlighting its merits. A high performance divide-by-2 circuit has been designed in a digital sub-micron CMOS process, with success, both with functionality and target frequency band of operation, in the first cut of silicon. The circuit performs the role of a divide-by-2 function over a broad frequency range and with a reasonable power consumption if the input and output buffers are ignored. Lastly, the circuit has been shown to function as desired in an application, both coupled directly to a 10GHz VCO as well as integrated within a PLL for a state-of-the-art optical networking application.

Bibliography

[1] C.Vaucher and M.Apostolidou. A low power 20GHz static frequency divider with programmable input sensitivity. In *IEEE Radio Frequency Integrated Circuits Symposium*, pages 235–238, Seattle, USA, June 2002.

[2] L.Tiemeijer et al. Record q spiral inductors in standard cmos. In *IEEE International Electron Devices Meeting*, pages 40.7.1 – 40.7.3, December 2001.

[3] D.B.M. Klaassen R.M.D.A. Velghe and F.M.Klaassen. MOS model 9 - Unclassified report. NL-UR 003/94. Technical report, Philips Research Laboratories, Eindhoven, The Netherlands, 1995.

[4] R. van Beek et al. A 2.5-to-10GHz SONET compliant clock multiplier unit in a 0.18μm CMOS technology. In *IEEE International Solid-State Circuits Conference*, page 10.3, February 2003.

[5] W.Redman-White and D.Leenaerts. $\frac{1}{f}$ noise in passive CMOS mixers for low and zero IF integrated receivers. In *27th European Solid-State Circuits Conference*, pages 68–71, Villach, Austria, September 2001.

Chapter 7

Dual modulus 16/17 divider in 0.18μm bulk CMOS

This chapter presents the design of a dual-modulus divider in bulk CMOS as part of the IEEE802.11a project. This chapter is neatly placed in this thesis as it is an important cell of the next block in the 5GHz frequency synthesiser. Based on the phase selector and state machine discussed in chapter 4, the following sections describe the design of the same divider architecture, but in bulk CMOS, hence without the series current stacking seen throughout chapter 5. Although there are no explicit measurements for the dual-modulus divider, the circuit was submitted for fabrication as part of a multi-modulus divider design to be discussed in the next chapter. This multi-modulus design has been verified to function correctly.

7.1 Circuit design

This block is intended to follow the output of the divide-by-2 circuit introduced in chapter 6. Although the choice of the division ratio is important in such an application specific project, its derivation will be tackled in the following chapter, which is concerned with the full programmable divider. It is therefore useful for the reader to concentrate on the design of the dual-modulus divider upon which the rest of the programmable divider is built.

The operation of the divider is exactly the same as in the SOI case, but with only 2 divide stages after the phase selector circuit. There still remain two asynchronous divide-by-2 stages before the phase selector, and the last divider stage in the 16/17 chain forms part of the 8 state FSM. Here, there is a need to have the complete divider running with an input frequency of at least 5.35GHz, and it should be designed and simulated

to run even faster to overcome any tolerances and modelling limitations. There is also a need for minimal power drain, as well as the restriction of having the gates of the input transistors no larger than 10μm in width. This restriction on gate width is set by the driving capability of the 12GHz divide-by-2 block described in the previous chapter, which is also loaded with the mixers

The striking difference now is the 'unstacking' of the divide-by-2 stages. With the common bulk connections between every NMOS transistor, the gate oxide of transistors higher up in the stack are stressed more than those closer to the lowest circuit potential. Hence, the circuit must now be kept within the 1.8V supply limits dictated by the process technology. With the bulk transconductances now active, the threshold voltages of transistors in the upper stages of the SCL dividers are higher. In circuits where the coupled source terminals of the differential pair are not at AC ground, the front gate drive must increase by some means, whether it be through a higher aspect ratio device or a larger swing on the input. The repercussions with the bulk implementation are that non-differential source followers must now be added to act as inter-stage couplers between pairs of dividers and to translate the bias between the output of one and the input of the next. These common-drain amplifier stages suffer from loading by the bulk transconductance, and are also a source of power drain and signal attenuation. These flaws only serve to emphasise the beauty of the SOI divider's circuit topology and the ease with which the designer can proceed with their high speed design.

Without unstacking, all blocks that existed within the divide stack in the SOI imple-

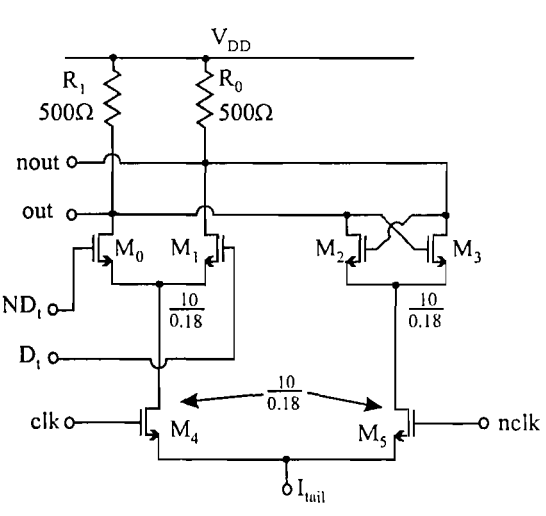

(a) Schematic of the first D-Latch circuit.

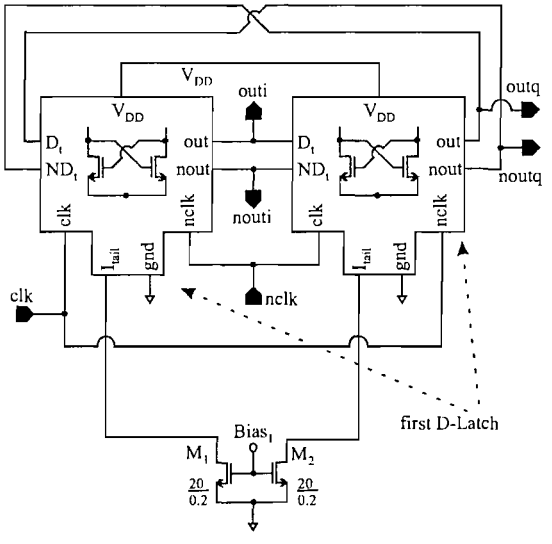

(b) Schematic of the first divide-by-2.

FIGURE 7.1: Schematics of the first stage.

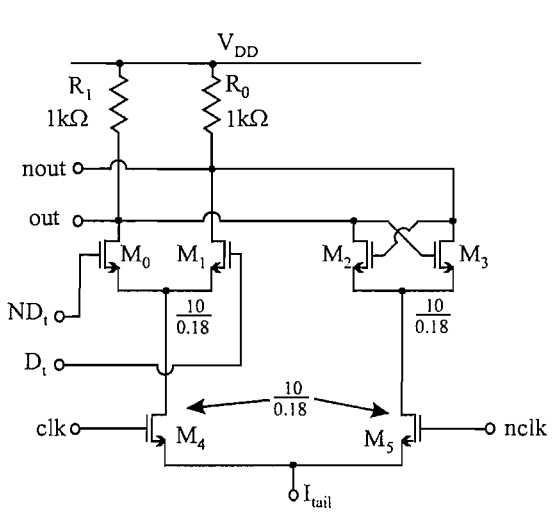

(a) Schematic of the second D-Latch circuit.

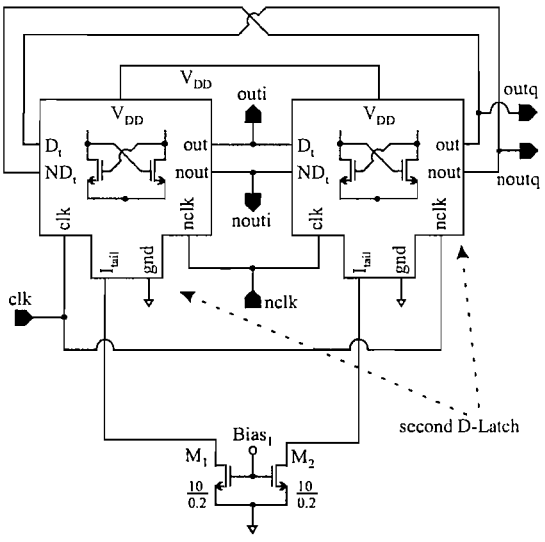

(b) Schematic of the second divide-by-2.

FIGURE 7.2: Schematics of the second stage.

mentation now have to sit on their own current sources, with the added penalty of a pair of clock transistors for each of its asynchronous dividers. At the point of coupling, the same three options exist as pointed out in section 3.1 as well as subsection 5.1.2: resistive divider, capacitive coupling and source followers. The source followers are naturally a drain on power, but the on-chip coupling capacitors sadly present a large parasitic capacitance to ground, as a large value is needed in order to keep the low frequency pole well away from the signal of interest. The resistive divider suffers from a proportional attenuation of the AC signal. One advantage with unstacking is the removal of the cumbersome voltage bias network that is seen in the stacked divider case.

7.1.1 Input dividers

Figures 7.1(a) and 7.1(b) are schematic diagrams of the first divider, showing the component parameters. The choice of clock transistor width is the maximum value specified by the quadrature generator. During the simulations of the quadrature generator, dummy loads consisting of simple transistor stages were used to emulate a subsequent capacitive load being a crude first order approximation of that found on the input of the passive mixer and the programmable divider. This divider is designed to run on 2.5mA, split between the master and slave latches, and together with the load resistor of 500Ω results in an output common-mode voltage sitting 300mV below the supply. This choice of current is found iteratively, as the input transistor size to this divider is constrained as mentioned before. The current may be reduced as simulations eventually showed the dual-modulus divider to operate beyond 5GHz, but has remained a conservative figure

in this first cut of silicon so as to increase the chances of a functional circuit after fabrication). The output of the divider is set to be approximately 600mV_{pp} centred around a 1.5V common mode voltage. This is the highest speed stage in this dual-modulus divider, and so the currents in the subsequent cells can drop accordingly, allowing the power drain to taper off, before the current consumption levels off per stage (there has to be a minimum static current in a SCL structure set by the fact that the gate width cannot be divided down to infinitesimal dimensions).

The pairs of schematics in Figs. 7.2, 7.3, 7.4, 7.5 and 7.6 show in detail the remaining divider stages in this chapter's design. After the third divider (inclusive), the currents through each D-Latch remain the same at just over 300μA. The output swing in each subsequent stage is kept the same, with the reluctance to go for lower current fueled by the need for first time right silicon.

At first, all the currents through all the dividers were kept the same, inevitably making it quite an inefficient but operational programmable divider. After functional simulation, the currents were reduced in a binary fashion, keeping the output voltage swing the same throughout, until the polysilicon resistors became too area consuming (high values of resistance imply a long length of resistive material and with it, attracts parasitic capacitance to the substrate). As would be expected in such an asynchronous chain, the first divider is the most power hungry amongst all the divide-by-2 blocks.

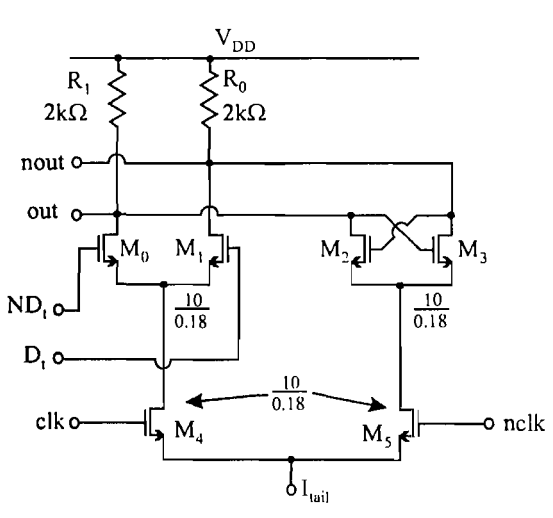

(a) Schematic of the third D-Latch circuit.

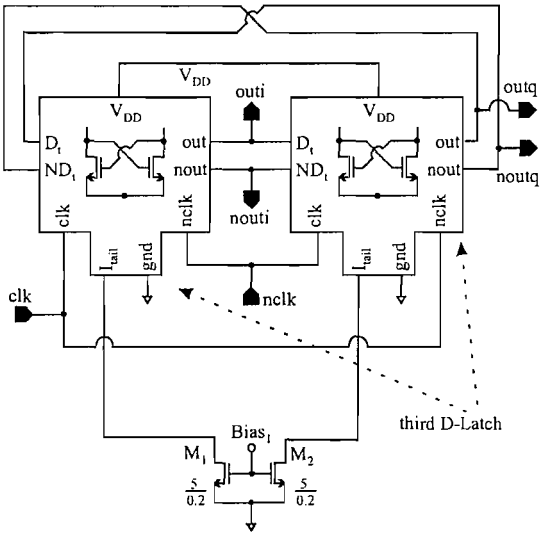

(b) Schematic of the third divide-by-2.

FIGURE 7.3: Schematics of the third stage.

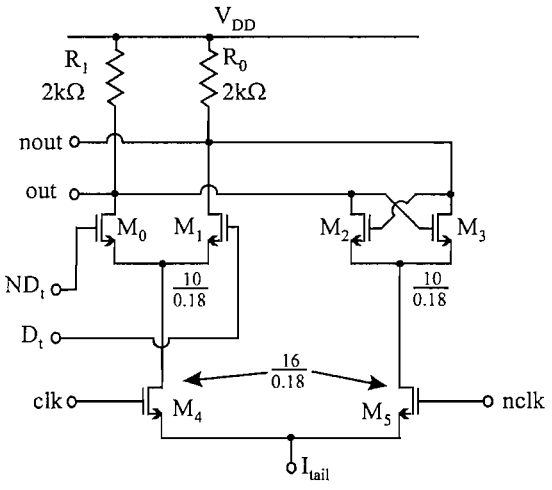

(a) Schematic of the fourth D-Latch circuit.

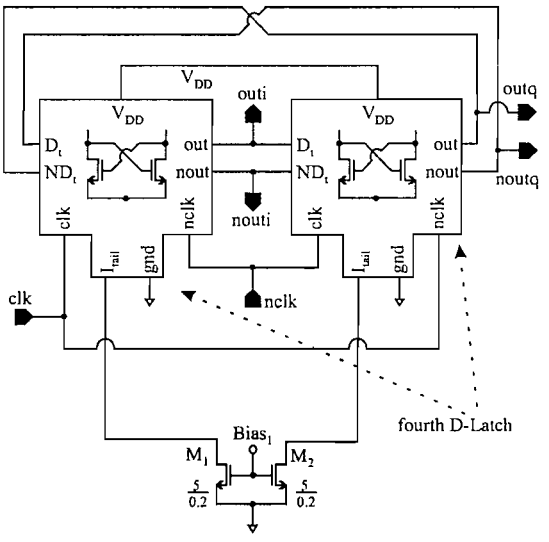

(b) Schematic of the fourth divide-by-2.

FIGURE 7.4: Schematics of the fourth stage.

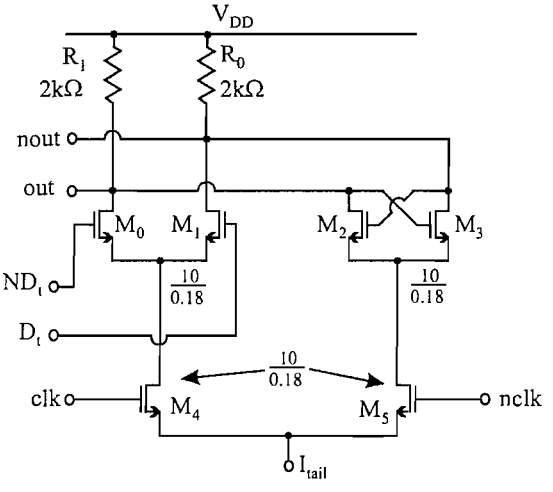

(a) Schematic of the fifth D-Latch circuit.

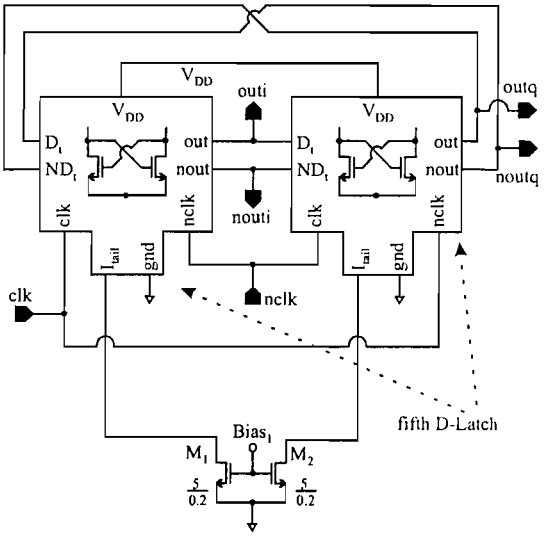

(b) Schematic of the fifth divide-by-2.

FIGURE 7.5: Schematics of the fifth stage.

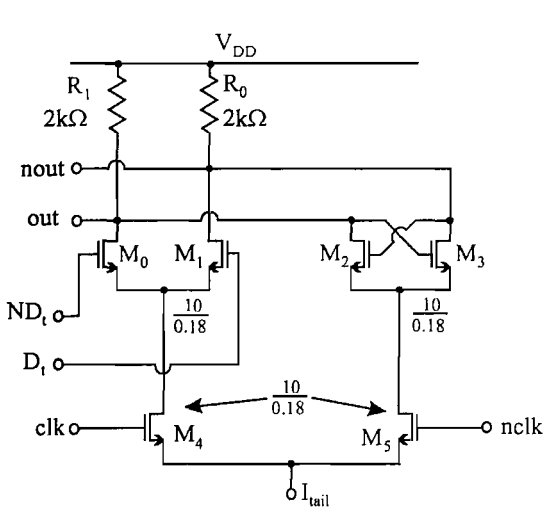

(a) Schematic of the sixth D-Latch circuit.

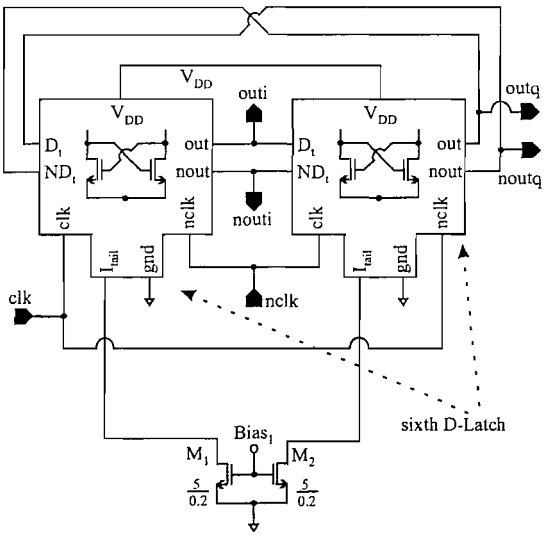

(b) Schematic of the sixth divide-by-2.

FIGURE 7.6: Schematics of the sixth stage.

7.1.2 Phase selector

The phase selector shown in Fig. 7.7(a) is no exception to the unstacking and requires both the clock transistors and the current source. This cell, although driven with a frequency one quarter of the input to the divider, poses quite a challenge with so many levels of transistors. The actual phase selector from the SOI divider can be placed with resistors, clock transistors and current source transistors, bringing the tally to four active devices plus a resistor and a 1.8V supply; this is quite a challenge. The alternative is to split the phase selector operation into two stages: 1) choosing whether to accept the quadrature inputs as they are or their inverses, 2) picking which quadrature to pass to the next divider. In the latter case, three active devices plus a load resistor must operate within the same 1.8V supply, but there is now a need for a set of interstage couplers between the split phase selector. However, it is preferable not to have these interstage couplers as they consume extra power, especially when a large drop is required. Simulations showed the complete phase selector to operate within a 1.8V supply and hence is chosen for the remainder of this design.

7.1.3 Level shifters

Regarding the source followers, those following the second divider have a greater drop than the rest of the level shifters (determined by the need to drop three levels of MOS devices as opposed to two,) but also suffer from the greatest degradation in high frequency

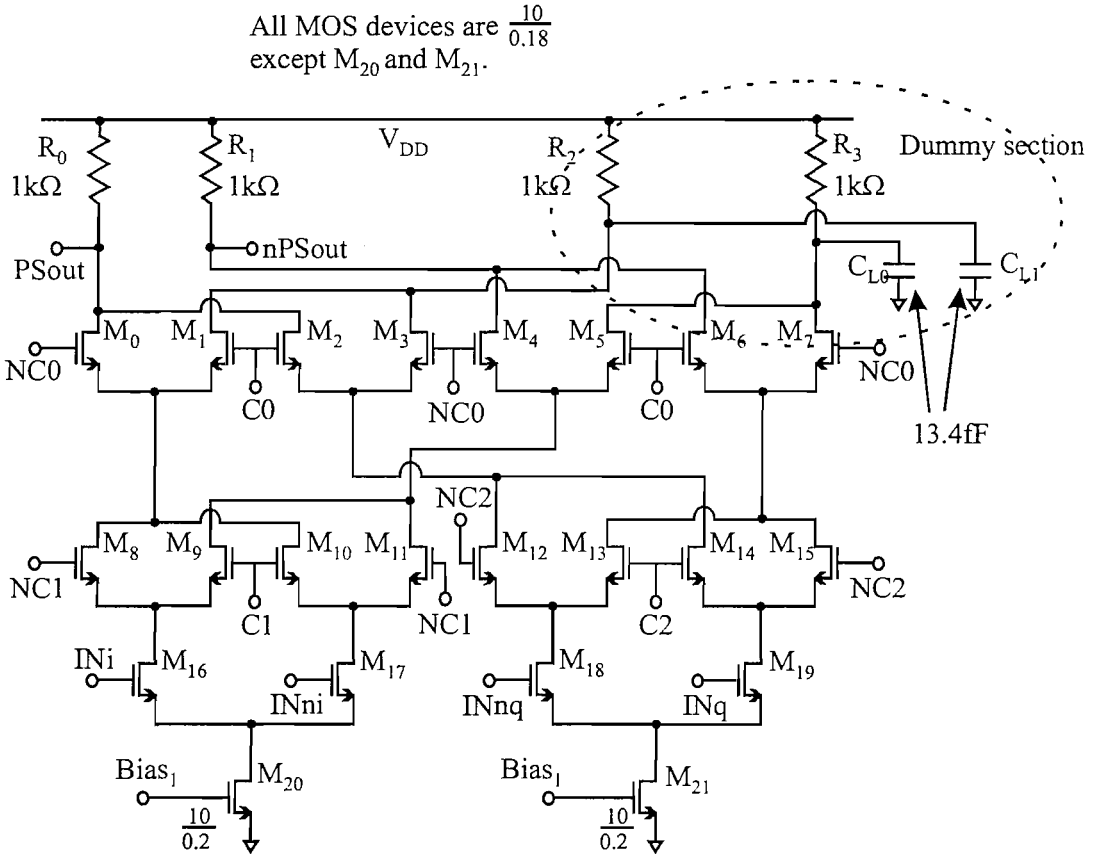

(a) Schematic of the phase selector.

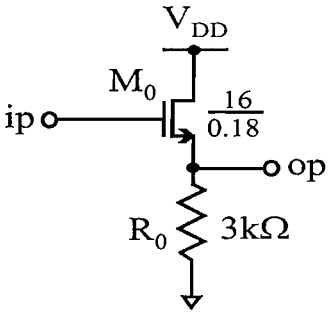

(b) Schematic of the source followers between the dividers.

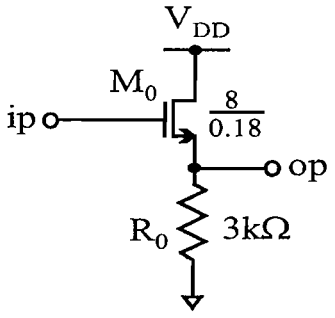

(c) Schematic of the source followers between the outputs of the second divide stage and phase selector inputs.

FIGURE 7.7: Schematics of various other subcircuits.

gain. Looking at the small signal gain of the source follower[1]:

$$\frac{v_o}{v_i} = \frac{g_m}{g_m + g_{mb} + \frac{1}{R_L}}, \tag{7.1}$$

where g_m is the front gate transconductance, R_L is the load resistance of the source follower circuit, and g_{mb} is the transconductance due to the body, any increase in the bulk transconductance value will lower the gain (assume a large load resistance). This term is almost inversely proportional to the bias across the source-bulk terminals (there is a square-root function in the denominator acting on this voltage). When the source-bulk voltage approaches zero, this transconductance approaches a value that is less than unity (voltage swing dependent gain). However, simulations showed the phase selector to have a high sensitivity at its inputs such that this was not a problem. These source followers required a DC drop of 0.8V with the other available amplifiers needing to drop 0.5V between input and output. If the output resistance was made lower, then the level shift would be even greater but at the expense of greater attenuation of the input signal, as explained above.

7.1.4 Modulus control

The modulus control circuit was initially implemented as in the SOI divider (Fig. 7.8(a)) using a differential pair NAND gate in the stacking close to the load resistors. The aim of this is to stop the state machine by forcing it into a 'particular' stable state, making sure that the divider's output prevents the subsequent dividers from toggling. Unfortunately, the problem arises with this particular stable state as the outputs are fixed to a certain value (see Wave3 in Fig. 7.9), when set **not** to divide by (N+1). This means that if the output of the divider is in any other state then the outputs will change when the divider modulus is changed; this leads to an erroneous change of outputs, as illustrated in Wave4 of Fig. 7.9. If the subsequent dividers in the FSM are sensitive to such an edge, then they also toggle, causing the phase selector to change phase at the wrong time and thus skip an input cycle. Instead, a transparent gate which holds onto the last known value of the clock level when set to do so by the modulus control is the solution (Fig. 7.8(b)). As Waves 5 and 6 in Fig. 7.9 show, the latch is transparent when dividing by (N+1), but holds the last state when the modulus control toggles. Although this may not be an issue in the case of a dual-modulus divider, it is fatal when included in a programmable divider loop (if just the dual-modulus divider were used and the subsequent circuit were to allow a settling time, then this error could be ignored.)

The SCL D-Latch in Fig. 7.10 implements this task, giving the desired operation. Obviously, this is a power hungry solution to the problem of modulus control, but nevertheless, it solves the problem. The response time of this latch must equal that of the

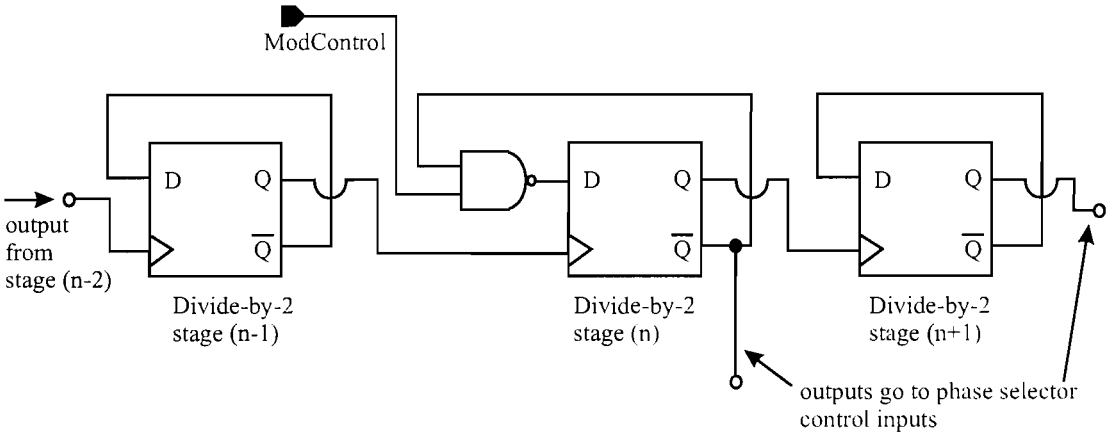

(a) NAND gate version.

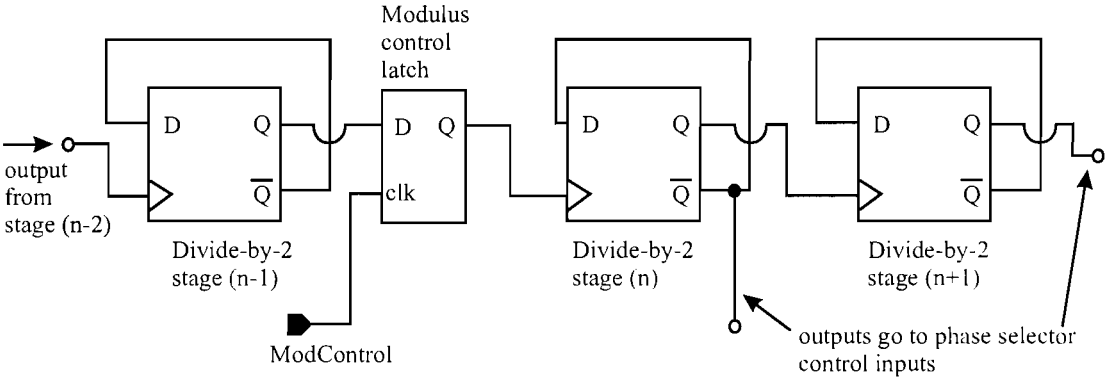

(b) version with latch.

FIGURE 7.8: Schematics related to the modulus control problem, where the upper diagram relates to the asynchronous modulus switching, leading to erroneous division in a programmable divider loop. The lower diagram is a solution to the mentioned problem.

last divider in the 16/17 divider chain in order for it to operate correctly. This observation allows its design to be no more power hungry than each of the latches in the divider stage preceding it, and also the capacitance it presents to the output of the divider. This latch is not as power efficient as one in CMOS topology, but it does remove the need for changing between CMOS to SCL logic levels with the signals running at a few hundred MHz (it should not be necessary to apply a conversion the other way, with only a DC level shift being required to centre the swing around the switching threshold of the latch.) This cell has been designed to consume 250μA but it should be feasible to reduce the current further in this simple cell.

Figure 7.11 shows a schematic of the CMOS-to-SCL converter that translates a CMOS rail-rail modulus control input into a differential SCL signal, which controls the aforementioned latch.

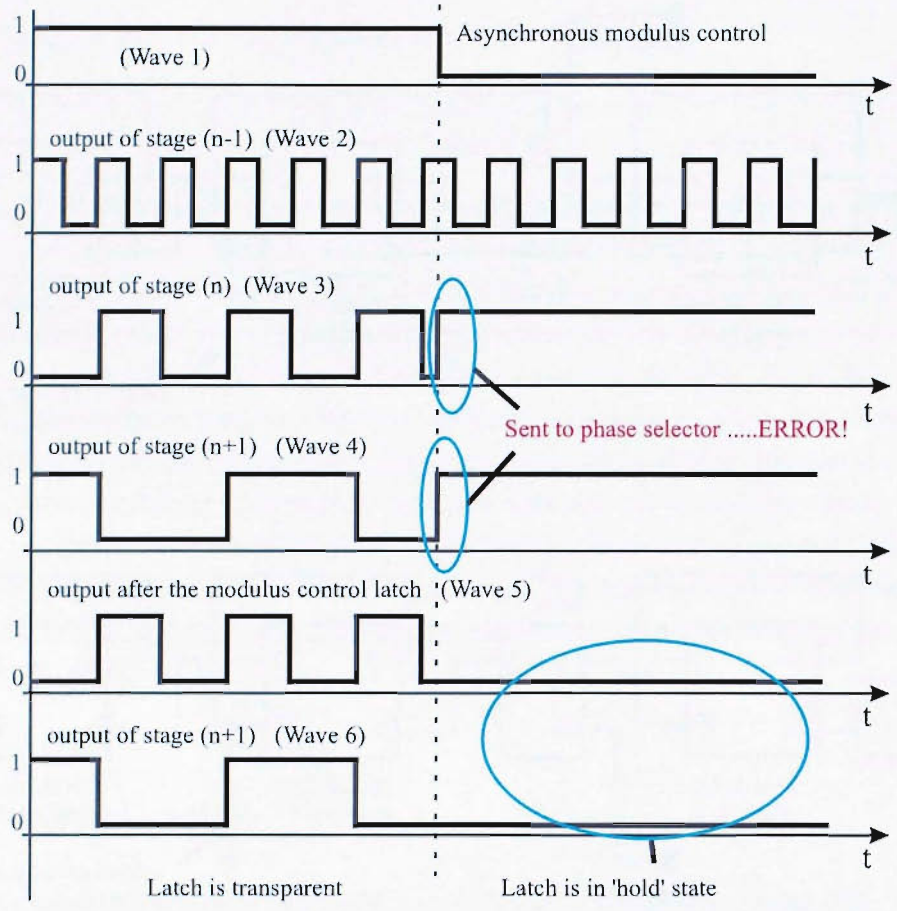

FIGURE 7.9: Time domain plots illustrating the problem with an embedded NAND gate in our implementation.

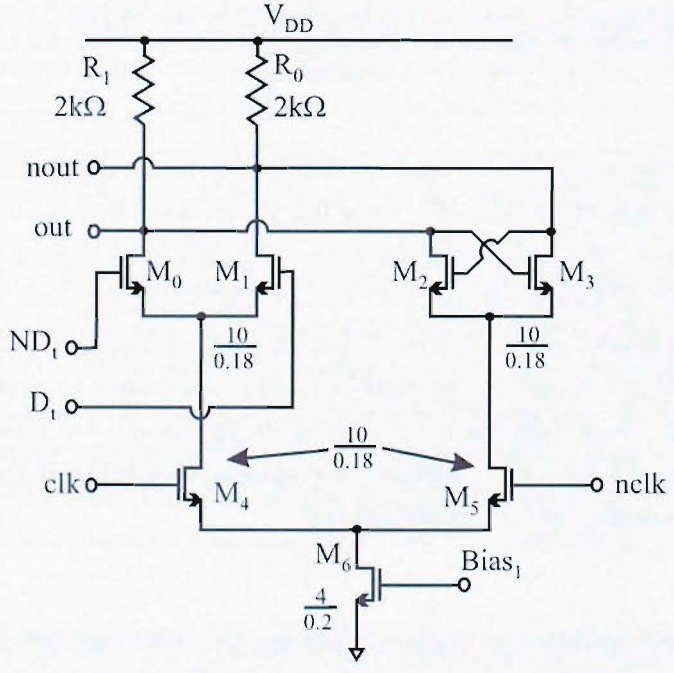

FIGURE 7.10: Schematic of the modulus control latch.

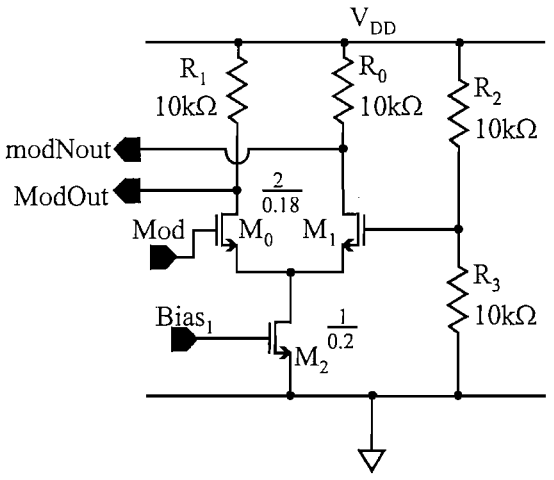

FIGURE 7.11: Schematic of the CMOS to SCL differential pair.

7.1.5 Complete circuit

The whole divider is represented schematically in Fig. 7.12. The current bias is set using a resistor between the input of a current mirror and the upper power rail. When used in the IEEE802.11a LO, a very stable current reference of 25μA is made available for the rest of the chip.

The observant reader should notice the presence of a great number of capacitors on the outputs of some of the dividers. They are deliberately placed in the design (including the layout) and their purpose is solely to balance the loads on the divider outputs, deemed critical for successful operation at high speeds. With most of the dividers in the division chain, capacitors have been placed on their redundant outputs, either the master or the slave. This concern to guard against unsuccessful division at high frequencies is taken to another extreme with a powered divider stage on the redundant pair of outputs of the first divide-by-2 stage. The redundant divider's outputs are then loaded with linear capacitors; the motivation behind this action is to have a balanced operation at the very high frequency end of the circuit. However, if one is to be thorough on this issue of matching, then the question can be raised as to why no source follower is present between the output of the first divider and the clock input of the redundant divider. The author acknowledges this as a mistake, but is satisfied with the effort to place some active load, as opposed to none at all. The measurements in the next chapter will show this oversight not to have marred the performance of this divider in the desired frequency band.

The schematic in Fig. 7.13 shows the connection of blocks after connecting to bond pads. A 50Ω output driver is placed after the dual-modulus divider to aid high frequency characterisation. There are also a considerable number of on-chip decoupling capacitors

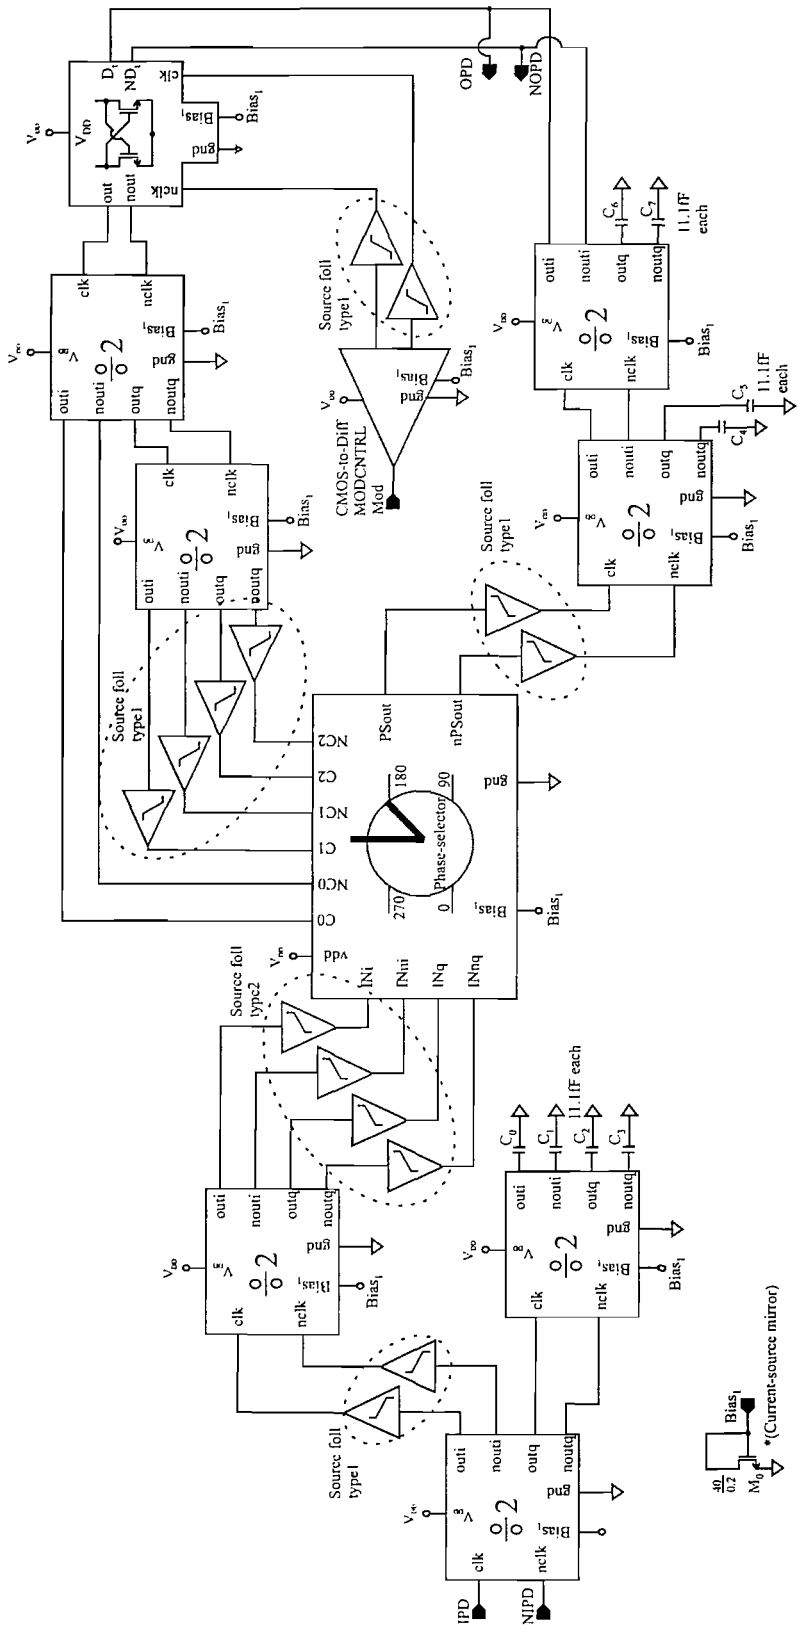

FIGURE 7.12: Schematic of the divide-by-16/17 core circuit.

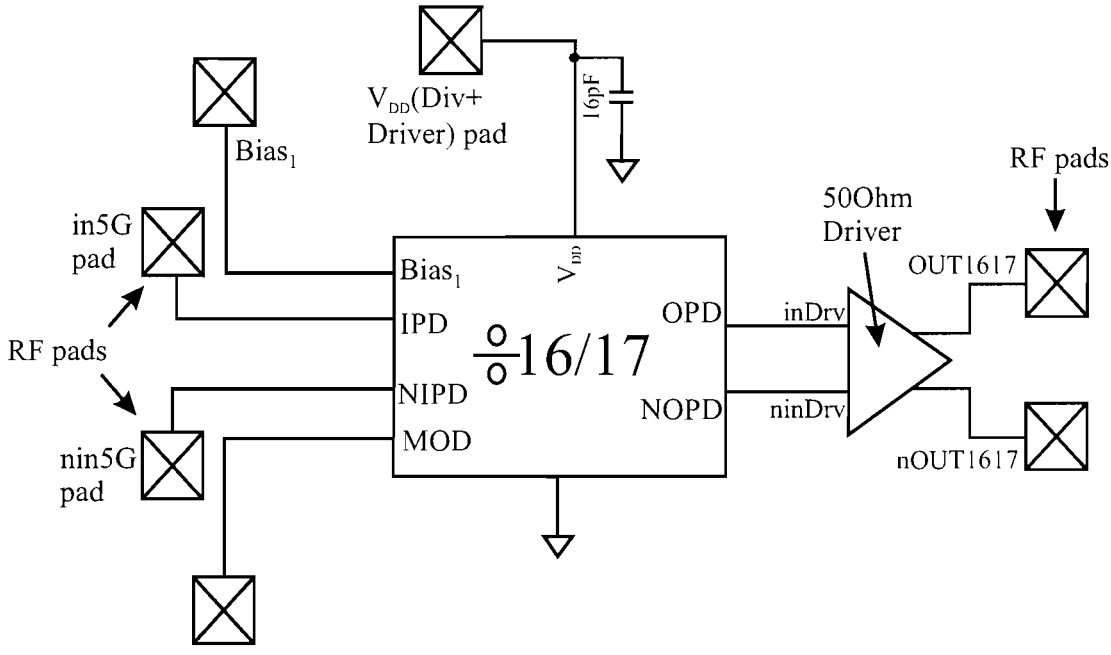

FIGURE 7.13: Schematic of the divide-by-16/17 fabricated circuit.

in order to supply the transient charge due to high frequency transitions, compensating for the effects of any bondwire inductance.

7.2 Layout

In order to characterise the dual-modulus divider, a stand-alone copy of this cell, outside the programmable loop, was inserted into a padring for characterisation and verification. The divider ratio is fairly low and hence a division of 5GHz by 16 drops the frequency down to around 300MHz, thus requiring a 50Ω driver. Looking at the individual layouts of the divider cells, it should be clear that the currents in the active regions travel in the same direction, either horizontally or vertically. This is also the case in the full dual-modulus divider core. The arrangement of the transistors and interconnect is largely the same for every cell, with the aspect ratio of the earlier dividers differing due to their larger power drain. The common feature amongst every divider cell is that the high frequency signal is confined to a small radius about the centre of each layout, hence justifying the position of the current source transistors. As all the CAD models used during simulation use the RF extension, it seems sensible to use a layout on which the model characterisation is based.

The phase selector pictured in Fig. 7.15(a) is a deviation from the rather similar divider cells pictured. The floorplan of this cell is dictated by the floorplan of the whole dual-

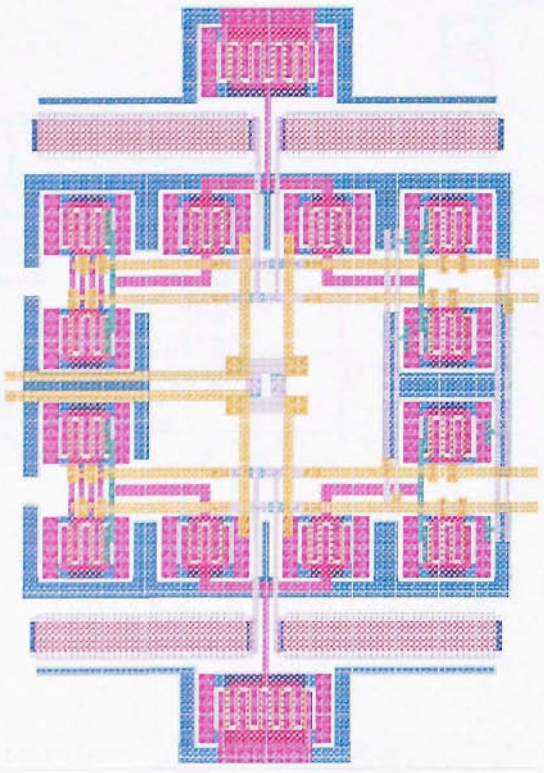

(a) First divider.

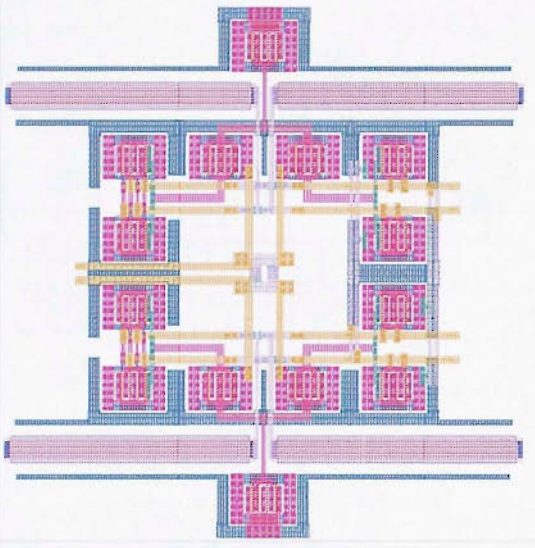

(b) Second divider.

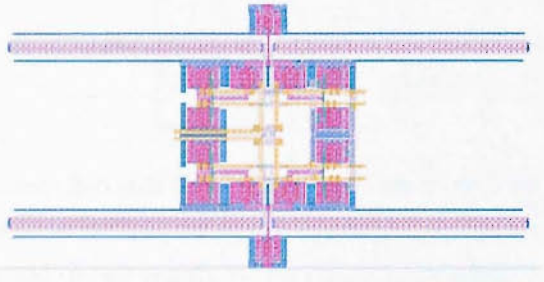

(c) Third divider.

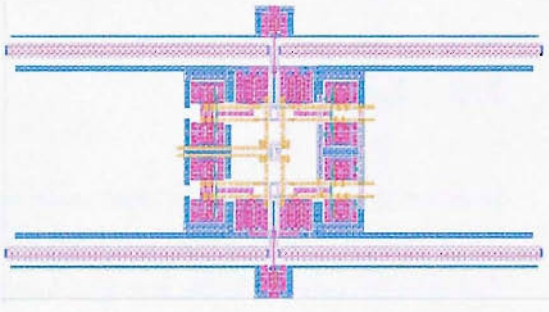

(d) Fourth, fifth and sixth dividers.

FIGURE 7.14: Layouts of the individual divider cells.

modulus divider. The arrangement of the transistors is similar to the circuit topology shown previously, with the redundant outputs capped with linear capacitive loads as well as polysilicon resistors. In such a critical cell, the need to balance parasitic capacitances is a high priority, so that the quadrature phase relationship between the inputs is not disturbed. Saying this, the frequency at which this phase selector cell operates has lower sensitivity to such parasitics, unlike in the case of the 11GHz VCO.

Figure 7.16 is the last layout in this section, showing the taped out cell. The high frequency signal arrives on the left, with the divided signal emerging on the right hand

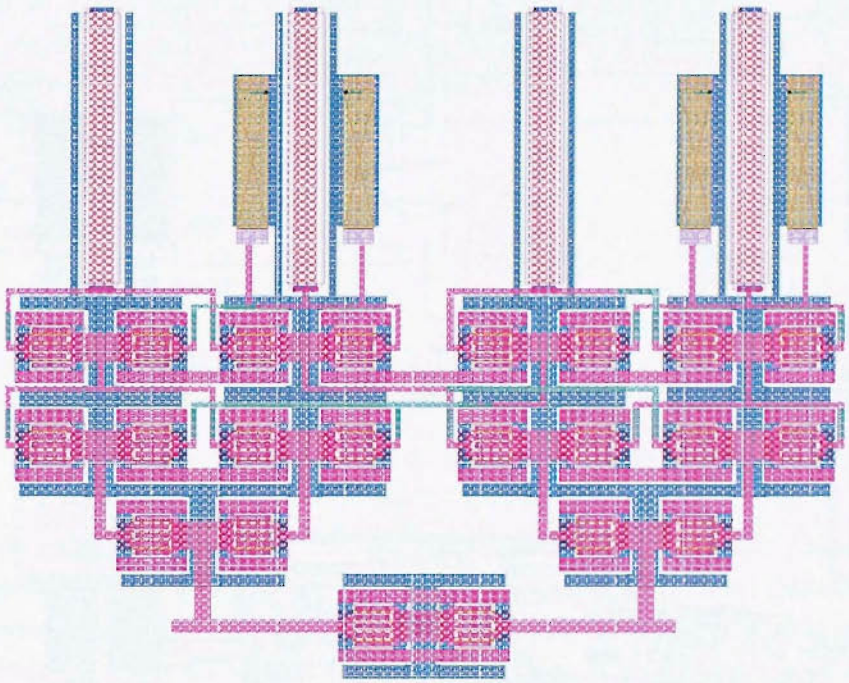

(a) Phase selector.

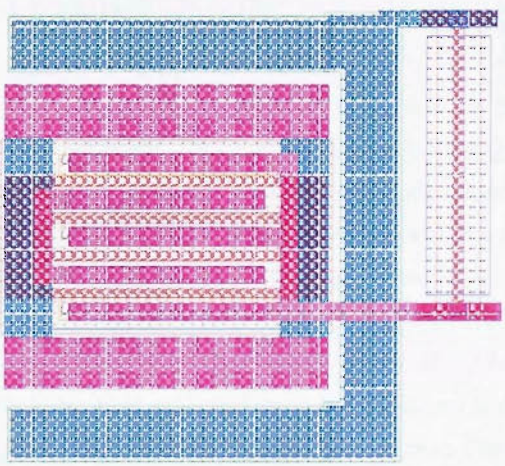

(b) Most common source follower.

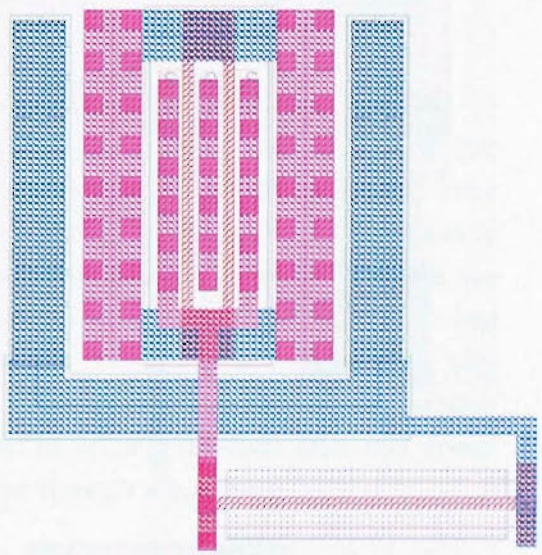

(c) Source followers between the second divider and phase selector inputs.

FIGURE 7.15: Layouts of miscellaneous cells.

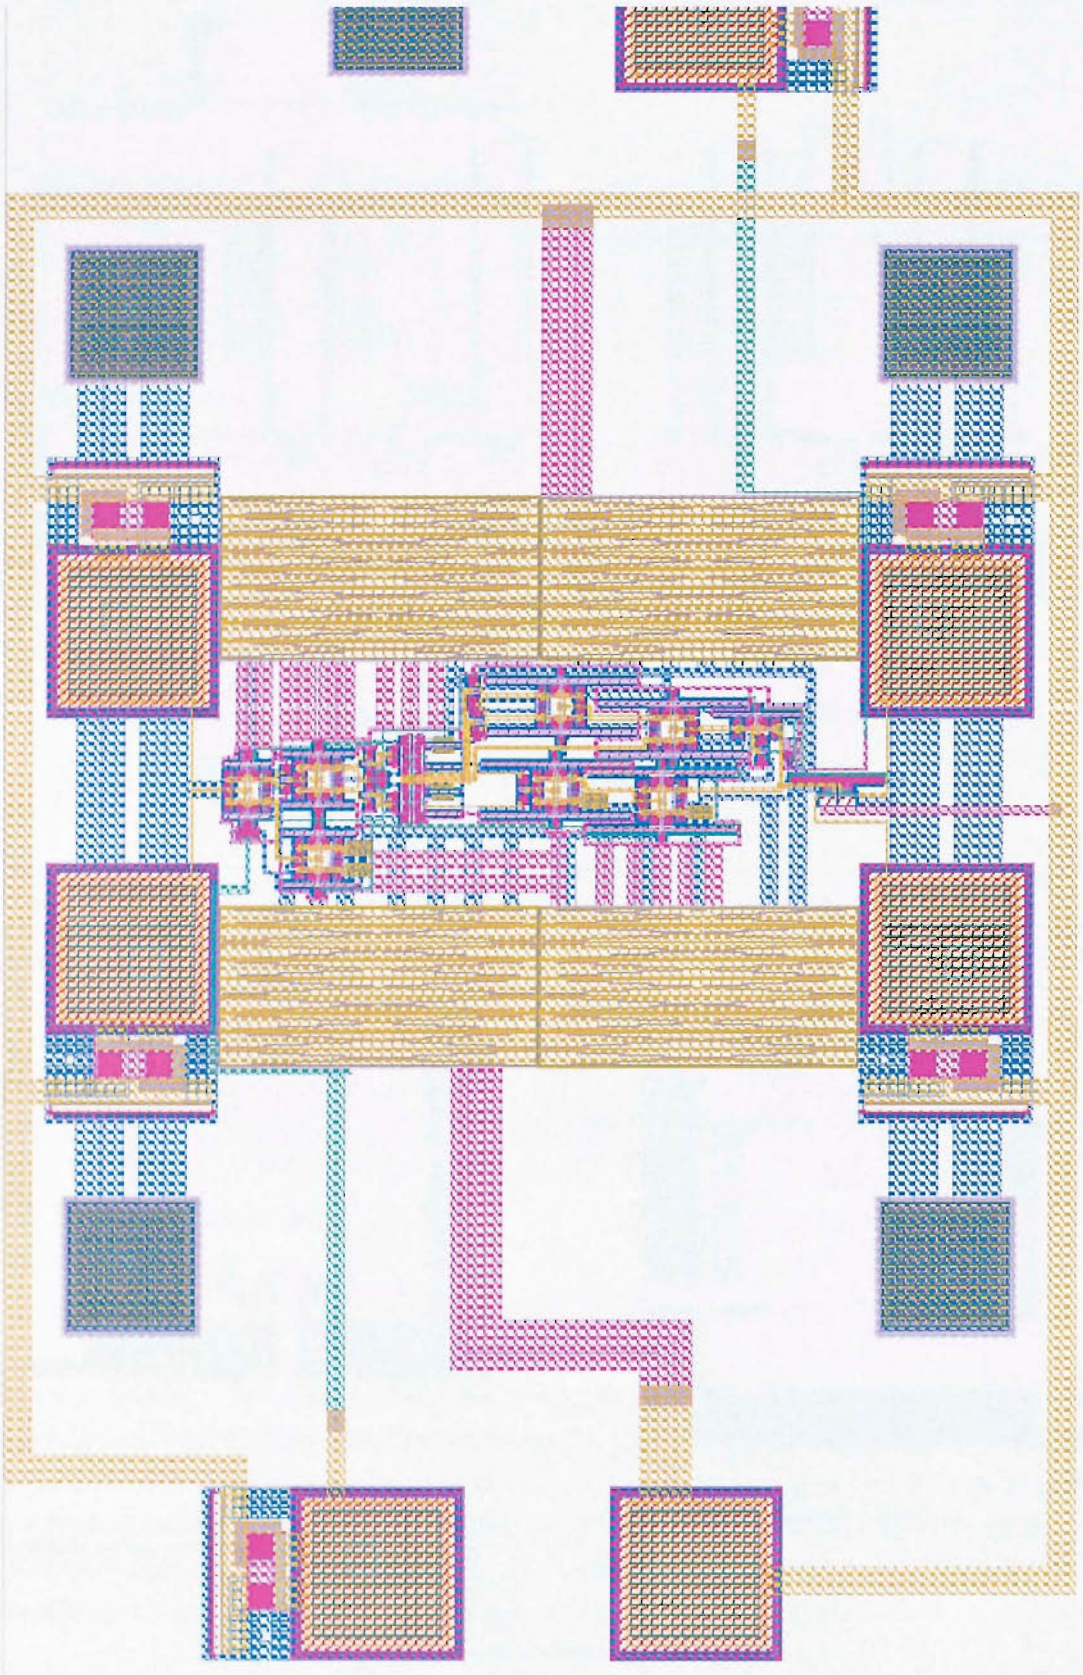

FIGURE 7.16: Layout of the divide-by-16/17 taped out circuit.

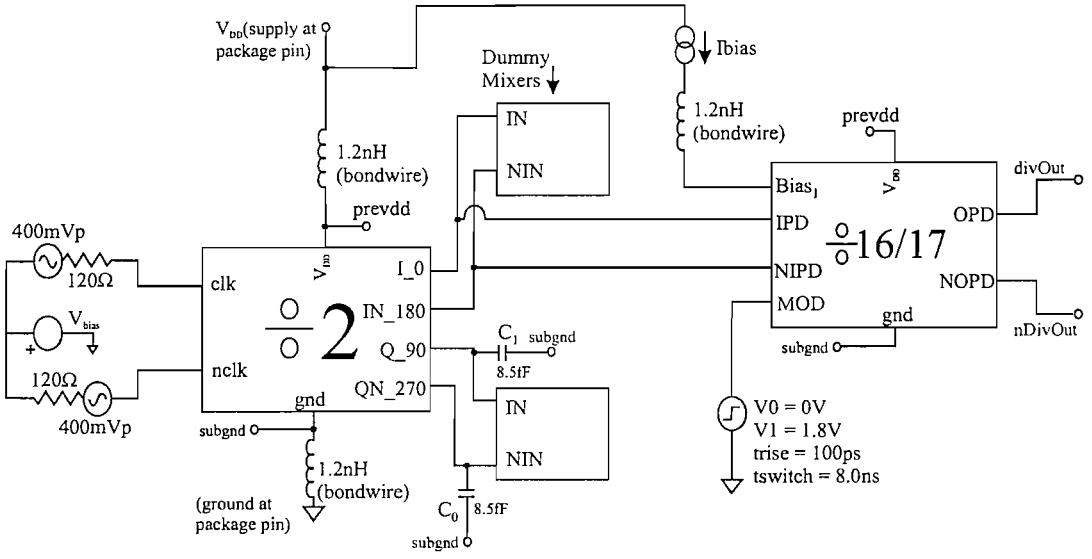

FIGURE 7.17: Schematic of the divide-by-16/17 in a testbench arrangement.

side. The flow of operation loops back from the output to the phase selector control inputs. Figures 7.15(b) and 7.15(c) show the simple active level shifters, that are placed symmetrically in pairs between the divider stages. There is also a 50Ω buffer placed before the *ground – signal – signal – ground* pads on the right hand side for matching and driving such low impedance without affecting the operation of the divider core.

7.3 Simulations

In this section, the simulation results are presented to show the dual-modulus divider running. Figure 7.17 shows the test arrangement, with the conscious decision to drive the circuit with the 12GHz quadrature generator. The quadrature generator incidentally drives the 'dummy' mixer loads as well as linear capacitors for the output without the dual-modulus divider. Although shown, the inductors are not relevant during the testing of this cell, and this is because the circuit is in fact probed and not wire-bonded to a package. However, it is useful to see the effects of bondwire inductance as the intention is to have this cell within a LO loop surrounded by other transceiver cells, and driven with the same power supply connection delivered through a bondwire.

Figure 7.18 presents the results after simulating the previous testbench with an 11GHz input signal on the input to the quadrature generator. The output of the divide-by-2 oscillates around 1.1V DC, which is slightly off the desired common-mode voltage. The peak-peak swing of the dual-modulus divider input is 750mV at 5.5GHz. On its output, a voltage swing of 700mV is produced. The upper graph in this figure shows the modulus control input, with it dividing by 17 when the signal is low. The effect of the bondwire

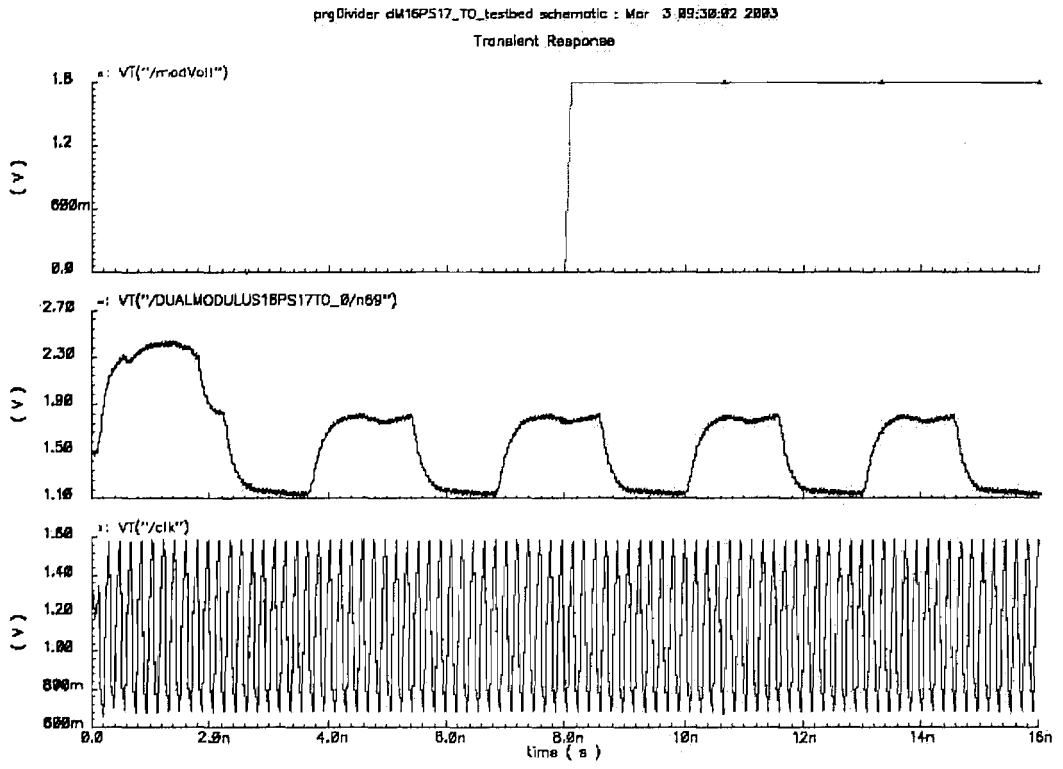

FIGURE 7.18: Simulations of the testbench with a 5.5GHz input. Top plot is the modulus control signal on the 'MOD' port of the divide-by-16/17 unit. The middle plot is the single-ended output, 'divOut.' The bottom plot is the single-ended clock input, 'IPD,' port of the divide-by-16/17 unit

inductance in series with the decoupling capacitors is shown in the output waveform; note that the supply to the core of the chip jumps to greater than 2.3V before settling back down to the nominal 1.8V. This is naturally a cause for concern, but a fault in the simulation is not having the power supply ramp up to the nominal value in the space of say a few milliseconds.

Although not captured, the circuit has been simulated to run as fast as 7GHz, exceeding the specifications of the IEEE802.11a standard. This is enough margin for unaccounted effects which could mar the performance of this cell and not allow the LO to operate up to the upper edge of the desired band.

7.4 Measurements

At the time of writing this thesis, no measurements existed solely for this dual-modulus divider. However, the design has been included in an integer-N programmable divider discussed in the next chapter. This design was found to function at the correct frequency and measurements will be included in that chapter.

7.5 Discussion

Although the primary aim of the chip discussed was to provide a circuit function that is suitable for 5GHz LO generation, a useful insight has been given into the bulk implementation of the topology and phase-select control scheme. The simulations have clearly shown the circuit to work as well as in the desired frequency range.

Looking back at the circuit, it should be obvious that designing such a circuit in bulk (schematically) is a challenge compared with the SOI design as one now has the back gate effect manifesting itself. From a layout point of view, in the case of body-tied PDSOI, the bulk design is far simpler, though with very high frequency designs, it is necessary to place a large number of ground substrate contacts. The low power supply does little to help the design, giving a low headroom for signal swing.

From a modelling point of view, highly accurate DC models are needed for short-channel devices, especially in a low supply voltage scheme where errors cannot be tolerated. Such models were made available for the project and the results in the next chapter provide evidence of their essential role in the design. With the use of source followers, the need for excellent RF modelling is now as important as ever, ensuring high frequency operation at the early stages of the divider, not seen in the SOI divider presented earlier. Without such modelling, it would be hard to quantify parameters such as phase shift due to the source followers and the attenuation in signal amplitude due to the capacitance (rather than the gain) of such a buffer.

Being in its own pad-ring partially masks the true performance of the chips, with pad parasitics and the output buffer. Saying this, by the time the input signal is divided down to its output, digital cells from a library should be acceptable for later stages. As in the SOI divider, it was decided to keep all dividers after the phase selector as static dividers, at the expense of a higher current consumption. The benefit is that when the quadrature outputs are needed, there is no need to convert back from CMOS to SCL in order to generate those outputs. However, this could well be considered a luxury and should ideally be simulated against a hybrid dual-modulus chain of CMOS and SCL

divide-by-2's.

The amendment to the modulus control function is implemented in this circuit and was prompted by the simulations of the full programmable divider discussed in the following chapter. The embedded NAND gate used in the SOI version, though simplistic, had a certain output state when signalling the divider not to swallow a pulse, causing a momentary but significant glitch. With the subsequent dividers sensitive to it, erroroneous division has been seen with such a scheme. The inclusion of the D-type latch acting as a sample-and-hold circuit for the clock signals (see Fig. 7.8) solves this problem, again at the expense of slightly higher current consumption. The operating characteristics of the dual-modulus divider before and after such amendment are quite distinct and really requires its inclusion in a programmable divider loop for the difference to make itself apparent.

Bibliography

[1] P.R.Gray and R.G.Meyer. *Analysis and design of analog integrated circuits (3rd edition)*. John Wiley, 1993.

Chapter 8

Programmable (513-544) divider in 0.18μm bulk CMOS

This chapter develops further the ideas presented earlier for the dual-modulus divide-by-16/17 design already discussed, in the design of a 513-544 divider in a 0.18μm bulk CMOS technology. The internal configuration of this programmable divider is what makes the PLL an integer-N PLL as opposed to a fractional-N PLL. The core of this programmable architecture is the divide-by-16/17 circuit which, with some additional sub-blocks, forms the heart of the synthesiser as part of a Philips Research project. The design has been fabricated and successfully measured, and is able to run with input frequencies up to 6.6GHz.

8.1 Circuit design

8.1.1 PLL architecture

The local oscillator used in the IEEE802.11a direct conversion receiver project (conducted at Philips Research Laboratories N.V, Eindhoven, The Netherlands) is a frequency synthesiser based on a PLL, whose frequency can be set to the centre frequencies of the appropriate channels. In order to perform this discrete channel selection, a programmable divider of some sort is required which, in our project, will come after the 12GHz divide-by-2 circuit discussed in chapter 6. The actual style and operation of this programmable divider is based on the chosen style of the frequency synthesiser. As introduced in chapter 1, one can have a fractional-N synthesiser as well as an integer-N synthesiser, each with its own merits and flaws. Although the fractional-N synthesiser is said to have a better phase noise performance and give the designer a greater degree of freedom with the choice of the reference frequency and locking time, no comparison

of either architecture has been performed prior to the design showing whether the spurious output of this style of frequency synthesiser would overshadow its benefits for this application. Both contain the same number of divider blocks, but differ in where the output is taken for phase comparison. In the fractional-N type synthesiser, the phase detector is required to operate at a higher frequency. Owing to better understanding and lack of time for further study, the integer-N PLL was the chosen style of frequency synthesis for this project, though this is certainly not necessarily the optimum choice.

The choice of channel spacing in an integer-N architecture is governed by the reference frequency arriving on one of the inputs of the phase detector. It is usual practice to have a higher frequency divided down to the desired step size, as the phase noise of a frequency divider is theoretically improved at its output. This follows the $20 \cdot log_{10}(N)$ rule [3], predicting the improvement at the output of any combined divider with respect to the input phase noise, assuming the noise floor of the divider isn't reached. The IEEE802.11a standard operates in the 5GHz unlicenced national information infrastructure (UNII) band. In this standard, an aggregate signal bandwidth of 300MHz is available, split over two bands (the lower has 200MHz). There are 8 channels in the lower band and 4 in the upper, each having a bandwidth of 20MHz. Within each OFDM (orthogonal frequency division multiplexing) channel, 52 subcarriers are present, each being a BPSK, QPSK, 16-QAM or 64-QAM signals and the channel has the capacity to allow a maximum data rate of 54Mb/s.

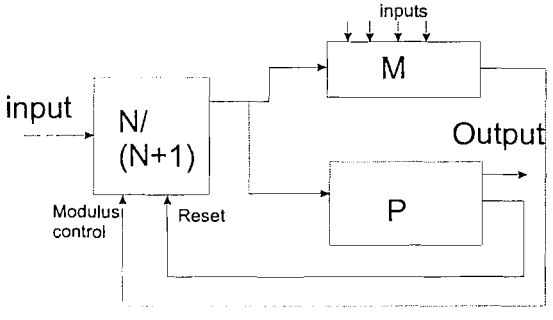

FIGURE 8.1: Block diagram for an integer-N programmable divider.

A step size of 20MHz would be an obvious choice for the synthesiser, but a conscious decision to have a 10MHz step size was made, giving a finer granularity to frequency selection if need. The drawback with this approach is that a larger set of division ratioes emerge because a greater frequency multiplier is needed to achieve channel selection in the same frequency band. This leads to a higher close-in phase noise at the output, as seen in equation (6.1), which is not attenuated by the low pass filtering of the loop bandwidth. Although the frequency divider circuit is said to reduce the phase noise between its input and output, the phase noise (close-in, within the loop bandwidth) at the output of the local oscillator goes up by the same factor.

Concentrating on the lower band, the range of integer values needed are derived as follows:

- the lower edge of the frequency band is 5.15GHz, which has to be compared with 10MHz $\therefore \frac{5.15\times10^9}{10\times10^7} = 515$
- the upper edge of the frequency band is 5.35GHz, which has to be compared with 10MHz $\therefore \frac{5.35\times10^9}{10\times10^7} = 535$

From this, the integer range of the divider is [515, 535]. To achieve this with dual-modulus dividers, one can use the architecture shown in Fig. 8.1 for integer division.

8.1.2 Division ratio

Referring to Fig. 8.1, the overall divider ratio is $P(N+1)-M$, where $P \geq M$ for a dual-modulus divider with ratio $N/(N+1)$, fixed divider, P, and a programmable counter with a value M [5]. The variables M, N and P can be any integer though there are the physical problems of implementing such a counter. The composition of the divider is much simpler when using 2^N division prescalers (with the programmable counter being an exception). With division ratios falling outside this requirement, the result is dividers that require additional combinatorial logic which can impede the speed performance of the circuit. With the required divisors and the range, the ratios were found iteratively as follows:

1. with a divide-by-2/3, the fixed divider can be 256, giving a ratio between 512 and 767 (or 513 and 768)
2. with a divide-by-4/5, the fixed divider can be 128, giving a ratio between 512 and 639 (or 513 and 640)
3. with a divide-by-8/9, the fixed divider can be 64, giving a ratio between 512 and 575 (or 513 and 576)
4. with a divide-by-16/17, the fixed divider can be 32, giving a ratio between 512 and 543 (or 513 and 544)
5. with a divide-by-32/33, the fixed divider can be 16, giving a ratio between 512 and 528 (or 513 and 527)

Following the decision to implement the phase select architecture of chapter 4 in bulk CMOS, the first two configurations were ruled out. The second of the two requires at

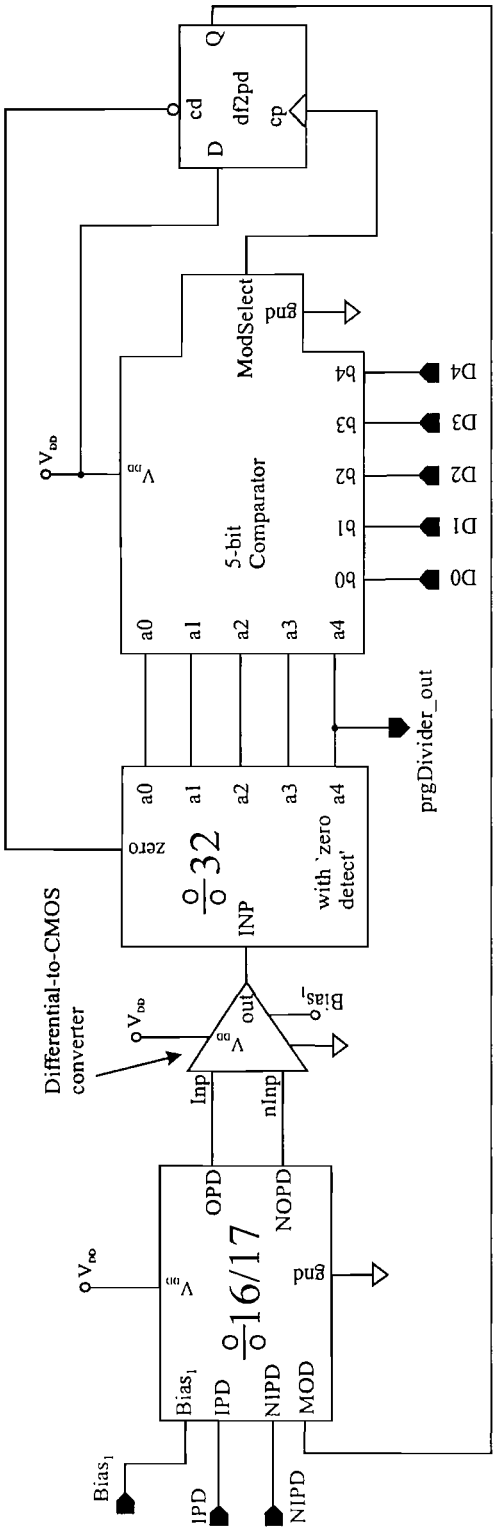

FIGURE 8.2: Schematic of the programmable divider core.

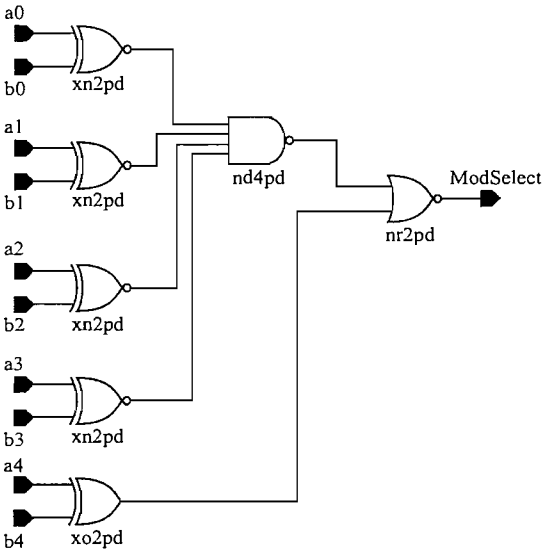

(a) Schematic of the digital comparator.

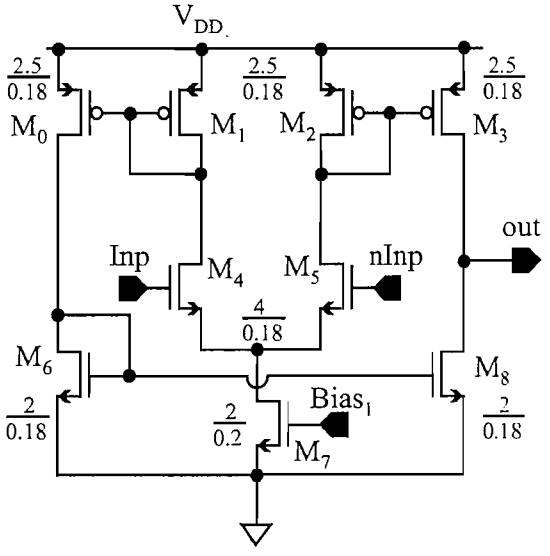

(b) Schematic of the SCL to CMOS translator.

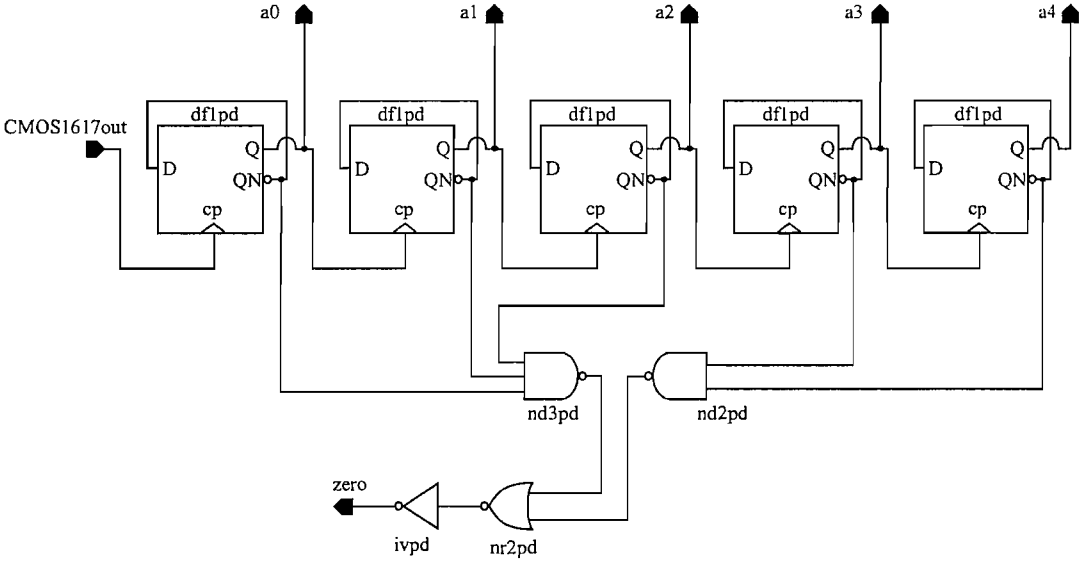

(c) Schematic of the extra divide by 32 circuit.

FIGURE 8.3: Schematics of various cells.

least one divide-by-2 stage (which will also be a part of the state machine) after the phase selector to avoid race conditions. The fifth proposal is also ruled out, with the range of divisors falling short of the target. The decision to choose the divide-by-16/17 over the divide-by-8/9 was based on the frequency at which the control signals to the phase-selector toggle. With the higher division ratio, the signals have more time to settle. The divide-by-8/9 option may insinuate more divider cells designed in a CMOS topology (in the overall programmable divider), but this would require custom designed latch cells because library cells would not be able to cope with such a high input frequency.

As mentioned before, the motivation to keep the signals in SCL in our divide-by-16/17 throughout the loop is based on the desire to have true differential signalling and to remove the overhead of translating the CMOS signal to an SCL one in order to control the phase-selector.

8.1.3 Transistor level design

Figure 8.2 shows the schematic of the programmable divider, with ratios 513 to 544 inclusive. The dual-modulus divider from Chapter 7 is used unchanged. Its SCL output signals are converted to rail-to-rail CMOS (using the circuit shown in Fig. 8.3(b)), before driving a fixed divide-by-32 cell (Fig. 8.3(c)) containing Philips 0.18μm CMOS library standard cells. The input frequency to this cell is over 300MHz and thus, due to the compactness of the design, should be simulated with extracted interconnect parasitics, together with MOS model 9 transistor models for the PMOS and NMOS transistors. The RF extensions to the MOS model 9 cannot be used for these cells, as the physical layout of the transistors in these digital cells do not follow the style of layout modelled by the RF extensions (for example, there is only one contact to the polysilicon gate and the lack of substrate/well taps per transistor increases the body-to-substrate terminal resistance.)

In order to toggle the modulus of the dual-modulus divider, a digital comparator checks the output of the fixed 5-bit counter with a programmed code designating the desired the division ratio. The choice of cells is based on the delay through each logic gate, with the fastest chosen for every gate, making sure the fanout isn't violated. The same is true for the fixed divider. The Boolean equation for the circuit in Fig. 8.3(a)is :

$$Q = \overline{(\overline{a0} \cdot \overline{b0} + a0{\cdot}b0) \cdot (\overline{a1} \cdot \overline{b1} + a1{\cdot}b1) \cdot (\overline{a2} \cdot \overline{b2} + a2{\cdot}b2) \cdot (\overline{a3} \cdot \overline{b3} + a3{\cdot}b3) + (\overline{a4}{\cdot}b4 + a4 \cdot \overline{b4})} \tag{8.1}$$

The output of the programmable divider is in fact the same as 'a4' marked on the divide-by-32 fixed counter (in Fig. 8.3(c)), and will generate one full cycle for every 513-544 cycles on the input of the dual-modulus divider, depending on the 5-bit word that is present on the inputs of the digital comparator.

The last cell in this loop is a D-type flip-flop which toggles the modulus control input. When the output of the digital comparator toggles (signalling a match), the transition at the output is sent to this flip-flop, causing its output to change from logic 0 to logic 1. When the fixed divider (counting down) reaches zero, the flip-flop is reset, and the dual-modulus divider begins to divide by the inital divisor at the beginning of the next cycle. The flip-flop's CMOS rail-rail signal is fed back to the dual-modulus

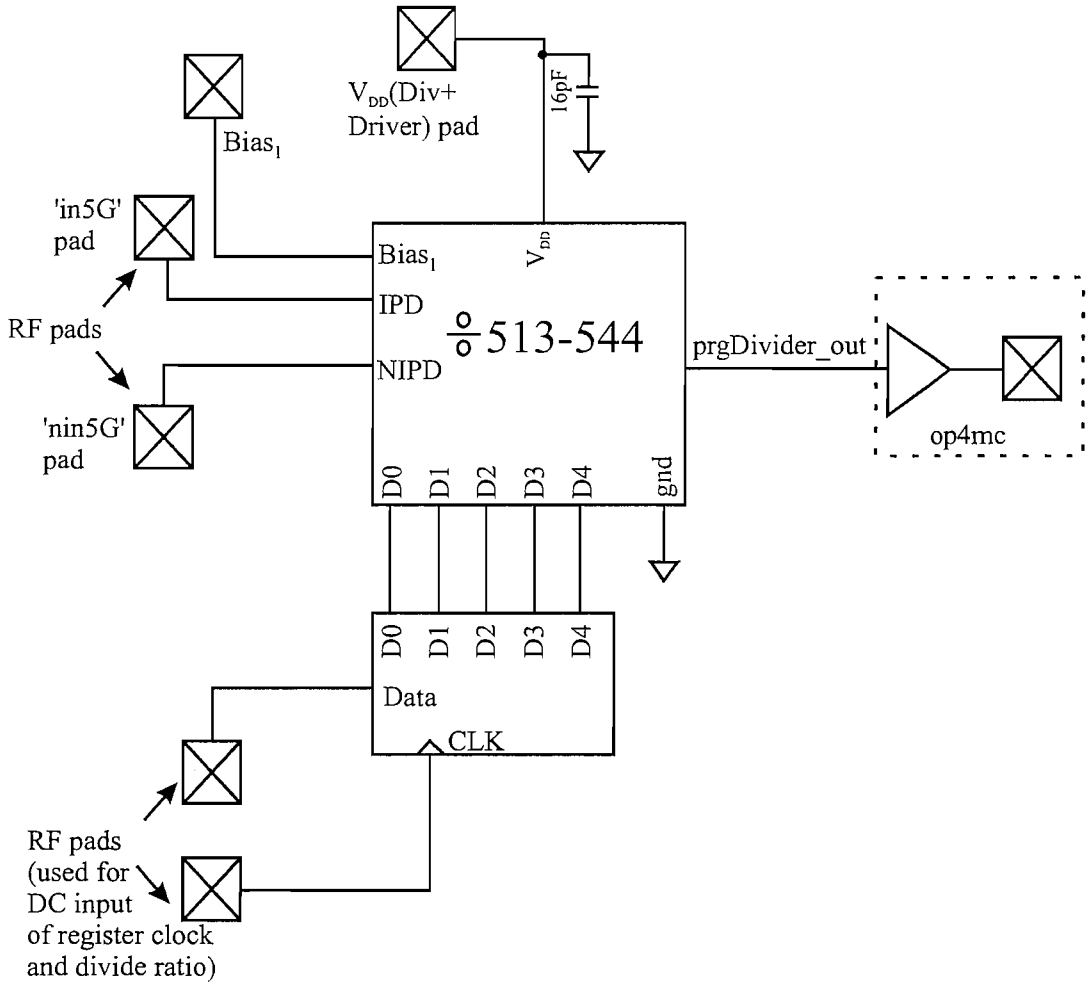

FIGURE 8.4: Schematic of the taped out programmable divider.

divider where the translator of Fig. 7.11 buffers the signal appropriately for internal use.

At the beginning of this chapter, we gave possible combinations of fixed- and dual-modulus dividers that, when connected together, resulted in the minimum range of divisors. For our choice of dual-modulus divider, there are 32 consecutive integer divisors, ranging from 513 to 544 inclusive. The 5-bit input word required to select the divisor ranges from '11111' for 513, to '00000' for 544. The dual-modulus divider's default modulus is 17, and the fixed counter is wired to count down in sequence, starting from '11111'. When the input (`b4,b3,b2,b1,b0`) is '00000', the single flip-flop remains in its default state and 32 output cycles emerge at the fixed counter output ('a4'), each with a period equivalent to 17 dual-modulus divider inputs, amounting to division-by-544. When the input is '11111', the flip-flop will instruct the dual-modulus divider to divide-by-16 immediately and will do so for every output cycle (31), until the output of the fixed counter (`a4,a3,a2,a1,a0`) reaches '00000'. At this instant, a 'zero' state is detected and the single flip-flop at the end of the programmable divider is reset back to

its original state to divide-by-17. The fixed counter will complete an output cycle with this input (17 dual-modulus input periods in length). Thus, the overall division will be 16*31 + 17*1 = 513.

Had the default divisor of the dual-modulus been 16, and the fixed counter counted 'up' in sequence, then the range of divisors would be 512 to 543. The expected 10MHz output of this circuit is a CMOS rail-to-rail voltage and is taken from the MSB of the fixed divider.

The final schematic shown in Fig. 8.4 shows the divider core surrounded by pad cells instances. For quick testing, it was decided that probing the wafer would retrieve results quicker. Considering the tight space around the probe station chuck where the probe holders sit, the number of pads that can be probed is limited so, in our design, where a 5-bit digital input is required, a serial input pad is placed in the pad ring to reduce this pad count. This pad feeds the input to a serial-in,parallel out register, clocked by the signal arriving on another pad (marked 'CLK'). The pad count is such that one '*ground – signal – signal – ground*' set accounts for the 5GHz input, 4 DC pads are used for circuit ground, current bias, circuit supply voltage, 10MHz CMOS output, and the finally another '*ground – signal – signal – ground*' pad set accounts for the divide ratio serial input and the corresponding register clock.

8.2 Layout

The bulk of the layout has already been described in Chapter 7, with the dual-modulus divider at the heart of the layout. The additional cells in the programmable divider loop have been added to the layout of the dual-modulus divider. These cells are collected in Fig. 8.5. Layouts constructed using RF layout transistors are consistent with the need to have all active area currents travelling in parallel directions. In the case of the library cells, this is not possible and so the matching issue is forgone with these cells.

With the circuits built around standard cells, these digital blocks have been cascaded where possible, making full use of their individual floorplan attributes and minimising additional routing.

The pad layout in Fig. 8.6 shows how little is changed between this circuit and the design described in the previous chapter. It should be pointed out that the pad assignments are different with the output found on only one pad. The GSSG pad arrangement on

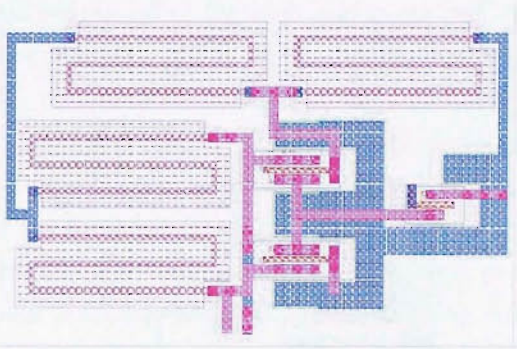

(a) Layout of the CMOS to SCL differential pair.

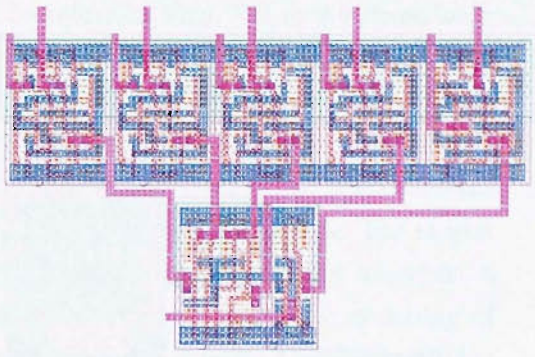

(b) Layout of the digital comparator with standard cells.

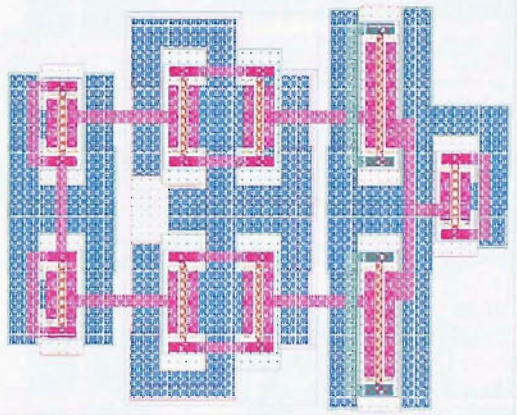

(c) Layout of the SCL to CMOS opamp.

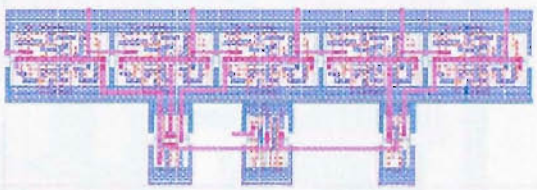

(d) Layout of the divide-by-32 cell using standard cells.

FIGURE 8.5: Layout of various cells.

the right hand side has one signal pad designated for the divider output and the other for clocking the division ratio register. The current bias is set on the corner pad at the top left, whilst division ratio is programmed into the internal register via the bottom right-hand pad. The signal pad in the top right corner resets the register. As before, considerable power supply decoupling is placed on chip close to the core of the divider.

8.3 Simulations

As part of the verification, the circuit in Fig. 8.2 is simulated before commencing layout. Being a large-signal circuit, performing small-signal ac analyses shows very little as the circuit is far from behaving linearly. During transient simulation, each divisor has been verified (with a >5GHz input,) usually in pairs to see if the divider hangs when changing divisor. As there are 32 possible divisors, it is not very informative to reproduce tran-

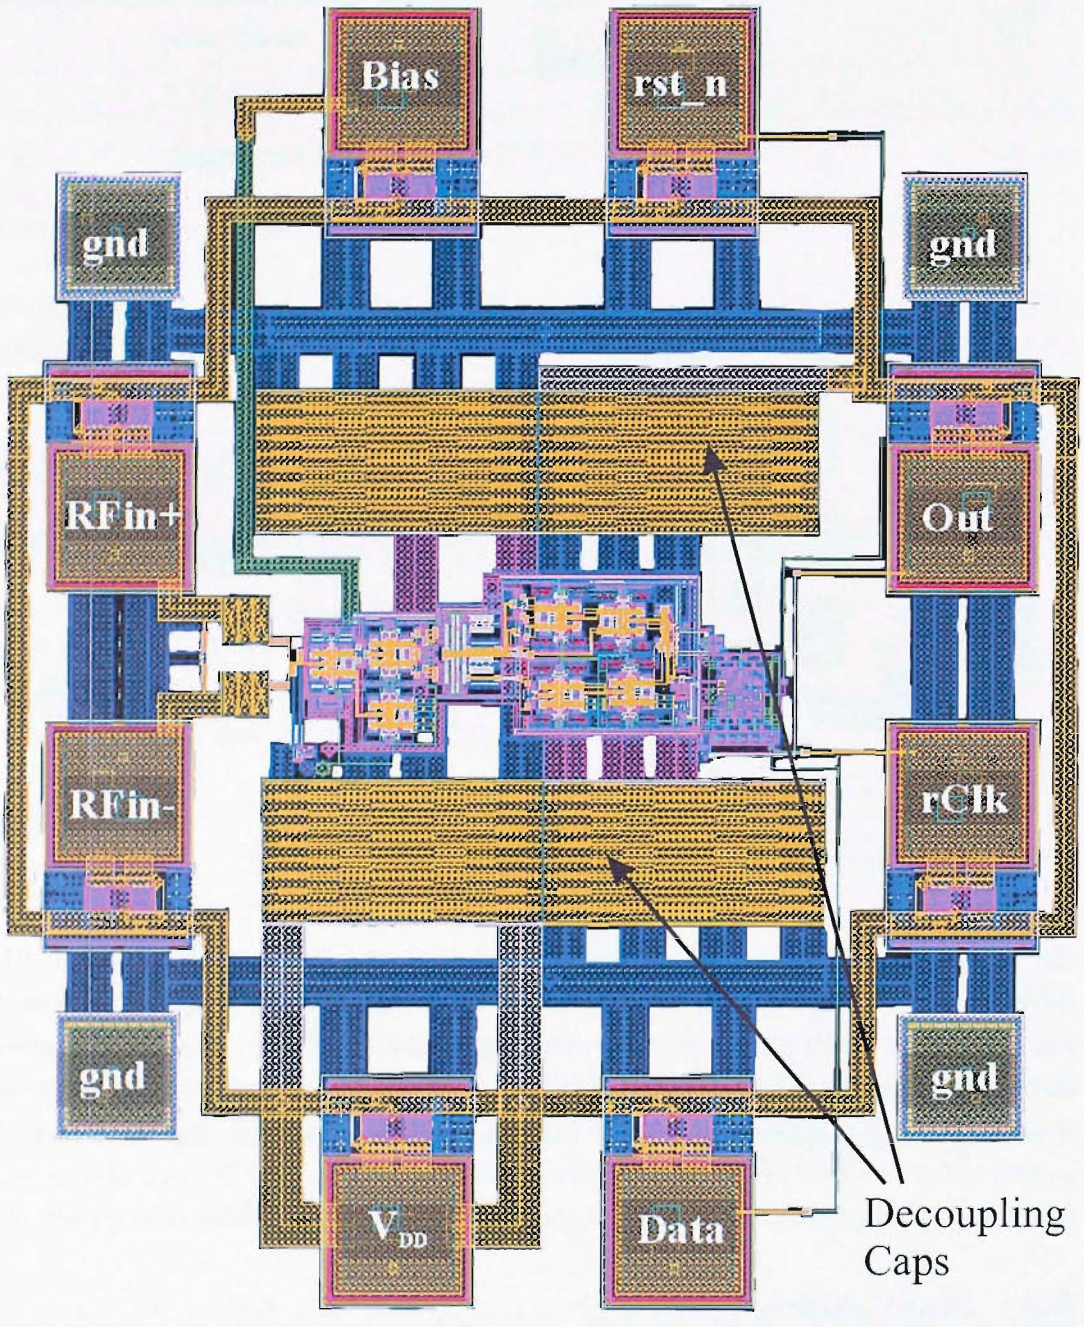

FIGURE 8.6: Layout of the taped out programmable divider.

sient output plots for each and every one, and therefore in Figs. 8.7 to 8.9 shows only those at the boundary of the group as well as one in the centre. As in the dual-modulus case, the programmable divider is driven by the 12 GHz divide-by-2 circuit described in Chapter 6 with a 10.7GHz clock signal on its input (thus translating to a 5.35GHz clock on the programmable divider input.) The input power level on its input remains the same as that from the corresponding simulation in the previous chapter. The reason behind this is that the additional circuit presented in this chapter does not introduce a global feedback loop. With the decision to change modulus affecting the switching of the circuit half-way into the dual-modulus circuit, the input stage to the dual-modulus divider and hence the programmable divider remains the same as in the previous chapter.

Though important, problems with the extraction tool hindered the manual extraction

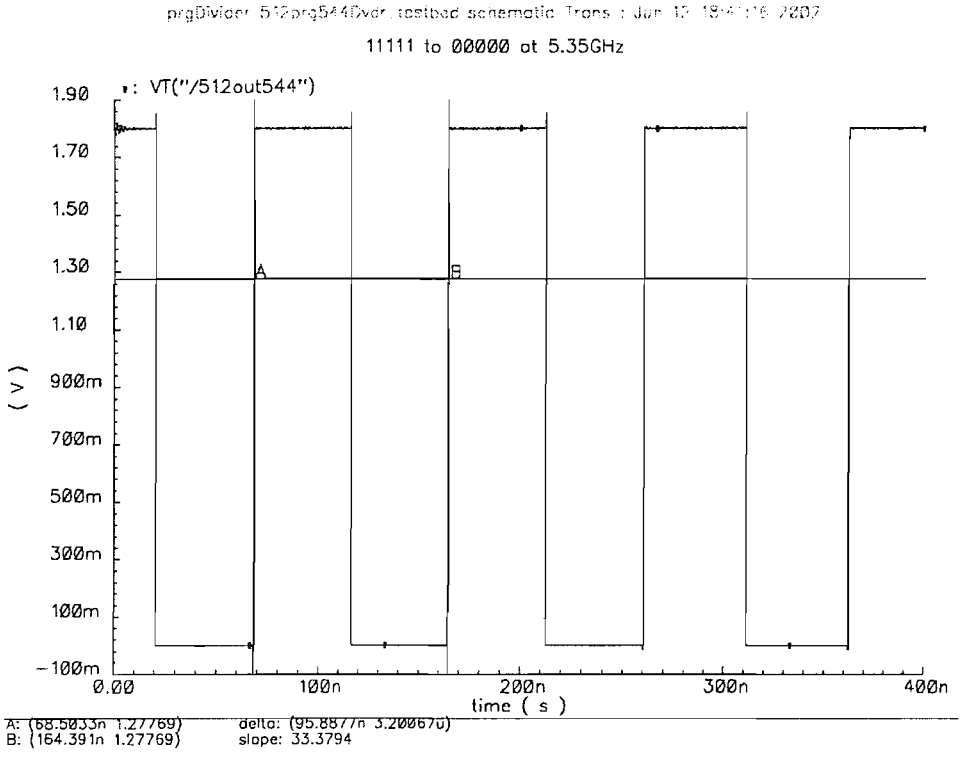

FIGURE 8.7: Screen captures output nodes of the complete divider (MSB output of the fixed divide-by-32 cell) after running transient simulations with a 5.35GHz input. Division by 513 (11111).

of parasitic capacitances of the divider circuit. To allow for the uncertainty resulting from this limitation, the dual-modulus divider has been designed with a margin such that the pre-layout simulations show the divider to operate at a higher frequency than the upper limit of the frequency band. Simulations with an input frequency of 7GIIz have been performed successfully.

Unfortunately, no phase-noise value has been extracted from simulations to check how

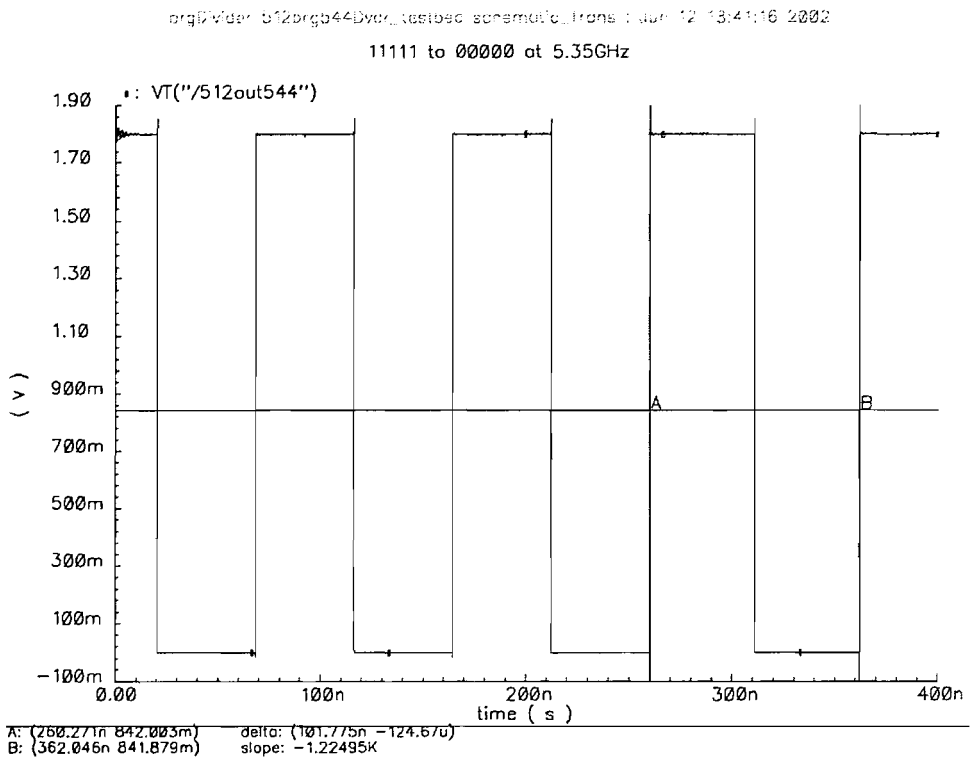

FIGURE 8.8: Screen captures output nodes of the complete divider (MSB output of the fixed divide-by-32 cell) after running transient simulations with a 5.35GHz input. Division by 544 (00000).

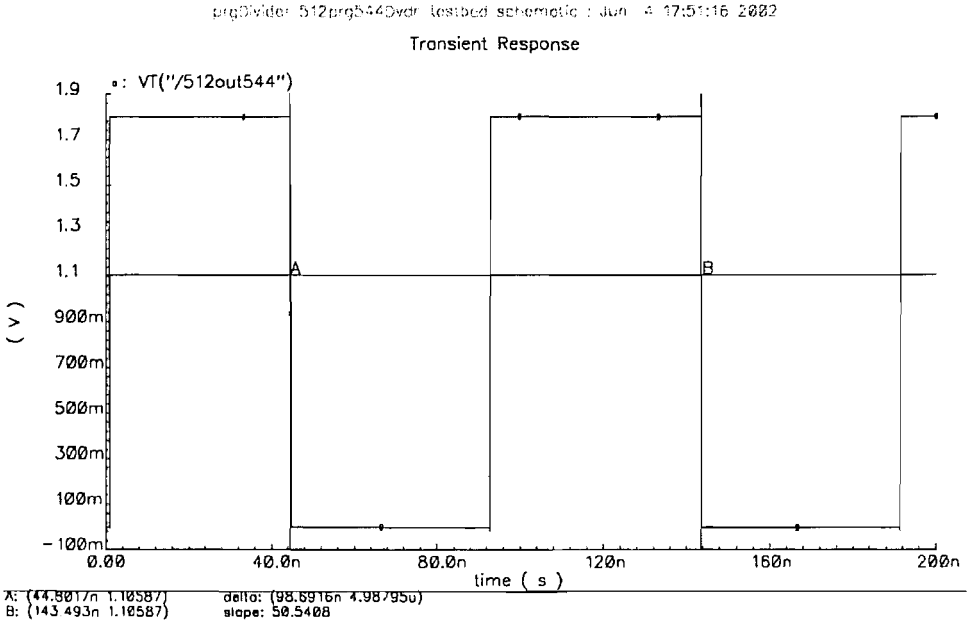

FIGURE 8.9: Screen captures output nodes of the complete divider (MSB output of the fixed divide-by-32 cell) after running transient simulations with a 5.35GHz input. Division by 528.

well it follows the expected relationship. The 'Periodic steady-state' analysis tool [1] [4] failed to converge at the time of the simulation. It is thought the mixture of models (plain MOS model 9 and MOS model 9 with RF extension) caused this failure, but this was difficult to conclude. One can simulate a pair of circuits with a common input and $\frac{\pi}{2}$ rad phase shift in one path before mixing the result down to baseband. The output transient simulation is then subjected to a n-point FFT in order to generate a discrete spectrum of the baseband result. Unfortunately, such simulation fails to show the true noise floor of such technology with equipment connected, resulting in an optimistic answer.

Another simulation that has been omitted is the input sensitivity plot. The algorithm used to extract such performance relies on identifying the correct minimum power level, with a slight quantisation error. For the divide-by-2 circuit, more than 24 hours were spent generating such data. With this programmable divider, the need to simulate long cycle times iteratively with a high resolution at a high frequency makes the simulation prohibitive. Without extracted parasitics of the taped out divider, any simulation would be dubious.

8.4 Measurements

The programmable divider was successfully fabricated in the Philips CMOS18 digital process and a photo of this chip is shown in Fig. 8.10. Although the divide-by-16/17 divider discussed in chapter 7 is the key to this programmable divider, there is no explicit output (nor is there a tap to the modulus control input) from that dual-modulus divider, in order to test its functionallity. Tapping the dual-modulus divider's output would have required a high speed buffer that was capable of driving a large capacitive load at a pad, whilst presenting a low capacitive load to that divider output. Though possible, time was a key factor in its omission.

The measurements on the returned samples have demonstrated that the programmable divider functions correctly, and within the frequency band of interest. From this, it is acceptable to infer that the dual-modulus divider is functioning correctly, though its performance (phase noise, power consumption, etc) is masked by the rest of the programmable divider.

The measurement setup for testing the die is shown in Fig. 8.11. The signal generator and spectrum analyser are equivalent to those used in Chapter 6. After generating the differential signals using a hybrid coupler, on-chip DC blocking capacitors prevent DC

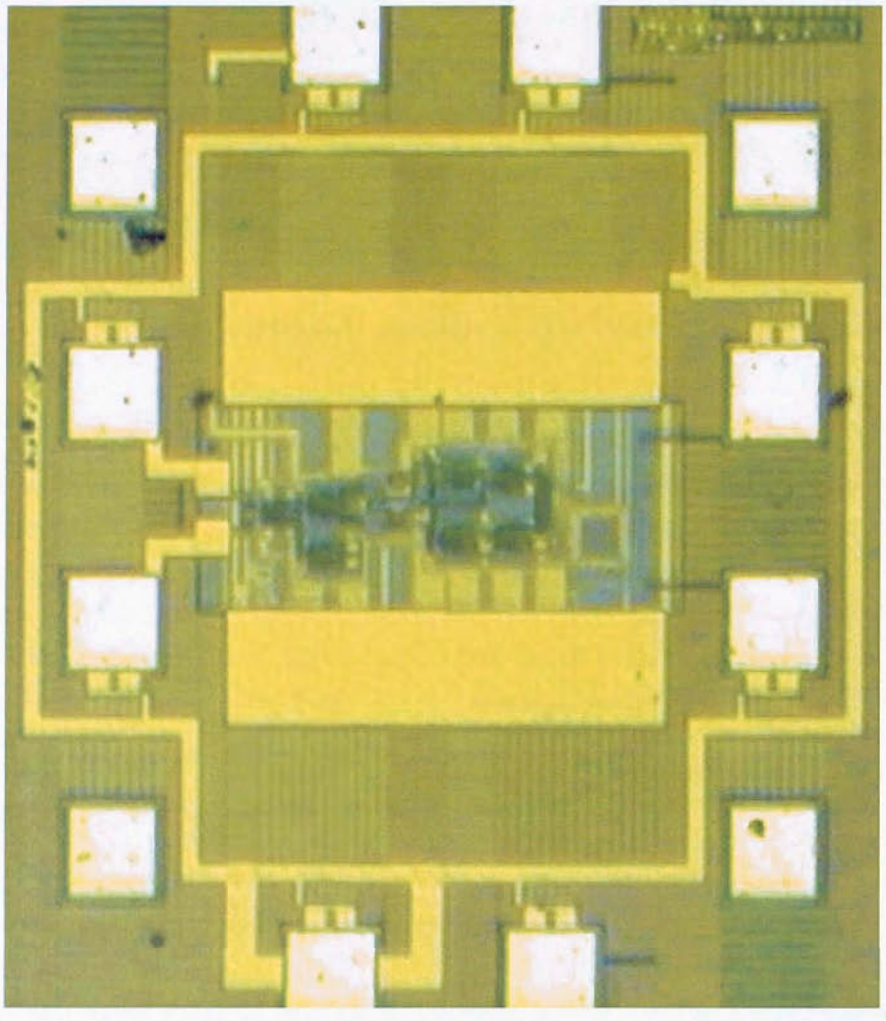

FIGURE 8.10: Chip photo of the fabricated programmable divider.

from the signal source arriving on the inputs to the divider core. The bias is set using the same arrangement as in Chapter 5. The input bias voltage to this programmable divider is equal to the output common-mode voltage of the bulk divide-by-2. With the passive mixer in the receiver chain dictating this common-mode voltage, the following measurements must adhere to this specifcation for the results to be credible within the context of the IEEE802.11a project. (NOTE: The DC blocking capacitors and input common-mode voltage generator have not been described explicitly in the design section of this chapter, as they were included immediately prior to tape out by engineers Nenad Pavlovic and Domine Leenaerts of *Philips Research Laboratories NV* to alleviate the need for external DC blocking capacitors and bias tees. The downside here, is that the input common-mode voltage is set via the master current bias, which also sets the current in all cells that use source-coupled logic.)

The division ratio is entered in a serial fashion using a DC source (mains-derived), as was the reset connection to the division ratio register. A battery powered data clock

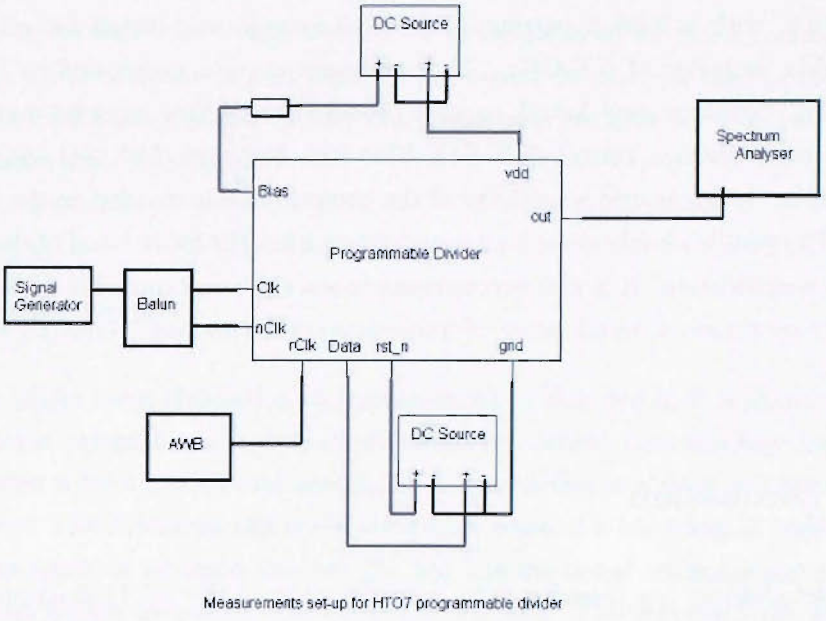

FIGURE 8.11: Measurement setup for checking the functionality of the divider.

generator provided our 1.8V single-ended CMOS clock for the register as it needed a 'clean' signal source. The mains DC power supply has been noted to be noisy, triggering the register to store the value on the 'DATA' pin more than once.

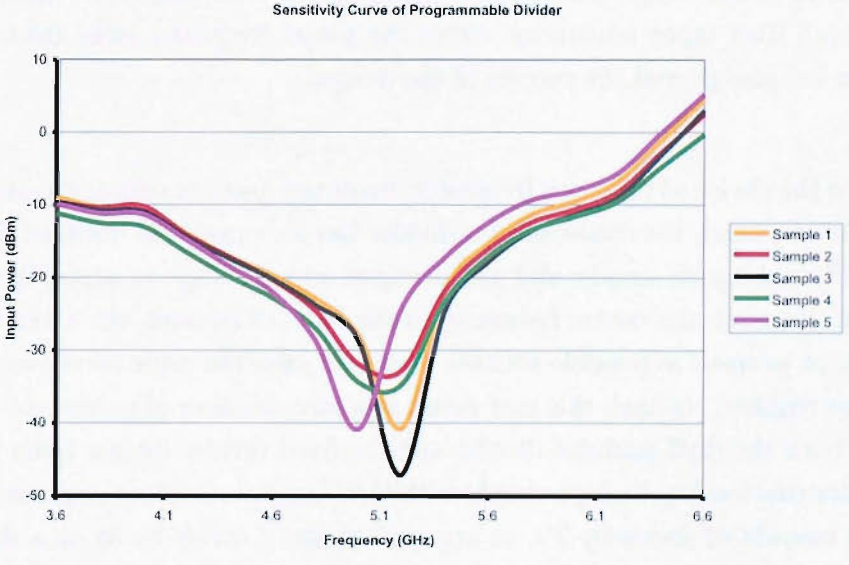

FIGURE 8.12: Input sensitivity curves for the programmable divider.

Although no explicit oscilloscope plots are shown of the low frequency output, a spectrum analyser was used to check the correct output frequency. Owing to bad probe contacts,

the power supply to the chip had to be raised, resulting in a current consumption of 12.6mA, with a 2.5mA current bias. One sample was tested for all divisors from 513 to 544 inclusive at 5.15GHz. The minimum current consumed by the divider for successful operation was found to be 11.8mA. Four other samples were checked for the following division ratios: 513, 514, 516, 520, 528, 529, 537, 541, 543, 544. Figure 8.12 graphs the measured sensitivity of the programmable divider, under the conditions given. The results clearly show high sensitivity within the lower band of the IEEE802.11a wireless specification. It is also encouraging to see the programmable divider functioning correctly over such a broad range of frequencies, with the top frequency being noted as 6.6GHz.

8.5 Discussion

This brief chapter has described the incorporation of the 16/17 dual-modulus divider within an integer-N divider loop. Together with the correction of the modulus control circuit, error-free operation has been simulated beyond 7GHz.

With the output frequency of the dual-modulus divider dropping to around 300MHz, the CMOS library cells have been used after converting the SCL signal to a CMOS logic output, resulting in a low current circuit as well as a quicker design process. The probed measurements have demonstrated the design to work as high as 6.6GHz, with a slight deviation in power supply and current consumption from the values used to simulate the design. High input sensitivity within the target frequency band from 5.15GHz to 5.35GHz has also proved the success of the design.

Owing to the choice of reference frequency, frequency band of operation, and centre frequency of the band, the dual-modulus divider has an 'open-loop' forward division ratio that is 2^N, making life simpler and not having to add glue logic to adjust the moduli. By virtue of the particular centre frequency of the band of interest, the extra fixed divider can be kept as small as possible without having to raise the value excessively in order to select the required channel; this may result in a large number of redundant divisors. By making both the dual-modulus divider and the fixed divider integer radix powers of 2, the implmentation can be kept simple with the complete divider comprising a straightforward cascade of divide-by-2's, as opposed to say a divide-by-35 or a dual-modulus divide-by-22/23.

The fabricated programmable divider yielded a fully functionally circuit, operating within the frequency band of interest. Initially, a conscious decision had been made to exclude the input common-mode stage and DC blocking capacitors. They were left

outside, in order to check the divider's tolerance to variation on the input common-mode voltage. One could designate a pad where it is possible to override the common-mode voltage (as in SOI dual-modulus divider case,) but this requires another pad and the chip can become difficult to probe.

8.5.1 Further work

Although no phase noise simulation or measurement is included in this chapter, a separate fabrication attempt of the dual-modulus divider would pave the way for a comparison between a fractional-N and an integer-N synthesiser in a deep sub-micron bulk CMOS process. The literature has noted the phase noise of a fractional-N implementation to be the superior amongst the two [2], but the fractional reference spurs around the synthesised centre frequency was cause for concern. The wireless LAN specifications were generous and thus sticking with the integer-N implementation may not prove to be fatal. It is agreed that simulations and studies should be performed to evaluate which is best. However, transistor-level simulations can be very time-consuming and may not even converge with the vast number of circuit nodes (each simulation of our programmable divider generating two output cycles required more than 4 hours). Assuming that powerful, dedicated computing resources, as well as time, were not an issue, then a $\Delta\Sigma$-modulator would still need to be designed and simulated for the fractional-N synthesiser.

Bibliography

[1] *Cadence Analog Artist.*

[2] B.Daly. Comparing integer-N and fractional-N synthesisers. *Microwaves & RF*, pages 93–133, September 2001.

[3] B.Razavi. *RF Microelectronics.* Prentice Hall, Inc, 1998.

[4] K.Kundert. Introduction to RF simulation and its application. *IEEE Journal of Solid-State Circuits*, 34(9):1298–1319, September 1999.

[5] W.Egan. *Frequency synthesis by phase-lock.* John Wiley and Sons (N.Y.), 1981.

Chapter 9

Summary, Conclusions and Further Work

This final chapter offers conclusions to the initial aims of the research, together with a summary of the results obtained in the latter half of this thesis. The central aim of the research has been to investigate a means for lowering the power consumption in high performance transceiver designs. More specifically, to look into very high speed, low power divider circuits for use within local oscillator designs found in radio transceiver architectures. The frequency divider is one block which runs at the highest signal frequency (and hence is a heavy consumer of power) in a phase locked loop for coherent indirect frequency synthesis [4]. The VCO is another example of a circuit which runs at the highest signal frequency and thus is a drain on the battery life in portable applications.

Chapters 1, 2 and 3 draw out the context in which our motivation is applied. The review material has set the tone of the work as well as focussing the attention on frequency dividers. With the circuit topology predominantly being source-coupled logic, a review is included with a discussion on the large signal DC and transient behaviour associated with this logic style. A ring-oscillator is central to the analysis of the logic's transient behaviour, with equations attempting to follow the trend of the circuit speed, based on a selection of MOS model parameters. However, these equations are far from complete, and deviate from the true operation where short chains based on short channel MOS devices suffer from a lack of second- or higher order modelling, leading to an error between the analytical expression and the results of the circuit simulator. A brief chapter follows this review, looking at some previously published frequency dividers other than the architecture used for our dual-modulus dividers, arguing their merits and flaws, and showing the variety of circuits investigated by other authors for specific applications.

In chapter 4, a new phase selector circuit is developed from a published idea that has been chosen because of its suitability to the problem of reducing the number of circuit blocks operating at the highest signal frequency. A new circuit topology is presented, together with a glitch-free controller, that is crucial to the implementation of two dual-modulus divider circuits in later chapters. In order to show its merits, the chapter relies on a specific set of diagrams justifying its suitability to the phase selector scheme, as well as verifying the controller's ability to switch between phases without fearing race conditions on its output lines.

Chapter 5 is the first design chapter in which a concept is taken forward and implemented in silicon, describing much of the core contributions of this study. A stacked SOI logic scheme is developed which exceeds the power supply limit usually given by foundries. This stacked arrangement also lends itself to the idea of current re-use, where circuit blocks close to the upper power rail are able to sink their bias currents into other circuit blocks; this follows the exact theme of this research, namely power reduction in high speed cells. The new phase-selector topology of Chapter 4 is 'adapted' for use within this stacked SOI logic scheme and the chosen circuit topology is a vast improvement (in terms of transistor count) on previous designs. The design evolves further with a new dual-modulus 64/65 divider circuit using the stacked SOI logic and current re-use, exploiting a technology feature. This divider is designed to operate above the maximum power supply quoted for this technology, without having to worry about back bias effects in the MOS transistors. The control unit described in Chapter 4 is implemented, so that the phase selector suffers no phase change transient. Together with the layout, the design flow used has been experimental and developed at Southampton University including parasitic capacitance extraction of the metal interconnect lines. The concept of stacking divider stages, along with the phase select topology and glitch-free controller are proven to work with actual measurements, illustrating competitive performance. In particular, the power consumption of the stack of dividers is an impressive figure, though this figure is masked by the current consumption of the rest of the circuit block. Overall, the architecture gives the least input load and has the fewest blocks running at the highest signal frequency relative to designs based on synchronous dividers sitting at the front of the dual-modulus design [3].

Staying on the topic of high speed frequency dividers for portable transceiver applications, chapter 6 presents the challenges of a very high speed quadrature generator based on a master-slave divide-by-2 circuit. This work has been targetted for a 0.18μm bulk CMOS process, predominantly optimised for large-scale integration (LSI) within digital applications. The circuit topology is again based on source-coupled logic, with the master-slave action giving the required quadrature outputs. A preceding VCO (not designed by this author) is presented with minimum load, and the divider's transistors

have been dimensioned to enable a passive mixer to be driven directly in a 5GHz wireless LAN application. This frequency divider has been fabricated and measured to run with input frequencies up to 12GHz. The chapter also collects results from its direct inclusion in a VCO-divider combination and a low jitter PLL for a 10GHz optical application. These application designs are crucial in drawing out the divider's true performance, as power hungry buffers become redundant and the actual power consumption of the divider at 10GHz and beyond shows its competitiveness.

Chapter 7 continues with the wireless LAN target application of chapter 6, though moves back to the design of a dual-modulus divider, with division ratios 16 and 17. The same phase-select architecture of chapter 5 is chosen, but has the important distinction of being implemented in a 0.18μm bulk CMOS process. This has led to the circuit topology reverting back to classical voltage-mode cascades of the divide-by-2 stages, with high speed level shifters playing the role of interstage coupling. The phase selector of chapter 4 is adapted to bulk CMOS technology and is designed within a restricted power supply, without having to 'unstack' the 4-to-1 line selector into 3 separate 2-to-1 multiplexors. The method of modulus control is another difference compared with the SOI divider, as the embedded NAND gate was found to be the culprit for erroneous division when incorporated in a programmable architecture. The remedy was to include a low power SCL latch in the signal path to block any further edges clocking the 8-state controller. Simulations have shown the circuit to function with input frequencies up to 7GHz.

The final design chapter incorporates the divide-by-16/17 circuit in a programmable architecture with integer moduli ranging from 513 to 544 inclusive. The dual-modulus divider forms the core of the circuit and is the only section of this circuit to receive the highest frequency input of 5.5GHz. Thus, with much of the crucial design work done in the dual-modulus divider, standard library cells have been added to complete the rest of the circuit, excluding one circuit which translates the SCL signal levels to rail-rail CMOS levels. The modulus is selected using a 5-bit binary input with CMOS logic levels, making it easy for integration with a digital controller designed with standard library cells. The full programmable divider has been fabricated in a 0.18μm bulk CMOS process and characterised to work up to 6.6GHz, as well as demonstrating high input sensitivity within the required 5.15-5.35GHz wireless band.

The designs presented in this work have relied heavily on robust compact transistor models for sub-micron CMOS technologies. The SOI dual-modulus divider was designed using the STAG SOI model with 0.35μm CMOS PDSOI parameters. At the layout stage, an extraction rules file was available and adapted to our parasitic extraction tool, though no RF modelling was performed on the devices with footprints specific to RF circuit design. This is one reason for the large discrepancy between the simulated

and measured maximum frequency of the SOI divider. In the case of the bulk CMOS dividers, the MOS models used have been developed with a considerable number of man-hours and resources, both in the DC model and the RF model. Characterisation of MOS devices with a special footprint suitable for RF circuit design is evident in the results for each divider. The lack of a qualified design flow in the SOI design has shown itself to be a prerequisite in high performance analogue IC design, not just for achieving functional silicon samples, but also for designs that need accurate biasing, signal swings, and top end frequency operation.

9.1 Further Work

Although the thesis has been completed, this does not mean the work on SCL frequency dividers is over. From the beginning of the work, it was always agreed by this author that a comprehensive analysis leading to a design method on how to design frequency dividers completely was needed. Once you have parameters for the input operating frequency range, power supply, technology, and/or output swing, as well as the circuit topology, there should be a set of equations showing whether the specifications can be met and how to dimension the transistors. The first order analysis of the SCL ring-oscillator shows gaps in its understanding at large-signal level. The published literature used in the analysis is far from giving useful design equations, constraining each stage to toggle continuously and not settle at one of the logic levels. Thus, more of an analysis on the SCL divide-by-2 stage is definitely required in order to obtain better analytical models. This is fuelled by the complexity of future designs in more advanced processes, with compact models becoming more and more thorough. Once found, such an equation would give an optimum operating point and, with more parameters, possibly a trajectory or trend as to where the fastest configuration occurs.

The SOI dual-modulus divider certainly has scope for optimisation and design improvement. As discussed in chapter 5, a more confident design flow would pave the way for lowering the current consumption amongst some of the lower speed divider stages in the design. The DC performance of the STAG MOS model should aid in lowering the redundant power consumption of the bias networks. Although the stacked SOI divider concept has been proven with measurements, it is by no means the best solution to the problem of frequency division if offered the choice of PDSOI CMOS technology. The phase selector circuit could benefit from optimisation, and would certainly prove worthy of more investigation. A better architecture, other than phase selecting, is certainly possible too, but long simulation times have prevented the author from exhausting all possibilities. With different architectures comes the possibility of different functions too (for example, a quad-modulus divider). Hence, there is much scope for developing SOI stacked logic styles in general. At the circuit level, attempts had been made by

the author to modify the design to have PMOS transistors for the divider loads, where their gate terminals would be clocked too, changing its behaviour. In recent times, this method has proved popular in high frequency designs [2].

Remaining with the SOI divider, a 'power up/down' strategy needs to be defined for the stacked circuits. As there is a strong dependency on voltage biasing within the divider, it seems logical that each of the bias levels be powered up and down in sequence to prevent fatal destruction of the circuit when first powered. The design is also in need of on-chip decoupling capacitors and very high speed *electrostatic discharge device* (ESD) pads that are strong enough to handle the likelihood of discharge around a laboratory bench. The test board on which the bare dice were mounted could also be redesigned to allow a second chip to be mounted and have high quality phase shifters and summers for phase noise measurements.

On the topic of the basic division-by-2, one circuit that could be investigated in parallel with the design of a VCO, is an injection-locked divider. Such circuits have proved themselves to be low power consumption circuits [1]. However, in the case of SOI technology, the lack of a broadband inductor and varactor model based on the SOI CMOS process made this difficult, given the time constraints. It was also dismissed as a choice for the quadrature generation in bulk CMOS design technology owing to the large silicon area occupied by the pairs of inductors. Ironically, scalable broadband inductor and varactor models existed in the bulk CMOS design flow, whereas silicon area was a little more abundant in the SOI tapeout, with more allocated than needed.

In the case of the bulk divider circuits, phase noise measurements of each would supplement the results in this thesis very well. In order to do this, two similar dividers (preferably adjacent to one another on a wafer) would have to be wire-bonded on to a high frequency substrate and driven with identical signals, with the output of one being phase shifted by $\frac{\pi}{2}$ rads and summed with the output of the other identical divider, before low-pass filtering and plotting the results on a spectrum analyser. As mentioned in the penultimate chapter, a study would be ideal to see whether design time and silicon area is justified for a fractional-N synthesiser over an integer-N implementation. Had the phase noise specifications of the local oscillator been more stringent, then the frequency synthesiser design may have swung in favour of the fractional-N implemention, with some clever spur rejection such as dithering pulse code modulation or high order loop filtering.

Finally, the dual-modulus dividers in this thesis, have centred around a phase selection architecture. Naturally, the SCL divide-by-2 stage based on a master-slave flip-flop lends itself to this method of dual-modulus division. At the architectural level, the phase selec-

tion scheme is not restricted to having two asynchronously cascaded divide-by-2 stages preceding the phase selector, with permutations of one (phase selector outputs separated by half an input cycle) and upwards entirely possible. This is irrespective of its implementation in either a bulk CMOS- or PDSOI CMOS technology. With respect to the phase selector transistor implementation, the current steering topology is difficult to implement in bulk technology, whereas the PDSOI CMOS technology has been exploited to make it suitable for integration in a submicron technology, even at very high frequencies. That is not to say, that the architecture itself is against the use of bulk CMOS technology, and examples have been given where non-current steering topologies have been designed to carry out the same procedure, albeit with less elegance than that achieved with the PDSOI CMOS technology presented in this dissertation.

Bibliography

[1] H.Rategh and T.H.Lee. Superharmonic injection-locked frequency dividers. *IEEE Journal of Solid-State Circuits*, 34(6):813–821, June 1999.

[2] H.Wang. A 1.8V 3mW 16.8GHz frequency divider in 0.25μm CMOS. In *IEEE International Solid-State Circuits Conference*, page TP12.1, 2000.

[3] R.Rogenmoser, Q.Huang, and F.Piazza. 1.57GHz asynchronous and 1.4GHz dual modulus 1.2 μm CMOS prescalers. In *IEEE Proc. Custom Integrated Circuits Conference*, pages 16.3.1–16.3.4., San Diego, May 1994.

[4] U.L.Rohde. *Digital PLL frequency synthesizers - Theory and design.* Prentice Hall, 1983.

List of Own Publications Related to This Work

1. K.Mistry, W.Redman-White, J.Benson and N.D'Halleweyn, *'Low power multi-gigahertz divider architectures in SOI CMOS employing series current reuse'*, in the Proceedings of European Conference on Circuit Theory and Design, ppIII.161-III.164, August, 2001

2. R. van Beek, C.S.Vaucher, D.M.W.Leenaerts, N.Pavlovic, K.Mistry, E.A.M.Klumperink and B.Nauta, *'A 2.5-to-10GHz SONET compliant clock multiplier unit in a 0.18μm CMOS technology'*, in IEEE International Solid-State Circuits Conference Digest of Technical Papers, paper 10.3, February 2003.

3. K.Mistry, W.Redman-White, J.Benson and N.D'Halleweyn, *'A high speed dual-modulus divider in SOI CMOS with stacked current steering phase selection architecture'*, in IEEE Radio Frequency Integrated Circuits (RFIC) Symposium Digest of Papers, pp.471-474, June 2003.

4. N.Pavlovic, J.Gosselin, K.Mistry, D.M.W.Leenaerts, *'A 10GHz frequency synthesiser for 802.11a in 0.18μm CMOS'*, to be presented at the European Solid-State Circuits Conference, Leuven, Belgium, pp.367-370, September 2004.

5. K.Mistry, W.Redman-White, J.Benson and N.D'Halleweyn, *'A CMOS-SOI dual-modulus 64/65 divider employing a stacked current-steering and reuse phase selection architecture.'*, journal paper in review.

6. K.Mistry, N.Pavlovic and D.M.W.Leenaerts, *'An IEEE802.11a compliant programmable divider in 0.18m CMOS, using current steering phase selection architecture'*, (journal paper to be submitted).

Appendix A

Design of test alumina board for dual modulus 64/65 frequency divider (SOI)

The design and configuration of the SOI chip testboard has been kept out of the main section of the thesis. For completeness, it will be set out in this chapter for the interested reader.

Figure A.1 shows the schematic of the test board. Owing to the high power supply, special arrangements had to be made such as with the regulators. The 6.8V power supply is supplied by a variable power supply and, as there are pads on the chip with signals above 3.3V (power for the technology), ESD protection diodes had to be placed on chip to protect these nodes. Hence, a variety of regulators are mounted on the tile too.

The 50Ω characteristic transmission lines are sized with the help of equations and manufacturing details. *ThickFilm Circuits Ltd* fabricated the tiles on a 0.635mm thick alumina (95% alumina, dielectric constant of 9.5) substrate, with a nominal 0.014mm thick conductor. As the wirebonder in the clean room at Southampton had specific requirements, two types of conductors were needed. For general soldering, a solderable silver-palladium (20-30mΩ/□) was layered and actually covers the majority of the tile (top- and underside.) Any location where wirebonding was needed, or placement of die, and a gold (3mΩ/□) paste was printed. The gold paste was actually layered and overlapped some of the silver-palladium lines, in order to make contact before being fired in a furnace.

As shown in Fig. A.2, the actual divider circuit with pads sits in a corner of the entire

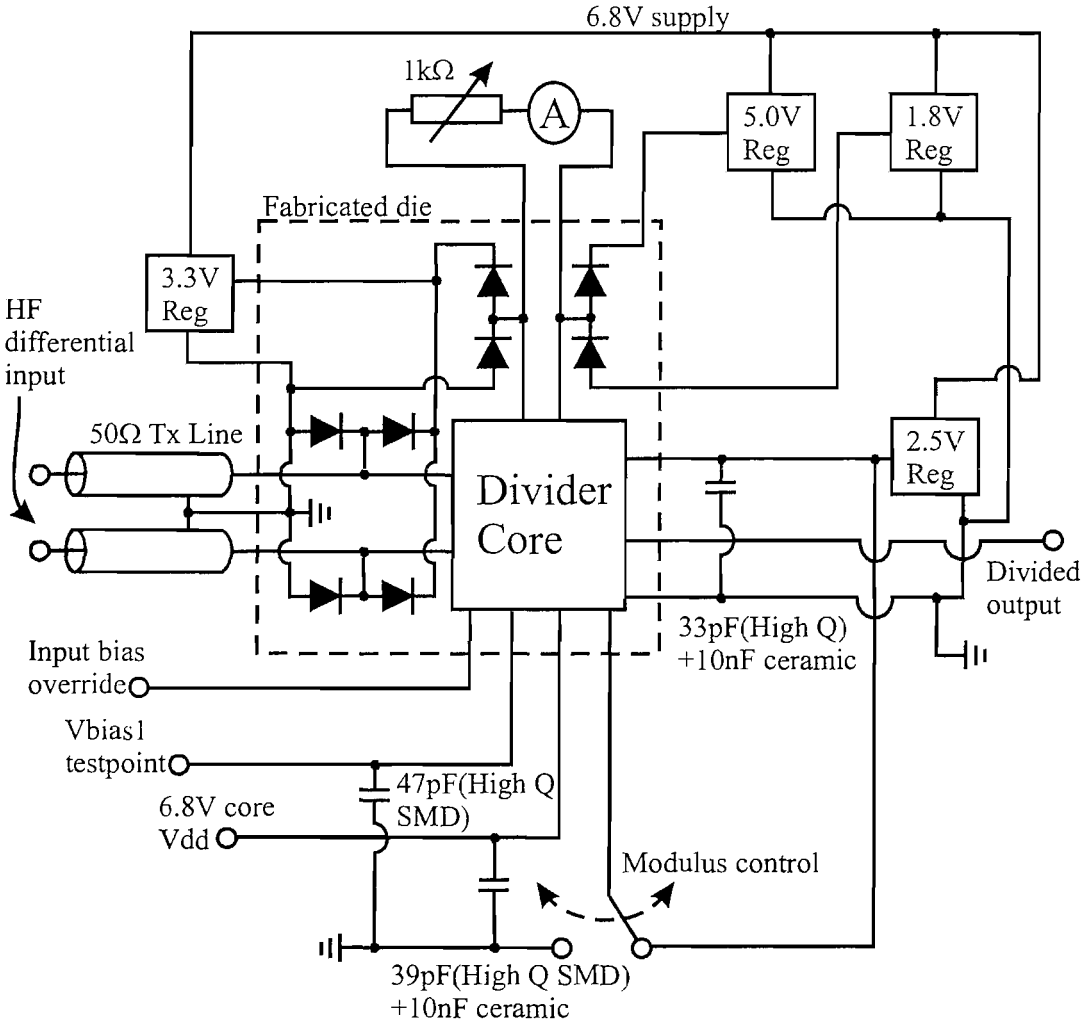

FIGURE A.1: Schematic of the testboard.

die. The die dimensions are 3.5mm by 3.2mm. A serious error in the picture and hence the fabricated design, is its position on the die. It would have been desirable to have the high frequency inputs as close to a sawn die edge as possible, thus leading to smaller inductance bond wires. The diagram in Fig. A.2 also shows the intended footprint of the gold conductor at the centre of the tile.

As the hope was to have 5GHz signals fed to the differential inputs, transmission lines were designed as a feed. The choice of coplanar lines (with a ground plane) and grounding regions along the length of the line implies that more of the EM energy can be contained within the dielectric. The dimensions of such a line are obtained with the following equations (refer to Fig. A.3):

$$Z_{ocp} = \frac{30 \cdot \pi}{\sqrt{\epsilon_{re}}} \cdot \frac{K'(k)}{K(k)} \tag{A.1}$$

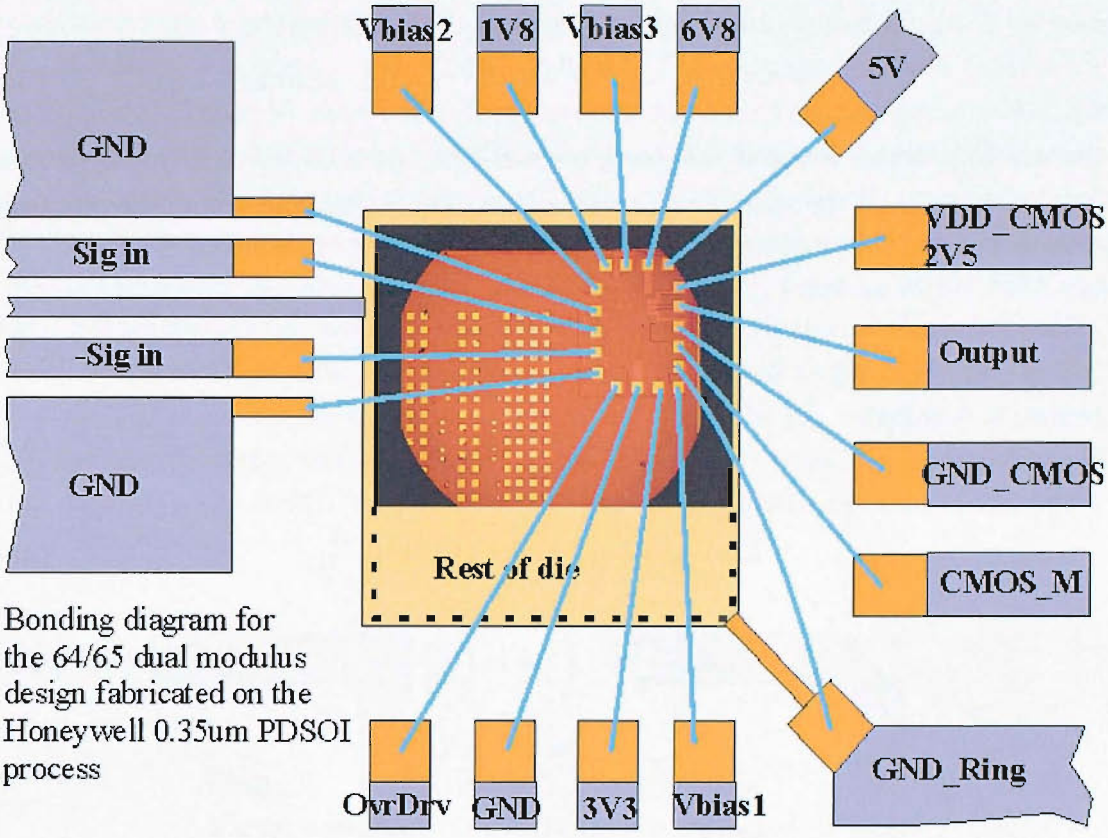

FIGURE A.2: Bond wire connection to the circuit on the fabricated die.

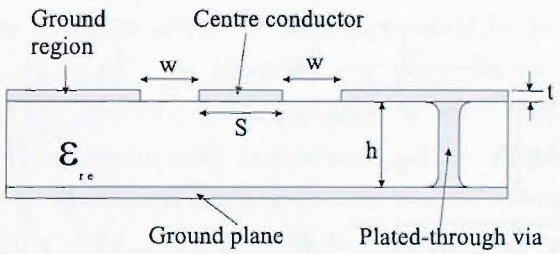

FIGURE A.3: Cross section of a coplanar transmission line, showing the various parameters used in the equations.

$$k = \frac{S}{S + 2 \cdot W} \tag{A.2}$$

$$k' = (1 - k^2)^{\frac{1}{2}} \tag{A.3}$$

$$\frac{K(k)}{K'(k)} = \frac{1}{\pi} \cdot ln\left(2 \cdot \frac{1 + \sqrt{k}}{1 - \sqrt{k}}\right) \quad for \quad 0.707 \leq k \leq 1 \tag{A.4}$$

$$\frac{K(k)}{K'(k)} = \frac{\pi}{ln\left(2 \cdot \frac{1+\sqrt{k'}}{1-\sqrt{k'}}\right)} \quad for \quad 0 \leq k \leq 0.707 \tag{A.5}$$

$$\epsilon_{re} = \frac{\epsilon_r + 1}{2} \cdot \left(tanh\left(1.785 \cdot log_{10}\left(\frac{h}{w}\right) + 1.75\right) \ldots \right.$$
$$\left. \ldots + \frac{k \cdot w}{h}\left(0.04 - 0.7 \cdot k + 0.01 \cdot (1 - 0.1 \cdot \epsilon_r)(0.25 + k)\right)\right). \quad \text{(A.6)}$$

where $\frac{K(k)}{K'(k)}$ is the '*complete integral of the first kind*,' ϵ_{re} is the effective relative permittivity, and Z_{ocp} is the characteristic impedance. These equations ignore the conductor thickness, and thus its inclusion would yield the following equations:

$$S_e = S + \Delta \quad \text{(A.7)}$$

$$W_e = W - \Delta \quad \text{(A.8)}$$

$$k_e = \frac{S_e}{S_e + 2 \cdot W_e} \approx k + (1 - k^2)\frac{\Delta}{2 \cdot W} \quad \text{(A.9)}$$

$$\Delta = \frac{1.25 \cdot t}{\pi} \cdot \left(1 + ln\left(\frac{4 \cdot \pi \cdot S}{t}\right)\right) \quad \text{(A.10)}$$

$$Z_{ocp} = \frac{30 \cdot \pi}{\sqrt{\epsilon_{re}{}^t}} \cdot \frac{K'(k_e)}{K(k_e)} \quad \text{(A.11)}$$

$$\epsilon_{re}{}^t = \epsilon_{re} - \frac{0.7 \cdot (\epsilon_{re} - 1) \cdot \frac{t}{w}}{\frac{K(k)}{K'(k)} + \frac{0.7 \cdot t}{w}} \quad \text{(A.12)}$$

A thorough treatment of these equations has been omitted, as their use was purely engineering and no modification was intended [2]. To check their accuracy, a field solver or advanced tool such as Agilent's ADS [1] could be used. Unfortunately, a lack of familiarity coupled with a shortage of time prevented this and, instead, a freeware application, 'AppCAD', was used to home in iteratively on the correct dimensions. By having S=0.4mm, W=0.3mm, h=0.635mm and ϵ_r=9.5, AppCAD yields Z_o=52.3Ω and ϵ_{eff}=5.31. Using the equations without strip thickness (equations (A.1)-(A.6)) gives k=0.4, k'=0.917, $\frac{K(k)}{K'(k)}$=0.694, Z_{ocp}=61.34Ω, ϵ_{re}=4.9. By including the conductor thickness (t=0.014mm), equations (A.7)-(A.12) give Δ=0.0383mm, S_e=0.4383mm, W_e=0.2617mm, k_e=0.4558, k'$_e$=0.89, $\frac{K(k_e)}{K'(k_e)}$=0.7428, $Z_{ocp}{}^t$=58.4Ω and $\epsilon_{ocp}{}^t$=4.72. It must be stated that the values described are not 'lucky' values and have actually been iterated to 'home' in on the planar dimensions of the transmission lines. With the board manufacturers stating the minimum track width and separation each to be 200μm, the values generated do not conflict with their process.

When laying down the lines, certain design rules have to be obeyed. Two factors governing this are printing resolution and the connectors acting as an interface between the microstrip transmission line and the coaxial transmission line connecting the source. With the aim to keep the input section symmetrical, a 'centre' line between the two input

pads on the die would coincidentally represent the line marking half the pitch between the two central conductors of the SMA conductor. To reduce the bondwire inductance, the differential pair of transmission lines carrying the high frequency input power has to be kept to very close together. On the other hand, the flange of the SMA conductors extends either side of the centre point of the connector. This forces the transmission line to change direction twice to bring the lines together. Two right-angled turns have been inserted into each line and there is a slight, but noticeable, tapering at the SMA-end of the transmission line to make soldering the central conductor to the board easier. The obvious choice of turns would be right-angled turns such as the one shown in Fig. A.4(a). Unfortunately, as shown in its accompanying model [3], inductance is present representing the disturbance in current. By chamfering the corner like the one shown in Fig. A.4(b) (where B≈0.57w), the inductance can be eliminated from the model [3].

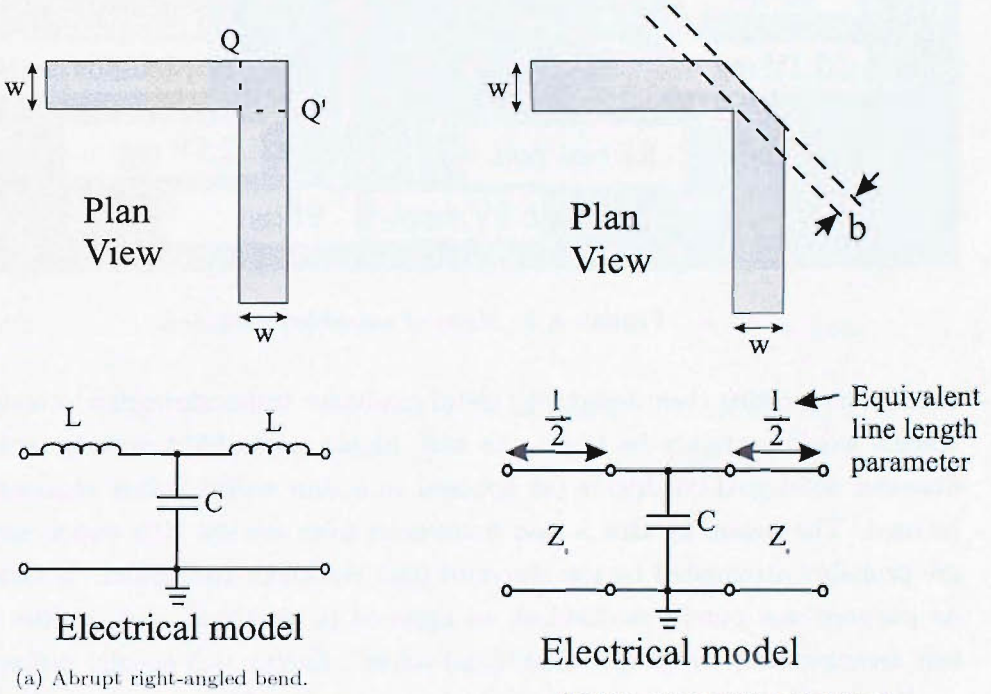

(a) Abrupt right-angled bend.

(b) Chamfer in right-angled bend.

FIGURE A.4: Ways to implement a bend in a microstrip.

The finished board with components has been captured and is shown in Fig. A.5. Surface mount device footprints have been used where necessary for certain passives and the regulators. With the large silver-palladium islands on the component side all connected to a common ground (as well as shorting to the back ground plane with the use of plated through vias), ceramic disc capacitors can be added to the tile as decoupling capacitors, though they would be more effective had they been integrated on chip.

One further thought on these boards is that if sufficient funds and time were available, then mounting the alumina tile on to a metal block with the high frequency SMA

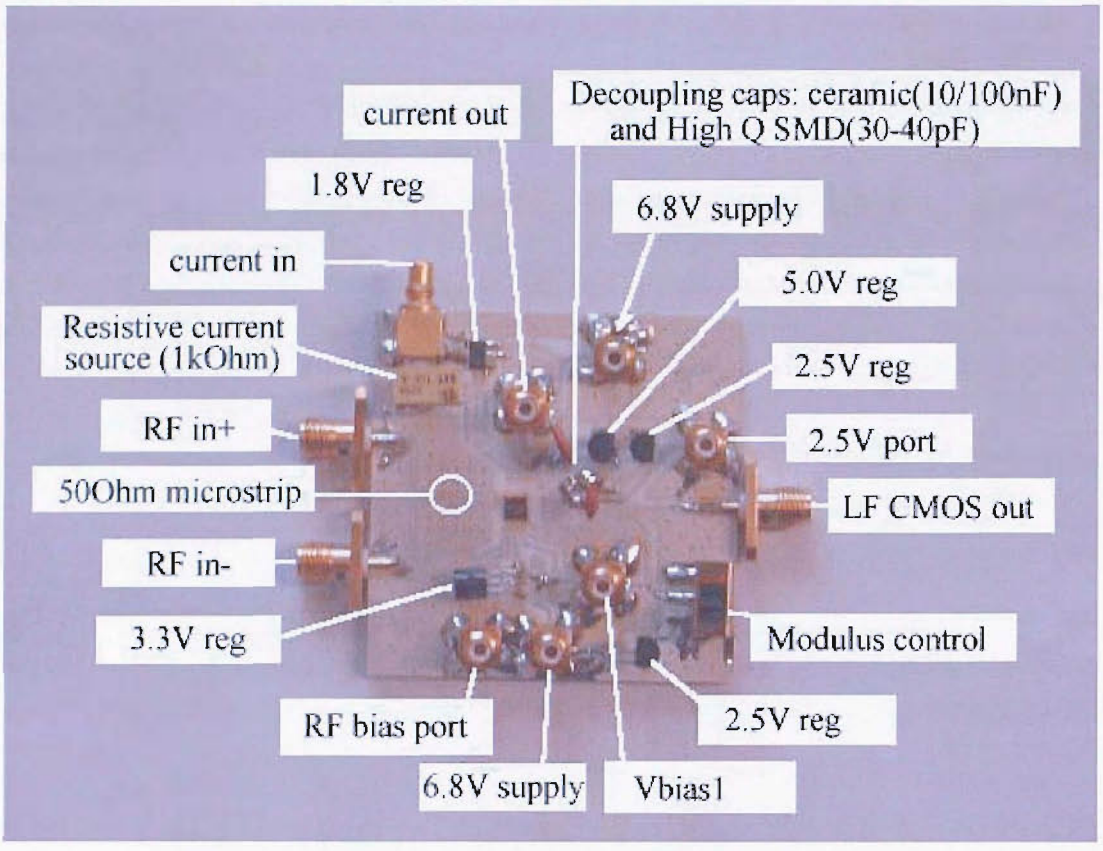

FIGURE A.5: Photo of assembled testboard.

conductors meeting their designated metal conductor tracks through a 'pressure contact' scheme would certainly be tried. As well, higher grade SMA sockets with a narrow diameter solid gold conductor (as opposed to a thin walled hollow conductor), would be used. The reason for this is that frequencies from the low GHz region and upwards, are probably attenuated by the standard lead/tin solder compound. It turns out that its purpose was purely mechanical, as opposed to electrical, and justifies why many test arrangements forego any traditional solder. Energy will usually radiate from the needle on this conductor and couple (electromagnetically) with the end point of the transmission line. Naturally, losses will occur as EM radiation passes through the air as opposed to some other low loss medium.

Bibliography

[1] Agilent. AppCAD software. Http://www.agilent.com/.

[2] K.C.Gupta, R.Garg, and I.J.Bahl. *Microstrip lines and slotlines*, chapter 7. Artech, 1979.

[3] T.C.Edward. *Foundations for microstrip circuit design*, pages 110–111. John Wiley and Sons, 1981.

Appendix B

SOI Dual Modulus Divider Schematics

This section collects the schematics for the design described in Chapter 5. The next few pages at times show a hierarchical view of the divider described in the body of this report. The last diagram shown is reserved for the netlists extracted from the final chip layout. Electrostatic protection diodes were manually placed because the extraction file was not configured to identify such devices. Also, one cannot view the insides of this block diagramatically, because the connectivity is in fact expressed in text form. A perfect balun is added onto the clock inputs, with low losses, but in fairness, should be replaced with an accurate model of the real balun which would include parasitics and loss parameters.

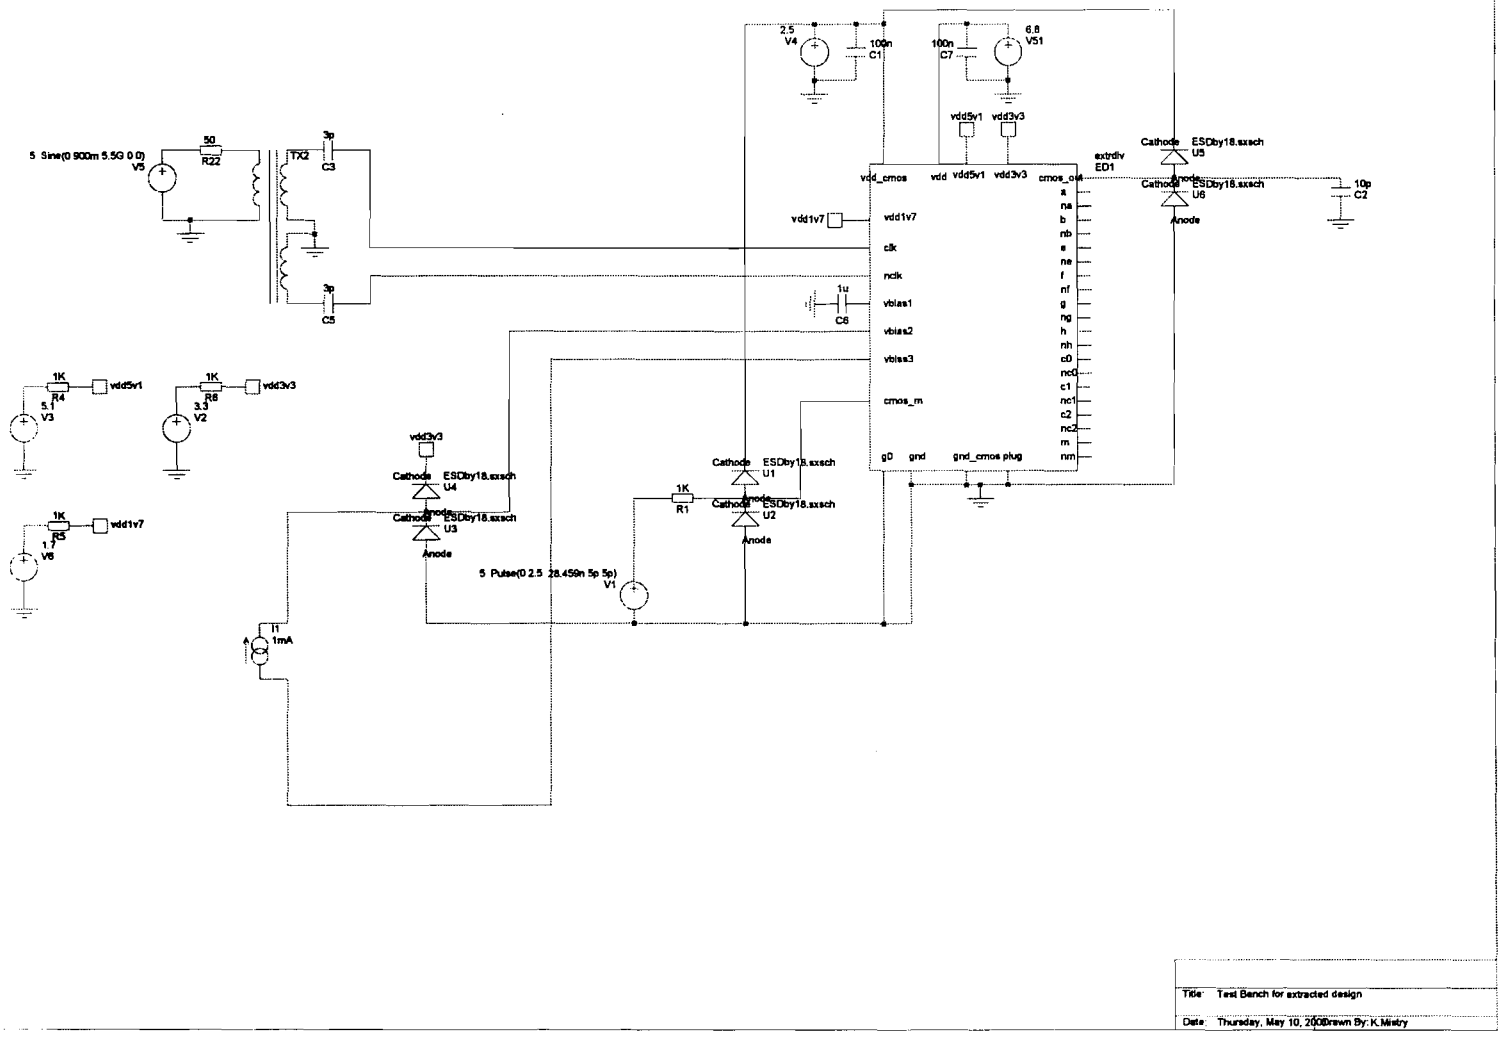
2.5
V4
100n
C1
100n
C7
6.8
V51
vdd5v1
vdd3v3
5 Sine(0 900m 5.5G 0 0)
V5
50
R22
TX2
3p
C3
3p
C5
vdd_cmos
vdd
vdd5v1
vdd3v3
cmos_out
extrdiv
ED1
Cathode ESDby18.sxsch
U5
Anode
Cathode ESDby18.sxsch
U6
Anode
10p
C2
vdd1v7
clk
nclk
vbias1
vbias2
vbias3
cmos_m
gD
gnd
gnd_cmos plug
a
na
b
nb
e
ne
f
nf
g
ng
h
nh
c0
nc0
c1
nc1
c2
nc2
m
nm
1u
C6
1K
R4
5.1
V3
vdd5v1
1K
R6
3.3
V2
vdd3v3
1K
R5
1.7
V6
vdd1v7
vdd3v3
Cathode ESDby18.sxsch
U4
Anode
Cathode ESDby18.sxsch
U3
Anode
1K
R1
Cathode ESDby18.sxsch
U1
Cathode ESDby18.sxsch
U2
Anode
5 Pulse(0 2.5 28.456n 5p 5p)
V1
I1
1mA
Title: Test Bench for extracted design
Date: Thursday, May 10, 2005 Drawn By: K.Mistry

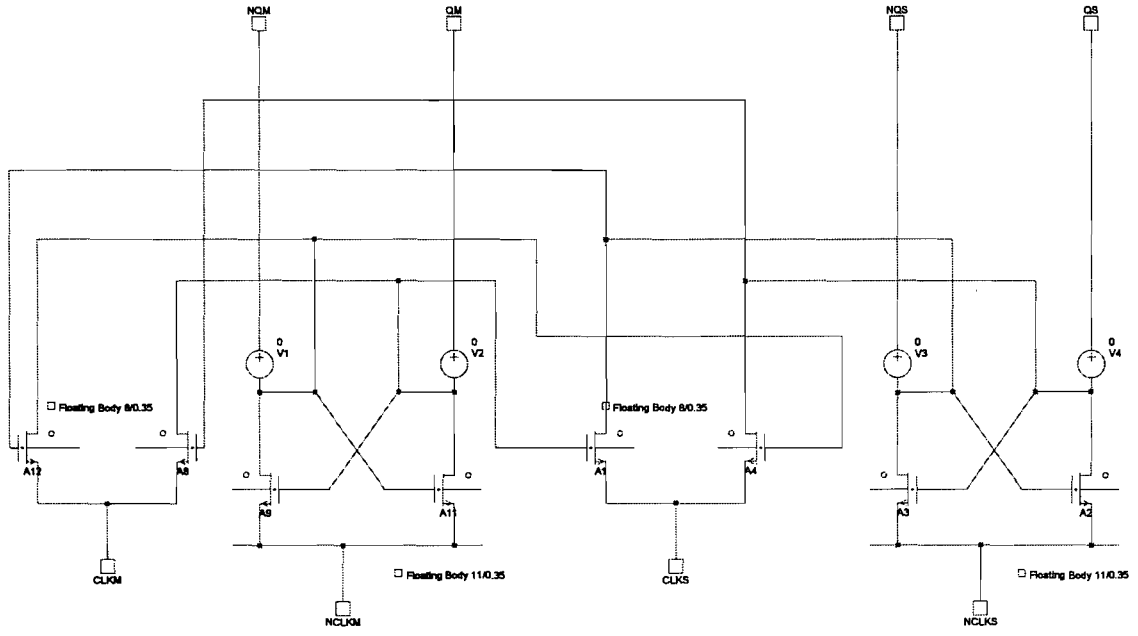

Title: First Divide by 2

Date: Thursday, May 10, 2005 Drawn By: K.Mistry

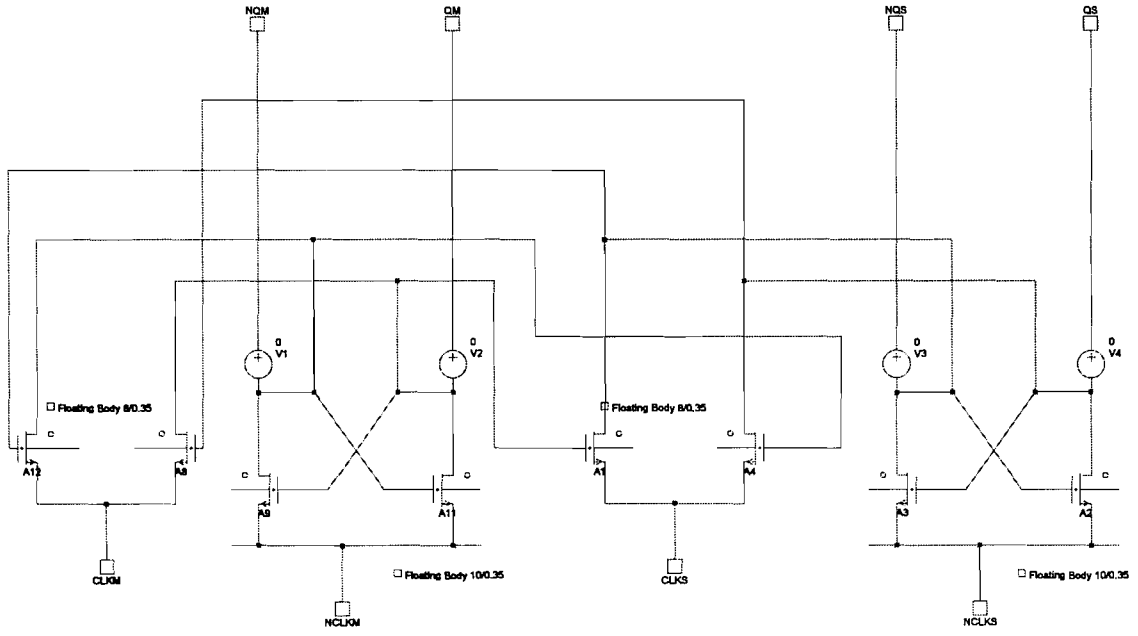

Title: Second Divide by 2

Date: Thursday, May 10, 2005 Drawn By: K.Mistry

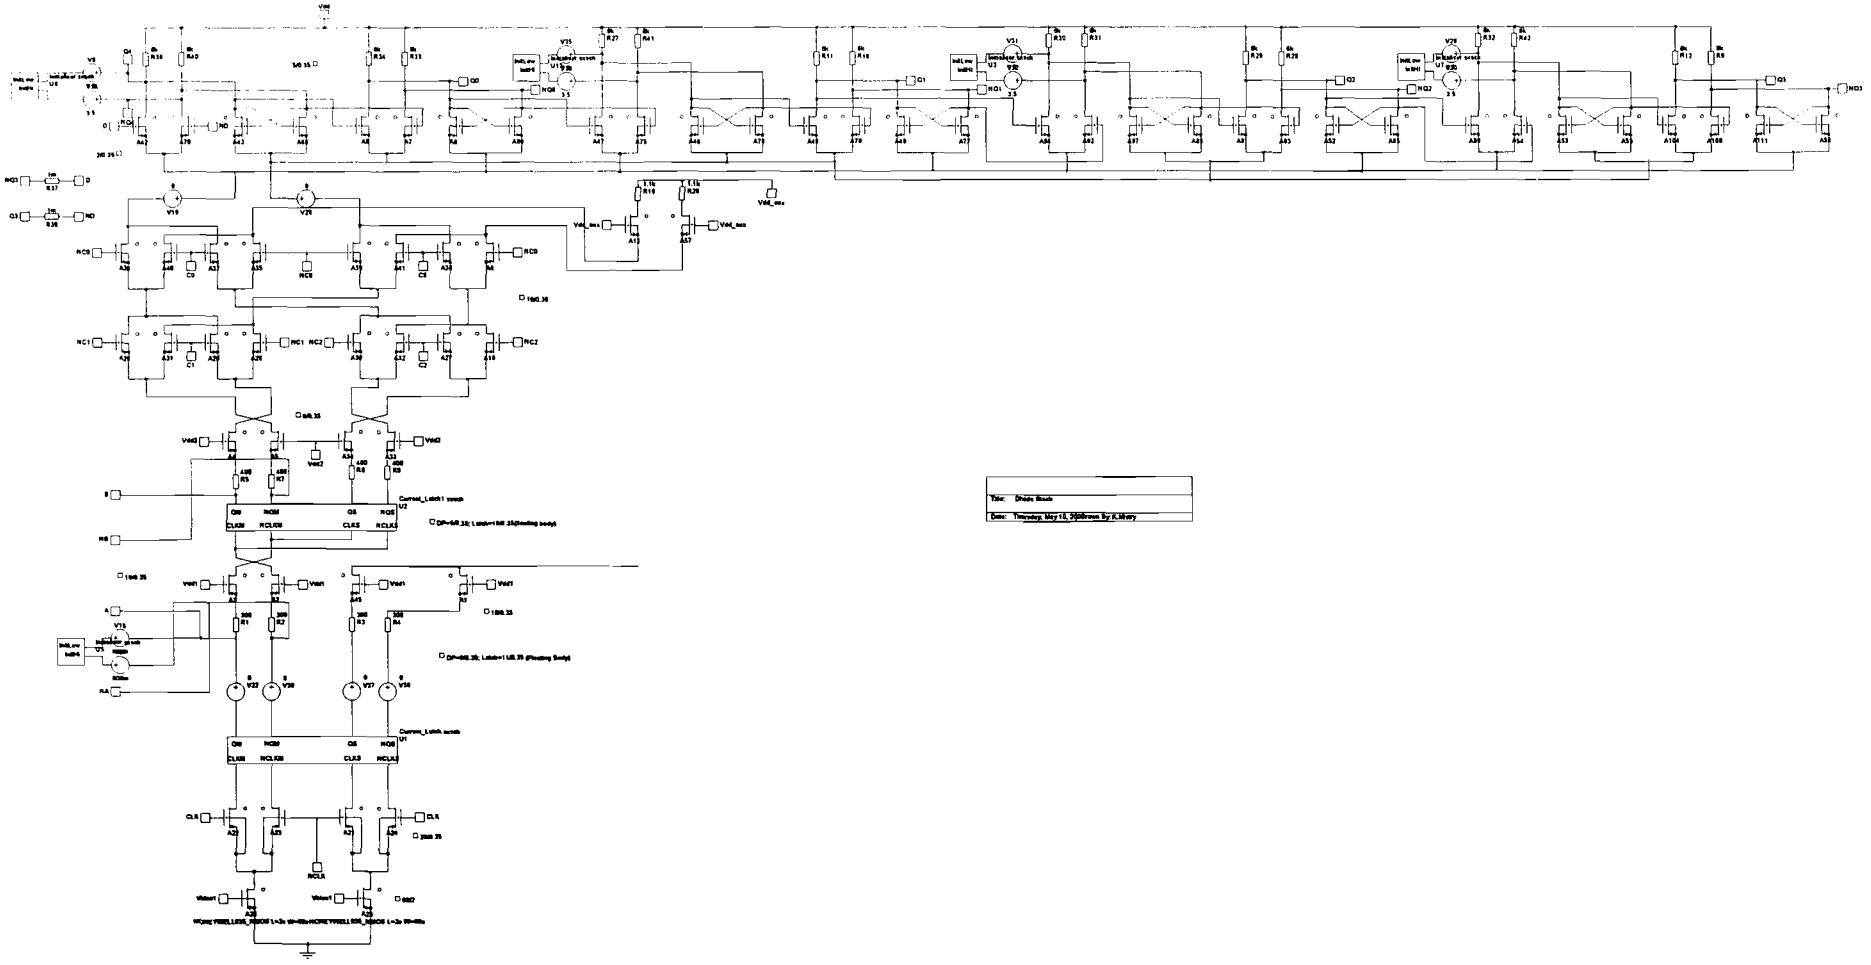

C:\texmf\Work\Full_Thesis\AppendixB\Schem4.sxsch - 27 May 2005

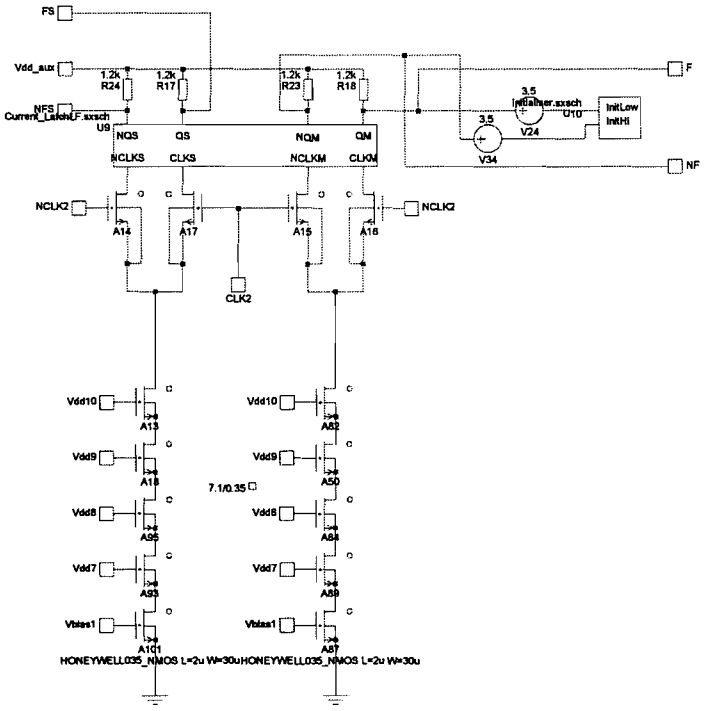

Title: First Off-stack Divide by 2

Date: Thursday, May 10, 2001 Drawn By: K.Mistry

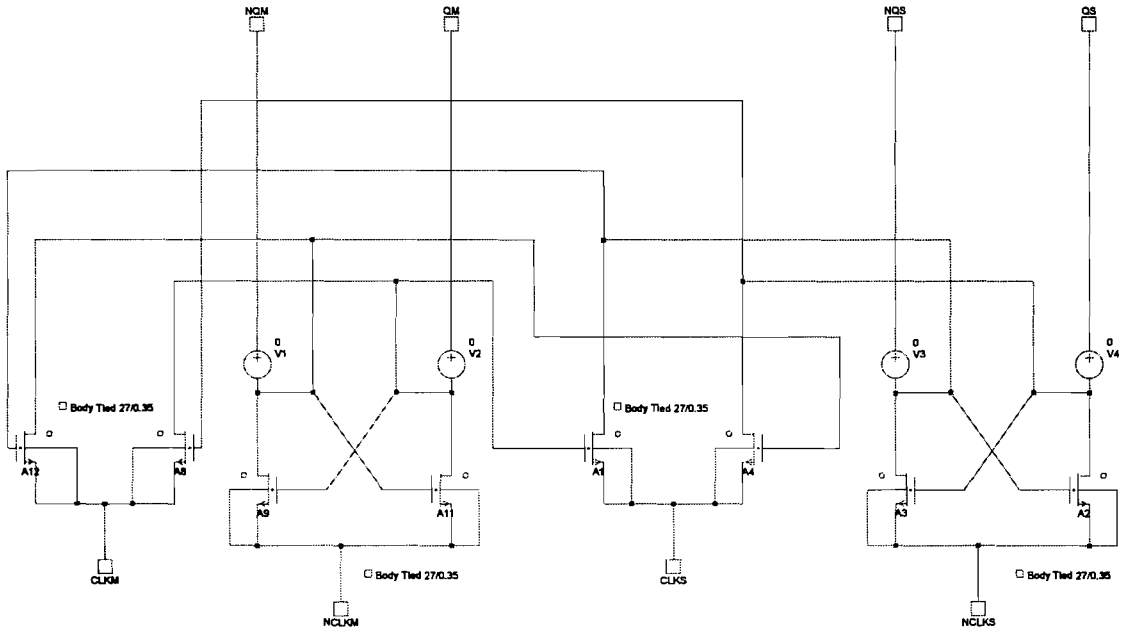

Title: Divide by 2 cells for the off-stack dividers

Date: Thursday, May 10, 2008 Drawn By: K.Mistry

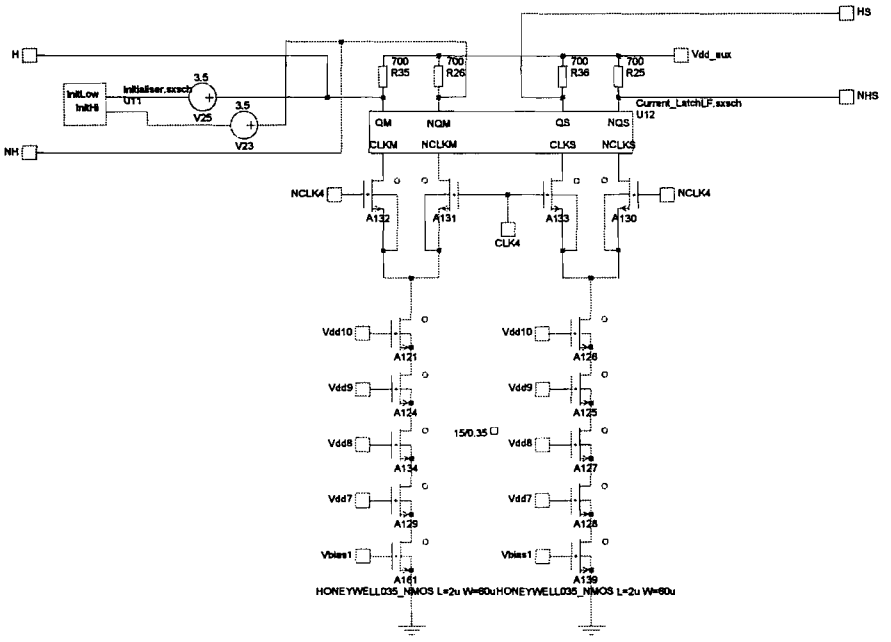

Title: Third Offstack Divide by 2

Date: Thursday, May 10, 2001 Drawn By: K.Mistry

Note: All Vbias transistors have been set to 50/2.

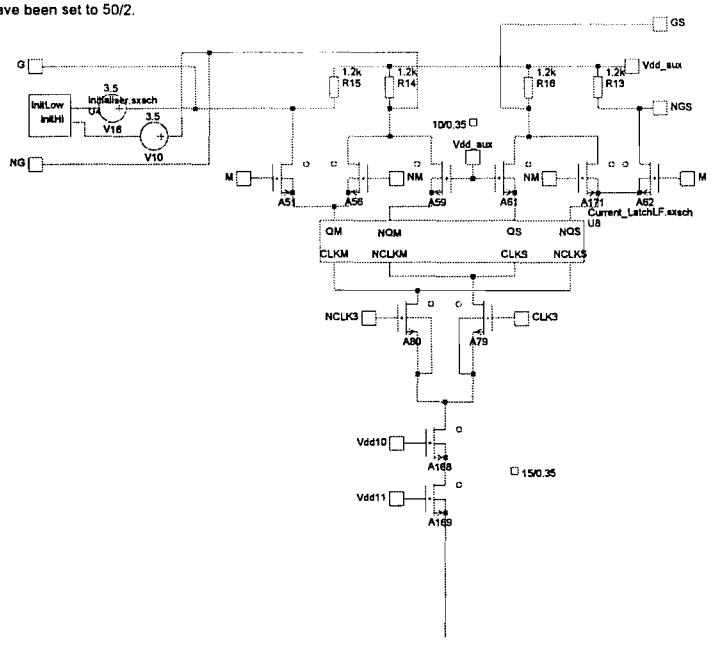

Although this unit sits on the Divide Stack, it does not receive a direct clock signal.

Title:	Second Off Stack Divide by 2
Date:	Thursday, May 10, 2001 Drawn By: K.Mistry

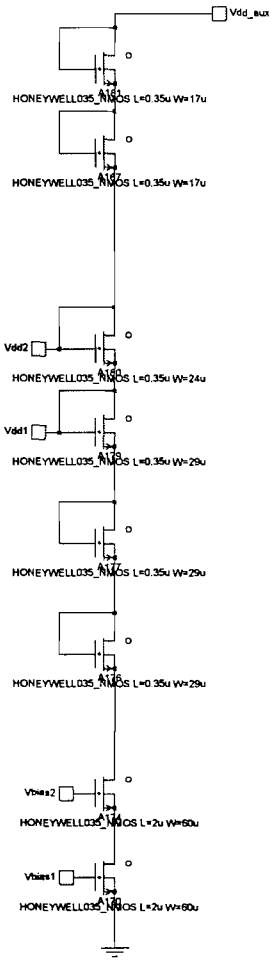

Bias Block 1, serving the divide stack

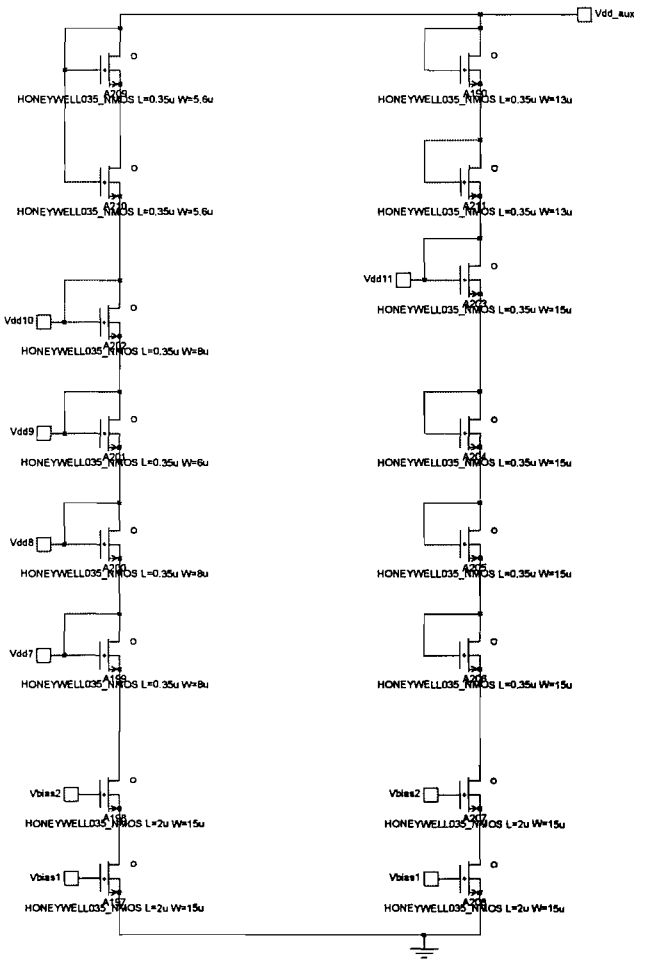

Bias Block 2 for the cascodes off-stack.

Title: Bias Blocks 1 & 2

Date: Thursday, May 10, 2005 Drawn By: K.Mistry

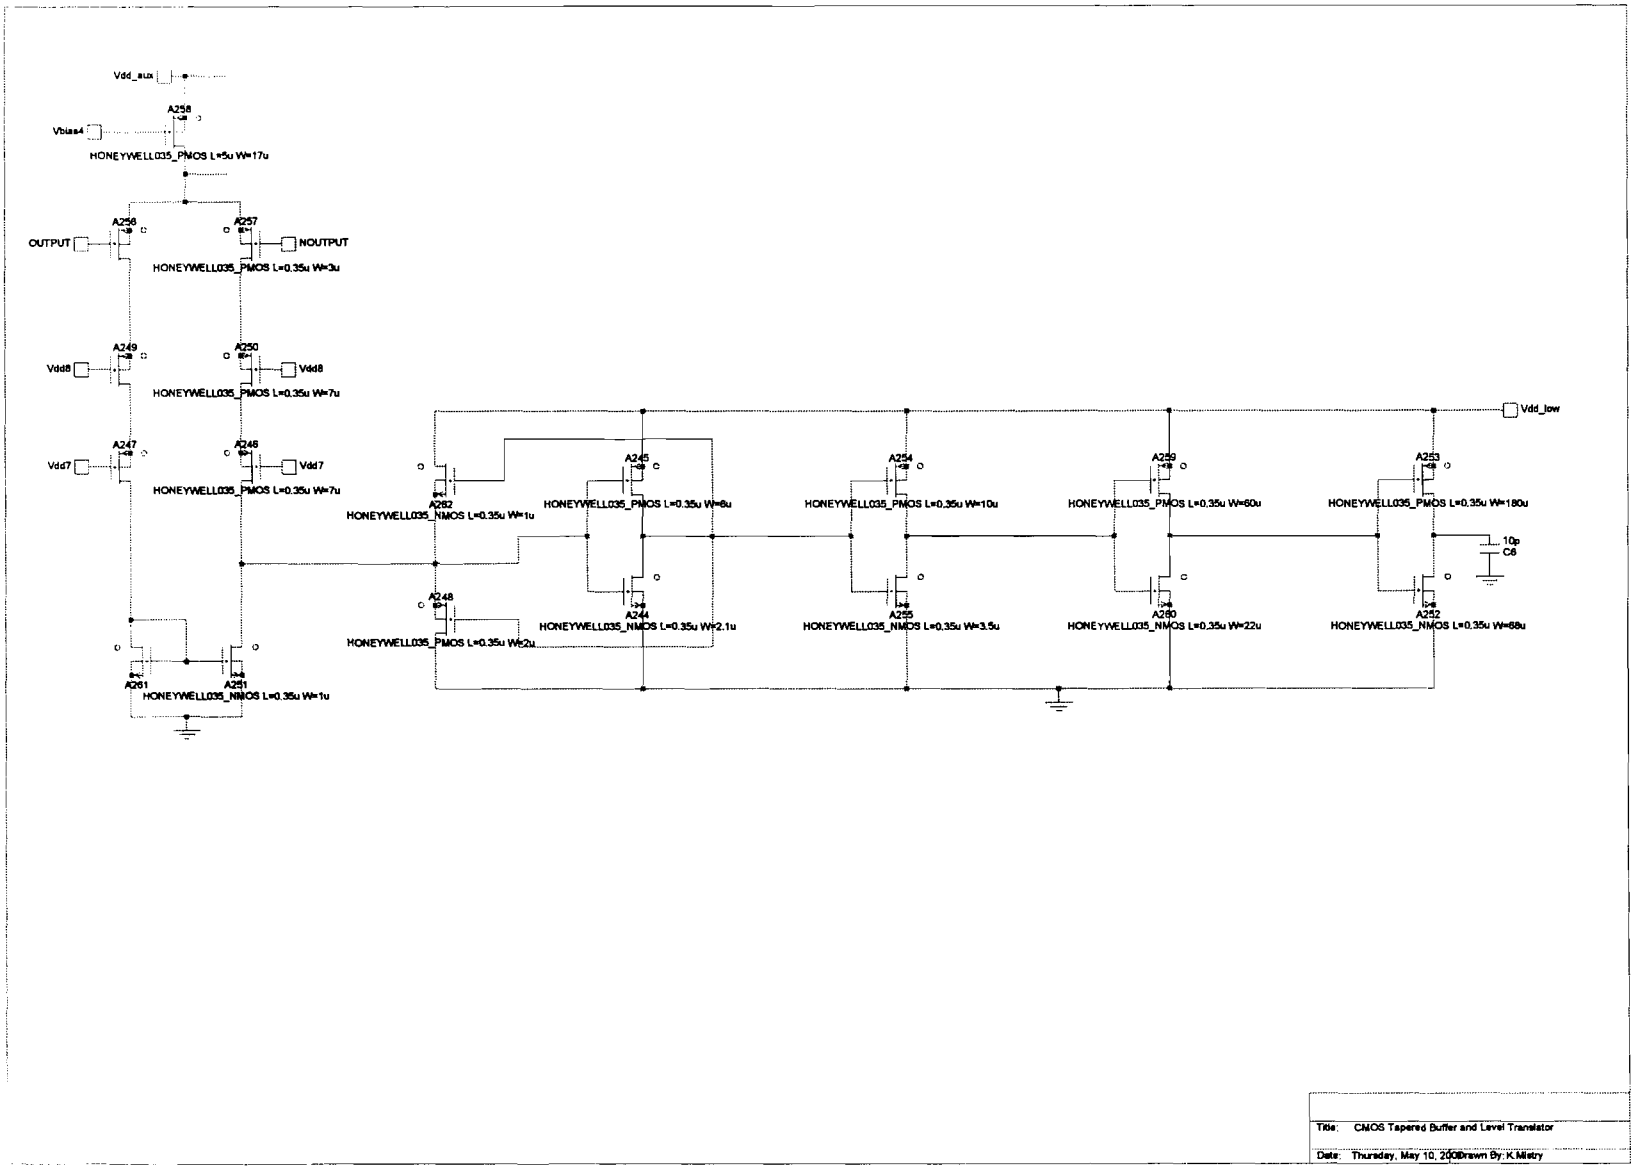
Vdd_aux
A258
Vbias4
HONEYWELL035_PMOS L=5u W=17u
A256
A257
OUTPUT
NOUTPUT
HONEYWELL035_PMOS L=0.35u W=3u
A249
A250
Vdd8
Vdd8
HONEYWELL035_PMOS L=0.35u W=7u
A247
A246
Vdd7
Vdd7
HONEYWELL035_PMOS L=0.35u W=7u
A261
A251
HONEYWELL035_NMOS L=0.35u W=1u
A262
HONEYWELL035_NMOS L=0.35u W=1u
A245
HONEYWELL035_PMOS L=0.35u W=6u
A254
HONEYWELL035_PMOS L=0.35u W=10u
A259
HONEYWELL035_PMOS L=0.35u W=60u
A253
HONEYWELL035_PMOS L=0.35u W=180u
Vdd_low
10p
C6
A248
HONEYWELL035_PMOS L=0.35u W=2u
A244
HONEYWELL035_NMOS L=0.35u W=2.1u
A255
HONEYWELL035_NMOS L=0.35u W=3.5u
A260
HONEYWELL035_NMOS L=0.35u W=22u
A252
HONEYWELL035_NMOS L=0.35u W=68u
Title: CMOS Tapered Buffer and Level Translator
Date: Thursday, May 10, 2005 Drawn By: K.Mistry

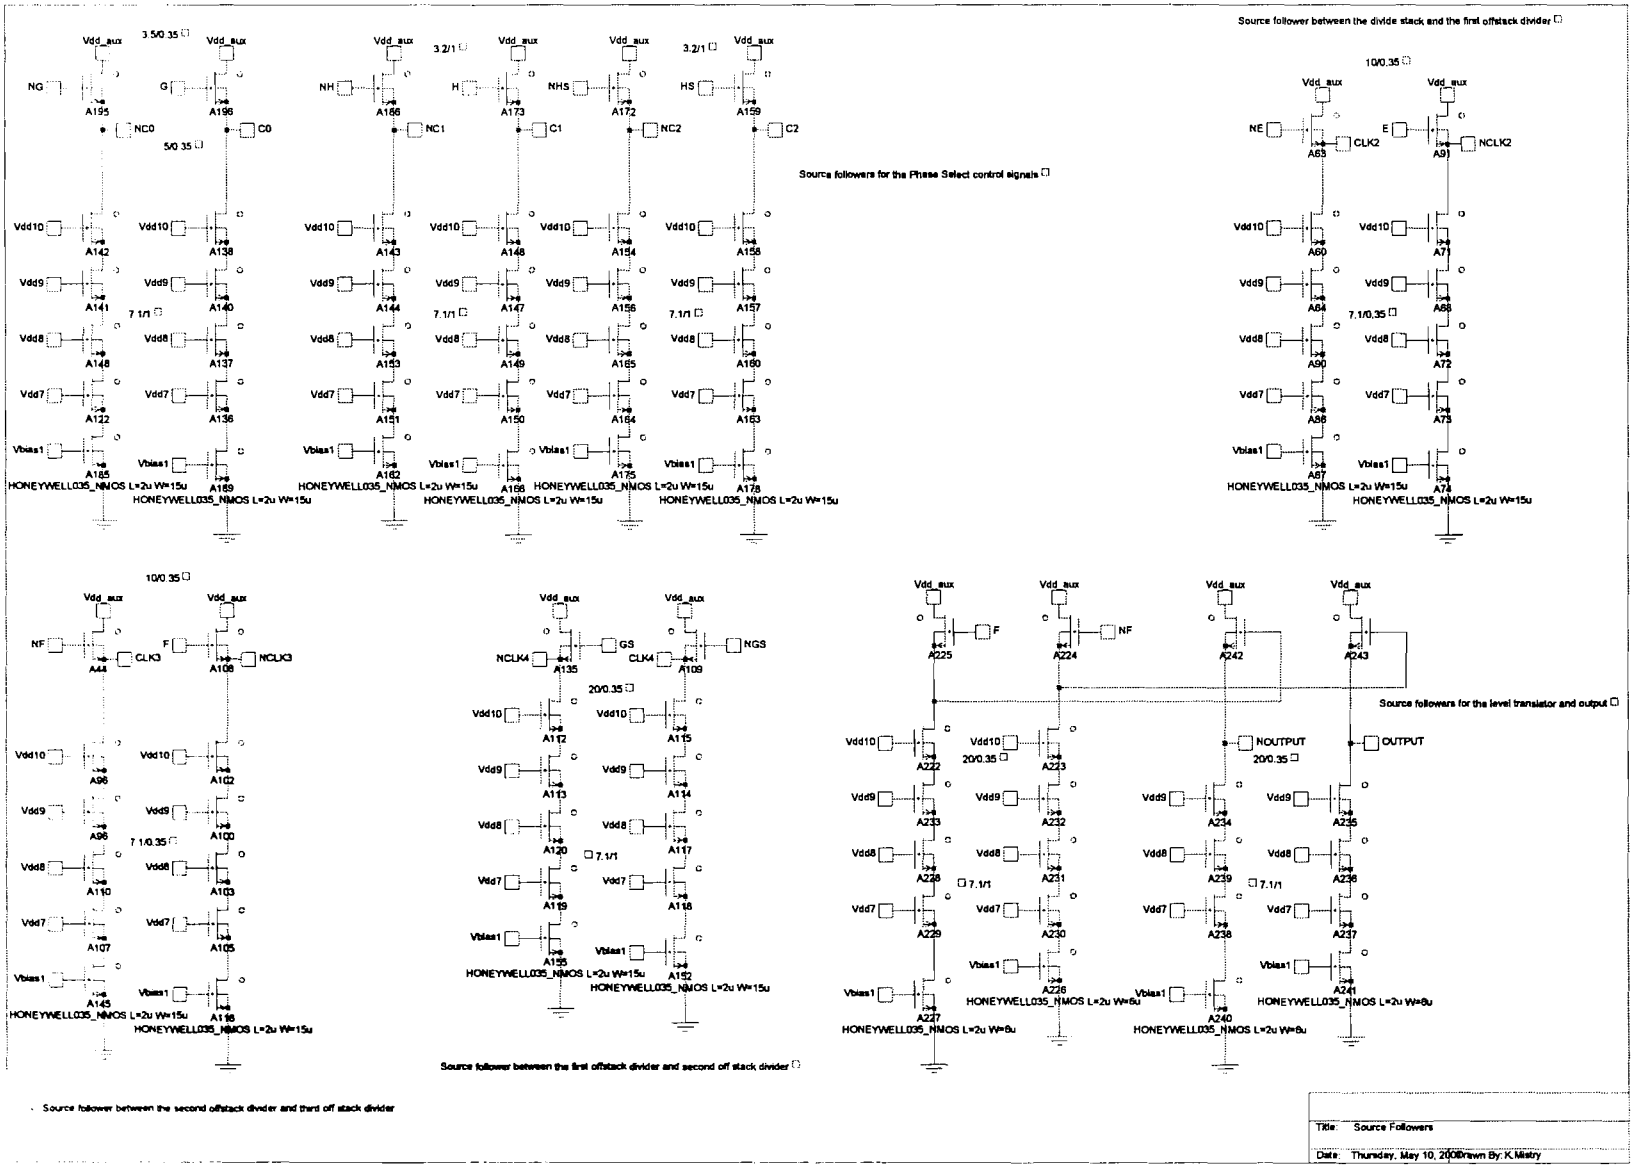

C:\texmf\Work\Full_Thesis\AppendixB\Schem11.sxsch - 27 May 2005

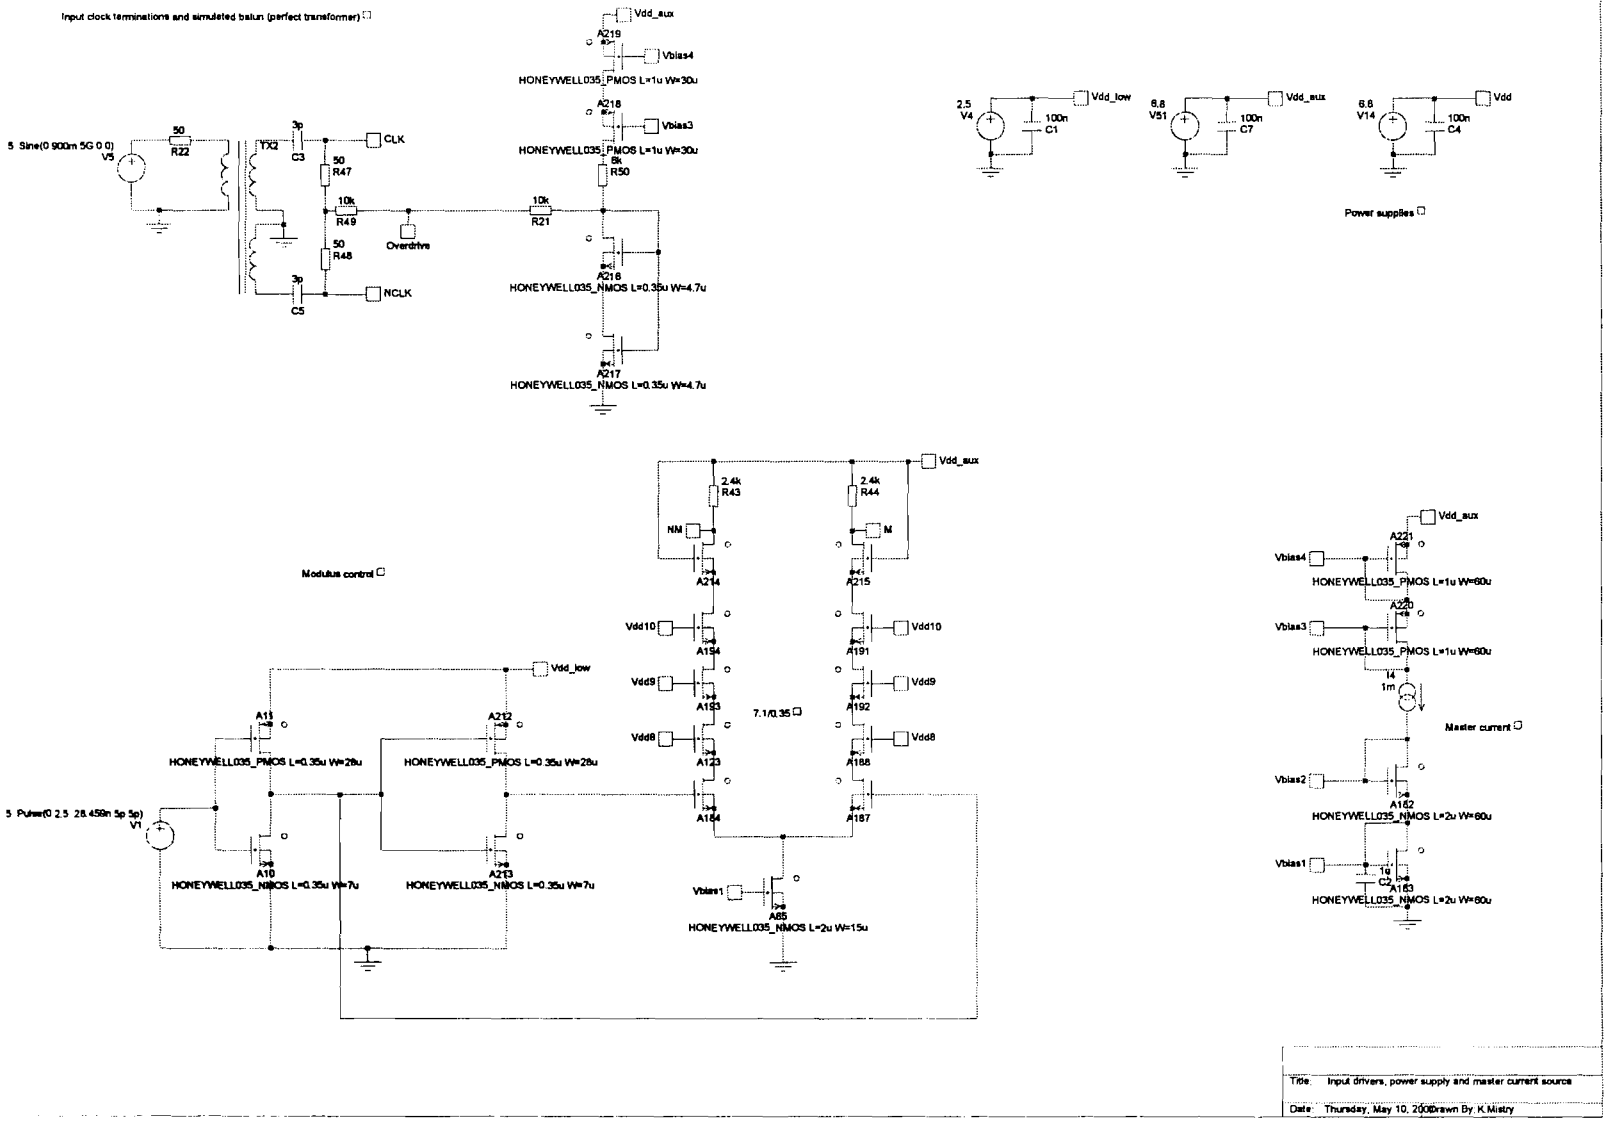
Input clock terminations and simulated balun (perfect transformer)
5 Sine(0 900m 5G 0 0)
V5
50
R22
TX2
3p
C3
50
R47
10k
R49
50
R48
3p
C5
CLK
NCLK
Overdrive
10k
R21
Vdd_aux
A219
Vbias4
HONEYWELL035_PMOS L=1u W=30u
A218
Vbias3
HONEYWELL035_PMOS L=1u W=30u
8k
R50
A216
HONEYWELL035_NMOS L=0.35u W=4.7u
A217
HONEYWELL035_NMOS L=0.35u W=4.7u
2.5
V4
Vdd_low
100n
C1
6.6
V51
Vdd_aux
100n
C7
6.6
V14
Vdd
100n
C4
Power supplies
Modulus control
2.4k
R43
2.4k
R44
NM
M
A214
A215
Vdd10
A154
A191
Vdd9
A193
A192
7.1/0.35
Vdd8
A123
A188
A184
A187
Vbias1
A65
HONEYWELL035_NMOS L=2u W=15u
Vdd_low
A11
A212
HONEYWELL035_PMOS L=0.35u W=28u
HONEYWELL035_PMOS L=0.35u W=28u
5 Pulse(0 2.5 28.458n 5p 5p)
V1
A10
A213
HONEYWELL035_NMOS L=0.35u W=7u
HONEYWELL035_NMOS L=0.35u W=7u
Vdd_aux
A221
Vbias4
HONEYWELL035_PMOS L=1u W=60u
A220
Vbias3
HONEYWELL035_PMOS L=1u W=60u
I4
1m
Master current
Vbias2
A182
HONEYWELL035_NMOS L=2u W=60u
Vbias1
1p
C2
A183
HONEYWELL035_NMOS L=2u W=60u
Title: Input drivers, power supply and master current source
Date: Thursday, May 10, 2005
Drawn By: K.Mistry